Faster relaxed multiplication*

Joris van der Hoeven

Dépt. de Mathématiques (Bât. 425)
CNRS, Université Paris-Sud
91405 Orsay Cedex
France
Email: joris@texmacs.org

April 13, 2012

In previous work, we have introduced several fast algorithms for relaxed power series multiplication (also known under the name on-line multiplication) up till a given order n. The fastest currently known algorithm works over an effective base field $\mathbb{K}$ with sufficiently many 2^p-th roots of unity and has algebraic time complexity $\mathcal{O}(n \log n \, e^{2\sqrt{\log 2 \log\log n}})$. In this note, we will generalize this algorithm to the cases when $\mathbb{K}$ is replaced by an effective ring of positive characteristic or by an effective ring of characteristic zero, which is also torsion-free as a $\mathbb{Z}$-module and comes with an additional algorithm for partial division by integers. We will also present an asymptotically faster algorithm for relaxed multiplication of p-adic numbers.

Keywords: power series, multiplication, on-line algorithm, FFT, computer algebra
A.M.S. subject classification: 68W30, 30B10, 68W25, 33F05, 11Y55, 42-04

1. Introduction

1.1. Relaxed resolution of recursive equations

Let $\mathbb{A}$ be an effective (possibly non commutative) ring. That is, we assume data structures for representing the elements of $\mathbb{A}$ and algorithms for performing the ring operations $+$, $-$ and $\times$. The aim of algebraic complexity theory is to study the cost of basic or more complex algebraic operations over $\mathbb{A}$ (such as the multiplication of polynomials or matrices) in terms of the number of operations in $\mathbb{A}$.

The algebraic complexity usually does not coincide with the bit complexity, which also takes into account the potential growth of the actual coefficients in $\mathbb{A}$. Nevertheless, understanding the algebraic complexity usually constitutes a first useful step towards understanding the bit complexity. Of course, in the special case when $\mathbb{A}$ is a finite field, both complexities coincide up to a constant factor.

One of the most central operations is polynomial multiplication. We will denote by $\mathsf{M}_{\mathbb{A}}(n)$ the number of operations required to multiply two polynomials of degrees $<n$ in $\mathbb{A}[x]$. If $\mathbb{A}$ admits primitive 2^p-th roots of unity for all p, then we have $\mathsf{M}_{\mathbb{A}}(x) = \mathcal{O}(n \log n)$ using FFT multiplication, which is based on the fast Fourier transform [CT65]. In general, it has been shown [SS71, CK91] that $\mathsf{M}_{\mathbb{A}}(n) = \mathcal{O}(n \log n \log\log n)$. The complexities of most other operations (division, Taylor shift, extended g.c.d., multipoint evaluation, interpolation, etc.) can be expressed in terms of $\mathsf{M}_{\mathbb{A}}(n)$. Often, the cost of such other operations is simply $\tilde{\mathcal{O}}(\mathsf{M}_{\mathbb{A}}(n)) = \tilde{\mathcal{O}}(n)$, where $\tilde{\mathcal{O}}(T(n))$ stands for $\mathcal{O}(T(n) \, (\log n)^{\mathcal{O}(1)})$; see [AHU74, BP94, GG02] for some classical results along these lines.

*. This work has been partly supported by the French ANR-09-JCJC-0098-01 MaGiX project, and by the Digiteo 2009-36HD grant of the Région Ile-de-France.

The complexity of polynomial multiplication is fundamental for studying the cost of operations on formal power series in $\mathbb{A}[[z]]$ up to a given truncation order n. Clearly, it is possible to perform the multiplication up to order n in time $\mathcal{O}(\mathsf{M}_{\mathbb{A}}(n))$: it suffices to multiply the truncated power series at order n and truncate the result. Using Newton's method, and assuming that $\mathbb{Q} \subseteq \mathbb{A}$, it is also possible to compute $\exp$, $\sin$, etc. up to order n in time $\mathcal{O}(\mathsf{M}_{\mathbb{A}}(n))$. More generally, it has been shown in [BK78, Hoe02, Sed01, Hoe10] that the power series solutions of algebraic differential equations with coefficients in $\mathbb{A}[[z]]$ can be computed up to order n in time $\mathcal{O}(\mathsf{M}_{\mathbb{A}}(n))$. However, in this case, the "$\mathcal{O}$" hides a non trivial constant factor which depends on the size of the equation that one wants to solve.

The *relaxed* approach for computations with formal power series makes it possible to solve equations in quasi-optimal time with respect to the sparse size of the equations. The idea is to consider power series $f \in \mathbb{A}[[z]]$ as streams of coefficients $f_0, f_1, \ldots$ and to require that all operations are performed "without delay" on these streams. For instance, when multiplying $h = f\,g$ two power series $f, g \in \mathbb{A}[[z]]$, we require that h_n is computed *as soon as* $f_0, g_0, \ldots, f_n, g_n$ are known. Any algorithm which has this property will be called a *relaxed* or *on-line* algorithm for multiplication.

Given a relaxed algorithm for multiplication, it is possible to let the later coefficients $f_{n+1}, g_{n+1}, f_{n+2}, g_{n+2}, \ldots$ of the input *depend* on the known coefficients $h_0, \ldots, h_n$ of the output. For instance, given a power series $f \in \mathbb{K}[z]$ with $f_0 = 0$, we may compute $g = \exp f$ using the formula

$$g \;=\; \int f'\, g. \tag{1}$$

Indeed, extraction of the coefficient of z^n in g and $\int f'\, g$ yields

$$g_n \;=\; \frac{1}{n}\,(f'\, g)_{n-1},$$

and $(f'\, g)_{n-1}$ only depends on $g_0, \ldots, g_{n-1}$. More generally, we define an equation of the form

$$f \;=\; \Phi(f) \tag{2}$$

to be *recursive*, if $\Phi(f)_n$ only depends on $f_0, \ldots, f_{n-1}$. Replacing $\mathbb{A}$ by $\mathbb{A}^r$, we notice that the same terminology applies to systems of r equations. In the case of an implicit equation, special rewriting techniques can be implied in order transform the input equation into a recursive equation [Hoe11, Hoe09, BL11].

Let $\mathsf{R}_{\mathbb{A}}(n)$ denote the cost of performing a relaxed multiplication up to order n. If Φ is an expression which involves s multiplications and t other "linear time" operations (additions, integrations, etc.), then it follows that (2) can be solved up to order n in time $\mathcal{O}(s\,\mathsf{R}_{\mathbb{A}}(n) + t\,n)$. If we had $\mathsf{R}_{\mathbb{A}}(n) = \mathcal{O}(\mathsf{M}_{\mathbb{A}}(n))$, then this would yield an optimal algorithm for solving (2) in the sense that the computation of the solution would essentially require the same time as its verification.

1.2. Known algorithms for relaxed multiplication

The naive $\mathcal{O}(n^2)$ algorithm for computing $h = f\,g$, based on the formula

$$h_n \;=\; f_0\, g_n + f_1\, g_{n-1} + \cdots + f_n\, g_0,$$

is clearly relaxed. Unfortunately, FFT multiplication is *not* relaxed, since $h_0, \ldots, h_n$ are computed simultaneously as a function of $f_0\, g_0, \ldots, f_n\, g_n$, in this case.

In [Hoe97, Hoe02] it was remarked that Karatsuba's $\mathcal{O}\left(n^{\log 3/\log 2}\right)$ algorithm [KO63] for multiplying polynomials can be rewritten in a relaxed manner. For small n, Karatsuba multiplication and relaxed multiplication thus require *exactly* the same number of operations. In [Hoe97, Hoe02], an additional fast relaxed algorithm was presented with time complexity

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(\mathsf{M}_{\mathbb{A}}(n)\log n). \tag{3}$$

We were recently made aware of the fact that a similar algorithm was first published in [FS74]. However, this early paper was presented in a different context of on-line (relaxed) multiplication of integers (instead of power series), and without the application to the resolution of recursive equations (which is quite crucial from our perspective).

