# Essentially optimal sparse polynomial multiplication 

Pascal Giorgi, Bruno Grenet, Armelle Perret Du Cray

## To cite this version:

Pascal Giorgi, Bruno Grenet, Armelle Perret Du Cray. Essentially optimal sparse polynomial multiplication. ISSAC: International Symposium on Symbolic and Algebraic Computation, Jul 2020, Kalamata, Greece. pp.202-209, 10.1145/3373207.3404026 . hal-02476609v1

## HAL Id: hal-02476609 https://hal.science/hal-02476609v1

Submitted on 13 May 2020 (v1), last revised 31 Aug 2020 (v2)

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

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

# Essentially optimal sparse polynomial multiplication 

Pascal Giorgi Bruno Grenet Armelle Perret du Cray<br>LIRMM, Univ. Montpellier, CNRS<br>Montpellier, France<br>\{pascal.giorgi,bruno.grenet,armelle.perret-du-cray\}@lirmm.fr

February 3, 2020


#### Abstract

In this article, we present a probabilistic algorithm to compute the product of two univariate sparse polynomials over a field with a number of bit operations that is quasi-linear in the size of the input and the output. Our algorithm works for any field of characteristic zero or larger than the degree. We mainly rely on sparse interpolation and on a new algorithm for verifying a sparse product that has also a quasi-linear time complexity. Using Kronecker substitution techniques we are able to extend our result to the multivariate case.


## 1 Introduction

Polynomial is one of the most basic object in computer algebra and the study of fast polynomial operations still remains one of the most challenging task. The most classical way of dealing with polynomial is to either use the dense representation that stores all the coefficients in a vector or the sparse one that keeps only the nonzero monomials. While quasi-optimal algorithms have been designed since decades for the dense representation, this not yet the case for the more compact sparse representation, mainly due to the specificity of the representation.

In the sparse representation, a polynomial $F=\sum_{i=0}^{D} f_{i} X^{i} \in R[X]$ is represented as a list of pairs $\left(e_{i}, f_{e_{i}}\right)$ such all the $f_{e_{i}}$ 's are nonzero. If $s$ is a bound on the size of the coefficients of $F, T$ its sparsity (that is its number of nonzero coefficients) and $D$ its degree, the size of the representation of $F$ is $\mathscr{O}(T(s+\log D)$ ) bits. It is classical to consider $s=\mathscr{O}\left(\log \left(\|F\|_{\infty}\right)\right)$ when $R=\mathbb{Z}$ and $s=\mathscr{O}(\log (q))$ when $R=\mathbb{F}_{q}$. This representation naturally extends to multivariate polynomials, where each exponent is replaced by a vector of exponents, thus giving a size of $\mathscr{O}(T(s+n \log D)$ where $n$ is the number of variables.

Several problems have been investigated to design fast algorithms for sparse polynomials, including arithmetic operations, interpolation and factorization (we refer to the excellent survey of Roche and the references therein [19]). Note that fast algorithms for sparse polynomials must have a logarithmic (or poly-logarithmic) dependency on the degree, contrary to the dense case. Unfortunately, as shown by several NP-hardness results, such fast algorithms might not even exist unless $P \neq N P$. This is for instance the case for GCD computations [16].

In this paper, we are interested in a more basic problem that is the product of two polynomials for which we provide the first algorithm being quasi-optimal in the input and the output size.

### 1.1 Previous work

The main difficulty and most interesting aspect of the sparse polynomial multiplication is the fact that the size of the output does not exclusively depend on the size of the inputs, contrary to the dense case. Indeed, the product of two polynomials $F$ and $G$ has at most $\# F \# G$ nonzero coefficients (where $\# F$ and $\# G$ denote their respective sparsities), but it may have as few as 2 nonzero coefficients.
Example 1. Let $F=X^{14}+2 X^{7}+2, G=3 X^{13}+5 X^{8}+3$ and $H=X^{14}-2 X^{7}+2$. Then $F G=3 X^{27}+5 X^{22}+$ $6 X^{20}+10 X^{15}+3 X^{14}+6 X^{13}+10 X^{8}+6 X^{7}+6$ has nine terms, while $F H=X^{28}+4$ has only two.

Given two polynomials of sparsity $T$, their product can be computed by generating the $T^{2}$ possible monomials, sorting them by increasing degree and merging the monomials of the same degree. Using radix sort, this algorithm takes $\mathscr{O}\left(T^{2}\left(\mathrm{M}_{R}+\log D\right)\right)$ bit operations, where $\mathrm{M}_{R}$ denotes the cost of one operation in $R$. A major drawback of this approach is its space complexity that exhibits a $T^{2}$, even if the result has less $T^{2}$
terms. Many improvements have been proposed to either lower down this space complexity while keeping the same time complexity or to extend this approach to the multivariate polynomials case and then to provide fast implementations in practice [12, 13, 14]. While all these results mainly improved the space complexity, none of them reduces the $T^{2}$ terms appearing in the time complexity.

Letting apart space complexity consideration, some other works have made the time complexity to be more output dependent when the inputs are sufficiently structured to know in advance that the output size will be small [18], especially in the multivariate case [23, 22], or when the support of the output is known in advance [9].

Since the output size could be quadratic in $T$ in the worst case, no improvement should be expected in terms of the inputs size only. But there could be less than $T^{2}$ terms in the output for two reasons. First, there may be some exponent collisions. If $F=\sum_{i=1}^{T} f_{i} X^{\alpha_{i}}$ and $G=\sum_{j=1}^{T} g_{j} X^{\beta_{j}}$, there may exist two distinct pairs of indices ( $i_{1}, j_{1}$ ) and ( $i_{2}, j_{2}$ ) such that $\alpha_{i_{1}}+\beta_{j_{1}}=\alpha_{i_{2}}+\beta_{j_{2}}$. Such collisions decrease the number of terms of the result. The second source of decrease comes from the coefficients. In the previous case, the resulting monomial is $\left(f_{i_{1}} g_{j_{1}}+f_{i_{2}} g_{j_{2}}\right) X^{\alpha_{i_{1}}+\beta_{j_{1}}}$ which could be zero depending on the coefficient values.

In [2] Arnold and Roche take into account the exponent collision by defining the structural support and the structural sparsity of a sparse polynomial product $F G$. The structural support of $F G$ is the sumset of the exponents appearing in $F$ and $G$ that is $\left\{\alpha_{i}+\beta_{j}: 1 \leq i, j \leq T\right\}$ and the structural sparsity its size. Let $F, G \in \mathbb{Z}[X]$ with coefficient of bitsize at most $B, H=F G$ and $S$ be the structural sparsity of $F G$. Arnold and Roche's algorithm computes $H$ in $\tilde{O}(S \log D+\# H B)$ bit operations. They also extend their algorithm to work over any finite fields and also to the multivariate polynomials case. Actually, Cole and Hariharan described an algorithm requiring $\tilde{\mathscr{O}}\left(S \log ^{2} D+\log ^{3} D\right)$ bit operations for polynomials with nonnegative integer coefficients of bitsize $\mathscr{O}(\log D)$ in the context of sparse wildcard matching [5].

Note that Arnold and Roche's algorithm is quasi-optimal when the structural sparsity is close, or even equal, to the sparsity of the result. Let $H=F G$ where $F$ and $G$ are two polynomials of sparsity $T \geq 2$, and $S$ be the structural sparsity of $F G$. Then by definition we have $2 \leq \# H \leq S \leq T^{2}$. Even though $S$ and \#H can be close, their difference can reach $\mathscr{O}\left(T^{2}\right)$ as shown by the next example, which means that their algorithm can be far from the optimal.
Example 2. Let $F=\sum_{i=0}^{T-1} X^{i}, G=\sum_{i=0}^{T-1}\left(X^{T i+1}-X^{T i}\right)$ and $H=F G$. Then $\# F=T$, $\# G=2 T$ and the structural sparsity of $F G$ is $T^{2}+1$ while $H=X^{T^{2}}-1$ has sparsity 2 .

An attempt to avoid the dependency on the structural sparsity in the complexity for the case of integer polynomials is presented in [15], in the word RAM model with large word size. Nevertheless, the algorithm translated to a more standard bit complexity model has non quasi-linear complexity $\tilde{\mathscr{O}}\left(\left(\# H \log ^{2} D+\right.\right.$ $\left.\log ^{3} D\right)(\log D+B)$, and it always returns polynomials with nonnegative coefficients.

While in the dense case the quasi-optimal algorithms rely on an evaluation-interpolation scheme, this is not yet the case for the fastest one in the sparse setting. Even if sparse interpolation has received a lot of attention, from the early results of Prony [17] and Ben-Or and Tiwari [4] to the most recent result of Huang [10], it does not resemble to the dense interpolation problem. Nevertheless, Arnold explores in his PhD thesis the use of sparse interpolation to perform sparse polynomial multiplication [1] and it reveals fruitful in obtaining the algorithm with Roche in [2]. Our work continues this approach a step further and we show that sparse interpolation can indeed be used to obtain a quasi-optimal multiplication algorithm for sparse polynomials.

