

Generic reductions for in-place polynomial multiplication

Pascal Giorgi* Bruno Grenet* Daniel S. Roche†

May 13, 2019

Abstract

The polynomial multiplication problem has attracted considerable attention since the early days of computer algebra, and several algorithms have been designed to achieve the best possible time complexity. More recently, efforts have been made to improve the space complexity, developing modified versions of a few specific algorithms to use no extra space while keeping the same asymptotic running time.

In this work, we broaden the scope in two regards. First, we ask whether an arbitrary multiplication algorithm can be performed in-place generically. Second, we consider two important variants which produce only part of the result (and hence have less space to work with), the so-called middle and short products, and ask whether these operations can also be performed in-place.

To answer both questions in (mostly) the affirmative, we provide a series of reductions starting with any linear-space multiplication algorithm. For full and short product algorithms these reductions yield in-place versions with the same asymptotic time complexity as the out-of-place version. For the middle product, the reduction incurs an extra logarithmic factor in the time complexity only when the algorithm is quasi-linear.

Keywords— arithmetic, polynomial multiplication, in-place algorithm, self reduction

1 Introduction

1.1 Polynomial multiplication

Polynomial multiplication is a fundamental problem in mathematical algorithms. It forms the basis (and key bottleneck) for other fundamental problems such as division with remainder, GCD computation, evaluation/interpolation, resultants, factorization, and structured linear algebra (see, e.g., [10, §8–15] and [3, §2–7,10,12]).

As such, significant effort has gone to improving the time to multiply two size- n polynomials, mainly following ideas from integers multiplication most notably with Karatsuba’s algorithm [17], Toom-Cook multiplication [9], and Schönhage-Strassen [22]. More recent results have improved the complexity further with dedicated polynomial approaches over finite field [13, 15] or over any algebra [7], yielding the best known quasi-linear time algorithms.

*LIRMM, Université de Montpellier, CNRS, Montpellier, France. `firstname.lastname@lirmm.fr`

†United States Naval Academy, Annapolis, Maryland, USA. `roche@usna.edu`

1.2 Space complexity

After minimizing the runtime, an important question both in theory and in practice is how much extra *space* these algorithms require. While the classical algorithm can be made to use only a constant number of temporary values, all the faster algorithms mentioned above require $O(n)$ space to multiply two size- n polynomials. In fact, proven time-space trade-offs in the algebraic circuit and branching program models indicate that space at least polynomial in n is required for any sub-quadratic multiplication algorithm [21, 1].

But in a model where the output space admits both random writes *and* reads, these time-space lower bounds can be broken. [20] developed a variant of Karatsuba’s algorithm using only $O(\log n)$ space. Later, an FFT-based multiplication algorithm using $O(n \log n)$ time and constant space was developed for the case that the coefficient ring contains a suitable root of unity [14]. Space-saving versions of Karatsuba’s algorithm can also be found in [24, 6, 23, 8].

1.3 Short and middle products

Besides the usual *full product* computation, two variants have also been extensively studied: the *short product* which truncates the output to the first n terms, and the *middle product* which truncates the result on both ends. These variants are important especially for power series, and specific variants of Karatsuba’s algorithm and others have been developed, usually gaining a constant factor compared to a full product followed by a truncation [19, 12, 11]. For space efficiency, dedicated algorithms are even mandatory as using full product increases the output size whence the space complexity.

In some specific model of straight line programs, [4] shows that transposing full product yields a middle product with the same time and space complexities. Nevertheless, such a result does not directly apply to the more general context of algorithms, since their model allows computations with fixed size inputs only.

1.4 Our work

In this paper, we develop reductions which transform any multiplication algorithm which uses $O(n)$ extra space into full, short, and middle product algorithms which use only $O(1)$ extra space. For this we provide two kinds of reductions: time-space preserving reductions that link the different product variants; and self-reductions for each variant that reduce space complexity to $O(1)$. The time complexity for full and short product is the same as that of the original, while that for middle product incurs an additional $\log n$ factor. This improves the $O(\log n)$ space of the most space-constrained Karatsuba algorithm [20], and implies for the first time: in-place versions of Toom-Cook multiplication; in-place FFT-based multiplication even when the ring does not contain a root of unity; in-place subquadratic short and middle product algorithms.

We begin by carefully stating our space complexity model and then defining the multiplications problems in Sections 2 and 3. Our time-space preserving reductions and some equivalences are presented next in Section 4, followed by the critical self-reductions in Section 5 which prove our main results.

2 Complexity model

We use the model of an *algebraic-RAM* that is equipped with two kinds of registers: the standard registers store integers as in the classical Word-RAM model, whereas the *algebraic registers* store elements from

the base field \mathbb{K} of coefficients. As in Word-RAM, we assume that the standard registers can store integers of size $O(\log n)$ where n is the number of coefficients in the inputs.

Word-RAM machines are a classical model in computational complexity, in particular for *fine-grained complexity* that classifies the difficulty of polynomial-time problems [25]. We use it in order to distinguish between the space needed to store indices (that is thus hidden in the standard registers) from the space needed to store elements from the base field.

Time complexity As mentioned, we use the number of arithmetic operations as the time complexity measure since the cost of the operations on indices is negligible with respect to arithmetic operations. Formally, we assume that any ring operation on the algebraic registers has cost 1.

Space complexity We divide the registers into three categories: the input space is made of the (algebraic) registers that store the inputs, the output space is made of the (algebraic) registers where the output must be written, and the work space is made of (algebraic and non-algebraic) registers that are used as extra space during the computation. The space complexity is then the maximum number of work registers used simultaneously during the computation. An algorithm is said to be “in-place” if its space complexity is $O(1)$, and “out-of-place” otherwise.

One can then distinguish different models depending on the read/write permissions on the input and output registers:

1. Input space is read-only, output space write-only;
2. Input space is read-only, output space is read/write;
3. Input and output spaces are both read/write.

The first model is the classical one from complexity theory [2]. Despite its theoretical interest, it does not reflect low-level computation where output is typically in some DRAM or Flash memory on which reading is no more costly than writing. Furthermore, polynomial multiplication here has a quadratic lower bound for time \times space [1], limiting the possibility for meaningful improvements.

The second model has been used in the context of in-place polynomial multiplication [20, 14]. This is a very reasonable model since it matches the paradigm of parallel computing with shared memory. This is the model in which we develop our algorithms.

The third model has been used in the context of straight line programs to provide a generic approach for preserving memory in Tellegen’s transposition principle [4]. However, this memory model allows to erase the input registers, which can be problematic especially for recursive algorithms that re-use their operands. Hence, we will not use this too-permissive model.

Notation The output space in our algorithms is denoted by R and registers are indexed from 0 to $n-1$. We write $R_{[k..\ell]}$ to denote the registers of indices k to $\ell-1$.

3 Polynomial multiplications

Define the *size* of a univariate polynomial as the number of coefficients in its (dense) representation; a polynomial of size n has degree at most $n-1$. Importantly, we allow zero padding: a size- n polynomial could have degree strictly less than $n-1$; the size indicates only how it is represented.

Let $f = \sum_{i=0}^{n-1} f_i X^i$ and $g = \sum_{i=0}^{n-1} g_i X^i$ be two size- n polynomials. Their product $h = fg$ is a polynomial of size $2n - 1$, what we call a *balanced full product*. More generally, if f has size m and g has size n , their product has size $m + n - 1$. We call this case the *unbalanced full product* of f and g .

