On p-adic differential equations with separation of variables

Pierre Lairez, Tristan Vaccon

On *p*-adic Differential Equations with Separation of Variables

Pierre Lairez*
Technische Universität
Berlin, Germany
pierre@lairez.fr

Tristan Vaccon
JSPS-Rikkyo University
Tokyo, Japan
vaccon@rikkyo.ac.jp

ABSTRACT

Several algorithms in computer algebra involve the computation of a power series solution of a given ordinary differential equation. Over finite fields, the problem is often lifted in an approximate p-adic setting to be well-posed. This raises precision concerns: how much precision do we need on the input to compute the output accurately? In the case of ordinary differential equations with separation of variables, we make use of the recent technique of differential precision to obtain optimal bounds on the stability of the Newton iteration. The results apply, for example, to algorithms for manipulating algebraic numbers over finite fields, for computing isogenies between elliptic curves or for deterministically finding roots of polynomials in finite fields. The new bounds lead to significant speedups in practice.

CCS Concepts

•**Computing methodologies** → **Algebraic algorithms;** *Number theory algorithms;* •**Mathematics of computing** → **Ordinary differential equations;**

Keywords

Ordinary differential equation; p-adic numbers; Newton iteration; numerical stability; differential precision

1. INTRODUCTION

We study, in a p-adic context, the loss of precision occuring during the computation of a power series solution of a certain class of differential equations. We use the method of differential precision that relies on a first-order analysis.

1.1 The *p*-adic context

Let $\mathbb{Z}_p$ be the ring of p-adic integer, for a given prime p, and $\mathbb{Q}_p$ its field of fractions. The p-adic valuation on $\mathbb{Q}_p$ is denoted v_p, and the p-adic norm is defined by $|a| = p^{-v_p(a)}$, for $a \in \mathbb{Q}_p$. For $a \in \mathbb{Q}_p$ and ε a positive real number, let $a + O(\varepsilon)$ denote the set of all $b \in \mathbb{Q}_p$ such that $|a - b| \leqslant \varepsilon$. For $a, b \in \mathbb{Z}_p$, we have $a = b + O(\varepsilon)$ if and only if $a \equiv b \pmod{p^{-\lfloor \log_p \varepsilon \rfloor}}$.

A computer can handle p-adic numbers given with bounded precision: a p-adic number a is approximately represented by a rational number a' and a radius ε such that $a \in a' + O(\varepsilon)$. This leads to a *ball arithmetic* over $\mathbb{Q}_p$. The ultrametric nature of $\mathbb{Q}_p$ makes ball arithmetic particularly convenient since errors do not propagate when adding two numbers:

$$(a + O(\varepsilon)) + (b + O(\eta)) = a + b + O(\max(\varepsilon, \eta)),$$

or multiplying them:

$$(a + O(\varepsilon))\,(b + O(\eta)) = ab + O(\max(\eta\,|a|\,, \varepsilon\,|b|)),$$

or dividing them:

$$\frac{a + O(\varepsilon)}{b + O(\eta)} = \frac{a}{b} + O(\max(\eta\,|a|\,|b|^{-2}\,, \varepsilon\,|b|^{-1})).$$

These formulae are optimal: the equalities are set equalities and not only left-to-right inclusions. When considering an algorithm performing additions, multiplications and divisions over p-adic numbers, it is possible to track the precision during all the intermediate steps. Thus, it is possible to run the algorithm on inputs given approximately as balls, and to return the result as a ball with the guarantee that whatever the *exact* values of the input are, the *exact* result lies in the ball returned. However, even if for every single operation the formulae above give the optimal precision of the result, the optimality does not compose. This is the well-known *dependency problem*. It is a major obstacle to the application of ball arithmetic over $\mathbb{Q}_p$, in the same way as it constricts interval arithmetic over $\mathbb{R}$.

For example, let us consider the computation of the determinant of a matrix with p-adic integer coefficients, given at precision ε. Since the determinant is an integral polynomial function of the coefficients, the determinant is also known at precision at least ε. However, if it is computed through a Gaussian elimination, and if at some point of the computation, one of the pivots has a positive valuation, then basic precision tracking will indicate that the result is only correct at precision less than ε. The *intrinsic* loss of precision is null, or even negative (Caruso, Roe, and Vaccon 2015), while the *algorithmic* loss, that depend on the algorithm, may be positive.

Caruso, Roe, and Vaccon (2014) have shown, in a p-adic setting, how to use first order analysis to obtain rigorous and optimal precision bounds on both the intrinsic and the

*Partially funded by the DFG grant BU 1371/2-2.

