Composition modulo powers of polynomials

Joris van Der Hoeven, Grégoire Lecerf

► To cite this version:

Joris van Der Hoeven, Grégoire Lecerf. Composition modulo powers of polynomials. 2017. hal-01455722v2

HAL Id: hal-01455722

<https://hal.science/hal-01455722v2>

Preprint submitted on 27 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Composition modulo powers of polynomials

Joris van der Hoeven and Grégoire Lecerf

Laboratoire d'informatique de l'École polytechnique, CNRS UMR 7161

École polytechnique

91128 Palaiseau Cedex, France

{vdhoeven,lecerf}@lix.polytechnique.fr

ABSTRACT

Modular composition is the problem to compose two univariate polynomials modulo a third one. For polynomials with coefficients in a finite field, Kedlaya and Umans proved in 2008 that the theoretical bit complexity for performing this task could be made arbitrarily close to linear. Unfortunately, beyond its major theoretical impact, this result has not led to practically faster implementations yet. In this paper, we study the more specific case of composition modulo the power of a polynomial. First we extend previously known algorithms for power series composition to this context. We next present a fast direct reduction of our problem to power series composition.

Keywords

Modular composition; Series Composition; Algorithm; Complexity

1. INTRODUCTION

Let $\mathbb{K}$ be an effective field and let f, g, h be polynomials in $\mathbb{K}[x]$ of degrees $\leq n$. The problem of *modular composition* is to compute $f \circ g$ modulo h . This is a fundamental problem in complexity theory because of its applications to polynomial factorization [16, 17, 18].

From a theoretical point of view, Kedlaya and Umans achieved a major breakthrough by showing that modular compositions over a finite field $\mathbb{F}_q$ can be computed in time $(n \log q)^{1+o(1)}$. Unfortunately, the practical impact of these results has been limited so far. The best current implementations still admit nearly quadratic complexities in n . More precisely, over a field $\mathbb{K}$, denoting by $M_{\mathbb{K}}(n)$ the cost to multiply two polynomials in $\mathbb{K}[x]$ of degrees $\leq n$, Brent and Kung [3] gave an algorithm with cost $O(\sqrt{n} M_{\mathbb{K}}(n) + n^2)$. It uses the *baby-step giant-step* technique due to Paterson–Stockmeyer [21] and even yields a sub-quadratic cost when using fast linear algebra (see [15, p. 185] and our Theorem 2 below).

In the special case when $h = x^n$, composition modulo h corresponds to the composition of truncated power series at order n . Better practical algorithms are known for this particular situation. Brent and Kung gave an algorithm [3] with complexity $O(\sqrt{n} M_{\mathbb{K}}(n) \log^{1/2} n)$ under the conditions that $g'(0)$ is invertible and that the characteristic is at least $\lceil \sqrt{n \log n} \rceil$. A variant proposed by van der Hoeven [11, Section 3.4.3] removes the condition on $g'(0)$. For fields of small positive characteristic, Bernstein [1] proposed an algorithm that is softly linear in the precision n but linear in the characteristic. Series with integer, rational or floating point coefficients can often be composed in quasi-linear time as well in suitable bit complexity models, as shown by Ritzmann [22]; see also [12].

In this paper, we focus on the intermediate case between power series and general modular compositions. More precisely, we assume h to be a power $\tilde{h}^m$ of some polynomial $\tilde{h}$. Composition modulo $\tilde{h}^m$ can also be regarded as a special “base” case of composition modulo a polynomial h that can be factored into $h = h_1 \cdots h_m$. Using Chinese remaindering, one may reduce composition modulo such h to compositions modulo powers of irreducible polynomials. In a separate paper [13], we exploit this observation for the design of fast practical (although not fully general) algorithms for modular composition over finite fields.

We develop two main approaches for fast composition modulo $h = \tilde{h}^m$. In section 3, we generalize Brent and Kung’s and Bernstein’s algorithms. We achieve a complexity exponent close to $3/2$, but only under the condition that n is much larger than $\deg \tilde{h}$. In section 4, we directly reduce composition modulo h to truncated power series composition. Our main result is Theorem 11: one composition modulo h is essentially reduced to $m + 1$ compositions modulo $\tilde{h}$ and one power series composition in $\mathbb{K}[z] / (\tilde{h}(z))[y] / (y^m)$. An important ingredient of this reduction is a softly optimal algorithm for converting between $\mathbb{K}[x] / (h(x))$ and $\mathbb{K}[z] / (\tilde{h}(z))[y] / (y^m)$ where x is sent to $y + z$. Potential applications of these conversions might also concern algebraic algorithms for plane curves such as integral basis computations and Puiseux expansions, which usually require several Taylor expansions at roots of discriminants.

Our paper is organized as follows. In section 2, we introduce basic notations and recall several well known complexity results. We also specify the complexity model that we work with. Further technical complexity results are gathered in the appendix. In sections 3 and 4, we present our main algorithms. For testing purposes, our algorithms are implemented in the MATHEMAGIX language within the FACTORIX library (www.mathemagix.org, revision 10538).

2. PRELIMINARIES

2.1 Costs of elementary operations

Recall that an *effective ring* is a ring $\mathbb{A}$ with unity whose elements can be represented on a computer and such that we have algorithms for performing the ring operations. Effective fields $\mathbb{K}$ and effective algebras over an effective ring are defined similarly.

Given an effective ring $\mathbb{A}$, *algebraic complexity models* express running times in terms of the number of ring operations in $\mathbb{A}$. Unless otherwise stated, we will analyze the costs of the algorithms in this paper in this way. More precisely, our results both apply for the straight-line program and computation tree models [4, Chapter 4].

Let $\mathbb{A}$ be an effective ring with unity, let $n \in \mathbb{N}$, and denote

$$\mathbb{A}[x]_{<n} = \{P \in \mathbb{A}[x] : \deg P < n\}.$$

We write $M_{\mathbb{A}} : \mathbb{N} \rightarrow \mathbb{R}^{\geq}$ for a cost function such that two polynomials in $\mathbb{A}[x]_{<n}$ can be multiplied using $M_{\mathbb{A}}(n)$ operations in $\mathbb{A}$. It will be convenient to assume that $M_{\mathbb{A}}(n)/n$ is a nondecreasing function in n . Notice that this assumption implies the *super-additivity* of $M_{\mathbb{A}}$, namely $M_{\mathbb{A}}(n_1) + M_{\mathbb{A}}(n_2) \leq M_{\mathbb{A}}(n_1 + n_2)$ for all $n_1 \geq 0$ and $n_2 \geq 0$.

The schoolbook algorithm allows us to take $M_{\mathbb{A}}(n) = O(n^2)$. The fastest currently known algorithm [5] yields $M_{\mathbb{A}}(n) = O(n \log n \log \log n)$. If $\mathbb{K}$ is a field of finite characteristic, then we even have $M_{\mathbb{K}}(n) = O(n \log n 8^{\log^* n})$, where $\log^* n = \min \{i \in \mathbb{N} : (\log \circ \dots \circ \log)(n) \leq 1\}$ [10].

The constant $\omega > 2$ represents a *feasible exponent* for the multiplication cost of matrices: two square matrices of size $n \times n$ can be multiplied using $O(n^\omega)$ operations in their coefficient ring. The constant $\varpi > 1.5$ is defined similarly but for multiplying a $\sqrt{n} \times \sqrt{n}$ matrix by a $\sqrt{n} \times n$ rectangular one. The best known bound $\omega < 2.3729$ is due to Le Gall [19]. This naturally yields $\varpi \leq (\omega + 1)/2 < 1.6845$. However the latter bound does not improve upon the earlier bound $\varpi < 1.6667$ due to Huang and Pan [14, Theorem 10.1].

