

FASTER INTEGER MULTIPLICATION USING SHORT LATTICE VECTORS

DAVID HARVEY AND JORIS VAN DER HOEVEN

ABSTRACT. We prove that n -bit integers may be multiplied in $O(n \log n 4^{\log^* n})$ bit operations. This complexity bound had been achieved previously by several authors, assuming various unproved number-theoretic hypotheses. Our proof is unconditional, and depends in an essential way on Minkowski's theorem concerning lattice vectors in symmetric convex sets.

1. INTRODUCTION

Let $M(n)$ denote the number of bit operations required to multiply two n -bit integers, where “bit operations” means the number of steps on a deterministic Turing machine with a fixed, finite number of tapes [17] (our results also hold in the Boolean circuit model). Let $\log^* x$ denote the iterated natural logarithm, i.e., $\log^* x := \min \{j \in \mathbb{N} : \log^{\circ j} x \leq 1\}$, where $\log^{\circ j} x := \log \cdots \log x$ (iterated j times). The main result of this paper is an algorithm achieving the following bound.

Theorem 1.1. *We have $M(n) = O(n \log n 4^{\log^* n})$.*

The first complexity bound for $M(n)$ of the form $O(n \log n K^{\log^* n})$ was established by Fürer [7, 8], for an unspecified constant $K > 1$. His algorithm reduces a multiplication of size n to many multiplications of size exponentially smaller than n , which are then handled recursively. The number of recursion levels is $\log^* n + O(1)$, and the constant K measures the “expansion factor” at each recursion level.

The first explicit value for K , namely $K = 8$, was given by Harvey, van der Hoeven and Lecerf [12]. Their method is somewhat different to Fürer, but still carries out an exponential size reduction at each recursion level. One may think of the constant $K = 8$ as being built up of three factors of 2, each coming from a different source.

The first factor of 2 arises from the need to perform both forward and inverse DFTs (discrete Fourier transforms) at each recursion level. This is a feature common to all of the post-Fürer algorithms, suggesting that significantly new ideas will be needed to do any better than $O(n \log n 2^{\log^* n})$.

The second factor of 2 arises from coefficient growth: a product of polynomials with k -bit integer coefficients has coefficients with at least $2k$ bits. This factor of 2 also seems difficult to completely eliminate, although Harvey and van der Hoeven have recently made some progress [11]: they achieve $K = 4\sqrt{2} \approx 5.66$ by arranging that, in effect, the coefficient growth only occurs at every second recursion level. This was the best known unconditional value of K prior to the present paper.

Harvey was supported by the Australian Research Council (grants DP150101689 and FT160100219).

The final factor of 2 occurs because the algorithm works over \mathbb{C} : when multiplying complex coefficients with say β significant bits, the algorithm first computes a full 2β -bit product, and then truncates to β bits. More precisely, after splitting the β -bit inputs into m exponentially smaller chunks, and encoding them into polynomials of degree m , the algorithm must compute the full product of degree $2m$, even though essentially only m coefficients are needed to resolve β significant bits of the product. Again, this factor of 2 has been the subject of a recent attack: Harvey has shown [9] that it is possible to work modulo a polynomial of degree only m , at the expense of increasing the working precision by a factor of $3/2$. This leads to an integer multiplication algorithm achieving $K = 6$.

Another way of attacking this last factor of 2 is to replace the coefficient ring \mathbb{C} by a finite ring $\mathbb{Z}/q\mathbb{Z}$ for a suitable integer q . By choosing q with some special structure, it may become possible to convert a multiplication modulo q directly into a polynomial multiplication modulo some polynomial of degree m , rather than $2m$. Three algorithms along these lines have been proposed.

First, Harvey, van der Hoeven and Lecerf suggested using *Mersenne primes*, i.e., primes of the form $q = 2^k - 1$, where k is itself prime [12, §9]. They convert multiplication in $\mathbb{Z}/q\mathbb{Z}$ to multiplication in $\mathbb{Z}[y]/(y^m - 1)$, where m is a power of two. Because k is not divisible by m , the process of splitting an element of $\mathbb{Z}/q\mathbb{Z}$ into m chunks is somewhat involved, and depends on a variant of the Crandall–Fagin algorithm [6].