ISSAC '16, July 19–22, 2016, Waterloo, ON, Canada

DOI: http://dx.doi.org/10.1145/2930889.2930912

algorithmic loss of precision. The method relies on the following fact:

LEMMA 1 (Caruso, Roe, and Vaccon 2014). *Let $\varphi : \mathbb{Q}_p^n \to \mathbb{Q}_p^m$ be a differentiable map and let $x \in \mathbb{Q}_p^n$. If the differential $\mathrm{d}_x\varphi$ at x is surjective, then $\varphi(x+B) = \varphi(x) + \mathrm{d}_x\varphi(B)$ for any zero-centered small enough ball B.*

In other words, if the precision on the input x is B, then the precision on the output $\varphi(x)$ is $\mathrm{d}_x\varphi(B)$; this is the inclusion of $\varphi(x+B)$ in $\varphi(x)+\mathrm{d}_x\varphi(B)$. And conversely, every pertubation of $\varphi(x)$ up to $\mathrm{d}_x\varphi(B)$ comes from a pertubation from the input x up to B; this is the converse inclusion of $\varphi(x)+\mathrm{d}_x\varphi(B)$ in $\varphi(x+B)$. What *small enough* is can be expressed in terms of the norms of the higher differentials of φ at x, in the case that φ is analytic at x.

1.2 Main result

We study here the computation of a power series with p-adic coefficients solution of a given first-order ordinary differential equation with separation of variables. Let g and h be power series in $\mathbb{Z}_p[\![t]\!]$ such that $h(0) = 1$ and $g(0) \neq 0$. We consider the following differential equation:

$$y' = g \cdot h(y), \quad y(0) = 0. \tag{E}$$

It has a unique solution $y \in \mathbb{Q}_p[\![t]\!]$. We make the assumption that this solution has integer coefficients, that is $y \in \mathbb{Z}_p[\![t]\!]$. Given the n first coefficients of g and h, at bounded precision, how far can we approximate y? At which computational cost?

The case of a general initial condition $y(0) = c$ reduces to the equation above by changing $h(t)$ in $h(t + c)$. In particular, the linear equation $y' = g \cdot y$, with $y(0) = 1$ can be written $z' = g \cdot (1 + z)$, with $z(0) = 0$, thanks to the transformation $z = y - 1$. Nevertheless, the initial condition $y(0) = 0$ ensures that the composition $h(y)$ is well defined when h is a general power series.

DEFINITION. For ε a positive real number and n a positive integer, an *approximation modulo* (p^κ, t^n) of a power series $f = \sum_{i\geqslant 0} a_i t^i \in \mathbb{Z}_p[\![t]\!]$ is a power series $\bar{f} = \sum_{i\geqslant 0} b_i t^i \in \mathbb{Z}_p[\![t]\!]$ such that $a_i = b_i \pmod{p^\kappa}$ for all $i < n$.

The complexity of computing y depends on the complexity of the power series multiplication and the composition $h(f)$ for a general power series f. Let $M_{\mathbb{Z}}(p^\lambda, n)$ the number of bit operations required to compute the product of two polynomials of degree n with coefficients in $\mathbb{Z}/p^\lambda\mathbb{Z}$. Let $C_h(p^\lambda, n)$, the number of bit operations needed to compute an approximation modulo (p^λ, t^n) of $h(f)$ given an approximation modulo (p^λ, t^n) of a power series $f \in \mathbb{Z}_p[\![t]\!]$. We assume that $C_h(p^\lambda, 2n) \geqslant 2C_h(p^\lambda, n)$. In the general case, Kedlaya and Umans (2011) proved the quasi-optimal bound $C_h(p^\lambda, n) = O\big((n\lambda \log p)^{1+o(1)}\big)$. In practice, h is given as a procedure that computes the composition $h(f)$ modulo (p^λ, t^n) for any $f \in \mathbb{Z}_p[\![t]\!]$ given modulo (p^λ, t^n). This composition is easy to compute in most applications: h is often a rational function of small degree or a radical of such a rational function so that $C_h(p^\lambda, n) = O(M_{\mathbb{Z}}(p^\lambda, n))$. We may regard h as known with infinite precision, but the computations depends only on a suitable approximation of h.

Our main result is then the following:

THEOREM 2. *Let $n > 0$, $\kappa > 0$ (or $\kappa > 1$ if $p = 2$) and let $\lambda = \kappa + \lfloor \log_p n \rfloor$. One can compute an approximation modulo (p^κ, t^{n+1}) of the solution y of* (E) *given approximations modulo (p^λ, t^n) of g and h, using $O\left(M_{\mathbb{Z}}(p^\lambda, n) + C_h(p^\lambda, n)\right)$ bit operations.*

This result was already known in the linear case: Bostan et al. (2005) gave the first proof and then Grenet, Hoeven, and Lecerf (2015) gave a simpler one. In the non-linear case, Lercier and Sirvent (2008) obtained a weaker bound: they showed that an approximation modulo (p^κ, t^n) can be computed from approximations modulo $(p^{\kappa+O(\log(n)^2)}, t^n)$ of g and h. A preliminary version of the present work appeared in Vaccon's PhD thesis (Vaccon 2015). Naturally, the result also holds over unramified extensions of $\mathbb{Q}_p$, see §2.