### 1.2 Our contributions

Our main result is summarized in Theorem 1.1. Here, and throughout the article, $\tilde{\mathscr{O}}(f(n))$ denotes $\mathscr{O}\left(f(n) \log ^{k_{1}}(f(n))\right)$ for some constant $k_{1}>0$. We extend the complexity notations to $\mathscr{O}_{\epsilon}$ and $\tilde{\mathscr{O}}_{\epsilon}$ for hiding some polynomial factors in $\log \left(\frac{1}{\epsilon}\right)$. For a polynomial $F=\sum_{i=1}^{T} f_{i} X^{e_{i}}$, we write $\|F\|_{\infty}=\max _{i}\left|f_{i}\right|$ for its height, $\# F$ for its number of nonzero terms and $\operatorname{supp}(F)=\left\{e_{1}, \ldots, e_{T}\right\}$.

Theorem 1.1. Given two sparse polynomials $F$ and $G$ over $\mathbb{Z}$, Algorithm SparseProduct computes $H=$ $F G$ in $\tilde{\sigma}_{\epsilon}(T(\log D+\log C))$ bit operations with probability at least $1-\epsilon$, where $D=\operatorname{deg}(H), C=$ $\max \left(\|F\|_{\infty},\|G\|_{\infty},\|H\|_{\infty}\right)$ and $T=\max (\# F, \# G, \# H)$. The algorithm extends naturally to finite fields with characteristic larger than $D$ with the same complexity where $C$ denotes the cardinality.

This result is based on two main ingredients. On the one hand, we adapt Huang's sparse interpolation algorithm [10] to be able to interpolate $F \times G$ in quasi-linear time. Note that the original algorithm does not reach such complexity. On the other hand, any sparse interpolation algorithms, including Huang's one, require a bound on the sparsity of the result. We circumvent the latter difficulty by designing a quasi-linear
time verification algorithm for sparse product which allows us to replace such an a priori knowledge by an $a$ posteriori check.

Theorem 1.2. Given three sparse polynomials $F, G$ and $H$ over $\mathbb{F}_{q}$ or $\mathbb{Z}$, Algorithm VERIFYSP tests whether $F G=H$ in $\tilde{O}_{\epsilon}(T(\log D+B))$ bit operations, where $D=\operatorname{deg}(H), B$ is a bound on the bitsize of the coefficients of $F, G$ and $H$, and $T=\max (\# F, \# G, \# H)$. The answer is always correct if $F G=H$, and the probability of error is at most $\epsilon$ otherwise.

Finally, we show that using Kronecker substitution our sparse polynomial multiplication algorithm also extends to the multivariate case with a quasi-linear bit complexity $\tilde{\sigma}_{\epsilon}(T(n \log d+B))$ where $n$ is the number of variables and $d$ is the maximal partial degree on each variable. Nevertheless, over finite fields our approach requires an exponentially large characteristic. Using the randomized Kronecker substitution [3] we derive a fast algorithm that works for characteristic that remains polynomial in the input size. Though not quasi-linear our bit complexity is $\tilde{\mathscr{O}}_{\epsilon}(n T(\log d+B))$ and it improves the known complexity for this case.

## 2 Preliminaries

We denote by $\mathrm{I}(n)=\mathscr{O}(n \log n)$ the bit complexity of the multiplication of two integers of at most $n$ bits [8]. Similarly, we denote by $\mathrm{M}_{q}(D)=\mathscr{O}\left(D \log q \log (D \log q) 4^{\log ^{*} D}\right)$ the bit complexity of the multiplication of two dense polynomials of degree at most $D$ over $\mathbb{F}_{q}$ where $q$ is prime [7]. The cost of multiplying two elements of $\mathbb{F}_{q^{s}}$ is $\mathrm{M}_{q}(s)$. The cost of multiplying two dense polynomials over $\mathbb{Z}$ of heights at most $C$ and degrees at most $D$ is $\mathrm{M}_{\mathbb{Z}}(D, C)=\mathrm{I}(D(\log C+\log D))[24$, Chapter 8].

Since our algorithms use reductions modulo $X^{p}-1$ for some prime number $p$, we first review useful related results.
Theorem 2.1 ([20]]). If $\lambda \geq 21$, there are at least $\frac{3}{5} \lambda / \ln \lambda$ prime numbers in $[\lambda, 2 \lambda]$.
Proposition 2.2. There exists an algorithm $\operatorname{RandomPrime}(\lambda, \epsilon)$ that returns an integer $p$ in $[\lambda, 2 \lambda]$, such that $p$ is prime with probability at least $1-\epsilon$. Its bit complexity is $\tilde{\sigma}_{\epsilon}\left(\log ^{k} \lambda\right)$ for some $k \leq 7$.

We need two distinct properties for the reductions modulo $X^{p}-1$. The first one is classical in sparse interpolation to bound the probability of exponent collision in the residue (see [2] Lemma 3.3]).

Proposition 2.3. Let $H$ be a polynomial of degree at most $D$ and sparsity at most $T, 0<\epsilon<1$ and $\lambda=$ $\max \left(21, \frac{10}{3 \epsilon} T^{2} \ln D\right)$. Then with probability at least $1-\epsilon$, RandomPrime $\left(\lambda, \frac{\epsilon}{2}\right)$ returns a prime number $p$ such that $H \bmod X^{p}-1$ has the same number of terms as $H$, that is no collision of exponents occurs.

The second property allows to bound the probability that a polynomial vanishes modulo $X^{p}-1$.
Proposition 2.4. Let $H$ be a nonzero polynomial of degree at most $D$ and sparsity at most $T, 0<\epsilon<1$ and $\lambda=\max \left(21, \frac{10}{3 \epsilon} T \ln D\right)$. Then with probability at least $1-\epsilon$, RandomPrime $\left(\lambda, \frac{\epsilon}{2}\right)$ returns a prime number $p$ such that $H \bmod X^{p}-1 \neq 0$.

Proof. It is sufficient, for $H \bmod X^{p}-1$ to be nonzero, that there exist one exponent $e$ of $H$ that is not congruent to any other exponents $e_{j}$ modulo $p$. In other words, it is sufficient that $p$ does not divide any of the $T-1$ differences $\delta_{j}=e_{j}-e$.

Noting that $\delta_{j} \leq D$, the number of primes in $[\lambda, 2 \lambda]$ that divide at least one $\delta_{j}$ is at most $\frac{(T-1) \ln D}{\ln \lambda}$. Since there exists $\frac{3}{5} \lambda / \ln \lambda$ primes in this interval, the probability that a prime randomly chosen from it divides at least one $\delta_{j}$ is at most $\epsilon / 2$. RandomPrime $(\lambda, \epsilon / 2)$ returns a prime in $[\lambda, 2 \lambda]$ with probability at least $1-\epsilon / 2$, whence the result.

The following two propositions will be useful to either reduce integer coefficients modulo some prime number or to construct an extension field.

Proposition 2.5. Let $H \in \mathbb{Z}[X]$ be a nonzero polynomial, $0<\epsilon<1$ and $\lambda \geq \max \left(21, \frac{10}{3 \epsilon} \ln \|H\|_{\infty}\right)$. Then with probability at least $1-\epsilon$, RandomPrime $\left(\lambda, \frac{\epsilon}{2}\right)$ returns a prime $q$ such that $H \bmod q \neq 0$.
Proof. Let $h_{i}$ be a nonzero coefficient of $H$,
a random prime from $[\lambda, 2 \lambda]$ divides $h_{i}$ with probability at most $\frac{5}{3} \ln \|H\|_{\infty} / \lambda \leq \epsilon / 2$. Since $\operatorname{RandomPrime}(\lambda, \epsilon / 2)$ returns a prime in $[\lambda, 2 \lambda]$ with probability at least $1-\epsilon / 2$ the result follows.

Proposition 2.6 ([21, Chapter 20]). There exists an algorithm that, given a finite field $\mathbb{F}_{q}$, an integer $s$ and $0<\epsilon<1$, computes a degree-s polynomial in $\mathbb{F}_{q}[X]$ that is irreducible with probability at least $1-\epsilon$. Its bit complexity is $\tilde{\sigma}_{\epsilon}\left(s^{3} \log q\right)$.

## 3 Sparse polynomial product verification

Verifying a product $F G=H$ of dense polynomials over an integral domain $R$ simply falls down to testing $F(\alpha) G(\alpha)=H(\alpha)$ for some random point $\alpha \in R$. This approach exhibits an optimal linear number of operations in $R$ but it is not deterministic (optimal deterministic algorithm does not yet exists). When $R=\mathbb{Z}$ or $\mathbb{F}_{q}$, the complexity becomes quasi-linear in the input size as the best algorithm using a divide and conquer approach requires $\tilde{O}(D B)$ where $B$ bounds the bitsize of the coefficients.

For sparse polynomials with $T$ nonzero coefficients, evaluation is not quasi-linear since the input size is only $\mathscr{O}(T(\log D+B))$. Indeed, computing $\alpha^{D}$ requires $\mathscr{O}(\log D)$ operations in $R$ which implies a bit complexity of $\tilde{O}(\log D \log q)$ when $R=\mathbb{F}_{q}$. Applying this computation to the $T$ non zero monomial gives a bit complexity of $\tilde{O}(T \log D \log q)$. We shall mention that the latter approach can be improved to $\tilde{O}(\log D \log q(1+T / \log \log D))$ using Yao's result [25] on simultaneous exponentiation. When $R=\mathbb{Z}$ the complexity is even worse as $\alpha^{D}$ has size $\mathscr{O}(D)$.

Our approach to obtain a quasi-linear complexity is to perform the evaluation modulo $X^{p}-1$ for some random prime $p$. This requires to evaluate the polynomial $\left[(F G) \bmod X^{p}-1\right]$ on $\alpha$ without computing it.