We now define precisely the short product, middle product, and half-additive full product.

Definition 3.1. Let f and g be two size- n polynomials. Their low short product is the size- n polynomial defined as

$$\text{SP}_{\text{lo}}(f, g) = (f \cdot g) \bmod X^n$$

and their high short product is the size- $(n - 1)$ polynomial defined as

$$\text{SP}_{\text{hi}}(f, g) = (f \cdot g) \text{quo} X^n.$$

The low short product is actually the meaningful notion of product for truncated power series. Note also that the definition of the high short product that we use implies that the result does not depend on all the coefficients of f and g . The rationale for this choice is to have the identity $fg = \text{SP}_{\text{lo}}(f, g) + X^n \text{SP}_{\text{hi}}(f, g)$.

Definition 3.2. Let f and g be two polynomials sizes $n + m - 1$ and n , respectively. Their middle product is the size- m polynomial made of the central coefficients of the product fg , that is

$$\text{MP}(f, g) = ((f \cdot g) \text{quo} X^{n-1}) \bmod X^m.$$

If $f = \sum_{i < n+m-1} f_i X^i$ and $g = \sum_{j < n} g_j X^j$, then

$$\text{MP}(f, g) = \sum_{n-1 \leq i+j < n+m-1} f_i g_j X^{i+j-n+1}.$$

The middle product, most commonly in the special case $n = m$, arises naturally in several algorithms manipulating polynomials or power series which are based on Newton's iteration, such as division or square root [11].

The most efficient algorithms for middle product are obtained by *Tellegen's transposition principle* from full product algorithms [11, 4], saving constant factor in time complexity compared to the naive approach. It is not yet known if this transposition or similar one can preserve space complexity. This question has been established as an open problem in [16], and only answered partially in [4] in a restricted model.

Definition 3.3. Let f and g be two size- n polynomials, and h be a size- $(n - 1)$ polynomial. The (low-order) half-additive full product of f and g given h is $\text{FP}_{\text{lo}}^+(f, g, h) = h + fg$. Similarly, their high-order half-additive full product is $\text{FP}_{\text{hi}}^+(f, g, h) = X^n h + fg$. An in-place half-additive full product algorithm is an algorithm computing a half-additive full product with h initially stored in the output space.

This variant of the full product which has a partially-initialized output space will be useful to derive other in-place algorithms.

3.1 Multiplications as linear maps

For ease of explanation, we will use the linear property of polynomial multiplications when an operand is fixed. Let $f = \sum_{i=0}^{n-1} f_i X^i$ and $g = \sum_{i=0}^{n-1} g_i X^i$ be two size- n polynomials. If f is fixed, the product

$h = fg$ can be described as a linear map from \mathbb{K}^n to \mathbb{K}^{2n-1} . The matrix, denoted $\mathfrak{M}_{\text{FP}(f)}$, for this map is to a Toeplitz matrix built from the coefficients of f , and the product fg corresponds to the following matrix-vector product:

$$\underbrace{\begin{pmatrix} f_0 & & & \\ \vdots & \ddots & & \\ f_{n-1} & & f_0 & \\ & \ddots & \vdots & \\ & & & f_{n-1} \end{pmatrix}}_{\mathfrak{M}_{\text{FP}(f)}} \times \underbrace{\begin{pmatrix} g_0 \\ \vdots \\ g_{n-1} \end{pmatrix}}_{\vec{g}} = \underbrace{\begin{pmatrix} h_0 \\ \vdots \\ h_{2n-1} \end{pmatrix}}_{\vec{h}} \quad (1)$$

where $\mathfrak{M}_{\text{FP}(f)} \in \mathbb{K}^{(2n-1) \times n}$, $\vec{g} \in \mathbb{K}^n$ and $\vec{h} \in \mathbb{K}^{2n-1}$.

The low and high short products being defined as part of the result of the full product, their corresponding linear maps are from \mathbb{K}^n to \mathbb{K}^n and from \mathbb{K}^n to \mathbb{K}^{n-1} respectively, given by submatrices of $\mathfrak{M}_{\text{FP}(f)}$ as follows:

$$\underbrace{\begin{pmatrix} f_0 & & & \\ \vdots & \ddots & & \\ f_{n-1} & \cdots & f_0 \end{pmatrix}}_{\mathfrak{M}_{\text{SP}_{\text{lo}}(f)}} \quad \underbrace{\begin{pmatrix} 0 & f_{n-1} & \cdots & f_1 \\ \vdots & & \ddots & \vdots \\ 0 & & & f_{n-1} \end{pmatrix}}_{\mathfrak{M}_{\text{SP}_{\text{hi}}(f)}} \quad (2)$$

Lastly, the middle product corresponds to a linear map from \mathbb{K}^n to \mathbb{K}^m with larger operand fixed, given by the $m \times n$ Toeplitz matrix

$$\underbrace{\begin{pmatrix} f_{n-1} & f_{n-2} & \cdots & f_1 & f_0 \\ f_n & f_{n-1} & & f_2 & f_1 \\ \vdots & \vdots & & \vdots & \vdots \\ f_{n+m-2} & f_{n+m-3} & \cdots & f_{m-2} & f_{m-1} \end{pmatrix}}_{\mathfrak{M}_{\text{MP}(f)}} \quad (3)$$

4 Time-space preserving reductions

In this section, we compare the relative difficulties of the full product, the half-additive full product, the low and high short products, and the balanced middle product, in the framework of time and space efficient algorithms. To this end, we define a notion of *time and space preserving reduction* between problems.

We say that a problem A is TISP-reducible to a problem B if, given an algorithm for B that has time complexity $t(n)$ and space complexity $s(n)$, one can deduce an algorithm for A that has time complexity $O(t(n))$ and space complexity $s(n) + O(1)$. We write $A \leq_{\text{TISP}} B$ if A is TISP-reducible to B and $A \equiv_{\text{TISP}} B$ if both $A \leq_{\text{TISP}} B$ and $B \leq_{\text{TISP}} A$. Note that the TISP-reduction is transitive. The reduction we use can be defined using oracles and is an adaptation of the notion of *fine-grained reduction* [25, Definition 2.1] adapted to time-space fine-grained complexity classes [18]. In this section, MP denotes the balanced middle product.