2.2 Problems related to modular composition

If f and g are polynomials in $\mathbb{A}[x]$, then $f \bmod g$ represents the remainder in the division of f by g . Let $\mathbb{A}$ be an effective ring and $\mathbb{B}$ be an effective $\mathbb{A}$ -algebra. We introduce the following cost functions:

- $S_{\mathbb{A}}(n)$: the cost of computing the *power series composition* $f \circ g \bmod x^n$, where $f, g \in \mathbb{A}[x]_{<n}$.
- $S_{\mathbb{B}/\mathbb{A}}(n)$: the cost of computing the *power series composition* $f \circ g \bmod x^n$ in terms of operations in $\mathbb{A}$, where $f, g \in \mathbb{B}[x]_{<n}$.
- $C_{\mathbb{A}}(n)$: the cost of computing the *modular composition* $f \circ g \bmod h$, where $h \in \mathbb{A}[x]$ is a monic polynomial of degree n and $f, g \in \mathbb{A}[x]_{<n}$.
- $Q_{\mathbb{A}}(n)$: the cost of computing the *characteristic polynomial* χ of $g \in \mathbb{A}[x]_{<n}$ modulo a monic polynomial $h \in \mathbb{A}[x]$ of degree n . This is the characteristic polynomial $\chi \in \mathbb{A}[x]$ of the multiplication endomorphism by $g \bmod h$ in $\mathbb{A}[x]/(h(x))$.

- $P_{\mathbb{A}}(n)$: the cost of *modular power projections*, i.e. the cost to compute $\varphi(1), \varphi(g), \dots, \varphi(g^{n-1} \bmod h)$, where $h \in \mathbb{A}[x]$ is monic of degree n and φ is a linear form on $\mathbb{A}[x]_{<n}$.

Let us recall a few known results about these functions. For $a \in \mathbb{R}$, we denote $\lfloor a \rfloor = \max \{b \in \mathbb{Z} : b \leq a\}$ and $\lceil a \rceil = \min \{b \in \mathbb{Z} : b \geq a\}$.

THEOREM 1. *Let $\mathbb{K}$ be a field of characteristic p .*

- a) *[3] If $p=0$ or $p > \lceil \sqrt{n \log n} \rceil$, then we may take $S_{\mathbb{K}}(n) = O(M_{\mathbb{K}}(n) \sqrt{n \log n})$.*
- b) *[1] If $p > 0$, then $S_{\mathbb{K}}(n) = O\left(pn + \frac{p}{\log p} M_{\mathbb{K}}(n) \log n\right)$.*

THEOREM 2. *Let h be a monic polynomial of degree n over a ring $\mathbb{A}$ and let $f, g \in \mathbb{A}[x]_{<n}$. The composed polynomial $f \circ g \bmod h$ may be computed with*

- a) *$O(n M_{\mathbb{A}}(n))$ operations in $\mathbb{A}$, or*
- b) *$O(n^\varpi + n^{1/2} M_{\mathbb{A}}(n))$ or $O(n^\varpi)$ operations in $\mathbb{A}$.*

PROOF. The first bound is immediate. The proof of the second bound is detailed in [8, Section 12.2]. The bound $O(n^\varpi)$ follows from $\varpi > 1.5$ by taking $M_{\mathbb{A}}(n) = n^{1+o(1)}$. $\square$

For a fixed monic polynomial h in $\mathbb{A}[x]$ of degree n , the modular composition $f \circ g \bmod h$ is a linear operation in f , for f and g in $\mathbb{A}[x]_{<n}$. The corresponding transposed application is precisely the operation of *modular power projections*: it corresponds to computing $\varphi(1), \varphi(g), \dots, \varphi(g^{n-1} \bmod h)$, where φ is a linear form on $\mathbb{A}[x]_{<n}$. If a modular composition algorithm with cost $C_{\mathbb{A}}(n)$ can be transposed in the sense of [2], then this leads to a power projection algorithm with cost $P_{\mathbb{A}}(n) = C_{\mathbb{A}}(n) + O(n)$.

THEOREM 3. *Let h be a monic polynomial of degree n over a ring $\mathbb{A}$ and let $g \in \mathbb{A}[x]_{<n}$. The characteristic polynomial χ of g modulo h can be computed using*

- a) *$O(M_{\mathbb{A}}(n^2) \log n + n M_{\mathbb{A}}(n) \log^2 n)$ operations in $\mathbb{A}$, including divisions in $\mathbb{A}$ (the partial division in $\mathbb{A}$ is supposed to be implemented), or*
- b) *$O(M_{\mathbb{A}}(n^2) \log n)$ operations in $\mathbb{A}$, if $\mathbb{A}$ is a field, or*
- c) *$O(n M_{\mathbb{A}}(n) \log n)$ operations in $\mathbb{A}$, if $\mathbb{A}$ is a field with $>n$ elements, or*
- d) *$P_{\mathbb{A}}(n) + M_{\mathbb{A}}(n)$ operations in $\mathbb{A}$, if there exist given inverses of $2, 3, \dots, n$ in $\mathbb{A}$.*

PROOF. The first two bounds directly rely on the formula $\chi(x) = \text{Res}_z(g(z) - x, h(z))$ by using [20, Corollary 29]. The third bound is obtained by computing $\text{Res}_z(g(z) - a, h(z))$ for $n+1$ values of a in $\mathbb{A}$, from which χ is interpolated using $O(M_{\mathbb{A}}(n) \log n)$ additional operations in $\mathbb{A}$. We refer to Appendix A.1 for the fourth bound. $\square$

For the particular kind of moduli $h = h^m$ that interest us in here, the problem of computing characteristic polynomials quite easily reduces to the case when $m = 1$:

PROPOSITION 4. *Let h be of degree d in $\mathbb{K}[x]$, let $h = h^m$ with $n = m d$, and let g be in $\mathbb{K}[x]_{<n}$. Then the characteristic polynomial χ of g modulo h can be computed using $Q_{\mathbb{K}}(d) + O(M_{\mathbb{K}}(n))$ operations in $\mathbb{K}$.*

PROOF. The characteristic polynomial χ is simply the m -th power of the characteristic polynomial of g modulo $\tilde{h}$. $\square$

3. EXTENDING KNOWN ALGORITHMS

From now on, we assume the modulus h to be the m -th power of a polynomial $\tilde{h}$ of degree d . In this section we extend the two fast known algorithms for power series composition from Theorem 1 to composition modulo h . We directly state and prove the generalized algorithms. For more explanations we refer to the original papers [1, 3] and [11, Section 3.4.3].

3.1 Extension of Bernstein's algorithm

If $\mathbb{K}$ has positive characteristic $p > 0$, then Bernstein's algorithm [1] exploits the Frobenius map in order to recursively reduce the degrees of the polynomials to be composed by a factor p .

Algorithm 1

Input. $\tilde{h}, h \in \mathbb{K}[x]$, $f, g \in \mathbb{K}[x]_{<n}$, where $n = \deg h$.

Output. $f \circ g \bmod h$.

Assumptions. $\text{char } \mathbb{K} = p$, $h = \tilde{h}^m$, $\deg \tilde{h} = d = n/m$.

1. If $m = 1$, then return $f \circ g \bmod h$, computed by a general modular composition algorithm of cost $\mathbb{C}_{\mathbb{K}}(d)$.
2. Split f into polynomials $\tilde{f}_0, \dots, \tilde{f}_{p-1}$ such that $f(x) = \sum_{i=0}^{p-1} \tilde{f}_i(x^p) x^i$. Let $\tilde{m} = \lceil m/p \rceil$ and $\tilde{g}, \tilde{h}$ be such that $\tilde{h}(x^p) = \tilde{h}(x)^p$ and $\tilde{g}(x^p) = g(x)^p \bmod \tilde{h}(x)^{\tilde{m}}$.
3. Recursively compute $\tilde{u}_i = \tilde{f}_i \circ \tilde{g} \bmod \tilde{h}^{\tilde{m}}$ for all $0 \leq i \leq p-1$.
4. Return $\sum_{i=0}^{p-1} \tilde{u}_i(x^p) g(x)^i \bmod h(x)$.

PROPOSITION 5. *Algorithm 1 is correct and takes $p m \mathbb{C}_{\mathbb{K}}(d) + O\left(\frac{p}{\log p} \mathbb{M}_{\mathbb{K}}(n) \log m\right)$ operations in $\mathbb{K}$.*

PROOF. Since $(d m - 1)/p < d \tilde{m}$, the degrees of $\tilde{g}$ and the $\tilde{f}_i$ are at most $d \tilde{m} - 1$. Now $\tilde{u}_i = \tilde{f}_i \circ \tilde{g} \bmod \tilde{h}^{\tilde{m}}$ implies

$$\tilde{u}_i(x^p) = \tilde{f}_i(\tilde{g}(x^p)) \bmod \tilde{h}(x^p)^{\tilde{m}} = \tilde{f}_i(g(x)^p) \bmod \tilde{h}(x)^{p \tilde{m}},$$

which ensures the correctness of the algorithm.