An interesting question remained: can the bound (3) be lowered further, be it by a constant factor? In [Hoe03], it was first noticed that an approximate factor of two can be gained if one of the multiplicands is known beforehand. For instance, if we want to compute $g = \exp f$ for a known series f with $f_0 = 0$, then the coefficients of f' are already known in the product $f'\,g$ in (1), so only one of the inputs depends on the output. An algorithm for the computation of $h = f\,g$ is said to be *semi-relaxed*, if h_n is written to the output as soon as $f_0, \ldots, f_n$ are known, but all coefficients of g are known beforehand. We will denote by $\mathsf{S}_{\mathbb{A}}(n)$ the complexity of semi-relaxed multiplication. We recall from [Hoe07] (see also section 3) that relaxed multiplication reduces to semi-relaxed multiplication

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(\mathsf{S}_{\mathbb{A}}(n)).$$

The first reduction of (3) by a non constant factor was published in [Hoe07], and uses the technique of *FFT blocking* (which has also been used for the multiplication of multivariate polynomials and power series in [Hoe02, Section 6.3], and for speeding up Newton iterations in [Ber00, Hoe10]). Under the assumption that $\mathbb{A}$ admits primitive 2^p-th roots of unity for all p (or at least for all p with $2^p \leqslant n$), we showed that

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(n \log n\, \mathrm{e}^{2\sqrt{\log 2 \log\log n}}). \tag{4}$$

Since this complexity will play a central role in the remainder of this paper, it will be convenient to abbreviate

$$\mathsf{R}_{*}(n) \;=\; n \log n\, \mathrm{e}^{2\sqrt{\log 2 \log\log n}}.$$

The function $\mathrm{e}^{2\sqrt{\log 2 \log\log n}}$ has slower growth than any strictly positive power of $\log n$. It will also be convenient to write $F(n) = \mathcal{O}^{\flat}(T(n))$ whenever $F(n) = \mathcal{O}(T(n)\,(\log n)^{\alpha})$ for all $\alpha > 0$. In particular, it follows that

$$\mathsf{R}_{*}(n) \;=\; \mathcal{O}^{\flat}(n \log n).$$

In section 3, we will recall the main ideas from [Hoe07] which lead to the complexity bound (4).

1.3. Improved complexity bounds

We recall that the characteristic of a ring $\mathbb{A}$ is the integer $k \in \mathbb{N}$ such that the canonical ring homomorphism $\mathbb{Z} \to \mathbb{A}$ has kernel $\mathbb{Z}/k\,\mathbb{Z}$. If $\mathbb{A}$ is torsion-free as a $\mathbb{Z}$-module, then we will say that $\mathbb{A}$ admits an effective partial division by integers if there exists an algorithm which takes $k \in \mathbb{Z}^*$ and $x \in k\,\mathbb{A}$ on input and which returns the unique $y \in \mathbb{A}$ with $x = k\,y$ on output. The main result of this paper is:

THEOREM 1. *Assume that one of the following two holds:*

- $\mathbb{A}$ *is an effective ring of characteristic zero, which is torsion-free as a* $\mathbb{Z}$*-module, and which admits an effective partial division by integers.*
- $\mathbb{A}$ *is an effective ring of positive characteristic.*

Then we have

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}^{\flat}(n \log n). \tag{5}$$

We notice that the theorem holds in particular if $\mathbb{A}$ is an effective field. In Section 7, we will also consider the relaxed multiplication of p-adic numbers, with $p \in \mathbb{N}$ and $p \geqslant 2$. If we denote by $\mathsf{I}(k)$ the bit complexity of multiplying two k-bit integers, then [FS74] essentially provided an algorithm of bit complexity $\mathcal{O}(\mathsf{I}(n) \log n)$ for the relaxed multiplication of 2-adic numbers at order $\mathcal{O}(2^n)$. Various algorithms and benchmarks for more general p were presented in [BHL11]. It is also well known [Too63, Coo66, SS71, Für07] that $\mathsf{I}(n) = \mathcal{O}^{\flat}(n \log n)$. Let $\mathsf{R}_p(n)$ denote the bit complexity of the relaxed multiplication of two p-adic numbers modulo p^n. In Section 7, we will prove the following new result:

THEOREM 2. *Let* $p \in \mathbb{N}$ *with* $p \geqslant 2$*. Then we have*

$$\mathsf{R}_p(n) \;=\; \mathcal{O}^{\flat}(n \log n \log p \log\log p).$$

The main idea which allows for the present generalizations is quite straightforward. In our original algorithm from [Hoe07], the presence of sufficiently many primitive 2^p-th roots of unity in $\mathbb{A}$ gives rise to an quasi-optimal evaluation-interpolation strategy for the multiplication of polynomials. More precisely, given two polynomials of degrees $<n$, their FFT-multiplication only requires $\mathcal{O}(n)$ evaluation and interpolation points, and both the evaluation and interpolation steps can be performed fast, using only $\mathcal{O}(n \log n)$ operations. Now it has recently been shown [BS05] that quasi-optimal evaluation-interpolation strategies still exist if we evaluate and interpolate at points in geometric progressions instead of roots of unity. This result is the key to our new complexity bounds, although further technical details have to be dealt with in order to make things work for various types of effective rings $\mathbb{A}$. We also notice that the main novelty of [BS05] concerns the interpolation step. Fast evaluation at geometric progressions was possible before using the so called *chirp transform* [RSR69, Blu70]. For effective rings of small characteristic $\mathbb{A}$, this would actually have been sufficient for the proving the bound (5).

Our paper is structured as follows. Since the algorithms of [BS05] were presented in the case when $\mathbb{A}$ is an effective field, Section 2 starts with their generalization to more general effective rings $\mathbb{A}$. These generalizations are purely formal and contain no essentially new ideas. In Section 3, we give a short survey of the algorithm from [Hoe07], but we recommentd the reader to download the original paper from our webpage for full technical details. In Section 4, we prove Theorem 1 in the case when $\mathbb{A}$ has characteristic zero. In Section 5, we turn our attention to the case when $\mathbb{A}$ has prime characteristic p. If the characteristic is sufficiently large, then we may find sufficiently large geometric progressions in $\mathbb{A} \cap \mathbb{Z}$ in order to generalize the results from Section 4. Otherwise, we have to work over $\mathbb{A} \otimes \mathbb{F}_{p^k}$ for some sufficiently large k. In Section 6, we complete the prove of Theorem 1; the case when p is a prime power is a refinement of the result from Section 5. The remaining case is done *via* Chinese remaindering. In Section 7, we will prove Theorem 2.

2. Multipoint evaluation and interpolation

Let $\mathbb{A}$ be an effective integral domain and let $\mathbb{K}$ be its quotient field. Assume that $\mathbb{K}$ has characteristic zero. Let $\mathbb{M}$ be an effective torsion-free $\mathbb{A}$-module and $\mathbb{V} = \mathbb{K} \otimes \mathbb{M}$. Elements of $\mathbb{K}$ and $\mathbb{V}$ are fractions x/s with $x \in \mathbb{A}$ (resp. $x \in \mathbb{M}$) and $s \in \mathbb{A}^*$, and the operations $-$, $+$, $\times$ on such fractions are as usual:

$$\begin{aligned} -\frac{x}{s} &= \frac{-x}{s} \\ \frac{x}{s} + \frac{y}{t} &= \frac{t\,x + s\,y}{s\,t} \\ \frac{x}{s}\,\frac{y}{t} &= \frac{x\,y}{s\,t} \end{aligned}$$

For $x/s \in \mathbb{K}^*$, we also have $(x/s)^{-1} = s/x$. It follows that $\mathbb{K}$ and $\mathbb{V}$ are effective fields and vector spaces. Moreover, all field operations in $\mathbb{K}$ (and all vector space operations in $\mathbb{V}$) can be performed using only $\mathcal{O}(1)$ operations in $\mathbb{A}$ (resp. $\mathbb{A}$ or $\mathbb{M}$).