Covanov and Thomé [5] later proposed using *generalised Fermat primes*, i.e., primes of the form $q = r^m + 1$, where m is a power of two and r is a small even integer. Here, multiplication in $\mathbb{Z}/q\mathbb{Z}$ is converted to multiplication in $\mathbb{Z}[y]/(y^m + 1)$. The splitting procedure consists of rewriting an element of $\mathbb{Z}/q\mathbb{Z}$ in base r , via fast radix-conversion algorithms.

Finally, Harvey and van der Hoeven [10] proposed using *FFT primes*, i.e., primes of the form $q = a \cdot 2^k + 1$, where a is small. They reduce multiplication in $\mathbb{Z}/q\mathbb{Z}$ to multiplication in $\mathbb{Z}[y]/(y^m + a)$ via a straightforward splitting of the integers into m chunks, where m is a power of two. Here the splitting process is trivial, as k may be chosen to be divisible by m .

These three algorithms all achieve $K = 4$, subject to plausible but unproved conjectures on the distribution of the relevant primes. Unfortunately, in all three cases, it is not even known that there are infinitely many primes of the required form, let alone that there exist a sufficiently high density of them to satisfy the requirements of the algorithm.

The main technical novelty of the present paper is a splitting procedure that works for an almost *arbitrary* modulus q . The core idea is to introduce an alternative representation for elements of $\mathbb{Z}/q\mathbb{Z}$: we represent them as expressions $a_0 + a_1\theta + \dots + a_{m-1}\theta^{m-1}$, where θ is some fixed $2m$ -th root of unity in $\mathbb{Z}/q\mathbb{Z}$, and where the a_i are small integers, of size roughly $q^{1/m}$. Essentially the only restriction on q is that $\mathbb{Z}/q\mathbb{Z}$ must contain an appropriate $2m$ -th root of unity. We will see that Linnik’s theorem is strong enough to construct suitable such moduli q .

In Section 2 we show that the cost of multiplication in this representation is only a constant factor worse than for the usual representation. The key ingredient is Minkowski’s theorem on lattice vectors in symmetric convex sets. We also give algorithms for converting between this representation and the standard representation. The conversions are not as fast as one might hope — in particular, we do not

know how to carry them out in quasilinear time — but surprisingly this turns out not to affect the overall complexity, because in the main multiplication algorithm we perform the conversions only infrequently.

Then in Sections 3 and 4 we prove Theorem 1.1, using an algorithm that is structurally very similar to [10]. We make no attempt to minimise the implied big- O constant in Theorem 1.1; our goal is to give the simplest possible proof of the asymptotic bound, without any regard for questions of practicality.

An interesting question is whether it is possible to combine the techniques of the present paper with those of [11] to obtain an algorithm achieving $K = 2\sqrt{2} \approx 2.83$. Our attempts in this direction have so far been unsuccessful. One might also ask if the techniques of this paper can be transferred to the case of multiplication of polynomials of high degree in $\mathbb{F}_p[x]$. However, this is not so interesting, because an unconditional proof of the bound corresponding to $K = 4$ in the polynomial case is already known [11].

Throughout the paper we use the following notation. We write $\lg n := \lceil \log_2 n \rceil$ for $n \geq 2$, and for convenience put $\lg 1 := 1$. We define $M_{SS}(n) = Cn \lg n \lg \lg n$, where $C > 0$ is some constant so that the Schönhage–Strassen algorithm multiplies n -bit integers in at most $M_{SS}(n)$ bit operations [18]. This function satisfies $nM_{SS}(m) \leq M_{SS}(nm)$ for any $n, m \geq 1$, and also $M_{SS}(dm) = O(M_{SS}(m))$ for fixed d . An n -bit integer may be divided by m -bit integer, producing quotient and remainder, in time $O(M_{SS}(\max(n, m)))$ [19, Ch. 9]. We may transpose an $n \times m$ array of objects of bit size b in $O(bnm \lg \min(n, m))$ bit operations [3, Appendix]. Finally, we occasionally use Xylouris’s refinement of Linnik’s theorem [20], which states that for any relatively prime positive integers a and n , the least prime in the arithmetic progression $p = a \pmod{n}$ satisfies $p = O(n^{5.2})$.