### 3.1 Modular product evaluation

Lemma 3.1. Let $F$ and $G$ be two sparse polynomials in $R[X]$ with $\operatorname{deg} F, \operatorname{deg} G \leq p-1$ and $\alpha \in R$. Then $(F G) \bmod X^{p}-1$ can be evaluated on $\alpha$ using $\mathscr{O}((\# F+\# G) \log p)$ operations in $R$.

Proof. Let $H=(F G) \bmod X^{p}-1$. The computation of $H$ corresponds to the linear map

$$
\underbrace{\left(\begin{array}{c}
h_{0} \\
h_{1} \\
\vdots \\
h_{p-1}
\end{array}\right)}_{\vec{h}}=\underbrace{\left(\begin{array}{cccc}
f_{0} & f_{p-1} & \cdots & f_{1} \\
f_{1} & f_{0} & \cdots & f_{2} \\
\vdots & \vdots & & \vdots \\
f_{p-1} & f_{p-2} & \cdots & f_{0}
\end{array}\right)}_{T_{F}} \underbrace{\left(\begin{array}{c}
g_{0} \\
g_{1} \\
\vdots \\
g_{p-1}
\end{array}\right)}_{\vec{g}}
$$

where $f_{i}$ (resp. $g_{i}, h_{i}$ ) is the coefficient of degree $i$ of $F$ (resp. $G, H$ ). Computing $H(\alpha)$ corresponds to the inner product $\vec{\alpha}_{p} \overrightarrow{\vec{h}}=\vec{\alpha}_{p} T_{F} \vec{g}$ where $\vec{\alpha}_{p}=\left(1, \alpha, \ldots, \alpha^{p-1}\right)$. This evaluation can be computed in $\mathscr{O}(p)$ operations in $R$ [6]. Here we reuse similar techniques in the context of sparse polynomials.

To compute $H(\alpha)$, we first compute $\vec{c}=\vec{\alpha}_{p} T_{F}$, and then the inner product $\vec{c} \vec{g}$. If $\operatorname{supp}(G)=\left\{j_{1}, \ldots, j_{\# G}\right\}$ with $j_{1}<\cdots<j_{\# G}<p$, we only need the corresponding entries of $\vec{c}$, that is all $c_{j_{k}}$ 's for $1 \leq k \leq \# G$. Since $c_{j}=\sum_{\ell=0}^{p-1} \alpha^{\ell} f_{(\ell-j) \bmod p}$, we can write $c_{j}=f_{p-j}+\alpha \sum_{\ell=0}^{p-2} \alpha^{\ell} f_{(\ell-j+1) \bmod p}$, that is $c_{j}=\alpha c_{j-1}+\left(1-\alpha^{p}\right) f_{p-j}$.

Applying this relation as many times as necessary, we obtain a relation to compute $c_{j_{k+1}}$ from $c_{j_{k}}$ :

$$
c_{j_{k+1}}=\alpha^{j_{k+1}-j_{k}} c_{j_{k}}+\left(1-\alpha^{p}\right) \sum_{\ell=j_{k}+1}^{j_{k+1}} \alpha^{\ell} f_{p-\ell}
$$

Each nonzero coefficient $f_{t}$ of $F$ appears in the definition of $c_{j_{k+1}}$ if and only if $p-j_{k+1} \leq t<p-j_{k}$. Thus, each $f_{t}$ is used exactly once to compute all the $c_{j_{k}}$ 's. Since for each summand, one needs to compute $\alpha^{\ell}$ for some $\ell<p$, the total cost for computing all the sums is $\mathscr{O}(\# F \log p)$ operations in $R$. Similarly, the computation of $\alpha^{j_{k+1}-j_{k}} c_{j_{k}}$ for all $k$ costs $\mathscr{O}(\# G \log p)$. The last remaining step is the final inner product which costs $\mathscr{O}(\# G)$ operations in $R$, whence the result.

The complexity is improved to $\mathscr{O}(\log p+(\# F+\# G) \log p / \log \log p)$ using again Yao's algorithm [25] for simultaneous exponentiation.

### 3.2 A quasi-linear time algorithm

Given three sparse polynomials $F, G$ and $H$ in $R[X]$, we want to assert that $H=F G$. Our approach is to take a random prime $p$ and to verify this assertion modulo $X^{p}-1$ through modular product evaluation. This method is explicitly described in the algorithm VERIFYSP that works over any large enough integral domain $R$. We further extend the description and the analysis of this algorithm for the specific cases $R=\mathbb{Z}$ and $R=\mathbb{F}_{q}$ in the next sections.

```
Algorithm 1 VERIFYSP
Input: \(H, F, G \in R[X] ; 0<\epsilon<1\).
Output: True if \(F G=H\), False with probability \(\geq 1-\epsilon\) otherwise.
    Define \(c_{1}>\frac{10}{3}\) and \(c_{2}>1\) such that \(\frac{10}{3 c_{1}}+\left(1-\frac{10}{3 c_{1}}\right) \frac{1}{c_{2}} \leq \epsilon\)
    \(D \leftarrow \operatorname{deg}(H)\)
    if \(\# H>\# F \# G\) or \(D \neq \operatorname{deg}(F)+\operatorname{deg}(G)\) then return False
    \(\lambda \leftarrow \max \left(21, c_{1}(\# F \# G+\# H) \ln D\right)\)
    \(p \leftarrow \operatorname{RANDOMPRIME}\left(\lambda, \frac{5}{3 c_{1}}\right)\)
    \(\left(F_{p}, G_{p}, H_{p}\right) \leftarrow\left(F \bmod X^{p}-1, G \bmod X^{p}-1, H \bmod X^{p}-1\right)\)
    Define \(\mathscr{E} \subset R\) of size \(>c_{2} p\) and choose \(\alpha \in \mathscr{E}\) randomly.
    \(\beta \leftarrow\left[\left(F_{p} G_{p}\right) \bmod X^{p}-1\right](\alpha) \quad \triangleright\) using Lemma 3.1
    return \(\beta=H_{p}(\alpha)\)
```

Theorem 3.2. If $R$ is an integral domain of size $\geq 2 c_{1} c_{2} \# F \# G \ln D$ VerifySP works as specified and it requires $\mathscr{O}_{\epsilon}\left(T \log (T \log D)\right.$ ) operations in $R$ plus $\mathscr{O}_{\epsilon}(T \mid(\log D))$ bit operations where $D=\operatorname{deg}(H)$ and $T=$ $\max (\# F, \# G, \# H)$.

Proof. Step 3 dismisses two trivial mistakes and ensures that $D$ is a bound on the degree of each polynomial.
If $F G=H$, the algorithm returns True for any choice of $p$ and $\alpha$. Otherwise, there are two sources of failure. Either $X^{p}-1$ divides $F G-H$, whence $(F G)_{p}(\alpha)=H_{p}(\alpha)$ for any $\alpha$. Or $\alpha$ is a root of the nonzero polynomial $(F G-H) \bmod X^{p}-1$. Since $F G-H$ has at most $\# F \# G+\# H$ terms, the first failure occurs with probability at most $\frac{10}{3 c_{1}}$ by Prop. 2.4. And since $(F G-H) \bmod X^{p}-1$ has degree at most $p-1$ and $\mathscr{E}$ has $c_{2} p$ points, the second failure occurs with probability at most $\frac{1}{c_{2}}$. Altogether, the failure probability is at most $\frac{10}{3 c_{1}}+\left(1-\frac{10}{3 c_{1}}\right) \frac{1}{c_{2}}$.

Let us remark that $c_{1}, c_{2}=\mathscr{O}\left(\frac{1}{\epsilon}\right)$ and $p=\mathscr{O}\left(\frac{1}{\epsilon} T^{2} \log D\right)$. Step 5 requires only $\tilde{\mathscr{O}}\left(\log ^{7}\left(\frac{1}{\epsilon} T \log D\right)\right)$ bit operations by Proposition[2.2. The operations in Step 6 are $T$ divisions by $p$ on integers bounded by $D$ which cost $\mathscr{\sigma}_{\epsilon}(T \mathrm{l}(\log D))$ bit operations, plus $T$ additions in $R$. The evaluation of $F_{p} G_{p} \bmod X^{p}-1$ on $\alpha$ at Step 8 requires $\mathscr{O}\left(T \log \left(\frac{1}{\epsilon} T \log D\right)\right.$ ) operations in $R$ by Lemma3.1. The evaluation of $H_{p}$ on $\alpha \operatorname{costs} \mathscr{O}(T \log (T \log D))$ operations in $R$. Other steps have negligible cost.

### 3.3 Analysis over finite fields

The first easy case is the case of large finite fields: If there are enough points for the evaluation, the generic algorithm has the same guarantee of success and a quasi-linear time complexity.

Corollary 3.3. Let $F, G$ and $H$ be three polynomials of degree at most $D$ and sparsity at most $T$ in $\mathbb{F}_{q}[X]$ where $q>2 c_{1} c_{2} \# F \# G \ln (D)$. Then Algorithm VERIFYSP has bit complexity $\mathscr{O}_{\epsilon}\left(n \log ^{2}(n) 4^{\log ^{*} n}\right)$ where $n=$ $T(\log D+\log q)$ is the input size.

Proof. By definition of $n$, the cost of step 6 is $\mathscr{O}_{\epsilon}(n \log n)$ bit operations. Each ring operation in $\mathbb{F}_{q}$ costs $\mathscr{O}\left(\log (q) \log \log (q) 4^{\log ^{*} q}\right)$ bit operations which implies that the bit complexity of step 8 is $\mathscr{O}_{\epsilon}\left(T \log (T \log D) \log (q) \log \log (q) 4^{\log ^{*} q}\right)$. Since $T \log q$ and $T \log D$ are bounded by $n, \log \log q \leq \log n$, the result follows.