We will say that $\mathbb{M}$ admits an effective partial division, if for every $s \in \mathbb{A}^*$ and $x \in s\,\mathbb{M}$, we can compute the unique $y \in \mathbb{M}$ with $x = s\,y$. From now on, we will assume that this is the case, and we will count any division of the above kind as one operation in $\mathbb{M}$. Given $n \in \mathbb{N}$, we define

$$\mathbb{M}[z]_n \;=\; \{P \in \mathbb{M}[z] : \deg P < n\}.$$

Given $P \in \mathbb{A}[z]_n$ and $Q \in \mathbb{M}[z]_n$, we will denote by $\mathsf{M}_{\mathbb{M}}(n)$ the number of operations in $\mathbb{A}$ and $\mathbb{M}$ which are needed in order to compute the product $PQ \in \mathbb{M}[z]$.

Lemma 3. *Let $\mathbb{A}$ be an effective and commutative integral domain and $\mathbb{M}$ an effective torsion-free $\mathbb{A}$-module with an effective partial division. There exists a constant K, such that the following holds: for any $n > 0$, $P \in \mathbb{M}[z]_n$ and $q \in \mathbb{A}^*$ such that $1, q, \ldots, q^{n-1}$ are pairwise distinct, we have:*

a) We may compute $P(1), \ldots, P(q^{n-1})$ from P using $K\,\mathsf{M}_{\mathbb{M}}(n)$ operations in $\mathbb{A}$ and $\mathbb{M}$.

b) We may reconstruct P from $P(1)$, ..., $P(q^{n-1})$ using $K\,\mathsf{M}_{\mathbb{M}}(n)$ operations in $\mathbb{A}$ and $\mathbb{M}$.

Proof. In the case when $\mathbb{A} = \mathbb{K}$ is a field of characteristic zero and $\mathbb{M} = \mathbb{V} = \mathbb{K}$, this result was first proven in [BS05]. More precisely, the conversions can be done using the algorithms **NewtonEvalGeom**, **NewtonInterpGeom**, **NewtonToMonomialGeom** and **MonomialToNewtonGeom** in that paper. Examining these algorithms, we observe that general elements in $\mathbb{A}$ are only multiplied with elements in $\mathbb{Q}[q]$ and divided by elements of the set $\{q, q-1, q^2-1, \ldots, q^{n-1}-1\}$. In particular, the algorithms can still be applied in the more general case when $\mathbb{A} = \mathbb{K}$ is a field of characteristic zero and $\mathbb{M} = \mathbb{V}$ a vector space.

If $\mathbb{A}$ is only an effective commutative integral domain of characteristic zero and $\mathbb{M}$ an effective torsion-free $\mathbb{A}$-module with an effective partial division, then we define the effective field $\mathbb{K}$ and the effective vector field $\mathbb{V}$ as above, and we may still apply the generalized algorithms for multipoint evaluation and interpolation in $\mathbb{V}$. In particular, both multipoint evaluation and interpolation can still be done using $\mathcal{O}(\mathsf{M}_{\mathbb{V}}(n))$ operations in $\mathbb{K}$ and $\mathbb{V}$, whence $\mathcal{O}(\mathsf{M}_{\mathbb{M}}(n))$ operations in $\mathbb{A}$ and $\mathbb{M}$. If we *know* that the end-results of these algorithms are really in the subspace $\mathbb{M}^n$ of $\mathbb{V}^n$ (or in the subring $\mathbb{M}[z]_n$ of $\mathbb{V}[z]_n$), then we use the partial division in $\mathbb{M}$ to replace their representations in $\mathbb{V}^n$ (or $\mathbb{V}[z]_n$) by representations in $\mathbb{M}^n$ (or $\mathbb{M}[z]_n$). □

LEMMA 4. *Let $\mathbb{A}$ be an effective and commutative integral domain and $\mathbb{M}$ an effective $\mathbb{A}$-module. There exists a constant K, such that the following holds: for any $n>0$, $P\in\mathbb{M}[z]_n$ and $q\in\mathbb{A}$ such that 1, q, ..., q^{n-1} are pairwise distinct and such that q and $q-1, q^2-1, ..., q^{n-1}-1$ are all invertible in $\mathbb{A}$, we have:*

a) We may compute $P(1),...,P(q^{n-1})$ from P using $K\,\mathsf{M}_{\mathbb{M}}(n)$ operations in $\mathbb{A}$ and $\mathbb{M}$.

b) We may reconstruct P from $P(1)$, ..., $P(q^{n-1})$ using $K\,\mathsf{M}_{\mathbb{M}}(n)$ operations in $\mathbb{A}$ and $\mathbb{M}$.

Proof. We again use straightforward generalizations of the algorithms in [BS05], using the fact that we only divide by elements in the set $\{q, q-1, q^2-1, ..., q^{n-1}-1\}$. □

3. SURVEY OF BLOCKWISE RELAXED MULTIPLICATION

Let $\mathbb{A}$ be an effective (possibly non commutative) ring and recall that

$$\mathbb{A}[z]_n \;=\; \{P\in\mathbb{A}[z]\colon \deg P<n\}.$$

Given a power series $f\in\mathbb{A}[[z]]$ and $i<j$, we will also use the notations

$$\begin{aligned} f_{i;j} &= f_i\,z^i+\cdots+f_{j-1}\,z^{j-1}\\ f_{;j} &= f_0+\cdots+f_{j-1}\,z^{j-1}\\ f_{i;} &= f_i\,z^i+f_{i+1}\,z^{i+1}+\cdots. \end{aligned}$$

The fast relaxed algorithms from [Hoe07] are all based on two main changes of representation: "blocking" and the fast Fourier transform. Let us briefly recall these transformations and how to use them for the design of fast algorithms for relaxed multiplication.

Blocking and unblocking. Given a block size $b>0$, the first operation of *blocking* rewrites a power series $f\in\mathbb{A}[[z]]$ as a series in $y=z^b$ with coefficients in $\mathbb{A}[z]_b$

$$\mathrm{B}_b(f) \;=\; \sum_i\sum_{j<b} f_{ib+j}\,z^j\,y^i \;\in\; \mathbb{A}[z]_b[[y]].$$

Given $f,g\in\mathbb{A}[[z]]$, we may then compute fg using

$$fg \;=\; \mathrm{B}_b^{-1}(\mathrm{B}_b(f)\,\mathrm{B}_b(g)),$$

where $\mathrm{B}_b(f)\,\mathrm{B}_b(g)\in\mathbb{K}[z]_{2b}[[y]]$ and

$$\begin{aligned} \mathrm{B}_b^{-1}\colon \mathbb{A}[z]_{2b}[[y]] &\to \mathbb{A}[[z]]\\ \sum_i P_i(z)\,y^i &\mapsto \sum_i P_i(z)\,z^{bi}. \end{aligned}$$

Discrete Fourier transforms. Assume now that $\mathbb{A}\ni 1/2$, that $b\in\{1,2,4,8,...\}$, and that $\mathbb{A}$ admits a primitive $2\,b$-th root of unity $\omega=\omega_{2b}$. Then the discrete Fourier transform provides us with an isomorphism

$$\begin{aligned} \mathrm{FFT}_\omega\colon \mathbb{A}[z]_{2b} &\to \mathbb{A}^{2b}\\ P &\mapsto (P(1),P(\omega),...,P(\omega^{2b-1})), \end{aligned}$$

and it is classical [CT65] that both FFT_ω and FFT_ω^{-1} can be computed using $\mathcal{O}(b\log b)$ operations in $\mathbb{A}$. The operations FFT_ω and FFT_ω^{-1} extend naturally to $\mathbb{A}[z]_{2b}[[y]]$ *via*

$$\mathrm{FFT}_\omega\Big(\sum_i f_i\,y^i\Big) \;=\; \sum_i \mathrm{FFT}_\omega(f_i)\,y^i.$$

This allows us to compute $f\,g$ using the formula

$$f\,g \;=\; \mathrm{B}_b^{-1}\big(\mathrm{FFT}_\omega^{-1}(\mathrm{FFT}_\omega(\mathrm{B}_b(f))\,\mathrm{FFT}_\omega(\mathrm{B}_b(g)))\big). \tag{6}$$

The first $k\,b$ coefficients of $f\,g$ can be computed using at most $2\,b\,\mathsf{M}_{\mathbb{A}}(k)+\mathcal{O}(k\,b\log b)$ operations in $\mathbb{A}$.