1.3 Applications

1.3.1 Newton sums

The problem studied by Bostan et al. (2005) is the recovery of a polynomial given its Newton sums. Let $f \in \mathbb{Z}_p[t]$ be a monic polynomial of degree d, and let ν_n be the nth Newton sum of f: if $\alpha_1, \ldots, \alpha_d$ are the roots of f in $\overline{\mathbb{Q}_p}$, then ν_n is the sum $\alpha_1^n + \cdots + \alpha_d^n$, it is an element of $\mathbb{Z}_p$. How can we recover f given Newton sums $\nu_0, \ldots, \nu_d$? Let g be the polynomial $x^d f(1/x)$, and H_f be the generating function $H_f(t) = \sum_{n\geqslant 0} \nu_{n+1} t^n$. Then $g' = -H_f g$, so that g is a solution of a first-order linear differential equation (Schönhage 1993). Therefore, knowing an approximation modulo (p^κ, t^d) of H_f makes it possible to recover each coefficient of g (and hence f) modulo $p^{\kappa - \lfloor \log_p n \rfloor}$.

An interesting application is the computation over $\mathbb{F}_p$ of composed products; composed sums can be treated similarly (Bostan et al. 2005). Let f and g be monic polynomials of $\mathbb{F}_p[t]$ of degree d and e respectively, with associated roots $(\alpha_i)_{1\leqslant 1\leqslant d}$ and $(\beta_j)_{1\leqslant j\leqslant e}$ in $\overline{\mathbb{F}_p}$. We define the composed product of f and g to be

$$f \otimes g = \prod_{i,j}(t - \alpha_i\beta_j) = \operatorname*{res}_y(y^d f(t/y), g(y)),$$

this is a polynomial in $\mathbb{F}_p[t]$ of degree de. Then $H_{f\otimes g}$ is the coefficient-wise product of the two power series H_f and H_g (also known as the Hadamard product). This gives a strategy to compute efficiently $f \otimes g$. Firsty, arbitrarily lift f and g as polynomials in $\mathbb{Z}_p[t]$, denoted $\bar{f}$ and $\bar{g}$. The composed product $\bar{f} \otimes \bar{g}$ is a polynomial in $\mathbb{Z}_p[t]$ and equals $f \otimes g$ modulo p. Secondly, compute approximations modulo (p^κ, t^{de}) of $H_{\bar{f}}$ and $H_{\bar{g}}$, with $\kappa = 1 + \lfloor \log_p(de) \rfloor$, and, with a coefficient-wise product, an approximation modulo (p^κ, t^{de}) of $H_{\bar{f}\otimes\bar{g}}$. Thirdly, compute an approximation modulo (p, t^{de+1}) of $\bar{f}\otimes\bar{g}$ using Theorem 2, and deduce the value of $f \otimes g$.

Another application of this procedure to root finding in finite fields is developed by Grenet, Hoeven, and Lecerf (2015). Our work does not improve the complexity given by Bostan et al. (2005), in the case of a linear equation, but it gives a simpler proof that generalizes to nonlinear differential equations.

1.3.2 Isogeny computation

To compute normalized isogenies between elliptic curves, Bostan et al. (2008) and Lercier and Sirvent (2008) studied the differential equation

$$y'^2 = g \cdot h(y), \tag{1}$$

where g and h are series in $\mathbb{Z}_p[\![t]\!]$. In their context, this differential equation is known to admit a solution in $\mathbb{Z}_p[\![t]\!]$. Like the

previous example, it comes from a lift of a problem over $\mathbb{F}_p$. This equation rewrites equivalently as $y' = \sqrt{g}\sqrt{h(y)}$, and, when $p \neq 2$, the series $\sqrt{g}$ and $\sqrt{h}$ are still in $\mathbb{Z}_p[\![t]\!]$, so we can apply Theorem 2. The study of this equation when $p = 2$ is still an open problem.

In order to compute an approximation of y modulo (p, t^{n+1}), we obtain that it is enough to have approximations of g and h modulo $(p^{1+\lfloor \log_p n \rfloor}, t^n)$. This improves upon the result of Lercier and Sirvent (2008) which requires approximations modulo $(p^{O(\log(n)^2)}, t^n)$.

Acknowledgements

We are grateful to Alin Bostan, Xavier Caruso, Luca De Feo, Reynald Lercier, Éric Schost and Kazuhiro Yokoyama for fruitful discussions.

2. THE ALGORITHM

We may consider the more general setting of an unramified finite extension K of $\mathbb{Q}_p$. This is useful, for example, for the computation of isogenies. Let $\mathcal{O}_K$ denote the ring of integers of K. For example, we may naturally consider $K = \mathbb{Q}_p$ and $\mathcal{O}_K = \mathbb{Z}_p$.