In step 2 we need: $O(n \log p)$ operations in $\mathbb{K}$ to compute all the necessary p -th powers, $O(\mathbb{M}_{\mathbb{K}}(d \tilde{m}))$ for $\tilde{h}^{\tilde{m}}$, and $O(\mathbb{M}_{\mathbb{K}}(n))$ for the division to obtain $\tilde{g}$. In step 4, by [8, Exercise 9.16], the computation of $\tilde{u}_i(x^p) \bmod h(x)$ for $i = 0, \dots, p-1$ may be done in time

$$\begin{aligned} O\left(p d \tilde{m} + \left\lceil \frac{p \tilde{m}}{m} \right\rceil \mathbb{M}_{\mathbb{K}}(n)\right) \\ = O\left(d(m+p) + \left(1 + \frac{p}{m}\right) \mathbb{M}_{\mathbb{K}}(n)\right) = O(p \mathbb{M}_{\mathbb{K}}(n)). \end{aligned}$$

The remaining calculations of step 4 amount to $O(p \mathbb{M}_{\mathbb{K}}(n))$ operations in $\mathbb{K}$.

Let $T(m)$ be the time complexity for precision m and let $e \in \mathbb{N}$ be such that $p^{e-1} < m \leq p^e$. By what precedes, we have

$$T(m) \leq p T(\lceil m/p \rceil) + c p \mathbb{M}_{\mathbb{K}}(d m),$$

for all $m > 1$ and a sufficiently large constant c . Unrolling this inequality $e-1 = O(\log_p m)$ times, we obtain

$$T(m) \leq p^e T(\lceil m/p^e \rceil) + c p \sum_{i=0}^{e-1} p^i \mathbb{M}_{\mathbb{K}}(d \lceil m/p^i \rceil).$$

Using $p^i \lceil m/p^i \rceil \leq m + p^i - 1$ and the super-additivity of $\mathbb{M}_{\mathbb{K}}$, we conclude that

$$T(m) = p m \mathbb{C}_{\mathbb{K}}(d) + O(p \mathbb{M}_{\mathbb{K}}(n) \log_p m). \quad \square$$

3.2 Extension of Brent and Kung's algorithm

We now extend Brent and Kung's series composition algorithm to composition modulo $h = \tilde{h}^m$. This method uses Taylor series expansion in order to reduce the degrees of the polynomials to be composed. For a given polynomial u , we write $u^{(l)}$ for its l -th derivative. Algorithm 2 is a sub-algorithm that is used when the arguments of the composition both have small degree.

Algorithm 2

Input. $u, g, h \in \mathbb{K}[x]$, where u has degree t .

Output. $u \circ g \bmod h$.

1. If $t \leq 0$ then return $u(0)$.
2. Split $u(x) = u_0(x) + x^{\lceil t/2 \rceil} u_1(x)$ with $\deg u_0 < \lceil t/2 \rceil$ and $\deg u_1 = t - \lceil t/2 \rceil$.
3. Recursively compute $v_i = u_i \circ g \bmod h$ for $i \in \{0, 1\}$,
4. Return $(v_1 + g^{\lceil t/2 \rceil} v_2) \bmod h$.

Algorithm 3

Input. $\tilde{h}, h \in \mathbb{K}[x]$, $f, g \in \mathbb{K}[x]_{<n}$, where $n = \deg h$.

Output. $f \circ g \bmod h$.

Assumption. $\mathbb{K}$ has characteristic $p \geq k$, where $l = \lceil \sqrt{m/(d \log n)} \rceil$ and $k = \lceil m/l \rceil$; $h = \tilde{h}^m$, $\tilde{h}$ is separable and irreducible, $\deg \tilde{h} = d = n/m$.

1. Compute $g_1 = g \bmod \tilde{h}^l$ and $g_2 = g - g_1$.
2. If $g_1 \neq 0$ then compute the valuation v of g_1' in $\tilde{h}$ and the modular inverse $u = (g_1'/\tilde{h}^v)^{-1} \bmod \tilde{h}^{m-l}$.
3. Compute $r_0 = f \circ g_1 \bmod h$ using Algorithm 2.
4. If $g_1 = 0$ then compute $r_i = f^{(i)}(0)$ for all i from 1 to $k-1$. Otherwise, for i from 0 to $k-2$, compute $r_{i+1} = u(r_i' \text{ quo } \tilde{h}^v) \bmod \tilde{h}^{m-(i+1)l}$.
5. Return $\sum_{i=0}^{k-1} \frac{r_i}{i!} g_2^i \bmod \tilde{h}^m$.

PROPOSITION 6. *Algorithm 3 is correct and takes $O(\sqrt{n} \mathbb{M}_{\mathbb{K}}(n) \log^{1/2} n)$ operations in $\mathbb{K}$ if $d = O(m/\log n)$.*

PROOF. The main ingredient of the proof is the following Taylor expansion:

$$f(g_1(x) + g_2(x)) = \sum_{i=0}^{k-1} \frac{f^{(i)}(g_1(x))}{i!} g_2^i \bmod \tilde{h}^m,$$

which makes use of $\text{val}_{\tilde{h}} g_2 \geq l$ and $k l \geq m$. If $g_1 = 0$ then $r_i = f^{(i)}(g_1(x))$ holds for all i , so the algorithm is correct. Otherwise we may write $g_1 = a(x) \tilde{h}^w \bmod \tilde{h}^{w+1}$ with $0 \leq w < l$ and $a \neq 0$ of degree $< d$, which implies $g_1'(x) = w a(x) \tilde{h}'(x) \tilde{h}^{w-1} \bmod \tilde{h}^w$. Since $p > w$ and $\tilde{h}$ separable, it follows that $0 \leq v \leq l-2$, whence u is well defined.

Let us show that $r_i = f^{(i)}(g_1(x)) \bmod \tilde{h}^{m-il}$, by induction on i . This is clear for $i=0$. Assume that the identity holds for some $i \geq 0$. From $f^{(i+1)} \circ g_1 = (f^{(i)} \circ g_1)' / g_1'$ we see that the valuation of $(f^{(i)} \circ g_1)'$ in $\tilde{h}$ is necessarily $\geq v$ and that $u(r_i' \text{ quo } \tilde{h}^v) = (f^{(i)} \circ g_1)' \bmod \tilde{h}^{m-il-v-1}$. The induction hypothesis is thus satisfied for $i+1$ because $i l + v + 1 \leq (i+1) l$. We are done with the correctness.

The $\hbar$ -adic expansion of g takes $O(M_{\mathbb{K}}(n) \log n)$ operations and dominates the cost of steps 1 and 2. Steps 4 and 5 amount to $O(\frac{m}{l} M_{\mathbb{K}}(n)) = O(\sqrt{n} M_{\mathbb{K}}(n) \log^{1/2} n)$ operations. As to step 3, we enter Algorithm 2 with $u = f$.

Let $T(t)$ be the cost function of Algorithm 2. The degrees of $u_1 \circ g_1$, $u_2 \circ g_1$ and $g_1^{\lceil t/2 \rceil}$ are $\leq dl \lceil t/2 \rceil$, so that

$$T(t) \leq 2T\left(\left\lceil \frac{t}{2} \right\rceil\right) + cM_{\mathbb{K}}\left(\min\left(d\left\lceil \frac{t}{2} \right\rceil, n\right)\right),$$

for some constant c . Let e be the largest integer such that $dl \lceil n/2^e \rceil > n$. Unrolling $e - 1$ times the inequality for $T(t)$ and $t = n$, we obtain

$$\begin{aligned} T(n) &\leq 2^e T(\lceil n/2^e \rceil) + c \sum_{i=0}^{e-1} 2^i M_{\mathbb{K}}(n) \\ &= 2^e T(\lceil n/2^e \rceil) + O(2^e M_{\mathbb{K}}(n)). \end{aligned}$$

In a similar way we obtain $T(\lceil n/2^e \rceil) = O(M_{\mathbb{K}}(n) \log n)$. Consequently, $T(n) = O(2^e M_{\mathbb{K}}(n) \log n)$. The conclusion follows from $2^e = O(dl) = O(\sqrt{n/\log n})$. $\square$

COROLLARY 7. *Let $\hbar$ be of degree d in $\mathbb{K}[x]$, let $h = \hbar^m$ be of degree $n = md$, and let f, g be in $\mathbb{K}[x]_{<n}$. If $d = O(m/\log n)$, then the modular composition $f \circ g \bmod h$ may be computed using $m\sqrt{n} C_{\mathbb{K}}(d) + O(\sqrt{n} M_{\mathbb{K}}(n) \log n)$ operations in $\mathbb{K}$.*

PROOF. If $p = O(\sqrt{n})$, then we may use Algorithm 1, which requires $m\sqrt{n} C_{\mathbb{K}}(d) + O(\sqrt{n} M_{\mathbb{K}}(n) \log n)$ operations in $\mathbb{K}$. Otherwise, we use Algorithm 3. $\square$

Notice that for fixed d , this corollary gives a cost $n^{1.5+o(1)}$, independently of the characteristic.