Relaxed multiplication. In formula (6) the k-th coefficient of the right hand side may depend on the $(k+b-1)$-th coefficients of f and g. In order to make (6) suitable for relaxed multiplication, we have to treat the first b coefficients of f and g apart. Indeed, the formula

$$\begin{aligned} f\,g \;=\;& f_{0;b}\,g_{0;b}+f_{0;b}\,g_{b;}+f_{b;}\,g_{0;b}+\\ & \mathrm{B}_b^{-1}\big(\mathrm{FFT}_\omega^{-1}(\mathrm{FFT}_\omega(\mathrm{B}_b(f_{b;}))\,\mathrm{FFT}_\omega(\mathrm{B}_b(g_{b;})))\big)\end{aligned}$$

allows for the relaxed computation of $f\,g$ at order $k\,b$ using at most

$$\mathsf{R}_{\mathbb{A}}(k\,b) \;\leqslant\; \mathsf{R}_{\mathbb{A}}(b)+2\,k\,\mathsf{S}_{\mathbb{A}}(b)+2\,b\,\mathsf{R}_{\mathbb{A}}(k)+\mathcal{O}(k\,b\log b)$$

operations in $\mathbb{A}$. Similarly, the formula

$$f\,g \;=\; f_{0;b}\,g+\mathrm{B}_b^{-1}\big(\mathrm{FFT}_\omega^{-1}(\mathrm{FFT}_\omega(\mathrm{B}_b(f_{b;}))\,\mathrm{FFT}_\omega(\mathrm{B}_b(g)))\big)$$

allows for the semi-relaxed computation of $f\,g$ at order $k\,b$ using at most

$$\mathsf{S}_{\mathbb{A}}(k\,b) \;\leqslant\; k\,\mathsf{S}_{\mathbb{A}}(b)+2\,b\,\mathsf{S}_{\mathbb{A}}(k)+\mathcal{O}(k\,b\log b)$$

operations in $\mathbb{A}$. For a given expansion order n, one may take $b\approx\sqrt{n}$, and use the above formula in a recursive manner. This yields [Hoe07, Theorem 11]

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(n\,(\log n)^{\log 3/\log 2}).$$

Remark 5. Since the block size b is chosen as a function of n, the above method really describes a relaxed algorithm for computing the product up to an order n_* which is specified in advance. In fact, such an algorithm automatically yields a genuine relaxed algorithm with the same complexity (up to a constant factor), by doubling the order n_* each time when needed.

Reducing relaxed to semi-relaxed multiplication. In the above discussion, we both provided bounds for $\mathsf{R}_{\mathbb{A}}(n)$ and $\mathsf{S}_{\mathbb{A}}(n)$. In fact, there exists a straightforward reduction of relaxed multiplication to semi-relaxed multiplication. First of all, the relaxed multiplication of two power series $f,\,g\in\mathbb{A}[[z]]$ up to order $\mathcal{O}(z^n)$ clearly reduces to the relaxed multiplication of the two polynomials $f_{0;n}$ and $g_{0;n}$ up to order $\mathcal{O}(z^{2n})$. Now the formula

$$f_{0;2n}\,g_{0;2n} \;=\; f_{0;n}\,g_{0;n}+f_{n;2n}\,g_{0;n}+f_{0;n}\,g_{n;2n}+f_{n;2n}\,g_{n;2n}$$

shows that a relaxed product of two polynomials $f_{0;2n}$ and $g_{0;2n}$ of degrees $<2\,n$ reduces to a relaxed product $f_{0;n}\;g_{0;n}$ of half the size, two semi-relaxed products $f_{n;2n}\;g_{0;n}$, $f_{0;n}\,g_{n;2n}$, and one non-relaxed product $f_{n;2n}\,g_{n;2n}$. Under the assumptions that $\mathsf{M}_{\mathbb{A}}(n)/n$ and $\mathsf{S}_{\mathbb{A}}(n)/n$ are increasing, a routine calculation thus yields

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(\mathsf{S}_{\mathbb{A}}(n)).$$

Multiple block sizes. Instead of taking a single block size b, we may write $n\approx n_1\cdots n_l$, take $l-1$ different block sizes $b_1=n_1<\cdots<b_{l-1}=n_1\cdots n_{l-1}$, and decompose f in l parts $f=f_{0;b_1}+f_{b_1;b_2}+\cdots+f_{b_{l-1};b_l}$, where $b_l=+\infty$. For semi-relaxed multiplication, this yields the formula

$$f\,g \;=\; f_{0;b_1}\,g+\sum_{i=1}^{l-1}\mathrm{B}_{b_i}^{-1}\Big(\mathrm{FFT}_{\omega_{2b_i}}^{-1}\big(\mathrm{FFT}_{\omega_{2b_i}}(\mathrm{B}_{b_i}(f_{b_i;b_{i+1}}))\,\mathrm{FFT}_{\omega_{2b_i}}(\mathrm{B}_{b_i}(g))\big)\Big) \tag{7}$$

and the complexity bound

$$\mathsf{S}_{\mathbb{A}}(n) \leqslant \frac{n}{n_1}\mathsf{S}_{\mathbb{A}}(n_1)+\frac{2\,n}{n_2}\mathsf{S}_{\mathbb{A}}(n_2)+\cdots+\frac{2\,n}{n_l}\mathsf{S}_{\mathbb{A}}(n_l)+\mathcal{O}(l\,n\log n). \tag{8}$$

For $n_1\approx n_2\approx\cdots\approx n_l$ and $l_p\approx \mathrm{e}^{\sqrt{\log 2\log p}}$ well chosen block sizes (depending on the expansion order n), the recursive application of this technique yields [Hoe07, Theorem 12]

$$\begin{aligned}\mathsf{R}_{\mathbb{A}}(n) &= \mathcal{O}(\mathsf{S}_{\mathbb{A}}(n)).\\ &\quad\ \mathcal{O}(n\log n\,\mathrm{e}^{2\sqrt{\log 2\log\log n}}).\end{aligned}$$

4. Relaxed multiplication in characteristic zero

Let us now consider the less favourable case when $\mathbb{A}$ is an effective ring which does not necessarily contain primitive 2^k-th roots of unity for arbitrarily high k. In this section, we will first consider the case when $\mathbb{A}$ is torsion-free as a $\mathbb{Z}$-module and also admits a partial algorithm for division by integers.

Given a block size $b\in\mathbb{N}$ and $q\in\mathbb{Z}\setminus\{-1,0,1\}$ (say $q=2$), we will replace the discrete Fourier transform FFT_ω at a $(2\,b)$-th primitive root of unity by multipoint evaluation at $1,\ldots,q^{2b-1}$. More precisely, we define

$$\begin{aligned}\mathrm{E}_{q,2b}\colon \mathbb{A}[z]_{2b} &\to \mathbb{A}^{2b}\\ P &\mapsto (P(1),P(q),\ldots,P(q^{2b-1}))\end{aligned}$$

and the inverse transform $\mathrm{E}_{q,2b}^{-1}\colon \operatorname{im}\mathrm{E}_{q,2b}\to\mathbb{A}[z]_{2b}$. By Lemma 3, these transforms can both be computed using $\mathcal{O}(\mathsf{M}(b))$ operations in $\mathbb{A}$. In a similar way as FFT_ω and FFT_ω^{-1}, we extend $\mathrm{E}_{q,2b}$ and $\mathrm{E}_{q,2b}^{-1}$ to power series in y.