Let $h \in \mathcal{O}_K[\![t]\!]$, a power series with integer coefficients, with $h(0) = 1$. For $g \in K[\![t]\!]$, let $Y(g)$ be the unique $y \in K[\![t]\!]$ such that $y(0) = 0$ and $y' = g \cdot h(y)$. Existence and uniqueness are clear because the differential equation rewrites equivalently into a well posed recurrence relation on the coefficients of y.

For $g \in K[\![t]\!]$, let N_g denote the Newton operator:

$$N_g(u) = u - h(u) \int \left(\frac{u'}{h(u)} - g \right),$$

where $\int f$, for $f \in K[\![t]\!]$, denotes the unique power series $F \in K[\![t]\!]$ such that $F' = f$ and $F(0) = 0$.

Proposition 3. *Let $g, u \in K[\![t]\!]$ be formal power series, and let $n > 0$. If $u = Y(g) \pmod{t^n}$ then*

$$N_g(u) = Y(g) \pmod{t^{2n}}.$$

Proof. Let $e = u'/h(u) - g$. Since $u = Y(g) \pmod{t^n}$, we have $e = 0 \pmod{t^{n-1}}$ and $\int e = 0 \pmod{t^n}$. With $v = N_g(u)$, we compute

$$v' - g\,h(v) = u' - u'h'(u) \int e - h(u)e - g\,h(v).$$

Then, the first-order expansion of $h(v)$ at $h(u)$ gives

$$h(v) = h(u) - h'(u)h(u) \int e \mod t^{2n},$$

and, using the equality $u' = h(u)e + g\,h(u)$, we obtain

$$v' - g\,h(v) = -e\,h(u)h'(u) \int e = 0 \mod t^{2n-1}.$$

This implies that $v = Y(g) \pmod{t^{2n}}$. □

The iteration of the Newton operator leads to Algorithm 1. In an exact setting, the correctness of this procedure would be clear, thanks to Proposition 3. In a p-adic setting, where the coefficients of the power series g and h are known with finite precision only, what can be obtained with Newton iteration is not clear because the operation $\int$ involves divisions.

Let us begin with a quick analysis of Algorithm 1. On input (g, h, n), it performs a recursive call and computes $u = \textsc{DSol}(g, h, m)$, with $m = \lceil \frac{n-1}{2} \rceil$. Let us assume that u is an approximation modulo (p^κ, t^{m+1}) of $Y(g)$, for some $\kappa > 0$.

Algorithm 1. The Newton iteration to solve a first-order differential equation.

Input. g and $h \in \mathcal{O}_K[\![t]\!]$ given modulo (p^λ, t^n) — that is, given as polynomials of degree less than n with coefficients in $\mathcal{O}_K/p^\lambda\mathcal{O}_K$.

Output. A power series $u \in \mathcal{O}_K[\![t]\!]$ given modulo (p^λ, t^{n+1}).

Specification. If $Y(g) \pmod{t^{n+1}}$ has integer coefficients and if $\lambda \geqslant \kappa + \lfloor \log_p n \rfloor$, then u is an approximation of $Y(g)$ modulo (p^κ, t^{n+1}).

```
function DSol(g, h, n)
    if n = 0 then
        return 0 (mod t)
    else
        u ← DSol(g, h, ⌈(n−1)/2⌉)
        return N_g(u) (mod t^{n+1})
                        ▷ Compute at fixed precision λ.
```

Then $N_g(u)$ is an approximation modulo $(p^{\kappa - \lfloor \log_p n \rfloor}, t^{n+1})$: Indeed, the computation of $N_g(u)$ involves divisions by the integers from 2 to n on distinct coefficients, so the loss of precision is at most the maximum valuation of these integers, which is $\lfloor \log_p n \rfloor$. Thus, if we define $\mu(0) = 0$ and $\mu(n) = \lfloor \log_p n \rfloor + \mu(\lceil \frac{n-1}{2} \rceil)$, we obtain that $\textsc{DSol}(g, h, n)$ is an approximation modulo $(p^{\lambda - \mu(n)}, t^{n+1})$ of $Y(g)$. We can check that $\mu(n) = O(\log(n)^2)$. Theorem 5 improves on that analysis and shows that the precision of the result is at least $\lambda - \lfloor \log_p n \rfloor$ and matches the intrinsic loss of precision.