We shall note that even if $q<2 c_{1} c_{2} \# F \# G \ln (D)$ we can make our algorithm work by using an extension field and this approach achieves the same complexity.

Theorem 3.4. One can adapt algorithm VerifySP to work over finite fields $\mathbb{F}_{q}$ such that $q<2 c_{1} c_{2} \# F \# G \ln (D)$. The bit complexity is $\mathscr{\sigma}_{\epsilon}\left(n \log (n) \log \log (n) 4^{\log ^{*} n}\right)$, where $n=T(\log D+\log q)$ is the input size.
Proof. To have enough elements in the set $\mathscr{E}$, we need to work over $\mathbb{F}_{q^{s}}$ where $q^{s}>c_{2} p \geq q^{s-1}$. An irreducible degree-s polynomial can be computed in $\tilde{\mathscr{O}}\left(s^{3} \log q\right)=\tilde{\mathscr{O}}(\log (T \log D) / \log q)$ by Proposition 2.6, Since $\alpha$ is taken in $\mathbb{F}_{q^{s}}$, the complexity becomes $\mathscr{O}_{\epsilon}\left(T \mathrm{I}(\log D)+T \log (T \log D) \mathrm{M}_{q}(s)\right)$ bit operations. Remarking that $T \leq D$ we have $T \log (T \log D) \leq T \log (D \log D)=\mathscr{O}(n)$. Since $s \log q=\mathscr{O}(\log (T \log D))=\mathscr{O}(\log n)$ we can obtain $\mathrm{M}_{q}(s)=\mathscr{O}\left(\log (n) \log \log (n) 4^{\log ^{*} n}\right)$ which implies that the second term of the complexity is $\mathscr{O}\left(n \log (n) \log \log (n) 4^{\log ^{*} n}\right)$. The first term is negligible since it is $\mathscr{O}(n \log n)$.

In order to achieve the same probability of success, we fix an error probability $1 / c_{3}<1$ for Proposition 2.6 and we take constants $c_{1}$ and $c_{2}$ in VERIFYSP such that $1-\left(1-\frac{10}{3 c_{1}}\right)\left(1-\frac{1}{c_{2}}\right)\left(1-\frac{1}{c_{3}}\right) \leq \epsilon$.

We note that for very sparse polynomials over some fields, the complexity is only dominated by the operations on the exponents.

Corollary 3.5. VerifySP has bit complexity $\mathscr{O}_{\epsilon}(n \log n)$ in the following cases :
(i) $s=1$ and $\log q=\mathscr{O}\left(\log ^{1-\alpha} D\right)$ for some constant $0<\alpha<1$,
(ii) $s>1$ and $T=\Theta\left(\log ^{k} D\right)$ for some constant $k$.

Proof. In both cases the cost of reducing the exponents modulo $p$ is $\mathscr{O}_{\epsilon}(n \log n)$ bit operations. In the first case, each multiplication in $\mathbb{F}_{q}$ costs $\mathscr{O}\left(\log (q) \log \log (q) 4^{\log ^{*} q}\right)=\mathscr{O}(\log D)$ bit operations as $\log \log (q) 4^{\log ^{*} q}=$ $\mathscr{O}\left(\log ^{\alpha} D\right)$. In the second case, $n=\mathscr{O}\left(\log ^{k+1} D\right)$ and $s \log q=\mathscr{O}_{\epsilon}\left(\log \left(T^{2} \log D\right)\right)=\mathscr{O}_{\epsilon}(\log \log D)$ which implies $\mathrm{M}_{q}(s)=\mathscr{\sigma}_{\epsilon}\left(s \log q \log (s \log q) 4^{\log ^{*} s}\right)=\mathscr{O}_{\epsilon}(\log D)$. In both cases, the algorithm performs $\mathscr{O}_{\epsilon}(T \log T \log D)=$ $\mathscr{O}_{\epsilon}(T \log n)$ operations in $\mathbb{F}_{q}$ (or in $\left.\mathbb{F}_{q^{s}}\right)$. Therefore the bit complexity is $\mathscr{O}_{\epsilon}(n \log n)$.

The following generalization will be used in our quasi-linear multiplication algorithm given in Section 4 ,
Corollary 3.6. Let $\left(F_{i}, G_{i}\right)_{0 \leq i<m}$ and $H$ be sparse polynomials over $\mathbb{F}_{q}$ of degree at most $D$ and sparsity at most $T$. We can verify if $\sum_{i=0}^{m-1} F_{i} G_{i}=H$, with error probability at most $\epsilon$ when they are different, in $\mathscr{O}_{\epsilon}(m(T \mid(\log D)+$ $\left.\left.T \log (m T \log D) \mathrm{M}_{q}(s)\right)\right)$ bit operations.

### 3.4 Analysis over the integers

In order to keep a quasi-linear time complexity over the integers, we must work over a prime finite field $\mathbb{F}_{q}$ to avoid the computation of too large integers. Indeed, $H_{p}(\alpha)$ could have a size of $p \log \alpha=\mathscr{O}_{\epsilon}\left(T^{2} \log D \log \alpha\right)$ which is not quasi-linear in the input size.

Theorem 3.7. One can adapt algorithm VERIFYSP to work over the integers. The bit complexity is $\mathscr{\sigma}_{\epsilon}(n \log n \log \log n)$, where $n=T(\log D+\log C)$ is the input size with $C=\max \left(\|F\|_{\infty},\|G\|_{\infty},\|H\|_{\infty}\right)$.
Proof. Before step 6, we choose a random prime number $q=\operatorname{RandomPrime}\left(\mu, \frac{5}{3 c_{2}}\right)$ with $\mu=$ $c_{2} \max \left(p, \ln \left(C^{2} T+C\right)\right)$ and we perform all the remaining step modulo $q$. Let us assume that the polynomial $\Delta=F G-H \in \mathbb{Z}[X]$ is nonzero. Our algorithm only fails in the following three cases: $p$ is such that $\Delta_{p}=\Delta \bmod X^{p}-1=0 ; q$ is such that $\Delta_{p} \equiv 0 \bmod q ; \alpha$ is a root of $\Delta_{p}$ in $\mathbb{F}_{q}$.

Using Proposition 2.4, $\Delta_{p}$ is nonzero with probability at least $1-\frac{10}{3 c_{1}}$. Actually, with the same probability, the proof of the proposition shows that at least one coefficient of $\Delta$ is preserved in $\Delta_{p}$. Since $\|\Delta\|_{\infty} \leq C^{2} T+C$, Proposition 2.5 ensures that $\Delta_{p} \not \equiv 0 \bmod q$ with probability at least $1-\frac{10}{3 c_{2}}$. Finally, $q$ has been chosen so that $\mathbb{F}_{q}$ has at least $c_{2} p$ elements whence $\alpha$ is not a root of $\Delta_{p} \bmod q$ with probability at least $1-\frac{1}{c_{2}}$. Altogether, taking $c_{1}, c_{2} \geq \frac{10}{3}$ such that $1-\left(1-\frac{10}{3 c_{1}}\right)\left(1-\frac{10}{3 c_{2}}\right)\left(1-\frac{1}{c_{2}}\right) \leq \epsilon$, our adaptation of VERIFYSP has an error probability at most $\epsilon$.

The reductions of $F, G$ and $H$ modulo $q$ add a term $\mathscr{O}(T I(\log C))$ to the complexity. Since operations in $\mathbb{F}_{q}$ have cost $\mathrm{I}(\log q)$, the complexity becomes $\mathscr{O}(T \mathrm{I}(\log D)+T \mathrm{I}(\log C)+T \log (T \log D) \mathrm{I}(\log q))$ bit operations. The first two terms are in $\mathscr{O}(n \log n)$. Moreover, $q=\mathscr{O}_{\epsilon}\left(\log \left(C^{2} T\right)+p\right)$ and $p=\mathscr{O}_{\epsilon}\left(T^{2} \log D\right)$, thus $\log q=$ $\mathscr{O}_{\epsilon}(\log (\log C+T \log D))=\mathscr{O}_{\epsilon}(\log n)$. Since $T \leq D, T \log (T \log D)=\mathscr{O}(n)$ and the third term in the complexity is $\mathscr{O}_{\epsilon}(n \log n \log \log n)$.

As over small finite fields, the complexity is actually better for very sparse polynomials.
Corollary 3.8. If $T=\Theta\left(\log ^{k} D\right)$ for some $k$, VERIFYSP has bit complexity $\mathscr{O}_{\epsilon}(n \log n)$.
Proof. If $T=\Theta\left(\log ^{k} D\right), T \log (T \log D)=\tilde{O}\left(\log ^{k} D\right)=o(n)$, thus the last term of the complexity in the proof of Theorem 3.7 becomes negligible with respect to the first two terms.

For the same reason as for finite fields, we extend the verification algorithm to a sum of product.
Corollary 3.9. Let $\left(F_{i}, G_{i}\right)_{0 \leq i<m}$ and $H$ be sparse polynomials of degree at most $D$, sparsity at most $T$, and height at most $C$. We can verify if $\sum_{i=0}^{m-1} F_{i} G_{i}=H$, with probability of error at most $\epsilon$ when they are different, in $\mathscr{O}_{\epsilon}(m T \mathbf{I}(\log D)+m T \mathbf{I}(\log C)+m T \log (m T \log D) \mathbf{I}(\log (m \log C+m T \log D)))$ bit operations.