Theorem 6. *Let $\mathbb{A}$ both be an effective ring and an effective torsion-free $\mathbb{Z}$-module. Then*

$$\begin{aligned}\mathsf{R}_{\mathbb{A}}(n) &= \mathcal{O}(\mathsf{R}_*(n)\sqrt{\log\log n})\\ &= \mathcal{O}^b(n\log n).\end{aligned}$$

Proof. It suffices to prove the complexity bound for semi-relaxed multiplication. We adapt the multiplication algorithm with l different block sizes from the end of the previous section, and replace (7) by

$$f\,g = f_{0;b_1}\,g+\sum_{i=1}^{l-1}\mathrm{B}_{b_i}^{-1}\big(\mathrm{E}_{q,2b_i}^{-1}(\mathrm{E}_{q,2b_i}(\mathrm{B}_{b_i}(f_{b_i;b_{i+1}}))\,\mathrm{E}_{q,2b_i}(\mathrm{B}_{b_i}(g)))\big). \tag{9}$$

This leads to the complexity bound

$$\mathsf{S}_{\mathbb{A}}(n) \leqslant \frac{n}{n_1}\mathsf{S}_{\mathbb{A}}(n_1)+\frac{2\,n}{n_2}\mathsf{S}_{\mathbb{A}}(n_2)+\cdots+\frac{2\,n}{n_l}\mathsf{S}_{\mathbb{A}}(n_l)+\mathcal{O}(l\,\mathsf{M}_{\mathbb{A}}(n)). \tag{10}$$

We now follow the proof of [Hoe07, Theorem 12]. Denote

$$U(p) = \frac{\mathsf{S}_{\mathbb{A}}(2^p)}{p\,2^p}$$

and take $n_1=\cdots=n_l$. Using that $\mathsf{M}_{\mathbb{A}}(n)=\mathcal{O}(n\log n\log\log n)$, this leads to

$$U(l\,p) \leqslant 2\,U(p)+\mathcal{O}(l\log(p\,l))$$

Applying this relation k times, we obtain

$$\begin{aligned} U(l^k) &\leqslant 2^k U(1) + \mathcal{O}\Big(2^k l \left(1 + \tfrac{2}{2} + \cdots + \tfrac{k}{2^k}\right) \log l\Big) \\ &= \mathcal{O}(2^k l \log l). \end{aligned}$$

Taking

$$\begin{aligned} k &= \frac{\log p}{\log l} \\ l &= \mathrm{e}^{\sqrt{\log 2 \log p}}, \end{aligned}$$

this yields

$$U(p) = \mathcal{O}(\sqrt{\log p}\, \mathrm{e}^{2\sqrt{\log 2 \log p}}).$$

Re-expressing this bound in terms of $\mathsf{S}_{\mathbb{A}}$ yields the desired result. □

5. Relaxed multiplication in prime characteristic

Let $\mathbb{A}$ now be an effective ring of prime characteristic p. For expansion orders $n \geqslant p$, the ring $\mathbb{A}$ does not necessarily contain n distinct points in geometric progression. Therefore, in order to apply Lemma 4, we will first replace $\mathbb{A}$ by a suitable extension, in which we can find sufficiently large geometric progressions.

Given n, let $k = 2\left\lceil \frac{\log(n+1)}{2 \log p} \right\rceil$ be even such that $p^k > n$. Let $P_{p^k} \in \mathbb{F}_p[z]$ be such that the finite field $\mathbb{F}_{p^k}$ is isomorphic to $\mathbb{F}_p[z]/(P_{p^k})$. Then the ring

$$\mathbb{B}_k := \mathbb{A}[z]/(P_{p^k})$$

has dimension k over $\mathbb{A}$ as a vector space, so we have a natural $\mathbb{A}$-linear bijection

$$\begin{aligned} \Lambda_k \colon \mathbb{A}[z]_k &\to \mathbb{B}_k \\ A &\mapsto A \bmod P. \end{aligned}$$

The ring $\mathbb{B}_k$ is an effective ring and one addition or subtraction in $\mathbb{B}_k$ corresponds to k additions or subtractions in $\mathbb{A}$. Similarly, one multiplication in $\mathbb{B}_k$ can be done using $\mathcal{O}(\mathsf{M}_{\mathbb{A}}(k))$ operations in $\mathbb{A}$.

In order to multiply two series $f, g \in \mathbb{A}[[z]]$ up to order $\mathcal{O}(z^n)$, the idea is now to rewrite f and g as series in $\mathbb{B}[[u]]$ with $u = z^{k/2}$. If we want to compute the relaxed product, then we also have to treat the first $k/2$ coefficients apart, as we did before for the blocking strategy. More precisely, we will compute the semi-relaxed product fg using the formula

$$fg = f_{0;k/2}\, g + \mathrm{B}_{k/2}^{-1}\big(\Lambda_k^{-1}(\Lambda_k(\mathrm{B}_{k/2}(f_{k/2;}))\, \Lambda_k(\mathrm{B}_{k/2}(g)))\big),$$

where we extended Λ_k to $\mathbb{A}[z]_k[[u]]$ in the natural way:

$$\Lambda_k\left(\sum_{i \geqslant 0} f_i\, u^i\right) = \sum_{i \geqslant 0} \Lambda_k(f_i)\, u^i.$$

From the complexity point of view, we get

$$\mathsf{S}_{\mathbb{A}}(n) \leqslant \frac{2\,n}{k}\, \mathsf{S}_{\mathbb{A}}\left(\frac{k}{2}\right) + \mathsf{S}_{\mathbb{B}}\left(\frac{2\,n}{k}\right) \mathcal{O}(\mathsf{M}_{\mathbb{A}}(k)). \tag{11}$$

Since $\mathbb{B}$ contains a copy of $\mathbb{F}_{p^k}$, it also contains at least $p^k - 1 \geqslant n$ points in geometric progression. For the multiplication up till order $2\,n/k$ of two series with coefficients in $\mathbb{B}$, we may thus use the blocking strategy combined with multipoint evaluation and interpolation.

THEOREM 7. *Let $\mathbb{A}$ be an effective ring of prime characteristic p. Then*

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(\mathsf{R}_*(n)\,(\log\log n)^{3/2}\log\log\log n)$$

Proof. With the notations from above, we may find a primitive (p^k-1)-th root of unity q in $\mathbb{F}_{p^k}\subseteq\mathbb{B}$. We may thus use formula (9) for the semi-relaxed multiplication of two series in $\mathbb{B}[[z]]$ up till order $2\,n/k\leqslant n$. In a similar way as in the proof of Theorem 6, we thus get

$$\mathsf{S}_{\mathbb{B}}\left(\frac{2\,n}{k}\right) \;=\; \mathcal{O}\left(\mathsf{R}_*(n)\,\frac{\sqrt{\log\log n}}{k}\right).$$

Using classical fast relaxed multiplication [Hoe97, Hoe02, FS74], we also have

$$\mathsf{S}_{\mathbb{A}}\left(\frac{k}{2}\right) \;=\; \mathcal{O}(k\log^2 k\log\log k),$$

whence (11) simplifies to

$$\mathsf{S}_{\mathbb{A}}(n) \;=\; \mathcal{O}\left(\mathsf{R}_*(n)\,\frac{\mathsf{M}_{\mathbb{A}}(k)}{k}\,\sqrt{\log\log n}\right).$$

Since $\mathsf{M}_{\mathbb{A}}(k)/k=\mathcal{O}(\log k\log\log k)$ and $k=\mathcal{O}(\log n)$, the result follows. □

Remark 8. In our complexity analysis, we have not taken into account the computation of the polynomial $P_{p^k}\in\mathbb{F}_p[z]$ with $\mathbb{F}_{p^k}=\mathbb{F}_p[z]/(P_{p^k})$. Using a randomized algorithm, such a polynomial can be computed in time $\tilde{\mathcal{O}}(k^2\log p)$; see [GG02, Corollary 14.44]. If $k=\mathcal{O}(\log n)$, then this is really a precomputation of negligible cost $\tilde{\mathcal{O}}(\log^2 n)$.