2. θ -REPRESENTATIONS

Throughout this section, fix an integer $q \geq 2$ and a power of two m such that

$$(2.1) \quad m \leq \frac{\log_2 q}{(\lg \lg q)^2}, \quad \text{or equivalently,} \quad q^{1/m} \geq 2^{(\lg \lg q)^2},$$

and assume we are given some $\theta \in \mathbb{Z}/q\mathbb{Z}$ such that $\theta^m = -1$.

For a polynomial $F = F_0 + F_1y + \dots + F_{m-1}y^{m-1} \in \mathbb{Z}[y]/(y^m + 1)$, define $\|F\| := \max_i |F_i|$. This norm satisfies $\|FG\| \leq m\|F\|\|G\|$ for any $F, G \in \mathbb{Z}[y]/(y^m + 1)$.

Definition 2.1. Let $u \in \mathbb{Z}/q\mathbb{Z}$. A θ -representation for u is a polynomial $U \in \mathbb{Z}[y]/(y^m + 1)$ such that $U(\theta) = u \pmod{q}$ and $\|U\| \leq mq^{1/m}$.

Example 2.2. Let $m = 4$ and

$$\begin{aligned} q &= 3141592653589793238462833, \\ \theta &= 2542533431566904450922735 \pmod{q}, \\ u &= 2718281828459045235360288 \pmod{q}. \end{aligned}$$

The coefficients in a θ -representation must not exceed $mq^{1/m} \approx 5325341.46$. Two examples of θ -representations for u are

$$\begin{aligned} U(y) &= -3366162y^3 + 951670y^2 - 5013490y - 3202352, \\ U(y) &= -4133936y^3 + 1849981y^2 - 5192184y + 1317423. \end{aligned}$$

By (2.1), the number of bits required to store $U(y)$ is at most

$$m(\log_2(mq^{1/m}) + O(1)) = \lg q + O(m \lg m) = \left(1 + \frac{O(1)}{\lg \lg q}\right) \lg q,$$

so a θ -representation incurs very little overhead in space compared to the standard representation by an integer in the interval $0 \leq x < q$.

Our main tool for working with θ -representations is the *reduction algorithm* stated in Lemma 2.9 below. Given a polynomial $F \in \mathbb{Z}[y]/(y^m + 1)$, whose coefficients are up to about twice as large as allowed in a θ -representation, the reduction algorithm computes a θ -representation for $F(\theta)$ (up to a certain scaling factor, discussed further below). The basic idea of the algorithm is to precompute a nonzero polynomial $P(y)$ such that $P(\theta) = 0 \pmod{q}$, and then to subtract an appropriate multiple of $P(y)$ from $F(y)$ to make the coefficients small.

After developing the reduction algorithm, we are able to give algorithms for basic arithmetic on elements of $\mathbb{Z}/q\mathbb{Z}$ given in θ -representation (Proposition 2.15), a more general reduction algorithm for inputs of arbitrary size (Proposition 2.17), and algorithms for converting between standard and θ -representation (Proposition 2.18 and Proposition 2.20).

We begin with two results that generate certain precomputed data necessary for the main reduction step.