Theorem 4.1. *Half-additive full products and short products are equivalent under TISP-reductions, that is*

$$\text{FP}_{\text{hi}}^+ \equiv_{\text{TISP}} \text{FP}_{\text{lo}}^+ \equiv_{\text{TISP}} \text{SP}_{\text{hi}} \equiv_{\text{TISP}} \text{SP}_{\text{lo}}.$$

Furthermore, if SP denotes either SP_{lo} or SP_{hi} ,

$$FP \leq_{TISP} SP \leq_{TISP} MP.$$

Proof. The equivalences $SP_{hi} \equiv_{TISP} SP_{lo}$ and $FP_{hi}^+ \equiv_{TISP} FP_{lo}^+$ are proved below in Lemmas 4.3 and 4.4. The equivalence $SP \equiv FP^+$ (where SP denotes any of SP_{lo} and SP_{hi} , and FP^+ any of FP_{lo}^+ and FP_{hi}^+) is proved in Section 4.2. The reduction $FP \leq_{TISP} SP$ simply amounts to the identity $FP(f, g) = SP_{lo}(f, g) + X^n SP_{hi}(f, g)$. The reduction $SP \leq_{TISP} MP$ follows from the following equalities where 0_n denotes the zero polynomial stored in size n :

$$\begin{aligned} SP_{lo}(f, g) &= MP(0_{n-1} + X^{n-1}f, g), \text{ and} \\ SP_{hi}(f, g) &= MP((f \text{ quo } X) + X^{n-1}0_{n-2}, g \text{ quo } X). \end{aligned}$$

Hence, one can compute the full product, the low and high short products of f and g simply by calling a middle product algorithm on f padded with zeroes and g . In our model of read-only inputs, an actual padding is not required. It is sufficient to use some kind of *fake padding* where the data structure storing f is responsible for returning 0 when needed. \square

The relative order of difficulty $FP \leq_{TISP} SP \leq_{TISP} MP$ makes intuitive sense based on the size of the output compared to the size of the inputs since the output can be used as work space: The full product maps $2n$ coefficients to $2n - 1$ coefficients, the short products map $2n$ coefficients to n coefficients and the middle product maps $3n$ coefficients to n coefficients. In Section 5, we shall give a partial converse to $SP \leq_{TISP} MP$: There exists a reduction from SP to MP which preserves space and either maintains the asymptotic complexity or increases it by a logarithmic factor.

4.1 Equivalences based on reverse polynomials

Definition 4.2. *The size- n reversal of a polynomial f is defined as $rev_n(f) = X^{n-1}f(1/X)$.*

We note that any algorithm whose input is a size- n polynomial f can be turned into a new algorithm that computes the same function with input $rev_n(f)$, simply by replacing a query to any coefficient with index i with one of index $n - i$, not affecting the number of ring operations. Let us now prove that $SP_{hi} \equiv_{TISP} SP_{lo}$.

Lemma 4.3. *Let f and g be two size- n polynomials. Then we get*

$$SP_{hi}(f, g) = rev_{n-1}(SP_{lo}(rev_{n-1}(f \text{ quo } X), rev_{n-1}(g \text{ quo } X))).$$

Proof. Let $\tilde{f} = rev_{n-1}(f \text{ quo } X)$ and $\tilde{g} = rev_{n-1}(g \text{ quo } X)$. Then

$$SP_{lo}(\tilde{f}, \tilde{g}) = \sum_{\substack{0 \leq i, j < n-1 \\ i+j < n-1}} f_{n-1-i} g_{n-1-j} X^{i+j},$$

whence

$$rev_{n-1}(SP_{lo}(\tilde{f}, \tilde{g})) = \sum_{\substack{0 \leq i, j < n-1 \\ i+j < n-1}} f_{n-1-i} g_{n-1-j} X^{n-2-(i+j)}.$$

One can change the indices of summation using $k = n - 1 - i$ and $\ell = n - 1 - j$. Then $n - 2 - (i + j) = k + \ell - n$ and the indices i and j such that $0 \leq i + j < n - 1$ are mapped to indices k and ℓ such that $2n - 1 > k + \ell \geq n$. In other words,

$$\text{rev}_{n-1}(\text{SP}_{\text{lo}}(\tilde{f}, \tilde{g})) = \sum_{\substack{0 < k, \ell \leq n-1 \\ n \leq k+\ell < 2n-1}} f_k g_\ell X^{k+\ell-n} = \text{SP}_{\text{hi}}(f, g). \quad \square$$

Similarly, we can prove that $\text{FP}_{\text{hi}}^+ \equiv_{\text{TISP}} \text{FP}_{\text{lo}}^+$.

Lemma 4.4. *Let f and g be two size- n polynomials and h be a size- $(n-1)$ polynomial. Then*

$$\text{FP}_{\text{hi}}^+(f, g, h) = \text{rev}_{2n-1}(\text{FP}_{\text{lo}}^+(\text{rev}_n(f), \text{rev}_n(g), \text{rev}_{n-1}(h))).$$

Proof. Let $f^* = \text{rev}_n(f)$, $g^* = \text{rev}_n(g)$ and $h^* = \text{rev}_{n-1}(h)$. First note that $\text{rev}_{2n-1}(h^*) = X^n h$ by definition. Since $\text{rev}_{2n-1}(f^* g^*) = \text{rev}_n(f^*) \text{rev}_n(g^*)$ we get $\text{rev}_{2n-1}(f^* g^* + h^*) = \text{rev}_n(f^*) \text{rev}_n(g^*) + \text{rev}_{2n-1}(h^*) = f g + X^n h = \text{FP}_{\text{hi}}^+(f, g, h)$. \square

4.2 Equivalence between short products and half-additive full products

Reduction from SP to FP^+ Let f and g be two size- n polynomials and h be a size- $(n-1)$ polynomial. The half-additive full product $\text{FP}_{\text{lo}}^+(f, g, h)$ equals $h + f g$. Note that $f g = \text{SP}_{\text{lo}}(f, g) + X^n \text{SP}_{\text{hi}}(f, g)$. This already proves that the non-additive full product can be computed using algorithms for low and high short products. For the half-additive full products, it is sufficient to store an intermediate result in the free registers of the output space.

Assuming $R_{[0..n-1]}$ holds the value of h , the following instructions reduce the computation of the half-additive full product $\text{FP}_{\text{lo}}^+(f, g, h)$ to two short products plus $(n-1)$ additions.

- 1: $R_{[n-1..2n-1]} \leftarrow \text{SP}_{\text{lo}}(f, g)$
- 2: $R_{[0..n-1]} \leftarrow R_{[0..n-1]} + R_{[n-1..2n-2]}$
- 3: $R_{n-1} \leftarrow R_{2n-1}$
- 4: $R_{[n..2n-1]} \leftarrow \text{SP}_{\text{hi}}(f, g)$

Reduction from FP^+ to SP Let f and g be two size- n polynomials. Splitting f and g by half such that $f = f_0 + X^{\lceil n/2 \rceil} f_1$ and $g = g_0 + X^{\lceil n/2 \rceil} g_1$, we have

$$\text{SP}_{\text{lo}}(f, g) = f_0 g_0 + X^{\lceil n/2 \rceil} (f_0 g_1 + f_1 g_0) \bmod X^n.$$

What is needed is the full product of f_0 and g_0 , and the low short products of f_0 and g_1 , and f_1 and g_0 . Actually, since f_0 is larger than g_1 when n is odd (and g_0 larger than f_1), one only needs the short products $\text{SP}_{\text{lo}}(f_0^-, g_1)$ and $\text{SP}(f_1, g_0^-)$ where $f_0^- = f \bmod X^{\lceil n/2 \rceil}$ and $g_0^- = g \bmod X^{\lceil n/2 \rceil}$.