If we insist on computing P_{p^k} in a deterministic way, then it is better to slightly change our algorithm, and systematically choose k to be a prime number minus one. Under this assumption, the cyclotomic polynomial $x^k+\cdots+1$ is irreducible over $\mathbb{F}_p$; see [GG02, Lemma 14.50]. For a fixed n, the first prime number q with $p^{q-1}>n$ still has size $\mathcal{O}(\log n)$, by the prime number theorem, and we may compute it in time $\tilde{\mathcal{O}}(\log n)$ using the sieve of Eratosthenes. This again shows that a suitable P_{p^k} can be precomputed with negligible cost.

6. RELAXED MULTIPLICATION IN MIXED CHARACTERISTIC

Let us now show that the technique from the previous section actually extends to the case when $\mathbb{A}$ is an arbitrary effective ring of positive characteristic. We first show that the algorithm still applies when the characteristic of $\mathbb{A}$ is a prime power. We then conclude by showing how to apply Chinese remaindering in our setting.

THEOREM 9. *Let $\mathbb{A}$ be an effective ring of prime power characteristic $s=p^r$. Then*

$$\mathsf{R}_{\mathbb{A}}(n) \;=\; \mathcal{O}(\mathsf{R}_*(n)\,(\log\log n)^{3/2}\log\log\log n).$$

Proof. Taking $k=2\left\lceil\frac{\log(n+1)}{2\log p}\right\rceil$, let $P=P_{p^k}$ be as in the previous section and pick a monic polynomial $\tilde{P}$ in $(\mathbb{Z}/s\,\mathbb{Z})[z]$ such that the reduction $\pi(\tilde{P})$ of $\tilde{P}$ modulo p yields P. Then we get a natural commutative diagram

$$\begin{array}{ccc}
(\mathbb{Z}/s\,\mathbb{Z})[z]/(\tilde{P}) & \hookrightarrow & \mathbb{A}[z]/(\tilde{P}) \\
\downarrow\pi & & \downarrow\pi \\
\mathbb{F}_p[z]/(P) & \hookrightarrow & \pi(\mathbb{A})[z]/(P),
\end{array}$$

where π stands for reduction modulo p. In particular, we have an epimorphism

$$\pi\colon (\mathbb{Z}/s\,\mathbb{Z})[z]/(\tilde{P}) \;\to\; \mathbb{F}_p[z]/(P) \cong \mathbb{F}_{p^k},$$

with $\ker \pi = (p)$.

Now let q be an element in $\mathbb{F}_{p^k}$ of order $p^k - 1$. Then any lift $\tilde{q} \in (\mathbb{Z}/s\,\mathbb{Z})[z]/(\tilde{P})$ of q with $\pi(\tilde{q}) = q$ has order at least $p^k - 1$. Moreover, $q - 1$, ..., $q^{p^k-2} - 1$ and q are all invertible. Consequently, $\tilde{q} - 1, ..., \tilde{q}^{p^k-2} - 1$ and $\tilde{q}$ do not lie in $\ker \pi = (p)$, whence they are invertible as well. It follows that we may still apply multipoint evaluation and interpolation in $\mathbb{A}[z]/(\tilde{P})$ at the sequence $1, \tilde{q}, ..., \tilde{q}^{p^k-2}$, whence Theorem 7 generalizes to the present case. □

Remark 10. For a fixed prime number p, we notice that the complexity bound is uniform in the following sense: there exists a constant K such that for all effective rings of characteristic p^r with $r \in \{1, 2, ...\}$, we have $\mathsf{R}_{\mathbb{A}}(n) \leqslant K\,\mathsf{R}_*(n)\,(\log\log n)^{3/2}\log\log\log n$. Indeed, the choice of k only depends on n and p, and any operation in $\mathbb{F}_{p^k}$ or $\mathbb{A}[z]/(P_{p^k})$ in the case $r = 1$ corresponds to exactly one lifted operation in $(\mathbb{Z}/s\,\mathbb{Z})[z]/(\tilde{P})$ or $\mathbb{A}[z]/(\tilde{P})$ in the general case.

THEOREM 11. *Let $\mathbb{A}$ be an effective ring of non zero characteristic s. Then*

$$\begin{aligned}\mathsf{R}_{\mathbb{A}}(n) &= \mathcal{O}(\mathsf{R}_*(n)\,(\log\log n)^{3/2}\log\log\log n)\\ &= \mathcal{O}^{\flat}(n\log n).\end{aligned}$$

Proof. We will prove the theorem by induction on the number of prime divisors of s. If s is a prime power, then we are done. So assume that $s = s_1\,s_2$, where s_1 and s_2 are relatively prime, and let $k_1, k_2 \in \mathbb{Z}$ be such that

$$k_1\,s_1 + k_2\,s_2 \;=\; 1.$$

Then we may consider the rings

$$\begin{aligned}\mathbb{A}_1 &= \mathbb{A}/s_1\,\mathbb{A}\\ \mathbb{A}_2 &= \mathbb{A}/s_2\,\mathbb{A}.\end{aligned}$$

These rings are effective, when representing their elements by elements of $\mathbb{A}$ and transporting the operations from $\mathbb{A}$. Of course, the representation of an element x of $\mathbb{A}_1$ (or $\mathbb{A}_2$) is not unique, since we may replace it by $x + y$ for any $y \in s_1\,\mathbb{A}$ (or $y \in s_2\,\mathbb{A}$). But this is not a problem, since our definition of effective ring did not require unique representability or the existence of an equality test.

Now let f, $g \in \mathbb{A}[[z]]$ and let $\pi_i(f)$, $\pi_i(g)$ be their projections in $\mathbb{A}_i[[z]]$, for $i = 1, 2$. Consider the relaxed products $\pi_i(f)\;\pi_i(g)$, for $i = 1, 2$. These products are represented by relaxed series h^1, $h^2 \in \mathbb{A}[[z]]$ *via* $\pi_i(h^i) = \pi_i(f)\;\pi_i(g)$, for $i = 1, 2$. By the induction hypotheses, we may compute h^1 and h^2 at order n using

$$\mathcal{O}(\mathsf{R}_*(n)\,(\log\log n)^{3/2}\log\log\log n)$$

operations in $\mathbb{A}$. The linear combination $h = k_2\,s_2\,h^1 + k_1\,s_1\,h^2 \in \mathbb{A}[[z]]$ can still be expanded up till order n with the same complexity. We claim that $h = fg$. Indeed,

$$\begin{aligned}k_2\,s_2\,h^1 - k_2\,s_2\,fg &\in k_2\,s_2\,s_1\,\mathbb{A} = \{0\}\\ k_1\,s_1\,h^2 - k_1\,s_1\,fg &\in k_1\,s_1\,s_2\,\mathbb{A} = \{0\}.\end{aligned}$$

Summing both relations, our claim follows. □

7. Relaxed multiplication of p-adic numbers

Let $p > 1$ be an integer, not necessarily a prime number, and denote $\mathbb{N}_p = \{0, ..., p-1\}$. We will regard p-adic numbers $a \in \mathbb{Z}_p$ as series $a_0 + a_1\, p + a_2\, p^2 + \cdots$ with $a_i \in \mathbb{N}_p$, and such that the basic ring operations $+$, $-$ and $\times$ require an additional carry treatment.

In order to multiply two relaxed p-adic numbers $a, b \in \mathbb{Z}_p$, we may rewrite them as series $\hat{a}, \hat{b} \in \mathbb{Z}[[z]]$, multiply these series $\hat{c} = \hat{a}\,\hat{b}$, and recover the product $c \in \mathbb{Z}_p$ from the result. Of course, the coefficients of $\hat{c}$ may exceed p, so some relaxed carry handling is required in order to recover c from $\hat{c}$. We refer to [BHL11, Section 2.7] for details. In particular, we prove there that c can be computed up to order $\mathcal{O}(p^n)$ using $\mathcal{O}(\mathsf{R}_{\mathbb{Z}}(n))$ ring operations in $\mathbb{Z}$ of bit size $\mathcal{O}(\log p + \log n)$.