Lemma 2.3. *In $q^{1+o(1)}$ bit operations, we may compute a nonzero polynomial $P \in \mathbb{Z}[y]/(y^m + 1)$ such that $P(\theta) = 0 \pmod{q}$ and $\|P\| \leq q^{1/m}$.*

Proof. We first establish existence of a suitable $P(y)$. Let $\bar{\theta}^i$ denote a lift of θ^i to \mathbb{Z} , and consider the lattice $\Lambda \subset \mathbb{Z}^m$ spanned by the rows of the $m \times m$ integer matrix

$$A = \begin{pmatrix} q & 0 & 0 & \cdots & 0 \\ -\bar{\theta} & 1 & 0 & \cdots & 0 \\ -\bar{\theta}^2 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \\ -\bar{\theta}^{m-1} & 0 & 0 & \cdots & 1 \end{pmatrix}$$

Every vector $(a_0, \dots, a_{m-1}) \in \Lambda$ satisfies the equation $a_0 + \dots + a_{m-1}\theta^{m-1} = 0 \pmod{q}$. The volume of the fundamental domain of Λ is $\det A = q$. The volume of the closed convex symmetric set $\Sigma := \{|a_i| \leq q^{1/m}\} \subset \mathbb{R}^m$ is $(2q^{1/m})^m = 2^m q$, so by Minkowski's theorem (see for example [13, Ch. V, Thm. 3]), there exists a nonzero vector (a_0, \dots, a_{m-1}) in $\Lambda \cap \Sigma$. The corresponding polynomial $P(y) := a_0 + \dots + a_{m-1}y^{m-1}$ then has the desired properties.

To actually compute $P(y)$, we simply perform a brute-force search. By (2.1) there are at most $(2q^{1/m} + 1)^m \leq (3q^{1/m})^m = 3^m q < q^{1+o(1)}$ candidates to test. Enumerating them in lexicographical order, we can easily evaluate $P(\theta) \pmod{q}$ in an average of $O(\lg q)$ bit operations per candidate. \square

Example 2.4. Continuing Example 2.2, the coefficients of $P(y)$ must not exceed $q^{1/m} \approx 1331335.36$. A suitable polynomial $P(y)$ is given by

$$P(y) = -394297y^3 - 927319y^2 + 1136523y - 292956.$$

Remark 2.5. The computation of $P(y)$ is closely related to the problem of finding an element of small norm in the ideal of the ring $\mathbb{Z}[\zeta_{2m}]$ generated by p and $\zeta_{2m} - \theta$, where ζ_{2m} denotes a primitive $2m$ -th root of unity.

Remark 2.6. The poor exponential-time complexity of Lemma 2.3 can probably be improved, by taking advantage of more sophisticated lattice reduction or shortest vector algorithms, but we were not easily able to extract a suitable result from the literature. For example, LLL is not guaranteed to produce a short enough vector [14], and the Micciancio–Voulgaris exact shortest vector algorithm [15] solves the problem for the Euclidean norm rather than the uniform norm. In any case, this has no effect on our main result.

Lemma 2.7. *Assume that $P(y)$ has been precomputed as in Lemma 2.3. Let r be the smallest prime exceeding $2m^2q^{1/m}$ such that $r \nmid q$ and such that $P(y)$ is invertible in $(\mathbb{Z}/r\mathbb{Z})[y]/(y^m + 1)$. Then $r = O(m^2q^{1/m})$, and in $q^{1+o(1)}$ bit operations we may determine r and compute a polynomial $J \in \mathbb{Z}[y]/(y^m + 1)$ such that $J(y)P(y) = 1 \pmod{r}$ and $\|J\| \leq r$.*

Proof. Let $R \in \mathbb{Z}$ be the resultant of $P(y)$ (regarded as a polynomial in $\mathbb{Z}[y]$) and $y^m + 1$. The primes r dividing R are exactly the primes for which $P(y)$ fails to be invertible in $(\mathbb{Z}/r\mathbb{Z})[y]/(y^m + 1)$. Therefore our goal is to find a prime $r > 2m^2q^{1/m}$ such that $r \nmid Rq$.