We assume the *fixed precision model* for computing with p-adic numbers: at precision λ, it amounts to work over the ring $\mathcal{O}_K/p^\lambda\mathcal{O}_K$. When a division a/b arises, with the approximation of a and $b \in \mathcal{O}_K$ given in $\mathcal{O}_K/p^\lambda\mathcal{O}_K$, three cases may arise:

- $v_p(b) = 0$, in which case b is invertible in $\mathcal{O}_K/p^\lambda\mathcal{O}_K$ and the division is well defined.
- $v_p(a) \geqslant v_p(b) > 0$, in which case a/b is in $\mathcal{O}_K$ but its approximation in $\mathcal{O}_K/p^\lambda\mathcal{O}_K$ is not fully determined by the approximations of a and b, so a/b is arbitrarily defined in this model to be the class in $\mathcal{O}_K/p^\lambda\mathcal{O}_K$ of the smallest integer $c \geqslant 0$ such that $a = bc \pmod{p^\lambda}$.
- $v_p(a) < v_p(b)$, in which case a/b is not an integer and an error is raised.

The main argument for the correctness of Algorithm 1 is the following proposition, proved in Section 3.

Proposition 4. *Let $n > 0$ and $\kappa > 0$ (or $\kappa > 1$ if $p = 2$) be integers, and let $g \in \mathcal{O}_K[\![t]\!]$ such that $Y(g) \pmod{t^{n+1}}$ has integer coefficients. For any $y \in K[\![t]\!]$ the following are equivalent:*

1. *$y = Y(\bar{g}) \pmod{t^{n+1}}$ for some power series $\bar{g} \in \mathcal{O}_K[\![t]\!]$ such that $\int(\bar{g} - g) = 0 \pmod{p^\kappa}$;*
2. *$y = Y(g) \pmod{p^\kappa, t^{n+1}}$.*

Theorem 5. *Algorithm 1 is correct: if $\kappa > 0$ (or $\kappa > 1$ if $p = 2$) and $\lambda \geqslant \lfloor \log_p n \rfloor + \kappa$, then for all $g \in \mathcal{O}_K[\![t]\!]$ such that $Y(g)$ has integer coefficients, the output of the procedure $\textsc{DSol}(g, h, n)$ equals $Y(g) \pmod{p^\kappa, t^{n+1}}$.*

Moreover, it performs $\mathcal{O}\left(M_{\mathcal{O}_K}(p^\lambda, n) + C_h(p^\lambda, n)\right)$ *bit operations, where* $M_{\mathcal{O}_K}(p^\lambda, n)$ *is the cost of computing the product of two polynomials of degree* n *with coefficients in* $\mathcal{O}_K/p^\lambda\mathcal{O}_K$.

Proof. We proceed by induction on n. The case $n = 0$ is trivial, so let us assume that $n > 0$. Let $g \in \mathcal{O}_K[\![t]\!]$ such that $Y(g)$ has integer coefficients, let $m = \lceil \frac{n-1}{2} \rceil$ and let $u \in K[\![t]\!]$ be the output of $\textsc{DSol}(g, h, m)$. By induction hypothesis, $u = Y(g) \pmod{p^\kappa, t^{m+1}}$. (In particular u has integer coefficients.) By Proposition 4, this implies that $y = Y(\bar{g}) \pmod{t^{m+1}}$ for some $\bar{g} \in \mathcal{O}_K[\![t]\!]$ such that $\int(g - \bar{g}) \pmod{p^\kappa}$. Proposition 3 gives that $Y(\bar{g}) = N_{\bar{g}}(u) \pmod{t^{n+1}}$ and Proposition 4 gives further that $Y(\bar{g}) = Y(g) \pmod{p^\kappa, t^{n+1}}$. We check that $N_g(u) = N_{\bar{g}}(u) - h(u)\int(g - \bar{g})$ and since $h(u)$ has integer coefficients, this implies that $N_g(u) = Y(g) \pmod{p^\kappa, t^{n+1}}$.

We now relate $N_g(u)$ to the output of the procedure $\textsc{DSol}(g, h, n)$. Let $e = u'/h(u) - g$. By definition, the output is $N_g(u) = u - h(u)\int e$, computed over $\mathcal{O}_K/p^\lambda\mathcal{O}_K$, in the fixed precision model. Let E be the primitive $\int e \pmod{t^{n+1}}$ computed in this model, so that the output is exactly $u - h(u)E$. Clearly $E = \int e + \int \eta \pmod{t^{n+1}}$ for some $\eta = 0 \pmod{p^\lambda}$ that reflects the indeterminacies in the divisions. Since $\lambda \geqslant \lfloor \log_p n \rfloor + \kappa$, $\int \eta = 0 \pmod{p^\kappa}$ and thus, the output $N_g(u) + \int \eta$ equals $Y(g) \pmod{p^\kappa, t^{n+1}}$. This concludes the proof of correctness.