We will only use this algorithm with $m=2$ and thus refer to it as $\operatorname{VERIFYSumSP}\left(H, F_{0}, G_{0}, F_{1}, G_{1}, \epsilon\right)$.

## 4 Sparse polynomial multiplication

Given two sparse polynomial $F$ and $G$, our algorithm aims at computing the product $H=F G$ through sparse polynomial interpolation. We avoid the difficulty of computing an a priori bound on the sparsity of $H$ needed for sparse interpolation by using our verification algorithm of Section3, Indeed, one can start with an arbitrary small sparsity and double it until the interpolated polynomial matches the product according to VERIFYSP.

The remaining difficulty is to be able to interpolate $H$ in quasi-optimal time given a sparsity bound, which is not yet achieved in the general case. In our case, we first analyze the complexity of Huang's sparse interpolation algorithm [10] when the input is a sum of sparse products. In order to obtain the desired complexity we develop a novel approach that interleaves two levels of Huang's algorithm.

### 4.1 Analysis of Huang's sparse interpolation

In [10] Huang proposes an algorithm that interpolates a sparse polynomial $H$ from its SLP representation, achieving the best known complexity for this problem, though it is not optimal. Its main idea is to use the dense polynomials $H_{p}=H \bmod X^{p}-1$ and $H_{p}^{\prime}=H^{\prime} \bmod X^{p}-1$ where $H^{\prime}$ is the derivative of $H$ and $p$ a small random prime. Indeed, if $c X^{e}$ is a term of $H$ that does not collide during the reduction modulo $X^{p}-1$, $H_{p}$ contains the monomial $c X^{e \bmod p}$ and $H_{p}^{\prime}$ contains $c e X^{e-1 \bmod p}$, hence $c$ and $e$ can be recovered by a mere division. Of course, the choice of $p$ is crucial for the method to work. It must be not too large to get a low complexity, but large enough for collisions to be sufficiently rare.

Lemma 4.1. There exists an algorithm FindTerms that takes as inputs a prime $p$, two polynomials $H_{p}=$ $H \bmod X^{p}-1, H_{p}^{\prime}=H^{\prime} \bmod X^{p}-1$, and bounds $D \geq \operatorname{deg}(H)$ and $C \geq\|H\|_{\infty}$ and it outputs an approximation $H^{*}$ of $H$ that contains at least all the monomials of $H$ that do not collide modulo $X^{p}-1$. Its bit complexity is $\mathscr{O}(T \mathrm{l}(\log C D))$, where $T=\# H$.

Proof. It is a straightforward adaptation of [10, Algorithm 3.4 (UTERMS)]. Here, taking $C$ as input allows us to only recover coefficients that are at most $C$ in absolute value and therefore we perform divisions with integers of bitlength at most $\log C D$.

Corollary 4.2. Let $H$ be a sparse polynomial such that $\# H \leq T$, $\operatorname{deg} H \leq D$ and $\|H\|_{\infty} \leq C$, and $0<\epsilon<1$. If $\lambda=\max \left(21, \frac{10}{3 \epsilon} T^{2} \ln D\right)$ and $p=\operatorname{RandomPRime}\left(\lambda, \frac{\epsilon}{2}\right)$, then with probability at least $1-\epsilon$, FindTerms ( $p, H \bmod X^{p}-1, H^{\prime} \bmod X^{p}-1, D, C$ ) returns $H$.
Proof. With probability at least $1-\epsilon$, no collision occurs in $H \bmod X^{p}-1$, and consequently neither in $H^{\prime}$ mod $X^{p}-1$, by Proposition 2.3. In this case FindTerms correctly computes $H$, according to Lemma 4.1.,

Theorem 4.3. There exists an algorithm InTERPSUMSP that takes as inputs $2 m$ sparse polynomials $\left(F_{i}, G_{i}\right)_{0 \leq i<m}$, three bounds $T \geq \# H, D>\operatorname{deg}(H)$ and $C \geq\|H\|_{\infty}$ where $H=\sum_{i=0}^{m-1} F_{i} G_{i}$, a constant $0<\mu<1$ and the list $\mathscr{P}$ of the first $2 N$ primes for $N=\max \left(1,\left\lfloor\frac{32}{5}(T-1) \log D\right\rfloor\right)$, and outputs $H$ with probability at least $1-\mu$.

Its bit complexity is $\tilde{\mathscr{O}}_{\mu}\left(m T_{1} \log D_{1} \log \left(C_{1} D_{1}\right)\right)$ where $T_{1}, D_{1}$ and $C_{1}$ are bounds on the sparsity, the degree and the height of $H$ and each $F_{i}$ and $G_{i}$.

Proof. It is identical to the proof of [10, Algorithm 3.9 (UIPoly)] taking into account that $H$ is not given as an SLP anymore but as $\sum_{i=0}^{m-1} F_{i} G_{i}$ where the polynomials $F_{i}$ and $G_{i}$ are given as sparse polynomials.

Remark 4.4. A finer analysis of algorithm InTERPSUMSP leads to a bit complexity $\mathscr{O}_{\mu}\left(m \log T_{1} \mathrm{M}_{\mathbb{Z}}\left(T_{1} \log D_{1} \log \left(T_{1} \log D_{1}\right), T_{1} C_{1} D_{1}\right)\right.$.

Remark 4.5. Even when InterpSumSP returns an incorrect polynomial, it has sparsity at most $2 T$, degree less than $D$ and coefficients bounded by $C$.

### 4.2 Multiplication

Our idea is to compute different candidates to $F G$ with a growing sparsity bound and to verify the result with VERIFYSP. Unfortunately, a direct call to InTERPSUMSP with the correct sparsity $T=\max (\# F, \# G, \#(F G))$ yields a bit complexity $\tilde{O}(T \log D \log C D)$ when the coefficients are bounded by $C$ and the degree by $D$. We shall remark that it is not nearly optimal since the input and output size are bounded by $T \log D+T \log C$.

To circumvent this difficulty, we first compute the reductions $F_{p}=F \bmod X^{p}-1$ and $G_{p}=G \bmod X^{p}-1$ of the input polynomials, as well as the reductions $F_{p}^{\prime}=F^{\prime} \bmod X^{p}-1$ and $G_{p}^{\prime}=G^{\prime} \bmod X^{p}-1$ of their
derivatives, for a random prime $p$ as in Corollary 4.2. The polynomials $H_{p}=F G \bmod X^{p}-1$ and $H_{p}^{\prime}=$ $(F G)^{\prime} \bmod X^{p}-1$ can be computed using InterpSumSP and VerifySP. Indeed, we first compute $F_{p} G_{p}$ by interpolation and then reduce it modulo $X^{p}-1$ to get $H_{p}$. Similarly for $H_{p}^{\prime}$ we first interpolate $F_{p}^{\prime} G_{p}+F_{p} G_{p}^{\prime}$ before its reduction. Finally we can compute the polynomial $F G$ from $H_{p}$ and $H_{p}^{\prime}$ using FindTerms according to Corollary 4.2. Our choice of $p$, which is polynomial in the input size, ensures that each call to InTERPSUMSP remains quasi-linear.

```
Algorithm 2 SpARSEPRODUCT
Input: \(F, G \in \mathbb{Z}[X] .0<\mu_{1}, \mu_{2}<1\) with \(\frac{\mu_{1}}{2} \leq \mu_{2}\).
Output: \(H \in \mathbb{Z}[X]\) s.t. \(H=F G\) with probability at least \(1-\mu_{1}\).
    \(t \leftarrow \max (\# F, \# G), D \leftarrow \operatorname{deg}(F)+\operatorname{deg}(G), C \leftarrow t\|F\|_{\infty}\|G\|_{\infty}\)
    \(\lambda \leftarrow \max \left(21, \frac{20}{3 \mu_{1}}(\# F \# G)^{2} \ln D\right), \mu^{*} \leftarrow \mu_{2}-\frac{\mu_{1}}{2}\)
    \(p \leftarrow \operatorname{RANDOMPRIME}\left(\lambda, \frac{\mu_{1}}{4}\right)\)
    \(F_{p} \leftarrow F \bmod X^{p}-1, G_{p} \leftarrow G \bmod X^{p}-1\)
    \(F_{p}^{\prime} \leftarrow F^{\prime} \bmod X^{p}-1, G_{p}^{\prime} \leftarrow G^{\prime} \bmod X^{p}-1\)
    repeat
        \(N \leftarrow \max \left(1,\left\lfloor\frac{32}{5}(t-1) \log D\right\rfloor\right)\)
        \(\mathscr{P} \leftarrow\{\) the first \(2 N\) primes in increasing order \(\}\)
        \(H_{1} \leftarrow \operatorname{InTERPSUMSP}\left(\left[\left(F_{p}, G_{p}\right)\right], t, 2 p, C, \frac{\mu^{*}}{2}, \mathscr{P}\right)\)
        \(H_{2} \leftarrow \operatorname{InterpSumSP}\left(\left[\left(F_{p}, G_{p}^{\prime}\right),\left(F_{p}^{\prime}, G_{p}\right)\right], t, 2 p, C D, \frac{\mu^{*}}{2}, \mathscr{P}\right)\)
        \(t \leftarrow 2 t\)
    until
        \(\operatorname{VERIFYSP}\left(H_{1}, F_{p}, G_{p}, \frac{\mu_{1}}{2}\right)\) and \(\quad \triangleright H_{1}=F_{p} G_{p}\)
        \(\operatorname{VERIFYSUMSP}\left(H_{2}, F_{p}, G_{p}^{\prime}, F_{p}^{\prime}, G_{p}, \frac{\mu_{1}}{2}\right) \quad \triangleright H_{2}=F_{p}^{\prime} G_{p}+F_{p} G_{p}^{\prime}\)
    \(H_{p} \leftarrow H_{1} \bmod X^{p}-1, H_{p}^{\prime} \leftarrow H_{2} \bmod X^{p}-1\).
    return FindTerms \(\left(p, H_{p}, H_{p}^{\prime}, D, C\right)\).
```

Lemmas 4.6 and 4.7 respectively provide the correctness and complexity bound of algorithm SPARSEPRODUCT. Together, they consequently form a proof of Theorem 1.1 by taking $\epsilon=\mu_{1}+\mu_{2}$. Note that this approach translates mutatis mutandis to the multiplication of sparse polynomials over $\mathbb{F}_{q}$ where the characteristic of $\mathbb{F}_{q}$ is larger than $D$.