4. REDUCTION TO SERIES COMPOSITION

Recall that the modulus h is the m -th power of a polynomial $\hbar$ of degree d . From now on, we assume that $\hbar$ is separable. So far, our algorithms for composition modulo h relied on extensions of known algorithms dedicated to power series. In this section, we investigate the other natural approach: we design algorithms that reduce one composition modulo $\hbar^m$ to one series composition of order m over a suitable extension of the ground field. We set $\mathbb{L} = \mathbb{K}[z]/(\hbar(z))$ and write α for the residue class of z in $\mathbb{L}$. The associated trace function is denoted by $\text{Tr}_{\mathbb{L}/\mathbb{K}}$. This function naturally extends to a map $\text{Tr}_{\mathbb{L}/\mathbb{K}}: \mathbb{L}[x] \rightarrow \mathbb{K}[x]$ in a coefficientwise fashion.

4.1 Case of small moduli

The first method we study is rather elementary. It works well for any characteristic, but it is only efficient when d is much smaller than n . We compute $(f \circ g)(y) \bmod (y - \zeta)^m$ for all roots ζ of $\hbar$ and recover $f \circ g \bmod \hbar^m$ by Chinese remaindering. The simultaneous computation with all roots of $\hbar$ is emulated *via* standard algebraic techniques.

Since $\hbar$ is not assumed to be irreducible, $\mathbb{L}$ is not necessarily a field. We will compute inverses in $\mathbb{L}$ using the technique of *dynamic evaluation*, which is briefly recalled in Appendix A.2. For convenience of the reader, we use the auxiliary variable y for power series and keep the variable x for polynomials.

Algorithm 4

Input. $\hbar, h \in \mathbb{K}[x]$, $f, g \in \mathbb{K}[x]_{<n}$, where $n = \deg h$.

Output. $f \circ g \bmod h$.

Assumptions. $h = \hbar^m$, $\hbar$ is separable, $\deg \hbar = d = n/m$.

1. Compute $G(y) = g(y + \alpha) - g(\alpha)$ in $\mathbb{L}[y]$.
2. Compute $F(y) = f(y + g(\alpha))$ in $\mathbb{L}[y]$.
3. Perform the power series composition $R(y) = (F \circ G)(y)$ in $\mathbb{L}[y]/(y^m)$.
4. Compute $H(x) = h(x)/(x - \alpha)^m$ in $\mathbb{L}[x]$.
5. With the algorithm underlying Corollary 20, compute the modular inverse $I(x) = H^{-1}(x) \bmod (x - \alpha)^m$.
6. Compute $Q(x) = H(x)I(x)$ and then $P(x) = R(x - \alpha)Q(x)$.
7. Return $\text{Tr}_{\mathbb{L}/\mathbb{K}}(P(x)) \bmod h(x)$.

PROPOSITION 8. *Algorithm 4 is correct and takes*

$$O(S_{\mathbb{L}/\mathbb{K}}(m) + M_{\mathbb{K}}(d) M_{\mathbb{K}}(n) \log d \log n)$$

operations in $\mathbb{K}$.

PROOF. For a root ζ of $\hbar$ in some algebraic closure $\bar{\mathbb{K}}$ of $\mathbb{K}$, we write σ_{ζ} for the homomorphism from $\mathbb{L}$ to $\bar{\mathbb{K}}$ that sends α to ζ . This homomorphism extends coefficientwise into a map $\sigma_{\zeta}: \mathbb{L}[x] \rightarrow \bar{\mathbb{K}}[x]$. One verifies that $\text{Tr}_{\mathbb{L}/\mathbb{K}}(P(x))$ equals

$$\begin{aligned} &\sum_{\hbar(\zeta)=0} \sigma_{\zeta}(R(x - \alpha)) \sigma_{\zeta}(Q(x)) \\ &= \sum_{\hbar(\zeta)=0} \left\{ ((f \circ g)(x) \bmod (x - \zeta)^m) \frac{h(x)}{(x - \zeta)^m} \right. \\ &\quad \left. \times \left(\left(\frac{h(x)}{(x - \zeta)^m} \right)^{-1} \bmod (x - \zeta)^m \right) \right\}, \end{aligned}$$

so $\text{Tr}_{\mathbb{L}/\mathbb{K}}(P(x)) \bmod (x - \zeta)^m$ equals $(f \circ g)(x) \bmod (x - \zeta)^m$ for all roots ζ of $\hbar$. This proves the correctness of the algorithm thanks to the Chinese remainder theorem.

For the complexity analysis, step 3 amounts to $S_{\mathbb{L}/\mathbb{K}}(m)$ operations in $\mathbb{K}$. In step 5, we use Corollary 20 to compute $I(x)$ with $O(M_{\mathbb{K}}(d) M_{\mathbb{K}}(n) \log d \log n)$ operations in $\mathbb{K}$. Steps 1, 2, 4, and 6 require $O(M_{\mathbb{L}}(n) \log n)$ additional ring operations in $\mathbb{L}$. In step 7, the row matrix of $\text{Tr}_{\mathbb{L}/\mathbb{K}}$ may be computed with $O(M_{\mathbb{K}}(d))$ operations in $\mathbb{K}$ by formula (20), so each trace takes $O(d)$ operations in $\mathbb{K}$. Therefore step 7 amounts to $O(M_{\mathbb{K}}(d) + M_{\mathbb{K}}(n) + dn)$ operations in $\mathbb{K}$, which is negligible. $\square$

When d is much smaller than m , we may directly benefit from fast power series composition. But overall this does not improve much on Corollary 7. In the rest of this section, we aim at decreasing the dependency in d in complexity bounds.

4.2 Fast reduction to series composition

In order to improve on Algorithm 4, we shall use a better way to expand $g(y + \alpha)$ and $f(y + g(\alpha))$ modulo y^m . For this purpose we introduce the $\mathbb{K}$ -algebra homomorphism

$$\begin{aligned} \pi_{\hbar, m}: \mathbb{K}[x]/(\hbar(x)^m) &\rightarrow \mathbb{L}[y]/(y^m) \\ u(x) &\mapsto u(y + \alpha). \end{aligned}$$

We regard $\pi_{h,m}$ and $\pi_{h,m}^{-1}$ as conversions that we respectively call the *untangling* and *tangling* mappings.

LEMMA 9. For all separable polynomial h and all $m \geq 1$, the map $\pi_{h,m}$ is a $\mathbb{K}$ -algebra isomorphism.

PROOF. $\pi_{h,m}$ is clearly a homomorphism between $\mathbb{K}$ -algebras of equal dimensions, so we need to prove its injectivity when h is separable. We consider the h -adic expansion $u_0(x) + u_1(x)h(x) + \dots + u_{m-1}(x)h(x)^{m-1}$ of $u(x)$ modulo $h(x)^m$, where all the u_i have degrees $< \deg h$, and we assume that $\pi_{h,m}(u) = 0$. It is immediate that $u_0 = 0$. We may thus suppose, by induction on $i \geq 1$, that $u_j = 0$ for all $0 \leq j \leq i-1$. From $h(y + \alpha) = h'(\alpha)y + O(y^2)$ we obtain $u(y + \alpha) = u_i(\alpha)h'(\alpha)^i y^i + O(y^{i+1})$. Since $h'(\alpha)$ is invertible, it follows that $u_i(\alpha) = 0$, whence $u_i = 0$. This completes the induction. $\square$

Remark. The separability assumption on h is necessary: if $h(x) = x^2$ and $m = 2$, then we have $\pi_{h,m}(x^3) = 0$.