To avoid any recursive call that would imply storing a call stack, we can actually use full products instead of short products: We first compute $f_0^- g_1 + f_1 g_0^-$ using a full product and a half-additive full product. Then we can forget about the higher order terms, and add $f_0 g_0$ to this sum using a second half-additive full product. The following instructions summarize this approach:

- 1: $R_{[0..2\lceil n/2 \rceil-1]} \leftarrow \text{FP}(f_0^-, g_1)$ ▷ half-additivity not needed
- 2: $R_{[0..2\lceil n/2 \rceil-1]} \leftarrow \text{FP}_{\text{lo}}^+(f_1, g_0^-, f_0^- g_1 \bmod X^{\lceil n/2 \rceil-1})$
- 3: $R_{[\lceil n/2 \rceil..n]} \leftarrow R_{[0..\lceil n/2 \rceil]}$ ▷ keep lower part of $f_0^- g_1 + f_1 g_0^-$
- 4: $R_{[0..2\lceil n/2 \rceil-1]} \leftarrow \text{FP}_{\text{hi}}^+(f_0, g_0, (f_0^- g_1 + f_1 g_0^-) \bmod X^{\lceil n/2 \rceil-1})$

The correctness is clear. The complexity of the algorithm is the cost of three full products in size approximately $n/2$: One non-additive full product in size $\lfloor n/2 \rfloor$ and two half-additive full products in size $\lfloor n/2 \rfloor$ and $\lceil n/2 \rceil$, respectively.

As direct consequence of Lemmas 4.3 and 4.4, one obtains the same reductions to SP_{hi} and from FP_{lo}^+ or FP_{hi}^+ .

4.3 From half-additive full product to unbalanced full product

The unbalanced full product can be computed using any algorithm for the (balanced) full product. Nevertheless, the space complexity increases since intermediate results must be stored. Given an algorithm for the balanced full product of space complexity $s(n)$, one obtains an algorithm with space complexity $s(n) + (n - 1)$ for the unbalanced full product. In this section, we prove that if the original full product algorithm is *half-additive*, the resulting unbalanced full product algorithm has the same space complexity.

Let f be a size- m polynomial and g be a size- n polynomial with $m > n$. Write $f = \sum_{k=0}^{\lceil m/n \rceil - 1} X^{kn} f_k$, where each sub-polynomial $f_0, \dots, f_{\lceil m/n \rceil - 1}$ has size at most n . The computation of $f \cdot g$ reduces to the computations of each $f_k \cdot g$. The following instructions prove that using half-additivity, the intermediate results $f_k \cdot g$ can be computed directly in the output space.

- 1: $\text{R}_{\lfloor \lceil m/n \rceil - 1 \rfloor n .. m+n} \leftarrow \text{FP}(f_{\lceil m/n \rceil - 1}, g)$
- 2: **for** k **from** $\lceil m/n \rceil - 2$ **down to** 0 **do**
- 3: $\text{R}_{\lfloor kn .. (k+2)n - 1 \rfloor} \leftarrow \text{FP}_{\text{hi}}^+(f_k, g, f_{k+1}g \bmod X^n)$

Note that at step 1, the size of $f_{\lceil m/n \rceil - 1}$ can be strictly smaller than n . To ensure a balanced product, fake padding can be done on this input and the output can be placed anywhere in the free output space. Thus, the time complexity is at most $\lceil m/n \rceil M(n)$ where $M(n)$ is the complexity of the half-additive full product.

5 In-place algorithms from out-of-place algorithms

In this section, we show how to obtain in-place algorithms from out-of-place algorithms. The theorem below summarizes the main results described in this section.

Theorem 5.1.

- (i) Given a full product algorithm with time complexity $M(n)$ and space complexity $\leq cn$, one can build an in-place algorithm for the half-additive full product with time complexity $\leq (2c + 7)M(n) + o(M(n))$.
- (ii) Given a (low or high) short product algorithm with time complexity $M(n)$ and space complexity $\leq cn$, one can build an in-place algorithm for the same problem with time complexity $\leq (2c + 5)M(n) + o(M(n))$.
- (iii) Given a middle product algorithm with time complexity $M(n)$ and space complexity $\leq cn$, one can build an in-place algorithm for the same problem with time complexity $\leq M(n) \log_{\frac{c+1}{c+2}}(n) + O(M(n))$ if $M(n)$ is quasi-linear, and $O(M(n))$ otherwise.

Actually, our reductions work for any space bound $s(n) \leq O(n)$. Smaller space bounds yield better time bounds though we do not have a general expression in terms of $s(n)$. Yet sublinear space bounds still imply an increase of the time complexity by a multiplicative constant for full and short products.




Figure 1: Tilings of the matrices $\mathfrak{M}_{\text{FP}(f)}$ (left), $\mathfrak{M}_{\text{SP}_0(f)}$ (center) and $\mathfrak{M}_{\text{MP}(f)}$ (right).

Formally, we give self-reductions for the three problems. That is, we use an out-of-place algorithm for the problem as building block of our in-place version. The general idea is similar in the three cases. In a first step, we use the out-of-place algorithm to compute some part of the output, using the unused output space as temporary work space. Then a recursive call finishes the work. The (constant) amount of space needed in our in-place algorithms correspond the space needed to process the base cases. As depicted in section 3.1, we aim to apply some specific matrix to a vector. Our general construction consists in first applying the top or bottom rows of the matrix to the vector using the out-of-place algorithm, and applying the remaining rows using a recursive call (*cf.* Fig. 1). Similar techniques have been already used for in-place fast matrix multiplication [5].

In the cases of full and short products, the diamond and triangular shapes of the corresponding matrices imply that the recursive call is made on two smaller inputs: For instance, to apply the first rows of a lower triangular matrix to a vector, one only needs to apply it to the first entries of the vector. For the middle product, the square shape imply that one input remains of the same size in the recursive call. This difference explains the difference in the time complexities in Theorem 5.1.

Our complexity analyses are based on the following Proposition which is proved in Section 5.4.

Proposition 5.2. *Let $T, M : \mathbb{N} \rightarrow \mathbb{R}_+$ such that M is super-linear.*

- (i) *If for all n , $T(n) \leq \sum_k a_k M(\lfloor \lambda_k n + \mu_k \rfloor) + bn + c + T(\lfloor \alpha n + \beta \rfloor)$ where $a_k, \lambda_k, \mu_k, b, c, \alpha, \beta \in \mathbb{R}_+^*$ with $\alpha < 1$ and $\lambda_k n + \mu_k < n$ for all k . Then*

$$T(n) \leq \sum_k \left(\frac{a_k \lambda_k}{1 - \alpha} \right) M(n) + \frac{bn}{1 - \alpha} + o(M(n)).$$

- (ii) *If for all $m \leq n$, $T(m) \leq (\lambda m/n + \mu/(m - \frac{1}{1-\alpha}) + 1)M(n) + T(\lfloor \alpha m + 1 \rfloor)$ where $\lambda, \mu, \alpha \in \mathbb{R}_+^*$ such that $\alpha < 1$. Then for $m = n$,*

$$T(n) \leq M(n) \log_{1/\alpha}(n) + \frac{\lambda + \mu\alpha}{1 - \alpha} M(n) + o(M(n)).$$

- (iii) *If for all $m \leq n$, $T(m) \leq \lambda \left(\frac{n}{(1-\alpha)^{m-1}} + 1 \right) ((1-\alpha)m)^\gamma + T(\lfloor \alpha m + 1 \rfloor)$ where $\alpha, \gamma, \lambda \in \mathbb{R}_+^*$ such that $\alpha < 1$ and $\gamma > 1$. Then for $m = n$, $T(n) \leq O(n^\gamma)$.*

5.1 In-place full product algorithm

Our aim is to build an in-place (low-order) half-additive full product algorithm iFP_{hi}^+ based on an out-of-place full product algorithm oFP that has space complexity cn . That is, we are given two size- n polynomials f and g in the input space and a size- $(n-1)$ polynomial h in the $(n-1)$ low-order registers of the output space \mathbb{R} and we aim to compute $fg + h$ in \mathbb{R} . The algorithm is based on a tiling of the matrix $\mathfrak{M}_{\text{FP}(f)}$ given in Eq. (1), see Fig. 1 (left).