Lemma 4.6. Let $F$ and $G$ be two sparse polynomials over $\mathbb{Z}$. Then algorithm SPARSEPRODUCT returns $F G$ with probability at least $1-\mu_{1}$.
Proof. Since $F G$ has sparsity at most $\# F \# G$, Corollary 4.2 implies that if $H_{p}=F G \bmod X^{p}-1$ and $H_{p}^{\prime}=$ $(F G)^{\prime} \bmod X^{p}-1$, the probability that FindTerms does not return $F G$ is at most $\frac{\mu_{1}}{2}$. The other reason for the result to be incorrect is that one of these equalities does not hold, which means that one of the two verifications fails. Since this happens with probability at most $\frac{\mu_{1}}{2}$, SPARSEPRODUCT returns $F G$ with probability at least $1-\mu_{1}$.

Lemma 4.7. Let $F$ and $G$ be two sparse polynomials over $\mathbb{Z}, T=\max (\# F, \# G, \#(F G)), D=\operatorname{deg}(F G)$, $C=\max \left(\|F\|_{\infty},\|G\|_{\infty},\|F G\|_{\infty}\right)$ and $\epsilon=\mu_{1}+\mu_{2}$. Then algorithm SPARSEPRODUCT has bit complexity $\tilde{\sigma}_{\epsilon}(T(\log D+\log C))$ with probability at least $1-\mu_{2}$. Writing $n=T(\log D+\log C)$, the bit complexity is $\mathscr{O}_{\epsilon}\left(n \log ^{2} n \log ^{2} T(\log T+\log \log n)\right)$.

Proof. In order to obtain the given complexity, we first need to prove that with high probability InTERPSUMSP never computes polynomials with a sparsity larger than 4\#(FG).

Let $T_{p}=\max \left(\#\left(F_{p} G_{p}\right), \#\left(F_{p} G_{p}^{\prime}+F_{p}^{\prime} G_{p}\right)\right)$. If $t \leq 2 T_{p}$ then the polynomials $H_{1}$ and $H_{2}$ satisfy $\# H_{1}, \# H_{2} \leq$ $4 T_{p}$ by Remark 4.5. Unfortunately, $T_{p}$ could be as large as $T^{2}$ and $t$ might reach values larger than $T_{p}$. We now prove that: (i) with probability at least $1-\mu^{*}$ the maximal value of $t$ during the algorithm is less than $2 T_{p}$; (ii) with probability at least $1-\frac{\mu_{1}}{2}, T_{p} \leq \#(F G)$. Together, this will prove that $\# H_{1}, \# H_{2} \leq 4 \#(F G)$ with probability at least $1-\mu^{*}-\frac{\mu_{1}}{2}=1-\mu_{2}$.
(i) As soon as $t \geq T_{p}$, Steps 9 and 10 compute both $F_{p} G_{p}$ and $F_{p} G_{p}^{\prime}+F_{p}^{\prime} G_{p}$ with probability at least $1-\mu^{*}$ by Theorem4.3. Since VErifySP never fails when the product is correct, the algorithm ends when $T_{p} \leq t<2 T_{p}$ with probability at least $1-\mu^{*}$.
(ii) Let us define the polynomials $\hat{F}_{p}$ and $\hat{G}_{p}$ obtained from $F_{p}$ and $G_{p}$ by replacing each nonzero coefficient by 1. The choice of $p$ in Step 3 ensures that with probability at least $1-\frac{\mu_{1}}{2}$ there is no collision in $\left(\hat{F}_{p} \hat{G}_{p}\right)$ mod
$X^{p}-1$ by applying Proposition 2.3 to the product $\hat{F}_{p} \hat{G}_{p}$. In that case, there is also no collision in $F_{p} G_{p} \bmod X^{p}-$ 1 and in $F_{p} G_{p}^{\prime}+F_{p}^{\prime} G_{p} \bmod X^{p}-1$ since $\operatorname{supp}\left(F_{p} G_{p}\right) \subset \operatorname{supp}\left(\hat{F}_{p} \hat{G}_{p}\right)$. Therefore, there are as many nonzero coefficients in $F_{p} G_{p}$ as in $F_{p} G_{p} \bmod X^{p}-1$, which is equal to $F G \bmod X^{p}-1$. Thus with probability at least $1-\frac{\mu_{1}}{2}$ we have $\#\left(F_{p} G_{p}\right)=\#(F G) \leq T$ and similarly $\#\left(F_{p}^{\prime} G_{p}+F_{p} G_{p}^{\prime}\right)=\#\left((F G)^{\prime}\right) \leq T$.

In the rest of the proof, we assume that the loop stops with $t \leq 2 T_{p}$ and that $T_{p} \leq T$. In particular, the number of iterations of the loop is $\mathscr{O}(\log T)$. Since $2 p=\mathscr{O}\left(\frac{1}{\epsilon} T^{4} \log D\right)$, Steps 9 and 10 have a bit complexity $\tilde{\sigma}_{\epsilon}(T \log p \log (p C D))=\tilde{\mathscr{O}}_{\epsilon}(T \log C D)$ by Theorem 4.3. Using Remark 4.5, VERIFYSP and VerifySumSP have polynomials of height at most $t C D$ as inputs. By Corollary 3.9, Step 12 has bit complexity $\mathscr{O}_{\epsilon}(T \log (T \log p) \mathrm{l}(\log C D))=\tilde{\mathscr{O}}_{\epsilon}(T \log C D)$. The list $\mathscr{P}$ can be computed incrementally, adding new primes when necessary. At the end of the loop, $\mathscr{P}$ contains $\mathscr{O}(T \log 2 p)$ primes, which means that it is computed in $\mathscr{\sigma}_{\epsilon}\left(T \log p \log ^{2}(T \log p) \log \log (T \log p)\right)$ bit operations [24, Chapter 18], that is $\tilde{\sigma}_{\epsilon}(T \log \log D)$ since $\log p=\mathscr{O}(\log (T \log D))$.

The total cost for the $\mathscr{O}(\log T)$ iterations of the loop is still $\tilde{\sigma}_{\epsilon}(T \log (C D))$. Step 14 runs in time $\mathscr{O}_{\epsilon}(T \mathrm{l}(\log C D))$ by Lemma 4.1 as the coefficients of $H_{p}^{\prime}$ are bounded by $2 T C^{2} D$ with $T \leq D$ and $\# H_{p}$, \# $H_{p}^{\prime} \leq$ \#H. Since other steps have negligible costs this yields a complexity of $\tilde{\mathscr{O}}_{\epsilon}(T(\log C+\log D)$ ) with probability at least $1-\mu_{2}$.

Using Remark 4.4 we can provide a more precise complexity for Steps 9 and 10 which is $\mathscr{O}_{\epsilon}\left(\log T \mathrm{M}_{\mathbb{Z}}(T \log p \log (T \log p), p D T C)\right)$ bit operations. It is easy to observe that the $\log T$ repetitions of these steps provide the dominant term in the complexity. A careful simplification yields a bit complexity $\mathscr{O}_{\epsilon}\left(n \log ^{2} n \log ^{2} T(\log T+\log \log n)\right.$ for SparseProduct where $n=T(\log D+\log C)$ bounds both input and output sizes.

### 4.3 Multivariate case

Using classical Kronecker substitution [24, Chapter 8] one can extend straightforwardly SparseProduct to multivariate polynomials. Let $F, G \in \mathbb{Z}\left[X_{1}, \ldots, X_{n}\right]$ with $\|F\|_{\infty},\|G\|_{\infty} \leq C$ and $\operatorname{deg}_{X_{i}}(F)+\operatorname{deg}_{X_{i}}(G)<d$. Writing $F_{u}(X)=F\left(X, X^{d}, \ldots, X^{d^{n-1}}\right)$ and $G_{u}(X)=G\left(X, X^{d}, \ldots, X^{d^{n-1}}\right)$, one can easily retrieve $F G$ from the univariate product $F_{u} G_{u}$. It is easy to remark that the Kronecker substitution preserve the sparsity and the height, and it increases the degree to $\operatorname{deg} F_{u}, \operatorname{deg} G_{u}<d^{n}$. Considering that $F$ and $G$ are sparse polynomial with at most $T$ nonzero terms, their sizes are at most $T(n \log d+\log C)$ which is exactly the sizes of $F_{u}$ and $G_{u}$. Since the Kronecker and inverse Kronecker substitutions cost $\tilde{\mathscr{O}}(T n \log d$ ) bit operations, one can compute $F_{u} G_{u}$ using SparseProduct within the following bit complexity.