Concerning the complexity, the last iteration involves a composition by h with cost $C_h(p^\lambda, n)$, a few multiplications with cost $M_{\mathcal{O}_K}(p^\lambda, n)$ and an inversion $1/h(u)$ with cost $\mathcal{O}\left(M_{\mathcal{O}_K}(p^\lambda, n)\right)$ too with a Newton iteration (Kung 1974). With the assumption that the cost of an iteration is greater than twice the cost of the previous one, it is well known that the cost of a Newton algorithm is dominated by the cost of the last iteration, which gives the result. □

The condition $\lambda \geqslant \lfloor \log_p n \rfloor + \kappa$ cannot be improved further: it matches the *intrinsic* loss of precision. This is shown, for example, by the differential equation $y' = at^{b-1}$, with $a \in K$, whose solution is $\frac{a}{b}t^b$. If we take $b = p^{\lfloor \log_p n \rfloor}$ and if a is known at precision λ then y is known at precision no more than $\lambda - \lfloor \log_p n \rfloor$.

3. DIFFERENTIAL PRECISION

We apply the method of Caruso, Roe, and Vaccon (2014) to study the loss in precision in the resolution of the differential equation (E) and give a proof of Proposition 4.

Let $n > 0$ and let E and F denote respectively the two n-dimensional K-vector spaces $K[\![t]\!]/(t^n)$ and $tK[\![t]\!]/(t^{n+1})$. Let $\mathcal{Y}$ be the polynomial map

$$\begin{aligned} \mathcal{Y} : E &\longrightarrow F \\ [u] &\longmapsto [Y(u)], \end{aligned}$$

which is well defined because the $n+1$ first coefficients of $Y(u)$ depend only on the n first coefficients of u. Let $g \in \mathcal{O}_K[\![t]\!]$ be such that $\mathcal{Y}(g)$ has integer coefficients in the monomial basis. Let $\mathrm{d}\mathcal{Y}$ denote the first differential of $\mathcal{Y}$: for any $g \in E$, $\mathrm{d}_g\mathcal{Y}$ is a linear map $E \to F$. Let $d^k\mathcal{Y}$ denote the higher differentials: for any $g \in E$, $\mathrm{d}_g^k\mathcal{Y}$ is a multilinear map $E^k \to F$.

Lemma 6. *For any* $w \in E$, $\mathrm{d}\mathcal{Y}(w) = h(\mathcal{Y})\int w$. *Moreover, for any* $k \geqslant 1$, *there exists a polynomial* $P_k \in \mathbb{Z}[u_0, \ldots, u_{k-1}]$ *such that for any* $w_1, \ldots, w_k \in E$,

$$\mathrm{d}^k\mathcal{Y}(w_1, \ldots, w_k) = P_k\left(h(\mathcal{Y}), h'(\mathcal{Y}), \ldots, h^{(k-1)}(\mathcal{Y})\right)\prod_{i=1}^{k}\int w_i. \quad (2)$$

Proof. Differentiating with respect to g the defining relation $\mathcal{Y}(g)' = g \cdot h(\mathcal{Y}(g)) \pmod{t^n}$ leads to

$$[\mathrm{d}_g\mathcal{Y}(w)]' = w \cdot h(\mathcal{Y}(g)) + g \cdot h'(\mathcal{Y}(g)) \cdot \mathrm{d}_g\mathcal{Y}(w),$$

which is a first-order inhomogeneous linear differential equation in $\mathrm{d}_g\mathcal{Y}(w)$. The initial condition $\mathrm{d}_g\mathcal{Y}(w)(0) = 0$ determines a unique solution, namely $h(\mathcal{Y}(g))\int w$.

The second claim follows by induction. Equation (2) holds for $k = 1$ with $P_1 = u_0$; differentiating it leads to the recurrence relation

$$P_{k+1}(u_0, \ldots, u_k) = u_0\sum_{i=0}^{k-1}\frac{\partial P}{\partial u_i}u_{i+1}. \qquad \square$$

The space F is endowed with the maximum norm in the monomial basis, denoted by $\|\cdot\|_F$. In particular, an element u of F has integer coefficients if and only if $\|u\|_F \leqslant 1$. The space E is endowed with the norm

$$\|u\|_E \stackrel{\text{def}}{=} \left\|\int u\right\|_F.$$

Let $\|\mathrm{d}_g^k\mathcal{Y}\|$ denote the operator norm of $\mathrm{d}_g^k\mathcal{Y}$, that is

$$\|\mathrm{d}_g^k\mathcal{Y}\| = \sup\left\{\|\mathrm{d}_g^k\mathcal{Y}(w_1, \ldots, w_k)\|_F \;\middle|\; \|\int w_i\|_F \leqslant 1\right\}. \quad (3)$$

Lemma 7. $\|\mathrm{d}_g^k\mathcal{Y}\| \leqslant 1$, *for any* $k \geqslant 1$.

Proof. Since h and $\mathcal{Y}(g)$ have integer coefficients, this follows easily from Lemma 6 and Equation (3). □

Proposition 8. *For any* $\varepsilon \leqslant \frac{1}{p}$ *(or* $\varepsilon \leqslant \frac{1}{4}$ *for* $p = 2$*),*

$$\mathcal{Y}(g + B_\varepsilon) = \mathcal{Y}(g) + \mathrm{d}_g\mathcal{Y}(B_\varepsilon),$$

where $B_\varepsilon = \{w \in E \;:\; \|w\|_E \leqslant \varepsilon\}$.