For some $k < n$ to be fixed later, let $f = \hat{f}X^k + f_0$ and $g = \hat{g}X^k + g_0$ where $\deg f_0, \deg g_0 < k$. Then we have

$$h + fg = h + f_0g + \hat{f}\hat{g}X^{2k} + \hat{f}g_0X^k + f_0\hat{g}X^k. \quad (4)$$

Recall that the output \mathbb{R} has size $2n-1$ with its $n-1$ lowest registers containing h . Then Eq. (4) can be evaluated as follows:

- 1: $\mathbb{R}_{[0..n+k-1[} \leftarrow h + f_0g$
- 2: $\mathbb{R}_{[k..n+k-1[} \leftarrow \mathbb{R}_{[k..n+k-1[} + \hat{f}g_0$
- 3: $\mathbb{R}_{[2k..2n[} \leftarrow \mathbb{R}_{[2k..2n[} + \hat{f}\hat{g}$

The first two steps correspond exactly to two *additive* unbalanced full products, that is unbalanced full products that must be added to some already filled output space. One can describe an algorithm oFP_{u}^+ for this task, based on a (standard) full product algorithm oFP : If f has size k and g has size n , $n > k$, we write $g = \sum_{i=0}^{\lceil n/k \rceil - 1} g_i X^{ki}$ with $\deg(g_i) < k$. Then $fg = \sum_i f g_i$: The algorithm computes the $\lceil n/k \rceil$ products $f g_i$ in $2k-1$ extra registers and adds them to the output. If oFP has time complexity $M(n)$ and space complexity cn , the time complexity of oFP_{u}^+ is $\lceil n/k \rceil (M(k) + 2k - 1)$ and its space complexity $(c+2)k-1$.

The last step computes $h + fg$ and corresponds to a half-additive full product on inputs of size $n-k$, since only the $n-k-1$ first registers of $\mathbb{R}_{[2k..2n[}$ are filled: Indeed, $\deg(h + f_0g + \hat{f}\hat{g}X^{2k}) < n+k-1$. This last step is thus a recursive call.

In order to make this algorithm run in place, k must be chosen so that the extra memory needed in the two calls to oFP_{u}^+ fits exactly in the unused part of \mathbb{R} . This is the case when

$$(c+2)k-1 \leq 2n-1-(n+k-1)$$

which gives $k \leq \frac{n+1}{c+3}$. The resulting algorithm is depicted below.

Algorithm 1 iFP_{hi}^+ from oFP

Input: f and g of size n in the input space, h of size $n-1$ in the output space \mathbb{R}

Output: \mathbb{R} contains $fg + h$

Required: Full product algorithm oFP of space complexity $\leq cn$

- 1: **if** $n < c+2$ **then**
 - 2: $\mathbb{R} \leftarrow \mathbb{R} + fg$ ▷ using a naive algorithm
 - 3: **else**
 - 4: $k \leftarrow \lfloor (n+1)/(c+3) \rfloor$
 - 5: $\mathbb{R}_{[0..n+k-1[} \leftarrow \text{oFP}_{\text{u}}^+(f_0, g, h)$ ▷ work space: $\mathbb{R}_{[n+k-1..2n[}$
 - 6: $\mathbb{R}_{[k..n+k-1[} \leftarrow \text{oFP}_{\text{u}}^+(f, g_0, h + f_0g)$ ▷ same work space
 - 7: $\mathbb{R}_{[2k..2n[} \leftarrow \text{iFP}_{\text{hi}}^+(\text{f quo } X^k, \text{g quo } X^k)$
-

Complexity analysis The algorithm uses two calls to oFP_u^+ with inputs of sizes (k, n) and $(n - k, k)$ respectively. The total complexity amounts to $\lceil n/k \rceil M(k) + (\lceil n/k \rceil - 1)M(k) + (2\lceil n/k \rceil - 1)(2k - 1)$ plus a recursive call in size $n - k$. Let $T(n)$ be the complexity of iFP_{hi}^+ , we thus have

$$T(n) = T(n - k) + (2\lceil n/k \rceil - 1)[M(k) + (2k - 1)].$$

Note that k depends upon n , this implies that the analysis must be done without k . Since $k = \lfloor (n + 1)/(c + 3) \rfloor$, $\lceil n/k \rceil \leq c + 4$ for $n \geq (c + 2)(c + 4)$. Therefore,

$$T(n) \leq T\left(\left\lfloor \frac{c+2}{c+3}(n+1) \right\rfloor\right) + (2c + 7)\left[M\left(\left\lfloor \frac{n+1}{c+3} \right\rfloor\right) + 2\frac{n}{c+3} - \frac{c+1}{c+3}\right].$$

Proposition 5.2(i) yields $T(n) \leq (2c + 7)M(n) + o(M(n))$.

5.2 In-place short product algorithm

Our goal is to describe an in-place (low) short product algorithm based on an out-of-place one. It corresponds to tiling $\mathfrak{M}_{\text{SP}_o(f)}$ given in Eq. (2), see Fig. 1 (center). Let $f = \sum_{i=0}^{n-1} f_i X^i$ and $g = \sum_{i=0}^{n-1} g_i X^i$, and let $h = \sum_{i=0}^{n-1} h_i X^i = \text{SP}_o(f, g)$. The idea is to fix some $k < n$ and to have two phases. The first phase corresponds to the bottom k rows of $\mathfrak{M}_{\text{SP}_o(f)}$ and computes h_{n-k} to h_{n-1} using the out-of-place algorithm on smaller polynomials. The second phase corresponds to the top $(n - k)$ rows and is a recursive call to compute h_0 to h_{n-k-1} : Indeed, $h \bmod X^{n-k} = \text{SP}_o(f \bmod X^{n-k}, g \bmod X^{n-k})$.

For the second phase, we remark that the bottom k rows can be tiled by $\lceil n/k \rceil$ lower triangular matrices (denoted $L_0, \dots, L_{\lceil n/k \rceil - 1}$ from the right to the left), and $\lceil n/k \rceil - 1$ upper triangular matrices (denoted $U_0, \dots, U_{\lceil n/k \rceil - 2}$). One can identify the matrices L_i and U_i as matrices of some low and high short products. More precisely, the coefficients that appear in the lower triangular matrix L_i are the coefficients of degree ki to $k(i + 1) - 1$ of f . Thus, $L_i = \mathfrak{M}_{\text{SP}_o(f_{ki, k(i+1)})}$ where $f_{ki, k(i+1)} = \sum_{j=ki}^{k(i+1)-1} f_j X^{j-ki}$. Similarly, $U_i = \mathfrak{M}_{\text{SP}_{hi}(f_{k(i+1), k(i+1)-1})}$. The matrices $L_{\lceil n/k \rceil - 1}$ and $U_{\lceil n/k \rceil - 2}$ must be padded if k does not divide n . Altogether, this proves that this part of the computation reduces to $\lceil n/k \rceil$ low short products and $\lceil n/k \rceil - 1$ high short products, in size k .