Corollary 4.8. There exists an algorithm that takes as inputs $F, G \in \mathbb{Z}\left[X_{1}, \ldots, X_{n}\right]$ and $0<\epsilon<1$, and computes $F G$ with probability at least $1-\epsilon$, using $\tilde{\sigma}_{\epsilon}(T(n \log d+\log C))$ bit operations where $T=\max (\# F, \# G, \#(F G))$, $d=\max _{i}\left(\operatorname{deg}_{X_{i}} F G\right)$ and $C=\max \left(\|F\|_{\infty},\|G\|_{\infty}\right)$.

Over a finite field $\mathbb{F}_{q^{s}}$ for some prime $q$, the previous technique requires that $q>d^{n}$ since SpARSEPRODUCT requires $q$ to be larger than the degree. The randomized Kronecker substitution method introduced by Arnold and Roche [3] allows to apply SPARSEPRODUCT to fields of smaller characteristic. The idea is to define univariate polynomials $F_{s}(X)=F\left(X^{s_{1}}, \ldots, X^{s_{n}}\right)$ and $G_{s}(X)=G\left(X^{s_{1}}, \ldots, X^{s_{n}}\right)$ for some random vector $\vec{s}=\left(s_{1}, \ldots, s_{n}\right)$ such that these polynomials have much smaller degrees than those obtained with classical Kronecker substitution. As a result, we obtain an algorithm that works for much smaller $q$ of order $\tilde{O}(n d \# F \# G)$.

Our approach is to first use some randomized Kronecker substitutions to estimate the sparsity of $F G$ by computing the sparsity of $H_{s}=F_{s} G_{s}$ for several distinct random vectors $\vec{s}$. With high probability, the maximal sparsity is close to the one of $F G$. Then, we use this information to provide a bound to some (multivariate) sparse interpolation algorithm. Note that our approach is inspired from [11] that slightly improves randomized Kronecker substitution.

Lemma 4.9. Let $H \in \mathbb{F}_{q^{s}}\left[X_{1}, \ldots, X_{n}\right]$ of sparsity $T$, and $\vec{s}$ be a vector chosen uniformly at random in $S^{n}$ where $S \subset \mathbb{N}$ is finite. The expected sparsity of $H_{s}(X)=H\left(X^{s_{1}}, \ldots, X^{s_{n}}\right)$ is at least $T\left(1-\frac{T-1}{\# S}\right)$.

Proof. If we fix two distinct exponent vectors $\vec{e}_{u}$ and $\vec{e}_{v}$ of $H$, they collide in $H_{s}$ if and only if $\vec{e}_{u} \cdot \vec{s}=\vec{e}_{v} \cdot \vec{s}$. Since $\vec{e}_{u} \neq \vec{e}_{v}$, they differ at least on one component, say $e_{u, j_{0}} \neq e_{v, j_{0}}$. The equality $\vec{e}_{u} \cdot \vec{s}=\vec{e}_{v} \cdot \vec{s}$ is then equivalent to

$$
s_{j_{0}}=\sum_{j \neq j_{0}} \frac{e_{v, j}-e_{u, j}}{e_{u, j_{0}}-e_{v, j_{0}}} s_{j}
$$

Writing $Y$ for the right-hand side of this equation we have

$$
\operatorname{Pr}\left[\vec{e}_{u} \cdot \vec{s}=\vec{e}_{v} \cdot \vec{s}\right]=\operatorname{Pr}\left[s_{j_{0}}=Y\right]=\sum_{y} \operatorname{Pr}\left[s_{j_{0}}=Y \mid Y=y\right] \operatorname{Pr}[Y=y]
$$

where the (finite) sum ranges over all possible values $y$ of $Y$. Since $s_{j_{0}}$ is chosen uniformly at random in $S$, $\operatorname{Pr}\left[s_{j_{0}}=Y \mid Y=y\right]=\operatorname{Pr}\left[s_{j_{0}}=y\right] \leq 1 / \# S$ and the probability that $\vec{e}_{u}$ and $\vec{e}_{v}$ collide is at most $1 / \# S$. This implies that the expected number of vectors that collide is at most $T(T-1) / \# S$.

Corollary 4.10. Let $H$ be as in Lemma 4.9 and $\vec{v}_{1}, \ldots, \vec{v}_{\ell} \in S^{n}$ be some vectors chosen uniformly and independently at random. Then $\operatorname{Pr}\left[\max _{i} \# H_{v_{i}} \leq T\left(1-2 \frac{T-1}{\# S}\right)\right] \leq 1 / 2^{\ell}$.
Proof. For each $\vec{v}_{i}$, the expected number of terms that collide in $H_{v_{i}}(X)$ is at most $T(T-1) / \# S$ by Lemma 4.9 , Using Markov's inequality, we have $\operatorname{Pr}\left[\# H_{v_{i}} \leq T-2 T(T-1) / \# S\right] \leq 1 / 2$. Since the vectors $\vec{v}_{i}$ are independent, the result follows.

```
Algorithm 3 SparsityEstimate
Input: \(F, G \in \mathbb{F}_{q^{s}}\left[X_{1}, \ldots, X_{n}\right], 0<\epsilon<1, \lambda>1\).
Output: An integer \(t\) such that \(t \leq \lambda \#(F G)\).
    \(N \leftarrow\left\lceil 2 \frac{\# F \# G-1}{1-1 / \lambda}\right\rceil, \ell \leftarrow\left\lceil\log \frac{2}{\epsilon}\right\rceil\).
    \(t^{\prime} \leftarrow 0, \mu \leftarrow \frac{\epsilon}{4 \ell}\).
    repeat \(\ell\) times
        \(\vec{s} \leftarrow\) random element of \(\{0, \ldots, N-1\}^{n}\).
        \(F_{s} \leftarrow F\left(X^{s_{1}}, \ldots, X^{s_{n}}\right), G_{s} \leftarrow G\left(X^{s_{1}}, \ldots, X^{s_{n}}\right)\)
        \(H_{s} \leftarrow \operatorname{SPARSEPRODUCT}\left(F_{s}, G_{s}, \mu, \mu\right)\)
        \(t^{\prime} \leftarrow \max \left(t^{\prime}, \# H_{s}\right)\)
    return \(\lambda t^{\prime}\).
```

Lemma 4.11. Algorithm SPARSITYEstimate is correct when $q \geq \frac{4 D \# F \# G}{1-1 / \lambda}$ where $D=\max (\operatorname{deg} F, \operatorname{deg} G)$. With probability at least $1-\epsilon$, it returns an integer $t \geq \#(F G)$ using $\tilde{\mathscr{O}}_{\epsilon}(T(n \log d+s \log q))$ bit operations where $T=\max (\#(F G), \# F, \# G)$ and $d=\max _{i}\left(\operatorname{deg}_{X_{i}} F G\right)$.

Proof. Since each polynomial $H_{s}$ has sparsity at most $\#(F G)$, SparsityEstimate returns an integer bounded by $\lambda \#(F G)$. SPARSEProduct can be used in step 5 since $\operatorname{deg} H_{s}=\operatorname{deg} F_{s}+\operatorname{deg} G_{s} \leq 2 N D \leq q$ by the definition of $N$. Assuming that SparseProduct returns no incorrect answer during the loop, Corollary 4.10 applied to the product $F G$ implies that $t^{\prime} \geq \#(F G)(1-2(\#(F G)-1) / N)$ with probability $\geq 1-\epsilon / 2$ at the end of the loop. By definition of $N$ and since $\# F \# G \geq \#(F G), t^{\prime} \geq \#(F G) / \lambda$. Taking into account the probability of failure of SPARSEPRODUCT, the probability that $\lambda t^{\prime} \geq \#(F G)$ is at least $1-\frac{3 \epsilon}{4}$.

The computation of $F_{s}$ and $G_{s}$ requires $\mathscr{O}(T n I(\log \max (d, N))+T s \log q)$ bit operations in Step 5. Since $\max \left(\# F_{s}, \# G_{s}, \# H_{s}\right) \leq T$ and $\operatorname{deg} H_{s}=\mathscr{O}\left(n d T^{2}\right)$ in Step 6, the bit complexity of each call to SparseProduct is $\tilde{\mathscr{O}}_{\mu}(T(\log (n d)+s \log q))$ with probability at least $1-\mu$ using Lemma 4.7. Therefore, SparsityEstimate requires $\tilde{\mathscr{O}}_{\epsilon}(T(n \log d+s \log q))$ bit operations with probability at least $1-\epsilon / 4$. Together with the probability of failure this concludes the proof.

Theorem 4.12. There exists an algorithm that takes as inputs two sparse polynomials $F$ and $G$ in $\mathbb{F}_{q^{s}}\left[X_{1}, \ldots, X_{n}\right]$ and $0<\epsilon<1$ that returns the product $F G$ in $\tilde{\sigma}_{\epsilon}(n T(\log d+s \log q))$ bit operations with probability at least $1-\epsilon$, where $T=\max (\# F, \# G, \#(F G)), d=\max _{i}\left(\operatorname{deg}_{X_{i}} F G\right), D=\operatorname{deg} F G$ and assuming that $q=\Omega(D \# F \# G+$ $D T \log D \log (T \log D))$.
Proof. The algorithm computes an estimate $t$ on the sparsity of $F G$ using $\operatorname{SparsityEstimate}\left(F, G, \frac{\epsilon}{2}, \lambda\right)$ for some constant $\lambda$. The second step interpolates $F G$ using Huang and Gao's algorithm [11, Algorithm 5 (MULPoLYSI)] which is parameterized by a univariate sparse interpolation algorithm. Originally, its inputs are a polynomial given as a blackbox and bounds on its degree and sparsity. In our case, the blackbox is replaced by $F$ and $G$, the sparsity bound is $t$ and the univariate interpolation algorithm is SparseProduct.