Given $k > 0$, let $\mathcal{Z}_k = \{i \in \mathbb{Z} \colon |i| < 2^{k-1}\}$, and consider two power series $f, g \in \mathcal{Z}_k[[z]]$. We will denote by $\mathsf{R}_{\mathbb{Z}}(n, k)$ (resp. $\mathsf{S}_{\mathbb{Z}}(n, k)$) the bit complexity of multiplying f and g up to order $\mathcal{O}(z^n)$ using a relaxed (resp. semi-relaxed) algorithm.

Lemma 12. *We have*

$$\mathsf{R}_{\mathbb{Z}}(n, k) \;=\; \mathcal{O}(\mathsf{R}_*(n)\, \mathsf{I}(k + \log n)\, (\log\log n)^{3/2} \log\log\log n).$$

Proof. Let $f, g \in \mathcal{Z}_k[[z]]$ and let p be a prime number (in practice, we recommend to take p to be a prime number which fits inside one machine word and such that $p - 1$ is divisible by a high power of two). Let $r = \lceil 2\, k \log(2\, n)/\log p \rceil$ be sufficiently large such that $n\, 2^{2k} < p^r$. Let $\hat{f}, \hat{g} \in (\mathbb{Z}/p^r\,\mathbb{Z})[[z]]$ be the reductions of f, g modulo p^r. Then $f\,g$ may be reconstructed up to order $\mathcal{O}(z^n)$ from the product $\hat{f}\,\hat{g}$. We thus get

$$\mathsf{R}_{\mathbb{Z}}(n, k) \;\leqslant\; \mathsf{I}(p^r)\, \mathsf{R}_{\mathbb{Z}/p^r\mathbb{Z}}(n)$$

By Theorem 9, we have

$$\mathsf{R}_{\mathbb{Z}/p^r\mathbb{Z}}(n) \;=\; \mathcal{O}(\mathsf{R}_*(n)\, (\log\log n)^{3/2} \log\log\log n).$$

By Remark 10, this bound is uniform in r. Since $p^r \sim n\, 2^{2k}$, the result follows. □

For the above strategy to be efficient, it is important that $\log n = \mathcal{O}(k)$. This can be achieved by combining it with the technique of p-adic blocking. More precisely, given a p-adic block size $b > 1$, then any p-adic number in $\mathbb{Z}_p$ can naturally be considered as a p^b-adic number in $\mathbb{Z}_{p^b}$, and *vice versa*. Assuming that numbers in $\mathbb{N}_p$ are written in base 2, the conversion is trivial if p is a power of two. Otherwise, the conversion involves base conversions and we refer to [BHL11, Section 4] for more details. In particular, the conversions in both directions up to order $\mathcal{O}(p^n)$ can be done in time $\mathcal{O}\big(\frac{n}{b}\, \mathsf{I}(b \log p) \log(b \log p)\big)$.

Let $\mathsf{R}_p(n)$ (resp. $\mathsf{S}_p(n)$) the complexity of relaxed (resp. semi-relaxed) multiplication in $\mathbb{Z}_p$ up till order $\mathcal{O}(p^n)$.

Theorem 13. *We have*

$$\begin{aligned} \mathsf{R}_p(n) \;&=\; \mathcal{O}\big(\mathsf{R}_*(n) \log p \log(\log n + \log p)(\log\log(n + \log p))^{\mathcal{O}(1)}\big) \\ &=\; \mathcal{O}^{\flat}(n \log n \log p \log\log p). \end{aligned}$$

Proof. Let $r = \lceil \log n/\log p \rceil$, so that

$$\begin{aligned} r \log p \;&\leqslant\; \log n + \log p \\ \mathsf{I}(r \log p) \;&=\; \mathcal{O}(r \log p \log(\log n + \log p) \log\log(n + \log p)) \\ \mathsf{I}(r\,(\log p + \log r)) \;&=\; \mathcal{O}(r\,(\log p + \log n) \log(\log n + \log p) \log\log(n + \log p)). \end{aligned}$$

Using the strategy of p-adic blocking, a semi-relaxed product in $\mathbb{Z}_p$ may then be reduced to one semi-relaxed product in $\mathbb{Z}_{p^r}$ and one relaxed multiplication with an integer in $\{0, ..., p^r - 1\}$. In other words,

$$\mathsf{S}_p(n) \leqslant \frac{n}{r}\mathsf{S}_p(r) + \mathsf{S}_{p^r}\Big(\Big\lceil \frac{n}{r} \Big\rceil\Big) + \mathcal{O}(n \log p),$$

where $\mathsf{S}_p(n)$ stands for the cost of semi-relaxed multiplication of two p-adic numbers in $\mathbb{Z}_p$ up till order $\mathcal{O}(p^n)$. By [BHL11, Proposition 4], we have

$$\begin{aligned} \mathsf{S}_p(r) &= \mathcal{O}(\mathsf{I}(r\,(\log p + \log r))\log r) \\ &= \mathcal{O}(r \log n\,(\log n + \log p) \log\,(\log n + \log p) \log\log\,(n + \log p)) \end{aligned}$$

By Lemma 12, we also have

$$\begin{aligned} \mathsf{S}_{p^r}\Big(\Big\lceil \frac{n}{r} \Big\rceil\Big) &= \mathcal{O}\Big(\mathsf{I}(r \log p + \log n)\,\frac{n}{r} \log n\,(\log\log n)^{3/2} \log\log\log n\,\mathrm{e}^{2\sqrt{\log 2 \log\log n}}\Big) \\ &= \mathcal{O}\Big(\mathsf{I}(r \log p)\,\frac{n}{r} \log n\,(\log\log n)^{3/2} \log\log\log n\,\mathrm{e}^{2\sqrt{\log 2 \log\log n}}\Big) \\ &= \mathcal{O}\Big(n \log n \log p \log\,(\log n + \log p)(\log\log\,(n + \log p))^{\mathcal{O}(1)}\,\mathrm{e}^{2\sqrt{\log 2 \log\log n}}\Big) \end{aligned}$$

In particular,

$$\frac{n}{r}\mathsf{S}_p(r) = \mathcal{O}\Big(\mathsf{S}_{p^r}\Big(\Big\lceil \frac{n}{r} \Big\rceil\Big)\Big),$$

which completes the proof of the theorem. □

Remark 14. The best previously known bound for relaxed multiplication in $\mathbb{Z}_p$ was

$$\begin{aligned} \mathsf{R}_p(n) &= \mathcal{O}(\mathsf{I}(n\,(\log p + \log n))\log n) \\ &= \begin{cases} \mathcal{O}^\flat(n \log^3 n) & \text{if } p = \mathcal{O}(n) \\ \mathcal{O}^\flat(n \log^2 n \log p \log\log p) & \text{if } n = \mathcal{O}(p) \end{cases} \end{aligned}$$

We thus improved the previous bound by a factor $\log n$ at least, up to sublogarithmic terms.

8. Final remarks

For the moment, we have not implemented any of the new algorithms in practice. Nevertheless, our old implementation of the algorithm from [Hoe07] allowed us to gain some insight on the practical usefulness of blockwise relaxed multiplication. Let us briefly discuss the potential impact of the new results for practical purposes.

Characteristic zero. In characteristic zero, our focus on algebraic complexity makes the complexity bounds more or less irrelevant from a practical point of view. In practice, two cases are of particular interest: floating point coefficients (which were already considered in [Hoe07]) and integer coefficients (rational coefficients can be dealt with similarly after multiplying by the common denominator).

In the case of integer coefficients, it is best to re-encode the integers as polynomials in $\mathbb{F}_p[x]$ for a prime number which fits into a machine word and such that $\mathbb{F}_p$ admits many 2^k-th roots of unity (it is also possible to take several primes p and use Chinese remaindering). After that, one may again use the old algorithm from [Hoe07]. Also, integer coefficients usually grow in size with n, so one really should see the power series as a bivariate power series in $\mathbb{F}_p[[x, z]]$ with a triangular support. One may then want to use TFT-style multiplication [Hoe04, Hoe05] in order to gain another constant factor.