In order for this algorithm to actually be in place, k must be small enough. If the out-of-place short product algorithm uses ck extra space, since we also need k free registers to store the intermediate results, k must satisfy $n - k \geq (c + 1)k$, that is $k \leq n/(c + 2)$.

Complexity analysis The algorithm performs $\lceil n/k \rceil$ low short products and $\lceil n/k \rceil - 1$ high short products plus one recursive call in size $n - k$. Let $M(k)$ be the complexity of a low short product algorithm. Then the high short product can be computed in time $M(k - 1)$. Let $T(n)$ be the complexity of the recursive algorithm. Then $T(n) = \lceil n/k \rceil M(k) + (\lceil n/k \rceil - 1)M(k - 1) + (2\lceil n/k \rceil - 1)k + T(n - k)$ (the linear time is for the additions). Since $k = \lfloor n/(c + 2) \rfloor$, $\lceil n/k \rceil \leq c + 3$ for $n \geq (c + 3)(c + 2)$ and $n - k \leq \frac{c+1}{c+2}n + 1$. Thus,

$$T(n) \leq (c + 3)M\left(\left\lfloor \frac{n}{c+2} \right\rfloor\right) + (c + 2)M\left(\left\lfloor \frac{n}{c+2} - 1 \right\rfloor\right) + 2n + T\left(\left\lfloor \frac{c+1}{c+2}n + 1 \right\rfloor\right).$$

Proposition 5.2(i) implies $T(n) \leq (2c + 5)M(n) + o(M(n))$.

5.3 In-place middle product algorithm

To build an in-place (unbalanced) middle product algorithm, we assume that we have an algorithm for the (balanced) middle product that uses cn extra space for inputs of size $2n - 1$ and n respectively.

Algorithm 2 $\text{iSP}_{1o_from_oSP}$

Input: f and g of size n **Output:** R contains $\text{SP}_{1o}(f, g)$ **Required:** Two short product algorithms oSP_{1o} and oSP_{hi} of space complexity $\leq cn$

```
1: if  $n < c + 2$  then
2:    $R \leftarrow \text{SP}_{1o}(f, g)$  ▷ using a naive algorithm
3: else
4:    $k \leftarrow \lfloor n/(c + 2) \rfloor$ 
5:   for  $i = 0$  to  $\lceil n/k \rceil - 1$  do ▷ work space:  $R_{[k..n-k]}$ 
6:      $R_{[0..k]} \leftarrow \text{oSP}_{1o}(f_{ki, k(i+1)}, g_{n-k(i+1), n-ki})$ 
7:      $R_{[n-k..n]} \leftarrow R_{[n-k..n]} + R_{[0..k]}$ 
8:   for  $i = 0$  to  $\lceil n/k \rceil - 2$  do ▷ same work space
9:      $R_{[0..k]} \leftarrow \text{oSP}_{hi}(Xf_{ki+1, k(i+1)-1}, g_{n-k(i+2), n-k(i+1)})$ 
10:     $R_{[n-k..n]} \leftarrow R_{[n-k..n]} + R_{[0..k]}$ 
11:   $R_{[0..n-k]} \leftarrow \text{iSP}_{1o\_from\_oSP}(f \bmod X^{n-k}, g \bmod X^{n-k})$ 
```

The in-place algorithm is based on a tiling of $\mathfrak{M}_{\text{MP}(f)}$ given in Eq. (3), see Fig. 1 (right) : The top k rows correspond to the matrix $\mathfrak{M}_{\text{MP}(f \bmod X^k)}$ and the bottom $m - k$ rows to the matrix $\mathfrak{M}_{\text{MP}(f \text{ quo } X^k)}$. The algorithm computes $\mathfrak{M}_{\text{MP}(f \bmod X^k)} \vec{g}$ using the out-of-place algorithm and then $\mathfrak{M}_{\text{MP}(f \text{ quo } X^k)} \vec{g}$ using a recursive call.

To make this algorithm work in place, the value of k has to be adjusted so that the work space is large enough. The result of a middle product in size k has degree $< k$ and needs ck extra work space by hypothesis. Therefore, if $m - k \geq (c + 1)k$, that is $k \leq m/(c + 2)$, the computation can be performed in place.

Algorithm 3 iMP_from_oMP

Input: f and g of size $n + m - 1$ and n respectively**Output:** R contains $\text{MP}(f, g)$ **Required:** Middle product algorithm oMP of space complexity $\leq cn$

```
1: if  $m < c + 2$  then
2:    $R \leftarrow \text{oMP}(f, g)$  ▷ using a naive algorithm
3: else
4:    $k \leftarrow \lfloor m/(c + 2) \rfloor$ 
5:    $R_{[0..k]} \leftarrow \text{oMP}(f \bmod X^{n+k}, g)$  ▷ work space:  $R_{[k..m]}$ 
6:    $R_{[k..m]} \leftarrow \text{iMP\_from\_oMP}(f \text{ quo } X^k, g)$  ▷ recursive call
```

Complexity analysis Let $M(k)$ be the cost of an out-of-place balanced middle product algorithm. The cost of an unbalanced middle product is thus $\lceil n/k \rceil M(k)$ for $k < n$. The in-place algorithm computes first a middle product using an out-of-place algorithm and then makes a recursive call on the remaining part. Note that n does not change during the algorithm and can be viewed as a large constant, while m is the parameter that varies. Then the cost of the algorithm verifies $T(m) \leq \lceil n/k \rceil M(k) + T(m - k)$. Since $k = \lfloor m/(c + 2) \rfloor$, $\lceil n/k \rceil < n(c + 2)/(m - c - 2) + 1$ and $m - k \leq (c + 1)m/(c + 2) + 1$. Furthermore,

$M(k) \leq m/n(c+2)M(n)$, thus $\lceil n/k \rceil M(k) \leq (m/(m-c-2) + m/n(c+2))M(n)$. That is,

$$T(m) \leq \left(\frac{m}{n(c+2)} + \frac{c+2}{m-c-2} + 1 \right) M(n) + T(\lfloor \frac{c+1}{c+2} m + 1 \rfloor).$$

Proposition 5.2(ii) implies $T(n) \leq M(n) \log_{\frac{c+2}{c+1}}(n) + O(M(n))$ for $m = n$, and $T(n) \leq O(M(n))$ if $M(n)$ is not quasi-linear.

If $M(n) \leq \lambda n^\gamma$ for some $\gamma > 1$, $M(k) \leq (\frac{m}{c+2})^\gamma$ and the recurrence becomes

$$T(m) \leq \left(\frac{n(c+2)}{m-c-2} + 1 \right) \lambda \left(\frac{m}{c+2} \right)^\gamma + T(\lfloor \frac{c+1}{c+2} m + 1 \rfloor).$$

Proposition 5.2(iii) implies $T(n) \leq O(n^\gamma)$ for $m = n$.