Algorithm 5

Input. $h, h \in \mathbb{K}[x]$, $f, g \in \mathbb{K}[x]_{<n}$, where $n = \deg h$.

Output. $f \circ g \text{ rem } h$.

Assumptions. $h = h^m$, h is separable, $\deg h = d = n/m$.

1. Compute $G = \pi_{h,m}(g) - g(\alpha)$.
2. Let $\bar{g} = g \text{ rem } h$ and compute the characteristic polynomial $\bar{\chi}$ of $\bar{g}$ modulo h .
3. Compute $\bar{F} = \pi_{\bar{\chi},m}(f)$.
4. Write $\bar{F}(y) = \bar{F}_0(\beta) + \bar{F}_1(\beta)y + \dots + \bar{F}_{m-1}(\beta)y^{m-1}$, where β is the class of z in $\mathbb{K}[z]/(\bar{\chi}(z))$, and the $\bar{F}_i$ have degrees $< d$. For each $i \in \{0, \dots, m-1\}$, compute $F_i = \bar{F}_i \circ \bar{g} \text{ rem } h$ and let $F(y) = F_0(\alpha) + F_1(\alpha)y + \dots + F_{m-1}(\alpha)y^{m-1} \in \mathbb{L}[y]$.
5. Perform the power series composition $R(y) = (F \circ G)(y)$ in $\mathbb{L}[y]/(y^m)$.
6. Return $\pi_{h,m}^{-1}(R)$.

From now on, in cost analyses, the expression $\Pi_{\mathbb{K}}(d, m)$ (resp. $\Pi_{\mathbb{K}}^*(d, m)$) represents a cost function for computing the *direct image* (resp. the *preimage*) of $\pi_{h,m}$. If good algorithms are known for $\pi_{h,m}$ and $\pi_{h,m}^{-1}$, for characteristic polynomials and for compositions modulo h , then we observe that Algorithm 5 boils down to one power series composition at precision m over $\mathbb{L}$.

PROPOSITION 10. Algorithm 5 is correct and it requires

$$H_{\mathbb{K}}(dm) = 2\Pi_{\mathbb{K}}(d, m) + \Pi_{\mathbb{K}}^*(d, m) + Q_{\mathbb{K}}(d) + m C_{\mathbb{K}}(d) + S_{\mathbb{L}/\mathbb{K}}(m) + O(dm)$$

operations in $\mathbb{K}$.

PROOF. The homomorphism $\mathbb{K}[z]/(\bar{\chi}(z)) \rightarrow \mathbb{L}$ that sends β to $g(\alpha)$ is well defined, so applying it coefficient-wise to $\bar{F}(y) = f(y + \beta) \in \mathbb{K}[z]/(\bar{\chi}(z))[y]$ we obtain $f(y + g(\alpha)) \bmod y^m = F(y)$. Therefore we have $(F \circ G)(y) = (f \circ g)(y + \alpha) \bmod y^m = \pi_{h,m}(f \circ g \text{ rem } h)$, which ensures the correctness of the algorithm. The cost directly follows from the definitions. $\square$

The rest of this section is devoted to algorithms with softly linear costs for computing $\pi_{h,m}$ and its inverse. More precisely, we will prove that we may take $\Pi_{\mathbb{K}}(d, m)$ and $\Pi_{\mathbb{K}}^*(d, m)$ to be $O(M_{\mathbb{K}}(d, m) \log^2 m + M_{\mathbb{K}}(d) \log d)$ in the latter proposition. This leads to our main theorem:

THEOREM 11. Let $h, h = h^m \in \mathbb{K}[x]$ and $f, g \in \mathbb{K}[x]_{<n}$, where $n = \deg h$, such that h is separable of degree $d = n/m$. Then, we can compute $f \circ g \text{ rem } h$ using

$$H_{\mathbb{K}}(d, m) = S_{\mathbb{L}/\mathbb{K}}(m) + m C_{\mathbb{K}}(d) + Q_{\mathbb{K}}(d) + O(M_{\mathbb{K}}(d, m) \log^2 m + M_{\mathbb{K}}(d) \log d)$$

operations in $\mathbb{K}$.

4.3 Untangling in characteristic zero

The following lemma is elementary, but we shall refer to it several times.

LEMMA 12. Given integers $0 \leq k < l$ and a polynomial $v \in \mathbb{K}[x]$, we have $(v h^l)^{(k)}(\alpha) = 0$.

PROOF. We have $(v h^l)^{(k)} = \sum_{i=0}^k \binom{k}{i} v^{(i)} (h^l)^{(k-i)}$ by Leibnitz formula, showing that $(v h^l)^{(k)}$ is divisible by h . $\square$

We begin with the easiest situation when the characteristic p of $\mathbb{K}$ is zero or sufficiently large and achieve softly linear time for $\Pi(d, m)$.

Algorithm 6

Input. $h, h \in \mathbb{K}[x]$, $u \in \mathbb{K}[x]_{<n}$.

Output. $u(\alpha), u'(\alpha), \dots, u^{(m-1)}(\alpha)$.

Assumptions. $h = h^m$, h is separable, $\deg h = d = n/m$.

1. If $m = 1$ then return $u(\alpha)$, otherwise let $l = \lfloor m/2 \rfloor$.
2. Recursively apply the algorithm to $u \text{ rem } h^l$, giving $u(\alpha), u'(\alpha), \dots, u^{(l-1)}(\alpha)$.
3. Recursively apply the algorithm to $u^{(l)} \text{ rem } h^{m-l}$, giving $u^{(l)}(\alpha), u^{(l+1)}(\alpha), \dots, u^{(m-1)}(\alpha)$.
4. Return $u(\alpha), u'(\alpha), \dots, u^{(m-1)}(\alpha)$.

PROPOSITION 13. Algorithm 6 is correct and takes $O(M_{\mathbb{K}}(n) \log m)$ operations in $\mathbb{K}$. If $p = 0$ or $p \geq m$, then, given $u \in \mathbb{K}[x]_{<n}$, one may compute $\pi_{h,m}(u)$ using $O(M_{\mathbb{K}}(n) \log m)$ operations in $\mathbb{K}$.

PROOF. The correctness follows from Lemma 12 which ensures that $u^{(k)}(\alpha) = (u \text{ rem } h^l)^{(k)}(\alpha)$ for all $0 \leq k < l$ and $u^{(l+k)}(\alpha) = (u^{(l)} \text{ rem } h^{m-l})^{(k)}(\alpha)$ for all $0 \leq k < m-l$. The cost analysis is straightforward. Finally, to deduce $\pi_{h,m}(u)$, thanks to the assumption on the characteristic, we may simply use the Taylor expansion $u(y + \alpha) = \sum_{i=0}^{m-1} u^{(i)}(\alpha) \frac{y^i}{i!} + O(y^m)$. $\square$

4.4 Untangling in positive characteristic

In order to handle the case $p < m$, we need to pay attention to vanishing coefficients when computing Taylor expansions using higher order derivatives. For this, we rely on the following adaptation of the Taylor expansion of u :

$$u(y + \alpha) = \sum_{i \geq 0} D^{(i)}(u)(\alpha) y^i.$$

The polynomial $D^{(l)}(u)$ is called the l -th order *Hasse derivative* of u . The operator $D^{(l)}$ is linear and we may compute it as follows: if $u = x^i$ is a monomial and $l > i$, then $D^{(l)}(x^i) = 0$; otherwise

$$D^{(l)}(x^i) = \binom{i}{l} x^{i-l}.$$

LEMMA 14. For all integers k and l , we have $D^{(k)} \circ D^{(l)} = \binom{k+l}{k} D^{(k+l)}$.