Finite fields. For large finite fields, it is easy to find large geometric progressions, so the algorithms of this paper can be applied without the need to consider field extensions. Moreover, for finite fields of the form $\mathbb{F}_{p^k}$ with p sufficiently large and $k>1$, it is possible to choose $q\in\mathbb{F}_p$, thereby speeding up evaluation and interpolation. For small finite fields of the form $\mathbb{F}_{p^k}$, it is generally necessary to make the *initial investment* of working in a larger ring with sufficiently large geometric progressions. Of course, instead of the ring extensions considered in Section 5, we may directly use field extensions of the form $\mathbb{F}_{p^l}$ with $l>k$.

Semi-relaxed multiplication. In principle, in the semi-relaxed case, it is possible to gain a factor 2 with respect to the fully relaxed case, using the technique from [Hoe03]. Unfortunately, the middle product is not always easy to implement. For instance, if we rely on Kronecker substitution for multiplications in $\mathbb{Z}[z]$, then we will need to implement an *ad hoc* analogue for the middle product. Since we did not use a fast algorithm for middle products for our benchmarks in [Hoe07, Section 5], the timings for the semi-relaxed product in were only about 25% instead of 50% better than the timings for the fully relaxed product. Nevertheless, modulo increased implementation efforts, we stress that a 50% gain should be achievable.

Cache friendliness. So far, we have not investigated the cache friendliness of blockwise relaxed multiplication, and it can be feared that a lot of additional work is required in order to make our algorithms really efficient from this point of view.

Bilinear maps. In order to keep the presentation reasonably simple, we have focussed on the case when $\mathbb{A}$ is an effective ring. In fact, a more general setting for relaxed multiplication is to consider a bilinear mapping $\mu\colon \mathbb{M}_1\times\mathbb{M}_2\to\mathbb{M}_3$, where $\mathbb{M}_1$, $\mathbb{M}_2$ and $\mathbb{M}_3$ are effective $\mathbb{A}$-modules, and extend it into a mapping $\hat{\mu}\colon \mathbb{M}_1[[z]]\times\mathbb{M}_2[[z]]\to\mathbb{M}_3[[z]]$ by $\hat{\mu}(f,\ g)=\sum_{i,j}\mu(f_i,\ g_j)\ z^{i+j}$. Under suitable hypothesis, the algorithms in this paper generalize to this setting.

Skew series. The relaxed approach can also be generalized to the case when the coefficients of the power series are operators which commute with monomials z^i in a non trivial way. More precisely, assume that we have an effective ring homomorphism $\phi\colon\mathbb{A}\to\mathbb{A}$ such that $z\,a=(\phi a)\,z$ for all $a\in\mathbb{K}$. For instance, one may take $\mathbb{A}=\mathbb{Q}[\delta]$ with $\delta=z\,\partial/\partial z$, so that $P(\delta)\,z=P(\delta+1)\,z$. Given a commutation rule of this kind, we define a skew multiplication on $\mathbb{A}[[z]]$ by

$$\left[\sum_{i\geqslant 0} f_i\,z^i\right]\left[\sum_{j\geqslant 0} g_j\,z^j\right] \;=\; \sum_{i,j\geqslant 0} f_i\,(\phi^i\,g_j)\,z^{i+j}.$$

If $\mathsf{M}_{\mathbb{A}}(n)$ denotes the cost of multiplying two polynomials $P\,z^i$ and $Q\ z^j$ in $\mathbb{A}[z]$ with $\deg P$, $\deg Q<n$, then the classical fast relaxed multiplication algorithm from [Hoe02, FS74] generalizes and still admits the time complexity $\mathcal{O}(\mathsf{M}_{\mathbb{A}}(n)\ \log\ n)$. However, the blockwise algorithm from this paper does not generalize to this setting, at least not in a straightforward way.

BIBLIOGRAPHY

[AHU74] A. Aho, J. Hopcroft, and J. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, Reading, Massachusetts, 1974.

[Ber00] D. Bernstein. Removing redundancy in high precision Newton iteration. Available from `http://cr.yp.to/fastnewton.html`, 2000.

[BHL11] J. Berthomieu, J. van der Hoeven, and G. Lecerf. Relaxed algorithms for p-adic numbers. *Journal de Théorie des Nombres de Bordeaux*, 23(3):541–577, 2011.

[BK78] R.P. Brent and H.T. Kung. Fast algorithms for manipulating formal power series. *Journal of the ACM*, 25:581–595, 1978.

[BL11] J. Berthomieu and R. Lebreton. Relaxed p-adic hensel lifting for algebraic systems. Work in preparation, 2011.

[Blu70] L.I. Bluestein. A linear filtering approach to the computation of the discrete fourier transform. *IEEE Trans. Electroacoustics*, AU-18:451–455, 1970.

[BP94] D. Bini and V.Y. Pan. *Polynomial and matrix computations. Vol. 1*. Birkhäuser Boston Inc., Boston, MA, 1994. Fundamental algorithms.

[BS05] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *Journal of Complexity*, 21(4):420–446, August 2005. Festschrift for the 70th Birthday of Arnold Schönhage.

[CK91] D.G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991.

[Coo66] S.A. Cook. *On the minimum computation time of functions*. PhD thesis, Harvard University, 1966.

[CT65] J.W. Cooley and J.W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Computat.*, 19:297–301, 1965.

[FS74] M.J. Fischer and L.J. Stockmeyer. Fast on-line integer multiplication. *Proc. 5th ACM Symposium on Theory of Computing*, 9:67–72, 1974.

[Für07] M. Fürer. Faster integer multiplication. In *Proceedings of the Thirty-Ninth ACM Symposium on Theory of Computing (STOC 2007)*, pages 57–66, San Diego, California, 2007.

[GG02] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2-nd edition, 2002.

[Hoe97] J. van der Hoeven. Lazy multiplication of formal power series. In W. W. Küchlin, editor, *Proc. ISSAC '97*, pages 17–20, Maui, Hawaii, July 1997.

[Hoe02] J. van der Hoeven. Relax, but don't be too lazy. *JSC*, 34:479–542, 2002.

[Hoe03] J. van der Hoeven. Relaxed multiplication using the middle product. In Manuel Bronstein, editor, *Proc. ISSAC '03*, pages 143–147, Philadelphia, USA, August 2003.

[Hoe04] J. van der Hoeven. The truncated Fourier transform and applications. In J. Gutierrez, editor, *Proc. ISSAC 2004*, pages 290–296, Univ. of Cantabria, Santander, Spain, July 4–7 2004.

[Hoe05] J. van der Hoeven. Notes on the Truncated Fourier Transform. Technical Report 2005-5, Université Paris-Sud, Orsay, France, 2005.

[Hoe07] J. van der Hoeven. New algorithms for relaxed multiplication. *JSC*, 42(8):792–802, 2007.

[Hoe09] J. van der Hoeven. Relaxed resolution of implicit equations. Technical report, HAL, 2009. `http://hal.archives-ouvertes.fr/hal-00441977/fr/`.

[Hoe10] J. van der Hoeven. Newton's method and FFT trading. *JSC*, 45(8):857–878, 2010.

[Hoe11] J. van der Hoeven. From implicit to recursive equations. Technical report, HAL, 2011. `http://hal.archives-ouvertes.fr/hal-00583125/fr/`.

[KO63] A. Karatsuba and J. Ofman. Multiplication of multidigit numbers on automata. *Soviet Physics Doklady*, 7:595–596, 1963.

[RSR69] L.R. Rabiner, R.W. Schafer, and C.M. Rader. The chirp z-transform algorithm and its application. *Bell System Tech. J.*, 48:1249–1292, 1969.

[Sed01] Alexandre Sedoglavic. *Méthodes seminumériques en algèbre différentielle ; applications à l'étude des propriétés structurelles de systèmes différentiels algébriques en automatique*. PhD thesis, École polytechnique, 2001.

[SS71] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. *Computing*, 7:281–292, 1971.

[Too63] A.L. Toom. The complexity of a scheme of functional elements realizing the multiplication of integers. *Soviet Mathematics*, 4(2):714–716, 1963.