Proof. We apply the result of Caruso, Roe, and Vaccon (2014, Corollary 3.16). Using their notations, we can use $C = 1$ because the closed ball of radius 1 in F (that is the set of elements with integer coefficients) is included in $\mathrm{d}_g\mathcal{Y}(B_1)$: Indeed for any $u \in F$ we have $\mathrm{d}_g\mathcal{Y}\left((u/\mathcal{Y}(g))'\right) = u$, and if $\|u\|_F \leqslant 1$ then $\|(u/\mathcal{Y}(g))'\|_E = \|(u/\mathcal{Y}(g))\|_F \leqslant 1$, because u and $\mathcal{Y}(g)$ have integer coefficients.

For $k \geqslant 2$, let M_k denote $\|\frac{1}{k!}\mathrm{d}_g^k\mathcal{Y}\|$. By Lemma 7, this is simply $|\frac{1}{k!}|$. Corollary 3.16 (ibid.), with $\rho = 1$ in their notations, implies that $\mathcal{Y}(g + B_\varepsilon) = \mathcal{Y}(g) + \mathrm{d}_g\mathcal{Y}(B_\varepsilon)$ as long as ε satisfies

$$\varepsilon < \exp\left(\inf_{k\geqslant 2}\frac{-\log M_k}{k-1}\right) = \inf_{k\geqslant 2} p^{-\frac{v_p(k!)}{k-1}}.$$

Let A denote the right-hand side. Legendre's formula for the p-adic valuation of $k!$ shows that $v_p(k!) \leqslant \frac{k}{p-1}$. Therefore $A \geqslant p^{-2/(p-1)}$. For $p \geqslant 5$, this bound gives $A > \frac{1}{p}$, which proves the claim. For $p = 3$, we have

$$A \geqslant \min\left(3^{-v_3(2!)}, \inf_{k\geqslant 3} 3^{-\frac{v_3(k!)}{k-1}}\right) \geqslant \min\left(1, 3^{-\frac{3}{4}}\right) > 3^{-1},$$

Figure 1. Timings in seconds, measured on a laptop, of Algorithm 1 run at precision λ_{old} (upper curve) and λ_{new} (lower curve) in order to compute an approximation modulo $(5, t^{4m+1})$ of the solution of Equation (4).

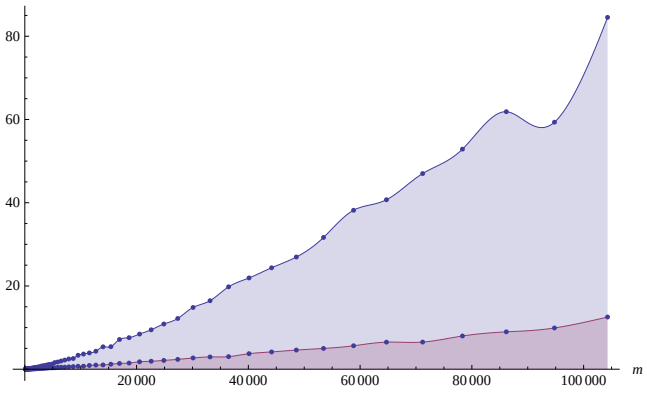

Figure 2. Practical speedup obtained with the new precision analysis compared with the theoretical improvement (m-axis in logarithmic scale).
(■) ratio $\lambda_{\text{old}}/\lambda_{\text{new}}$; (•) actual speedup.

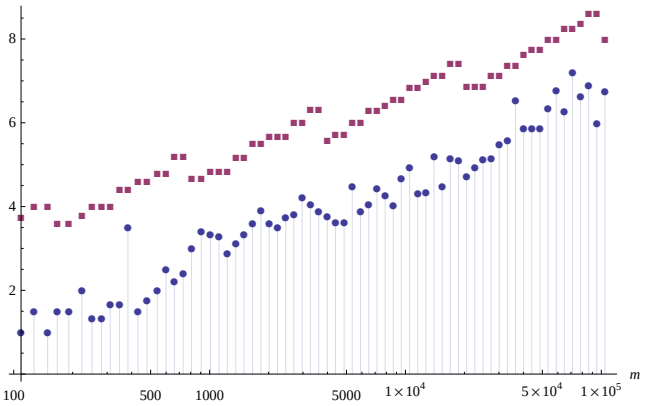

and for $p = 2$, we have

$$A \geqslant \min\left(2^{-v_2(2!)}, \inf_{k\geqslant 3} 2^{-\frac{v_2(k!)}{k-1}}\right) \geqslant \min\left(2^{-1}, 2^{-\frac{3}{2}}\right) > 2^{-2},$$

which concludes the proof. □

Proof of Proposition 4. Let $\varepsilon = p^{-\kappa}$. The norm of an element $v \in F$ is given by $p^{-\lambda}$ where λ is the largest integer such that $v = 0 \pmod{p^\lambda}$. Since $h(\mathcal{Y}(g))$ is invertible modulo (p^λ, t^{n+1}), for any $\lambda > 0$, this shows that $\|v\|_F = \|h(\mathcal{Y}(g))v\|_F$ for any $v \in F$. Therefore, with Lemma 6 and the definition of the norms,

$$\mathrm{d}_g\mathcal{Y}(B_\varepsilon) = \{v \in F \mid v = 0 \mod p^\kappa\}.$$