PROOF. Let $i \geq k + l$. A straightforward calculation in characteristic zero yields

$$\begin{aligned} D^{(k)} \circ D^{(l)}(x^i) &= \binom{i-l}{k} \binom{i}{l} x^{i-l-k} \\ &= \binom{i}{k+l}^{-1} \binom{i-l}{k} \binom{i}{l} D^{(k+l)}(x^i) \\ &= \binom{k+l}{k} D^{(k+l)}(x^i). \end{aligned}$$

Now if the characteristic $p > 0$ divides the integer $\binom{i}{k+l}$, then it also divides $\binom{i-l}{k} \binom{i}{l}$, and the lemma remains correct. $\square$

It is convenient to also introduce the *Pochhammer symbol*

$$(i)_j = i(i-1) \cdots (i-j+1).$$

The next lemma expresses conditions to compute Hasse derivatives from ones of lower orders, in positive characteristic.

LEMMA 15. Let $l \geq p$ be an integer power of $p > 0$, let s be an integer such that $\text{val}_p s > \text{val}_p l$, and let $a < b$ be two integers in $\{0, \dots, p-1\}$. Then, for any polynomial $u \in \mathbb{K}[x]$, we may compute $D^{(s+bl)}(u)$ from $D^{(s+al)}(u) = \sum_{i \geq 0} v_i x^i$ using the formula

$$D^{(s+bl)}(u) = \sum_{c \geq b-a} \frac{(c)_{b-a}}{(b)_{b-a}} \sum_{cl \leq i < (c+1)l} v_i x^{i-(b-a)l}. \quad (1)$$

PROOF. The previous lemma implies that $D^{(s+bl)} = \binom{s+bl}{(b-a)l}^{-1} D^{((b-a)l)} \circ D^{(s+al)}$. Now in characteristic zero, we have

$$\begin{aligned} &\binom{s+bl}{(b-a)l}^{-1} D^{((b-a)l)}(x^i) \\ &= \binom{s+bl}{(b-a)l}^{-1} \binom{i}{(b-a)l} x^{i-(b-a)l} \\ &= \frac{(i)_{(b-a)l}}{(s+bl)_{(b-a)l}} x^{i-(b-a)l} \end{aligned}$$

for all $i \geq (b-a)l$. In positive characteristic, we claim that the ratios $\frac{(i)_{(b-a)l}}{(s+bl)_{(b-a)l}}$ have non-negative valuations in p .

Let us first consider the case when $b = a + 1$. If $i = cl$, with $c \geq 1$, then $\text{val}_p i \geq \text{val}_p l = \text{val}_p(s+bl)$. If $0 < j < l$, then $\text{val}_p(i-j) = \text{val}_p j = \text{val}_p(s+bl-j)$. This shows that $(i)_l / (s+bl)_l = c/b \pmod p$ is well defined and may be computed in $\mathbb{K}$ whenever i is a multiple of l .

At a second stage, we make the induction hypothesis that $(i)_l / (s+bl)_l$ has been computed in $\mathbb{K}$ for $i \geq cl$ such that $i+1 < (c+1)l$. The value for $i+1$ may then be obtained via

$$\frac{(i+1)_l}{(s+bl)_l} = \frac{i+1}{i+1-l} \frac{(i)_l}{(s+bl)_l} = \frac{c}{b} \pmod p,$$

since $\text{val}_p(i+1) \geq \text{val}_p(i+1-l)$. This deals with the case when $b = a + 1$.

Finally, for any $b > a$, the ratio $(i)_{(b-a)l} / (s+bl)_{(b-a)l}$ is also well defined since it rewrites into

$$\begin{aligned} &\frac{(i)_{(b-a)l}}{(s+bl)_{(b-a)l}} \\ &= \frac{(i)_l}{(s+bl)_l} \frac{(i-l)_l}{(s+(b-1)l)_l} \cdots \frac{(i-(b-a-1)l)_l}{(s+(a+1)l)_l} \\ &= \frac{c(c-1) \cdots (c-(b-a-1))}{b(b-1) \cdots (a+1)} \pmod p, \end{aligned}$$

which concludes the proof. $\square$

We are now in a position to adapt the “divide and conquer” Algorithm 6 to the case when $\mathbb{K}$ has small positive characteristic p .

Algorithm 7

Input. $\tilde{h} \in \mathbb{K}[x]$ of degree d ; integers s, l, a, b ; $D^{(s+al)}(u) \text{ rem } \tilde{h}^{(b-a)l}$.

Output. $D^{(s+kl)}(u) \text{ rem } \tilde{h}^l$, for all $k \in \{a, \dots, b-1\}$.

Assumption. $\mathbb{K}$ has positive characteristic p ; l is a power of p , $1 \leq \text{val}_p l < \text{val}_p s$; $0 \leq a < b \leq p$; $u \in \mathbb{K}[x]_{<n}$.

1. If $b = a + 1$, then return $D^{(s+al)}(u) \text{ rem } \tilde{h}^l$.
2. Let $c = \lceil (a+b)/2 \rceil$.
3. Compute $D^{(s+al)}(u) \text{ rem } \tilde{h}^{(c-a)l}$.
4. By using formula (1), compute $D^{(s+cl)}(u) \text{ rem } \tilde{h}^{(b-c)l}$ from $D^{(s+al)}(u) \text{ rem } \tilde{h}^{(c-a)l}$.
5. Recursively apply the algorithm to $u, \tilde{h}, s, l, a, c$, and $D^{(s+al)}(u) \text{ rem } \tilde{h}^{(c-a)l}$. This yields $D^{(s+kl)}(u) \text{ rem } \tilde{h}^l$ for all $k \in \{a, \dots, c-1\}$.
6. Recursively apply the algorithm to $u, \tilde{h}, s, l, c, b$, and $D^{(s+cl)}(u) \text{ rem } \tilde{h}^{(b-c)l}$. This yields $D^{(s+kl)}(u) \text{ rem } \tilde{h}^l$, for all $k \in \{c, \dots, b-1\}$.
7. Return $D^{(s+kl)}(u) \text{ rem } \tilde{h}^l$, for all $k \in \{a, \dots, b-1\}$.

PROPOSITION 16. Algorithm 7 is correct and takes $O(M_{\mathbb{K}}((b-a)dl \log(b-a)))$ operations in $\mathbb{K}$.

PROOF. The correctness follows from Lemmas 12 and 15; the cost analysis is straightforward. $\square$

Algorithm 8

Input. $\tilde{h}, h \in \mathbb{K}[x]$ with $d = \deg \tilde{h}, \deg h = n = dm$; an integer s ; $D^{(s)}(u) \text{ rem } \tilde{h}^m$.

Output. $D^{(s)}(u)(\alpha), \dots, D^{(s+m-1)}(u)(\alpha)$.

Assumptions. $\mathbb{K}$ has positive characteristic p ; $h = \tilde{h}^m$; $p^{\text{val}_p s} \geq m$; $u \in \mathbb{K}[x]_{<n}$.

1. If $m \leq p$, then use Algorithm 6 to compute $\pi_{h,m}(D^{(s)}(u)) = \sum_{i=0}^{m-1} (D^{(i)} \circ D^{(s)})(u)(\alpha) y^i + O(y^m)$ and return $D^{(s)}(u)(\alpha), \dots, (D^{(m-1)} \circ D^{(s)})(u)(\alpha)$.
2. Let e be minimal with $p^e \geq m$, let $l = p^{e-1}$, and $k = \lfloor m/l \rfloor$.
3. Compute $D^{(s+il)}(u) \text{ rem } \tilde{h}^l$ for all $i \in \{0, \dots, k-1\}$ using Algorithm 7 (called with $s, l, a=0$ and $b=k$).
4. If $kl < m$, then compute $D^{(s+kl)}(u) \text{ rem } \tilde{h}^{m-kl}$ from $D^{(s)}(u) \text{ rem } \tilde{h}^m$ with formula (1).
5. For all $i \in \{0, \dots, k-1\}$, use the algorithm recursively with $D^{(s+il)}(u) \text{ rem } \tilde{h}^l$ in order to obtain $D^{(s+il)}(u)(\alpha), \dots, D^{(s+(i+1)l-1)}(u)(\alpha)$.
6. If $kl < m$ then use the algorithm recursively with $D^{(s+kl)}(u) \text{ rem } \tilde{h}^{m-kl}$ in order to obtain $D^{(s+kl)}(u)(\alpha), \dots, D^{(s+m-1)}(u)(\alpha)$.
7. Return $D^{(s)}(u)(\alpha), \dots, D^{(s+m-1)}(u)(\alpha)$.