Reduction from short products to middle product The middle product of f and g can be computed as the sum of the low short product of f quo X^n with g and the high short product of $f \bmod X^n$ with g . Yet this reduction does not preserve the space complexity since one needs to store the results of the two short products in two zones of size n before summing them. Actually, the reduction given above from oMP to iMP can easily be adapted to a reduction from SP to MP that is space-preserving. Yet, the complexity also worsens with a logarithmic factor. Thus, we cannot conclude that $\text{MP} \leq_{\text{TISP}} \text{SP}$.

5.4 Proof of Proposition 5.2

Lemma 5.3. *Let $T(n)$ be a function satisfying $T(n) \leq f(n) + T(\lfloor \alpha n + \beta \rfloor)$ for some $\alpha < 1$, β and non-decreasing f . Then*

$$T(n) \leq T(\lfloor n_K \rfloor) + \sum_{i=0}^{K-1} f(n_i)$$

where $n_i = \alpha^i n + \beta \frac{1-\alpha^{i+1}}{1-\alpha}$ and $K \leq \log_{1/\alpha}(n)$.

Proof. By definition of n_i , $n = n_0$ and $T(\lfloor n_i \rfloor) \leq f(n_i) + T(\lfloor n_{i+1} \rfloor)$. Then by recurrence, $T(n) \leq T(\lfloor n_{i+1} \rfloor) + \sum_{j=0}^i f(n_j)$. \square

Lemma 5.4. *Let $n_i = \alpha^i n + \beta \frac{1-\alpha^{i+1}}{1-\alpha}$. Then*

$$\sum_{i=0}^{K-1} n_i \leq \frac{n + \beta K}{1 - \alpha}.$$

Proof. Since $0 < \alpha < 1$, $\sum_{i=0}^{K-1} \alpha^i < \sum_{i=0}^{\infty} \alpha^i = 1/(1-\alpha)$. Also, $\sum_{i=0}^{K-1} (1 - \alpha^{i+1})/(1-\alpha) \leq \sum_{i=0}^{K-1} 1/(1-\alpha) = K/(1-\alpha)$. \square

Lemma 5.5. *Let $n_i = \alpha^i n + \beta \frac{1-\alpha^{i+1}}{1-\alpha}$. Then*

$$\sum_{i=0}^{K-1} \frac{1}{n_i - \beta/(1-\alpha)} = \frac{\alpha(\alpha^{-K} - 1)}{(1-\alpha)n - \alpha\beta}.$$

Proof. Since $n_i = \alpha^i(n - \beta\alpha/(1-\alpha)) + \beta/(1-\alpha)$, $n_i - \beta/(1-\alpha)$ is a multiple of α^i . Thus,

$$\sum_{i=0}^{K-1} \frac{1}{n_i - \beta/(1-\alpha)} = \frac{1}{n - \beta\alpha/(1-\alpha)} \sum_{i=0}^{K-1} \alpha^{-i}.$$

Then, $\sum_i \alpha^{-i} = (1 - \alpha^{-K})/(1 - 1/\alpha) = \alpha(\alpha^{-K} - 1)/(1 - \alpha)$, and $\sum_i 1/(n_i - \beta/(1-\alpha)) = \alpha(\alpha^{-K} - 1)/((1 - \alpha)n - \alpha\beta)$. \square

Lemma 5.6. *If $M(n)/n$ is non-decreasing, and $n_i = \alpha^i n + \beta(1 - \alpha^{i+1})/(1 - \alpha)$ for some $\alpha < 1$, then*

$$\sum_{i=0}^{K-1} M(\lfloor \lambda n_i + \mu \rfloor) = \frac{\lambda}{1 - \alpha} M(n) + o(M(n))$$

for $K \leq \log_{1/\alpha}(n)$ and any λ and μ such that $\lambda n_i + \mu \leq n$ for all n_i .

Proof. Since $M(n)/n$ is non-decreasing, we get $M(\lfloor \lambda n_i + \mu \rfloor) \leq \frac{\lambda n_i + \mu}{n} M(n)$. Thus, $\sum_i M(\lfloor \lambda n_i + \mu \rfloor) \leq M(n)/n \sum_i (\lambda n_i + \mu)$. By Lemma 5.4, $\sum_i M(\lfloor \lambda n_i + \mu \rfloor) \leq \lambda M(n)/(1 - \alpha) + \lambda \beta K M(n)/n(1 - \alpha) + \mu K M(n)/n$. Since $K = O(\log n)$, $K M(n)/n = o(M(n))$. \square

Proof of Proposition 5.2. Proof of (i): By hypothesis, we have $T(n) \leq \sum_k a_k M(\lfloor \lambda_k n + \mu_k \rfloor) + bn + c + T(\lfloor \alpha n + \beta \rfloor)$ with $\alpha < 1$ and $\lambda_k n + \mu_k < n$ for all k . By Lemma 5.3, $T(n) \leq T(\lfloor n_K \rfloor) + \sum_i f(n_i)$ with n_i defined as in the lemma and $f(n) = \sum_k a_k M(\lfloor \lambda_k n + \mu_k \rfloor) + bn + c$. Then

$$\begin{aligned} \sum_{i=0}^{K-1} f(n_i) &= \sum_k a_k \sum_{i=0}^{K-1} M(\lfloor \lambda_k n_i + \mu_k \rfloor) + b \sum_{i=0}^{K-1} n_i + Kc \\ &\leq \sum_k a_k \left(\frac{\lambda_k}{1 - \alpha} M(n) + o(M(n)) \right) + b \frac{n + \beta K}{1 - \alpha} + Kc \\ &= \sum_k \frac{a_k \lambda_k}{1 - \alpha} M(n) + \frac{bn}{1 - \alpha} + o(M(n)) \end{aligned}$$

since $K = o(M(n))$ and the sum over k is of fixed size.

Proof of (ii): Since $T(m) \leq (\lambda m/n + \mu/(m - \frac{1}{1-\alpha}) + 1)M(n) + T(\lfloor \alpha m + 1 \rfloor)$ with $\alpha < 1$ and $m \leq n$, Lemma 5.3 implies

$$T(m) \leq T(\lfloor m_K \rfloor) + M(n) \sum_i \left(\frac{\lambda m_i}{n} + \frac{\mu}{m_i - 1/(1-\alpha)} + 1 \right)$$

where $m_i = \alpha^i m + (1 - \alpha^{i+1})/(1 - \alpha)$. By Lemma 5.4, $\sum_i m_i \leq \frac{m+K}{1-\alpha}$ and by Lemma 5.5, $\sum_i 1/(m_i - \frac{1}{1-\alpha}) \leq \alpha^{-K+1}/((1-\alpha)m - \alpha)$. Altogether,

$$T(m) \leq T(\lfloor m_K \rfloor) + KM(n) + \frac{\lambda(m+K)}{n(1-\alpha)} M(n) + \frac{\mu\alpha}{1-\alpha} \cdot \frac{(1/\alpha)^K}{m-\alpha/(1-\alpha)} M(n).$$

If we plug $K = \log_{1/\alpha}(m)$ and fix $m = n$, we get

$$T(n) \leq T(\lfloor n_K \rfloor) + M(n) \log_{1/\alpha} n + \frac{\lambda + \mu\alpha}{1 - \alpha} M(n) + o(M(n)).$$