Moreover, $B_\varepsilon = \left\{u \in E \;\middle|\; \int u = 0 \mod p^\kappa\right\}$, and Proposition 4 now appears as a rewording of Proposition 8. □

4. EXPERIMENTS

Let us consider the differential equation

$$y' = \sqrt{\frac{1 + \frac{1}{4}m^2y^2 + m^6y^6}{1 + \frac{1}{4}t^2 + t^6}}, \quad y(0) = 0, \tag{4}$$

inspired from algorithms for computing isogenies (Bostan et al. 2008; Lercier and Sirvent 2008). Using an implementation in Magma (Bosma, Cannon, and Playoust 1997) of Algorithm 1, we computed the power series expansion of $y \pmod{5, t^{4m+1}}$ for several m. We compared (Figure 1) the CPU time spent on the computation when using on the one hand the precision $\lambda_{\text{new}} = 1 + \lfloor \log_5(4m) \rfloor$, following Theorem 5, and using on the other hand the precision $\lambda_{\text{old}} = 1 + \mu(4m) = \mathcal{O}(\log(m)^2)$ found by a straightforward precision analysis — see the discussion in §2 for the definition of μ. For example, with $m = 104281$, we compute $\lambda_{\text{old}} = 72$ and $\lambda_{\text{new}} = 9$. The number of arithmetic operations performed does not depend on the precision λ, only on m, but the number of bit operations does since the base ring for the computation is $\mathbb{Z}/5^\lambda\mathbb{Z}$. Thus, the expected speedup is $\lambda_{\text{old}}/\lambda_{\text{new}}$, which is close to what we observed (Figure 2). The implementation is available at

https://gist.github.com/lairez/d648b0d7b5392d0fef74.

References

W. Bosma, J. Cannon, and C. Playoust (1997). "The Magma algebra system. I. The user language". In: *J. Symbolic Comput.* 24.3-4, pp. 235–265. DOI: 10.1006/jsco.1996.0125.

A. Bostan, L. González-Vega, H. Perdry, and É. Schost (2005). "From Newton sums to coefficients: complexity issues in characteristic p". In: *Proc. of MEGA*. Ed. by A. Dickenstein, P. Gianni, and T. Recio. Porto Conte, Italy.

A. Bostan, F. Morain, B. Salvy, and É. Schost (2008). "Fast algorithms for computing isogenies between elliptic curves". In: *Math. Comput.* 77.263, pp. 1755–1778. DOI: 10.1090/S0025-5718-08-02066-8.

X. Caruso, D. Roe, and T. Vaccon (2014). "Tracking p-adic precision". In: *LMS J. Comput. Math.* 17.suppl. A, pp. 274–294. DOI: 10.1112/S1461157014000357.

– (2015). "p-Adic Stability In Linear Algebra". In: *Proc. of ISSAC*. Bath, United Kingdom. ACM, pp. 101–108. DOI: 10.1145/2755996.2756655.

B. Grenet, J. van der Hoeven, and G. Lecerf (2015). "Deterministic root finding over finite fields using Graeffe transforms". In: *Appl. Algebra Engrg. Comm. Comput.* DOI: 10.1007/s00200-015-0280-5.

K. S. Kedlaya and C. Umans (2011). "Fast polynomial factorization and modular composition". In: *SIAM J. Comput.* 40.6, pp. 1767–1802. DOI: 10.1137/08073408X.

H. T. Kung (1974). "On computing reciprocals of power series". In: *Numer. Math.* 22, pp. 341–348.

R. Lercier and T. Sirvent (2008). "On Elkies subgroups of ℓ-torsion points in elliptic curves defined over a finite field". In: *J. Théor. Nombres Bordeaux* 20, pp. 783–797.

A. Schönhage (1993). "Fast parallel computation of characteristic polynomials by Leverrier's power sum method adapted to fields of finite characteristic". In: *Automata, languages and programming (Lund, 1993)*. Vol. 700. LNCS. Berlin: Springer, pp. 410–417. DOI: 10.1007/3-540-56939-1_90.

T. Vaccon (2015). "Précision p-adique. Applications en calcul formel, théorie des nombres et cryptographie". PhD thesis. Université de Rennes I.