PROPOSITION 17. Algorithm 8 is correct and takes

$$\Pi_{\mathbb{K}}(d, m) = O(M_{\mathbb{K}}(dm) \log m)$$

operations in $\mathbb{K}$.

PROOF. Let us first examine the case when $m \leq p$. We need to show that $(D^{(j)} \circ D^{(s)})(u)(\alpha)$ coincides with $D^{(s+j)}(u)(\alpha)$ for all $j \in \{0, \dots, m-1\}$. From Lemma 14, we have $D^{(s+j)} = \binom{s+j}{j}^{-1} D^{(j)} \circ D^{(s)}$. Since $\text{val}_p s \geq 1$, it follows that $\binom{s+j}{j} \equiv 1 \pmod p$.

Now we suppose that $m > p$. We have $e \geq 2$, so l is a power of p , we have $k \leq p$, and the assumption $p^{\text{val}_p s} \geq m$ is preserved through recursive calls. The conditions of Algorithm 7 are satisfied, whence the correctness of step 3 by Proposition 16. Step 4 is a consequence of Lemmas 12 and 15. We are done with the correctness.

Step 1 costs $O(M_{\mathbb{K}}(dm) \log m)$ by Proposition 13. Step 3 costs $O(M_{\mathbb{K}}(k d l) \log p) = O(M_{\mathbb{K}}(d m) \log p)$ by Proposition 16 and step 4 takes $O(M_{\mathbb{K}}(dm))$ operations in $\mathbb{K}$. The total cost function $T(m)$ of the algorithm satisfies

$$T(m) \leq k T(l) + T(m - kl) + c M_{\mathbb{K}}(dm) \log p,$$

where c is a sufficiently large constant and $T(m) \leq c M_{\mathbb{K}}(d m) \log m$ when $m \leq p$. This leads to $T(m) = O(M_{\mathbb{K}}(dm) \log m + M_{\mathbb{K}}(d m) \log p \log_p m)$. $\square$

4.5 Tangling

The next algorithm is independent of the characteristic. It relies on the untangling algorithms from the previous subsections, by using the “divide and conquer” strategy. For a polynomial or series f in x and $a \leq b$, we denote $f_{a;b} = \sum_{a \leq i < b} f_i x^{i-a}$ and $f_{;b} = f_{0;b}$. For convenience, we assume that $\Pi_{\mathbb{K}}(d, m)/m$ is a nondecreasing function of m .

Algorithm 9

Input. $\tilde{h} \in \mathbb{K}[x]$ of degree d ; an integer $m \in 2^{\mathbb{N}}$; $\sigma \in \mathbb{L}[y] / (y^m)$, where $\mathbb{L} = \mathbb{K}[z] / (\tilde{h}(z))$; $(w(y)^{2^i})_{;2^i}$ for $1 \leq 2^i \leq m/2$ where $w(y) = (y/\tilde{h}(y + \alpha))_{;m}$.

Output. $\pi_{\tilde{h},m}^{-1}(\sigma)$.

1. If $m=1$, then let $u(z)$ be the preimage of $\sigma(0)$ in $\mathbb{K}[z]$ and return $u(x)$.
2. Let $l = m/2$ and compute $v_0(x) = \pi_{\tilde{h},l}^{-1}(\sigma_{;l})$, $v_1(x) = (\sigma - \pi_{\tilde{h},m}(v_0(x)))_{;l;m}$.
3. Return $v_0(x) + \tilde{h}(x)^l \pi_{\tilde{h},m-l}^{-1}((w(y)^l v_1(y))_{;m-l})$.

LEMMA 18. *Algorithm 9 is correct and takes time $\Pi_{\mathbb{K}}(d, m) \log m + M_{\mathbb{K}}(d m) \log m + M_{\mathbb{K}}(d) \log d$. In particular, for any m , we have*

$$\Pi_{\mathbb{K}}^*(d, m) = O(M_{\mathbb{K}}(dm) \log^2 m + M_{\mathbb{K}}(d) \log d).$$

PROOF. The computation is clear when $m=1$. Otherwise we verify that

$$\begin{aligned} & \pi_{\tilde{h},m}(v_0(x) + \tilde{h}(x)^l \pi_{\tilde{h},m-l}^{-1}((w(y)^l v_1(y))_{;m-l})) \\ &= \pi_{\tilde{h},m}(v_0(x)) + \pi_{\tilde{h},m}(\tilde{h}(x)^l (w(y)^l v_1(y))_{;m-l}) + O(y^m) \\ &= \pi_{\tilde{h},m}(v_0(x)) + \tilde{h}(y + \alpha)^l (w(y)^l v_1(y))_{;m-l} + O(y^m) \\ &= \pi_{\tilde{h},m}(v_0(x)) + y^l v_1(y) + O(y^m) \\ &= \sigma_{0;l}(y) + y^l \pi_{\tilde{h},m}(v_0(x))_{;l;m} + y^l (\sigma - \pi_{\tilde{h},m}(v_0(x)))_{;l;m} \\ &= \sigma(y). \end{aligned}$$

Kronecker substitution [8, Chapter 8, Section 4] allows us to multiply two series in $\mathbb{L}[[y]]$ at precision $O(y^m)$ using $O(M_{\mathbb{K}}(dm))$ operations in $\mathbb{K}$.

The computation of $w(y)$ requires one inversion in $\mathbb{L}$ that takes $O(M_{\mathbb{K}}(d) \log d)$ operations in $\mathbb{K}$, together with $O(M_{\mathbb{K}}(d m))$ operations for the series inversion. The other $(w(y)^{2^i})_{;2^i}$ amount to $O(M_{\mathbb{K}}(d m))$ more operations in $\mathbb{K}$. The cost function $T(m)$ satisfies

$$T(m) \leq 2 T(m/2) + \Pi_{\mathbb{K}}(d, m) + c M_{\mathbb{K}}(d m),$$

for some constant c , which leads to $T(m) = O((\Pi_{\mathbb{K}}(d, m) + M_{\mathbb{K}}(d m)) \log m + M_{\mathbb{K}}(d) \log d)$. The final bound of the lemma follows from Proposition 17. $\square$

APPENDIX A

A.1 Trace and Newton–Girard identities

Let h be a monic polynomial in $\mathbb{A}[x]$ of degree n , let φ be a linear form on $\mathbb{A}[x]_{<n}$, and let $g \in \mathbb{A}[x]_{<n}$. We study the *modular power projection* problem, which corresponds to computing $\varphi(1), \varphi(g), \dots, \varphi(g^{n-1} \text{rem } h)$.

Let $\mathbb{B}$ represent $\mathbb{A}[x] / (h(x))$. If we take $\varphi(u) = \text{Tr}_{\mathbb{B}/\mathbb{A}}(g u)$, where $\text{Tr}_{\mathbb{B}/\mathbb{A}}$ denotes the usual trace function of $\mathbb{B}$ over $\mathbb{A}$, then the latter power projections write as

$$\text{Tr}_{\mathbb{B}/\mathbb{A}}(g), \text{Tr}_{\mathbb{B}/\mathbb{A}}(g^2 \text{rem } h), \dots, \text{Tr}_{\mathbb{B}/\mathbb{A}}(g^n \text{rem } h).$$

It is well known that these traces are related to the characteristic polynomial of g modulo h by the *Newton–Girard identities*:

$$\begin{aligned} -\frac{\tilde{\chi}'(x)}{\tilde{\chi}(x)} &= \text{Tr}_{\mathbb{B}/\mathbb{A}}(g) + \text{Tr}_{\mathbb{B}/\mathbb{A}}(g^2 \text{rem } h) x + \dots \\ &\quad + \text{Tr}_{\mathbb{B}/\mathbb{A}}(g^n \text{rem } h) x^{n-1} + O(x^n), \end{aligned}$$

where $\tilde{\chi}(x) = x^n \chi(1/x)$ is the reverse polynomial of χ . If there exist given inverses of $2, 3, \dots, n$ in $\mathbb{A}$, then it is possible to integrate the latter differential equation in χ with $O(M_{\mathbb{A}}(n))$ operations in $\mathbb{A}$, which directly leads to an algorithm to compute χ from the power sums of the roots of χ (see for instance [9, Section 2] for details, which also cover the case when $\mathbb{A}$ is a finite field).