The algorithm MULPOLYSI requires $\mathscr{\sigma}_{\epsilon}\left(n \log t+\log ^{2} t\right)$ interpolation of univariate polynomials with degree $\tilde{\mathscr{O}}(t D)$ and sparsity at most $t$. Each interpolation with SPARSEPRODUCT is done with $\mu_{1}, \mu_{2}$ such that $\mu_{1}+\mu_{2}=$ $\epsilon / 4(n+1) \log t$, so that MULPOLYSI returns the correct answer in $\tilde{\mathscr{O}}_{\epsilon}(n T(\log d+s \log q)$ ) bit operations with probability at least $1-\frac{\epsilon}{2}$ [11, Theorem 6]. Altogether, our two-step algorithm returns the correct answer using $\tilde{\mathscr{O}}_{\epsilon}(n T(\log d+s \log q))$ bit operations with probability at least $1-\epsilon$. The value of $q$ is such that it bounds the degrees of the univariate polynomials returned by SPARSEPRODUCT during the algorithm.

### 4.4 Small characteristic

We now consider the case of sparse polynomial multiplication over a field $\mathbb{F}_{q^{s}}$ with characteristic smaller than the degree of the product $F G$ (or, in the multivariate case, smaller than the degree of the product after randomized Kronecker substitution). We can no more use Huang's interpolation algorithm since it uses the derivative to encode the exponents into the coefficients and thus it only keeps the value of the exponents modulo $q$. Our idea to circumvent this problem is similar to the one in [2] that is to rather consider the polynomials over $\mathbb{Z}$ before calling our algorithm SparseProduct.

The following proposition is only given for the multivariate case as it encompasses univariate's one. It matches exactly with the complexity result given by Arnold and Roche [2].

Proposition 4.13. There exists an algorithm that takes as inputs two sparse polynomials $F$ and $G$ in $\mathbb{F}_{q^{s}}\left[X_{1}, \ldots, X_{n}\right]$ and $0<\epsilon<1$ that returns the product $F G$ in $\tilde{\sigma}_{\epsilon}(S(n \log d+s \log q))$ bit operations with probability at least $1-\epsilon$, where $S$ is the structural sparsity of $F G$ and $d=\max _{i}\left(\operatorname{deg}_{X_{i}} F G\right)$.
Proof. If $s=1$, the coefficients of $F$ and $G$ map easily to the integers in $\{0, \ldots, q-1\}$. Therefore, the product $F G$ can be obtained by using an integer sparse polynomial multiplication, as the one in Corollary4.8, followed by some reductions modulo $q$. Unfortunately, mapping the multiplication over the integers implies that the cancellations that could have occur in $\mathbb{F}_{q}$ do not hold anymore. Consequently, the support of the product in $\mathbb{Z}$ before modular reduction is exactly the structural support of $F G$.

If $s>1$, the coefficients of $F$ and $G$ are polynomials over $\mathbb{F}_{q}$ of degree $s-1$. As previously, mapping $\mathbb{F}_{q}$ to integers, $F$ and $G$ can be seen as $F_{Y}, G_{Y} \in \mathbb{Z}[Y]\left[X_{1}, \ldots, X_{n}\right]$ where the coefficients are polynomials in $\mathbb{Z}[Y]$ of degree at most $s-1$ and height at most $q-1$.

If $T=\max (\# F, \# G)$, the coefficients of $F_{Y} G_{Y}$ are polynomials of degree at most $2 s-2$ and height at most $T s q^{2}$. Therefore, the product $F G \in \mathbb{F}_{q^{s}}$ can be computed by: (i) computing $F_{B}, G_{B} \in \mathbb{Z}\left[X_{1}, \ldots, X_{n}\right]$ by evaluating the coefficients of $F_{Y}$ and $G_{Y}$ at $B=T s q^{2}$ (Kronecker substitution); (ii) computing the product $H_{B}=F_{B} G_{B}$; (iii) writing the coefficients of $H_{B}$ in base $B$ to obtain $H_{Y}=F_{Y} G_{Y}$ (Kronecker segmentation); (iv) and finally mapping back the coefficients of $H_{Y}$ from $\mathbb{Z}[Y]$ to $\mathbb{F}_{q^{s}}$.

Similarly as the case $s=1, H_{B}$ and then $H_{Y}$ have at most $S$ nonzero coefficients. The Kronecker substitutions in (i) require $\tilde{O}(T s \log q)$ bit operations, while the Kronecker segmentations in (iii) need $\tilde{\mathscr{O}}(S s \log q)$ bit operations. In ( $i v$ ) we first compute $S s$ reductions modulo $q$ on integers smaller than B , and then $S$ polynomial divisions in $\mathbb{F}_{q}[Y]$ with polynomial of degree $\mathscr{O}(s)$. Thus, it can be done in $\tilde{\mathscr{O}}(S s \log q)$ bit operations. Finally the computation in (ii) is dominant and it requires $\tilde{\mathscr{O}}_{\epsilon}(S(n \log d+s \log q))$ bit operations with probability at least $1-\epsilon$ using Corollary 4.8.

## References

[1] Arnold, A. Sparse polynomial interpolation and testing. PhD thesis, University of Waterloo, 2016.
[2] Arnold, A., And Roche, D. Output-sensitive algorithms for sumset and sparse polynomial multiplication. In ISSAC'15 (2015), ACM, pp. 29-36.
[3] Arnold, A., and Roche, D. S. Multivariate sparse interpolation using randomized Kronecker substitutions. In ISSAC'14 (2014), ACM, pp. 35-42.
[4] Ben-Or, M., And Tiwari, P. A Deterministic Algorithm for Sparse Multivariate Polynomial Interpolation. In STOC'88 (1988), ACM, pp. 301-309.
[5] Cole, R., and Hariharan, R. Verifying candidate matches in sparse and wildcard matching. In STOC'02 (2002), ACM, pp. 592-601.
[6] Giorgi, P. A probabilistic algorithm for verifying polynomial middle product in linear time. Information Processing Letters 139 (2018), 30-34.
[7] Harvey, D., and van der hoeven, J. Faster polynomial multiplication over finite fields using cyclotomic coefficient rings. J. Complexity 54 (2019).
[8] Harvey, D., and van der Hoeven, J. Integer multiplication in time O(n $\log n$ ). manuscript, 2019.
[9] Hoeven, J. v. D., and Lecerf, G. On the bit-complexity of sparse polynomial and series multiplication. J. Symbolic Computation 50 (2013), 227-254.
[10] HUANG, Q.-L. Sparse polynomial interpolation over fields with large or zero characteristic. In ISSAC'19 (2019), ACM, pp. 219-226.
[11] Huang, Q.-L., And Gao, X.-S. Revisit Sparse Polynomial Interpolation Based on Randomized Kronecker Substitution. In CASC'19 (2019), Springer, pp. 215-235.
[12] Johnson, S. C. Sparse polynomial arithmetic. ACM SIGSAM Bulletin 8, 3 (1974), 63-71.
[13] Monagan, M., and Pearce, R. Parallel sparse polynomial multiplication using heaps. In ISSAC'09 (2009), ACM, p. 263.
[14] Monagan, M., and Pearce, R. Sparse polynomial division using a heap. J. Symbolic Computation 46, 7 (2011), 807-822.
[15] NАкоs, V. Nearly optimal sparse polynomial multiplication, 2019.
[16] Plaisted, D. A. New NP-hard and NP-complete polynomial and integer divisibility problems. Theoretical Computer Science 31, 1 (1984), 125-138.
[17] Prony, R. Essai expérimental et analytique sur les lois de la dilatabilité de fluides élastique et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de l'alkool, à différentes températures. J. de l'École Polytechnique 1 (Floréal et Prairial III (1795)), 24-76.
[18] Roche, D. S. Chunky and equal-spaced polynomial multiplication. J. Symbolic Computation 46, 7 (2011), 791-806.
[19] Roche, D. S. What can (and can't) we do with sparse polynomials? In ISSAC'18 (2018), ACM, pp. 2530.
[20] Rosser, J. B., And Schoenfeld, L. Approximate formulas for some functions of prime numbers. Illinois J. Math. 6, 1 (1962), 64-94.
[21] Shoup, V. A Computational Introduction to Number Theory and Algebra, second ed. Cambridge University Press, 2008.
[22] van der Hoeven, J., Lebreton, R., and Schost, É. Structured FFT and TFT: Symmetric and Lattice Polynomials. In ISSAC'13 (2013), ACM, pp. 355-362.
[23] Van der Hoeven, J., and Lecerf, G. On the Complexity of Multivariate Blockwise Polynomial Multiplication. In ISSAC'12 (2012), ACM, pp. 211-218.
[24] von zur Gathen, J., and Gerhard, J. Modern Computer Algebra, 3rd ed. Cambridge University Press, 2013.
[25] Yao, A. C.-C. On the Evaluation of Powers. SIAM Journal on Computing 5, 1 (1976), 100-103.