Proof of (iii): Let $T(m) \leq \lambda(n/((1-\alpha)m-1)+1)((1-\alpha)m)^\gamma + T(\lfloor \alpha m + 1 \rfloor)$. We claim that there exist constants μ and ν such that $T(m) \leq \mu m^{\gamma-1}n + \nu m^\gamma + o(m^{\gamma-1}n + m^\gamma)$ and prove it by induction. Using the recurrence relation and the induction hypothesis,

$$T(m) \leq \lambda n((1-\alpha)m)^{\gamma-1} + \lambda((1-\alpha)m)^\gamma + \mu(\alpha m)^{\gamma-1}n + \nu(\alpha m)^\gamma + o(m^{\gamma-1}n + m^\gamma)$$

The result follows as soon as $\lambda(1-\alpha)^{\gamma-1} + \mu\alpha^{\gamma-1} \leq \mu$ and $\lambda(1-\alpha)^\gamma + \nu\alpha^\gamma \leq \nu$. We can thus fix

$$\mu = \frac{\lambda(1-\alpha)^{\gamma-1}}{1-\alpha^{\gamma-1}} \text{ and } \nu = \frac{\lambda(1-\alpha)^\gamma}{1-\alpha^\gamma}.$$

Finally, taking $m = n$, we conclude that $T(n) \leq (\mu + \nu)n^\gamma + O(n^{\gamma-1})$. □

6 Perspectives

We have presented algorithms for polynomial products which are efficient in terms of both time and space. Our results show that any algorithm for the full and short products can be turned into another algorithm with the same asymptotic time complexity while using only $O(1)$ extra space. We obtain similar results for the middle product but only proved it for algorithms that do not have a quasi-linear time complexity. In the latter case, an increase of the time complexity by a logarithmic factor occurs. We provided analysis of our reductions that make their constants explicit. In particular, their values ensure that our reductions are practicable.

In a future work, we plan to address some remaining issues. By examining the constants in the already known algorithms, we can choose the algorithms to use as starting points of our reductions to optimize the complexity. For instance three variants of Karatsuba's algorithm with different time and space complexities are known [20, 24, 17]. Furthermore, it seems possible to improve on the complexity of low-space versions of Karatsuba's and Toom-Cook's algorithm, yielding faster in-place algorithms through our reductions. Another promising approach is to slightly relax the model of computation and work in model in which one can write on the input space as long as the original inputs are restored by the end of the computation. Preliminary results for Karatsuba's algorithm suggest that this could also yield a lower constant in the time complexity.

Finally, we have started to explore the design of in-place algorithms for a broader range of problems of polynomials, such as division or evaluation/interpolation. The use of in-place middle and short products becomes crucial since one needs to avoid any increase in the size of the intermediate results.

Acknowledgements

This work was begun while the last author was graciously hosted by the LIRMM at the Université Montpellier. This work was supported in part by the National Science Foundation under grants 1319994, 1618269; and also by the Région Occitanie.

References

- [1] K. Abrahamson. Time-space tradeoffs for branching programs contrasted with those for straight-line programs. In *27th Annual Symposium on Foundations of Computer Science (FOCS 1986)*, pages 402–409, 1986. DOI:10.1109/SFCS.1986.58.

- [2] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 1st edition, 2009.
- [3] A. Bostan, F. Chyzak, M. Giusti, R. Lebreton, G. Lecerf, B. Salvy, and E. Schost. *Algorithmes Efficaces en Calcul Formel*. 1 edition, Aug. 2017.
- [4] A. Bostan, G. Lecerf, and E. Schost. Tellegen’s principle into practice. In *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation*, ISSAC ’03, pages 37–44. ACM, 2003. DOI:10.1145/860854.860870.
- [5] B. Boyer, J.-G. Dumas, C. Pernet, and W. Zhou. Memory efficient scheduling of strassen-winograd’s matrix multiplication algorithm. In *Proceedings of the 2009 International Symposium on Symbolic and Algebraic Computation*, ISSAC ’09, pages 55–62. ACM, 2009. DOI:10.1145/1576702.1576713.
- [6] R. Brent and P. Zimmermann. *Modern Computer Arithmetic*. Cambridge University Press, New York, NY, USA, 2010.
- [7] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991. DOI:10.1007/BF01178683.
- [8] Y. Cheng. Space-efficient karatsuba multiplication for multi-precision integers, 2016. ARXIV: 1605.06760.
- [9] S. A. Cook. *On the minimum computation time of functions*. PhD thesis, Harvard University, May 1966.
- [10] J. v. z. Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.
- [11] G. Hanrot, M. Quercia, and P. Zimmermann. The middle product algorithm I. *Applicable Algebra in Engineering, Communication and Computing*, 14(6):415–438, Mar 2004.
- [12] G. Hanrot and P. Zimmermann. A long note on Mulders’ short product. *Journal of Symbolic Computation*, 37(3):391–401, 2004. DOI:10.1016/j.jsc.2003.03.001.
- [13] D. Harvey, J. van der Hoeven, and G. Lecerf. Faster polynomial multiplication over finite fields. *J. ACM*, 63(6):52:1–52:23, Jan. 2017. DOI:10.1145/3005344.
- [14] D. Harvey and D. S. Roche. An in-place truncated Fourier transform and applications to polynomial multiplication. In *ISSAC ’10: Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation*, pages 325–329. ACM, 2010.
- [15] D. Harvey and J. Van Der Hoeven. Polynomial multiplication over finite fields in time $O(n \log n)$. HAL:hal-02070816, Mar. 2019.
- [16] E. Kaltofen. Challenges of symbolic computation: my favorite open problems. *Journal of Symbolic Computation*, 29(6):891–919, 2000.
- [17] A. Karatsuba and Y. Ofman. Multiplication of Multidigit Numbers on Automata. *Soviet Physics-Doklady*, 7:595–596, 1963.

- [18] A. Lincoln, V. Vassilevska Williams, J. R. Wang, and R. R. Williams. Deterministic Time-Space Trade-Offs for k -SUM. In I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 58:1–58:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. DOI:10.4230/LIPIcs.ICALP.2016.58.
- [19] T. Mulders. On Short Multiplications and Divisions. *Applicable Algebra in Engineering, Communication and Computing*, 11(1):69–88, 2000. DOI:10.1007/s002000000037.
- [20] D. S. Roche. Space- and time-efficient polynomial multiplication. In *Proceedings of the 2009 International Symposium on Symbolic and Algebraic Computation, ISSAC '09*, pages 295–302. ACM, 2009.
- [21] J. Savage and S. Swamy. Space-time tradeoffs for oblivious integer multiplication. In H. Maurer, editor, *Automata, Languages and Programming*, volume 71 of *Lecture Notes in Computer Science*, pages 498–504. Springer Berlin / Heidelberg, 1979. DOI:10.1007/3-540-09510-1_40.
- [22] A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing*, 7:281–292, 1971. DOI:10.1007/BF02242355.
- [23] C. Su and H. Fan. Impact of Intel’s new instruction sets on software implementation of $GF(2)[x]$ multiplication. *Information Processing Letters*, 112(12):497–502, 2012. DOI:10.1016/j.ipl.2012.03.012.
- [24] E. Thomé. Karatsuba multiplication with temporary space of size $\leq n$. online, 2002.
- [25] V. Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *Proceedings ICM*, volume 3, pages 3431–3472, 2018.