First, let us explain how to compute φ efficiently. Recall that $\tilde{h}(x) = x^n h(1/x)$. In the basis $1, x, \dots, x^{n-1}$, the linear form $\text{Tr}_{\mathbb{B}/\mathbb{A}}$ may be computed as

$$\begin{aligned} -\frac{\tilde{h}'(x)}{\tilde{h}(x)} &= \text{Tr}_{\mathbb{B}/\mathbb{A}}(x) + \text{Tr}_{\mathbb{B}/\mathbb{A}}(x^2) x + \dots \\ &\quad + \text{Tr}_{\mathbb{B}/\mathbb{A}}(x^n) x^{n-1} + O(x^n), \end{aligned} \tag{2}$$

which uses $O(M_{\mathbb{A}}(n))$ operations in $\mathbb{A}$. If γ represents the multiplication endomorphism by g in $\mathbb{B}$, then we can write $\varphi = \text{Tr}_{\mathbb{B}/\mathbb{A}} \circ \gamma$ and therefore $\varphi^\top = \gamma^\top \circ \text{Tr}_{\mathbb{B}/\mathbb{A}}^\top$. By transposing the modular multiplication by g with the techniques from [2, Sections 4 and 5], the computation of $\varphi^\top$ requires $O(M_{\mathbb{A}}(n))$ operations in $\mathbb{A}$.

A.2 Modular inversions

Let $\mathbb{K}$ be a field. Often in computer algebra, we are led to compute in monogen algebras such as $\mathbb{L} = \mathbb{K}[z] / (\lambda(z))$, where λ is separable of degree d , but not necessarily irreducible. Ring operations in $\mathbb{L}$ require $O(M_{\mathbb{K}}(d))$ operations in $\mathbb{K}$. Algorithms such as the determinant are often slower over rings than their counterparts over fields. For this reason, we often like to conduct our computations as if $\mathbb{L}$ were a field. From the programming point of view, this is easy to achieve: whenever we need to invert an element $a \in \mathbb{L}$

or test whether $a=0$, we factor $\lambda = \lambda_0 \lambda_1$, where $a=0 \bmod \lambda_0$ and a is invertible modulo λ_1 . We then restart the relevant part of the computations with λ_0 and/or λ_1 in the role of λ . This approach is known as the *dynamic evaluation principle* in computer algebra [6, 7]. We borrow the following complexity result for extended gcds from Dahan *et al.* [6]:

PROPOSITION 19. *Let λ be a separable polynomial of degree d over $\mathbb{K}$ and let A and B be univariate polynomials over $\mathbb{L} = \mathbb{K}[z] / (\lambda(z))$ of degrees $\leq n$. Using $O(M_{\mathbb{K}}(d) M_{\mathbb{K}}(n) \log d \log n)$ operations in $\mathbb{K}$, one can compute a factorization $\lambda_1 \cdots \lambda_s$ of λ and triples of polynomials (G_i, U_i, V_i) respectively over $\mathbb{K}[z] / (\lambda_i(z))$, such that $\deg \lambda_i \geq 1$, G_i is monic, G_i generates the extension of the ideal (A, B) to $\mathbb{K}[z] / (\lambda_i(z))$, and the Bézout relation $G_i = A_i U_i + B_i V_i$ holds with $\deg U_i < \deg B_i$ and $\deg V_i < \deg A_i$, for all $i \in \{1, \dots, s\}$.*

PROOF. This statement rephrases [6, Propositions 2.4 and 4.1], where we replace the assumption “ λ is square-free over a perfect field” by “ λ is separable”. In fact all arguments of [6] apply in this setting *mutatis mutandis*. An alternative point of view is to apply [6, Propositions 2.4 and 4.1] over the algebraic closure of $\mathbb{K}$, while noticing that all actual computations are really done over $\mathbb{K}$. $\square$

COROLLARY 20. *Let λ be a separable polynomial of degree d over $\mathbb{K}$ and let A and B be univariate polynomials over $\mathbb{L} = \mathbb{K}[z] / (\lambda(z))$ of degrees $\leq n$, such that B is monic and A is invertible modulo B . Then the inverse of A modulo B may be computed using $O(M_{\mathbb{K}}(d) M_{\mathbb{K}}(n) \log d \log n)$ operations in $\mathbb{K}$.*

PROOF. The triples (G_i, U_i, V_i) of the previous proposition satisfy $G_i = 1 = A U_i \bmod B$ over $\mathbb{K}[z] / (\lambda_i(z))$. We recover the inverse U of A modulo B from the U_i using the Chinese remainder theorem. $\square$

5. REFERENCES

- [1] D. J. Bernstein. Composing power series over a finite ring in essentially linear time. *J. Symbolic Comput.*, 26(3):339–341, 1998.
- [2] A. Bostan, G. Lecerf, and É. Schost. Tellegen’s principle into practice. In Hoon Hong, editor, *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation*, ISSAC ’03, pages 37–44, New York, NY, USA, 2003. ACM.
- [3] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. *J. ACM*, 25(4):581–595, 1978.
- [4] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 1997.
- [5] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Infor.*, 28(7):693–701, 1991.
- [6] X. Dahan, M. Moreno Maza, É. Schost, and Yuzhen Xie. On the complexity of the D5 principle. In J.-G. Dumas, editor, *Proceedings of Transgressive Computing 2006: a conference in honor of Jean Della Dora*, pages 149–168. U. J. Fourier, Grenoble, France, 2006.
- [7] J. Della Dora, C. Dicrescenzo, and D. Duval. A new method for computing in algebraic number fields. In G. Goos and J. Hartmanis, editors, *Eurocal’85 (2)*, volume 174 of *Lect. Notes in Comp. Science*, pages 321–326. Springer, 1985.
- [8] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, New York, 2 edition, 2003.
- [9] B. Grenet, J. van der Hoeven, and G. Lecerf. Deterministic root finding over finite fields using Graeffe transforms. *Appl. Alg. Eng. Comm. Comp.*, 27(3):237–257, 2016.
- [10] D. Harvey, J. van der Hoeven, and G. Lecerf. Faster polynomial multiplication over finite fields. *J. ACM*, 63(6), 2017. Article 52.
- [11] J. van der Hoeven. Relax, but don’t be too lazy. *J. Symbolic Comput.*, 34(6):479–542, 2002.
- [12] J. van der Hoeven. Fast composition of numeric power series. Technical Report 2008-09, Université Paris-Sud, Orsay, France, 2008.
- [13] J. van der Hoeven and G. Lecerf. Modular composition via factorization. Technical report, CNRS & École polytechnique, 2017. <http://hal.archives-ouvertes.fr/hal-01457074>.
- [14] Xiaohan Huang and V. Y. Pan. Fast rectangular matrix multiplication and applications. *J. Complexity*, 14(2):257–299, 1998.
- [15] E. Kaltofen and V. Shoup. Fast polynomial factorization over high algebraic extensions of finite fields. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation*, ISSAC ’97, pages 184–188, New York, NY, USA, 1997. ACM.
- [16] E. Kaltofen and V. Shoup. Subquadratic-time factoring of polynomials over finite fields. *Math. Comp.*, 67(223):1179–1197, 1998.
- [17] K. S. Kedlaya and C. Umans. Fast modular composition in any characteristic. In *FOCS’08: IEEE Conference on Foundations of Computer Science*, pages 146–155, Washington, DC, USA, 2008. IEEE Computer Society.
- [18] K. S. Kedlaya and C. Umans. Fast polynomial factorization and modular composition. *SIAM J. Comput.*, 40(6):1767–1802, 2011.
- [19] F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ISSAC ’14, pages 296–303, New York, NY, USA, 2014. ACM.
- [20] G. Lecerf. On the complexity of the Lickteig-Roy sub-resultant algorithm. Technical report, CNRS & École polytechnique, 2017. <https://hal.archives-ouvertes.fr/hal-01450869>.
- [21] M. S. Paterson and L. J. Stockmeyer. On the number of nonscalar multiplications necessary to evaluate polynomials. *SIAM J. Comput.*, 2(1):60–66, 1973.
- [22] P. Ritzmann. A fast numerical algorithm for the composition of power series with complex coefficients. *Theoret. Comput. Sci.*, 44:1–16, 1986.