Since m is a power of two, $y^m + 1$ is a cyclotomic polynomial and hence irreducible in $\mathbb{Q}[y]$. Thus $y^m + 1$ and $P(y)$ have no common factor, and so $R \neq 0$. Also, we have $R = \prod_{\alpha} P(\alpha)$ where α runs over the complex roots of $y^m + 1$. These roots all lie on the unit circle, so $|P(\alpha)| \leq m\|P\| \leq mq^{1/m}$, and hence by (2.1) we obtain $|Rq| \leq (mq^{1/m})^m q = m^m q^2 < q^3$.

On the other hand, the prime number theorem implies that there exists an absolute constant $C > 2$ such that for any $x \geq 1$ we have $\sum_{2x \leq p \leq Cx} \log_2 p \geq 3x$ (sum taken over primes). Taking $x := m^2q^{1/m}$, by (2.1) again we get

$$\sum_{2m^2q^{1/m} \leq p \leq Cm^2q^{1/m}} \log_2 p \geq 3m^2q^{1/m} \geq 3 \cdot 2^{(\lg \lg q)^2} \geq 3 \lg q \geq \log_2(q^3).$$

In particular, there must be at least one prime in the interval $2m^2q^{1/m} \leq r \leq Cm^2q^{1/m}$ that does not divide Rq .

To find the smallest such r , we first make a list of all primes up to $Cm^2q^{1/m}$ in $(Cm^2q^{1/m})^{1+o(1)} < q^{1+o(1)}$ bit operations. Then for each prime r between $2m^2q^{1/m}$ and $Cm^2q^{1/m}$, we check whether r divides q in $(\lg q)^{1+o(1)}$ bit operations, and attempt to invert $P(y)$ in $(\mathbb{Z}/r\mathbb{Z})[y]/(y^m + 1)$ in $(m \lg r)^{1+o(1)} = (\lg q)^{1+o(1)}$ bit operations [19, Ch. 11]. \square

Example 2.8. Continuing Example 2.2, we have $r = 42602761$ and

$$J(y) = 17106162y^3 + 6504907y^2 + 30962874y + 8514380.$$

Now we come to the main step of the reduction algorithm, which is inspired by Montgomery’s method for modular reduction [16].

Lemma 2.9. *Assume that $P(y)$, r and $J(y)$ have been precomputed as in Lemma 2.3 and Lemma 2.7. Given as input $F \in \mathbb{Z}[y]/(y^m + 1)$ with $\|F\| \leq m^3(q^{1/m})^2$, we may compute a θ -representation for $F(\theta)/r \pmod{q}$ in $O(M_{\text{SS}}(\lg q))$ bit operations.*

Proof. We first compute the “quotient” $Q := FJ \pmod{r}$, normalised so that $\|Q\| \leq r/2$. This is done by means of Kronecker substitution [19, Ch. 8], i.e., we pack the polynomials $F(y)$ and $J(y)$ into integers, multiply the integers, unpack the result, and reduce the result modulo $y^m + 1$ and modulo r . The packed integers have

at most $m(\lg \|F\| + \lg r + \lg m)$ bits, where the $\lg m$ term accounts for coefficient growth in $\mathbb{Z}[y]$. By (2.1) and Lemma 2.7, this simplifies to $O(\lg q)$ bits, so the integer multiplication step costs $O(M_{\text{SS}}(\lg q))$ bit operations. This bound also covers the cost of the reductions modulo r .

Next we compute the product QP , again using Kronecker substitution, at a cost of $O(M_{\text{SS}}(\lg q))$ bit operations. Since $\|Q\| \leq r/2$ and $\|P\| \leq q^{1/m}$, we have $\|QP\| \leq \frac{1}{2}rmq^{1/m}$.

By construction of J we have $QP = F \pmod{r}$. In particular, all the coefficients of $F - QP \in \mathbb{Z}[y]/(y^m + 1)$ are divisible by r . The last step is to compute the ‘‘remainder’’ $G := (F - QP)/r$; again, this step costs $O(M_{\text{SS}}(\lg q))$ bit operations. Since $r \geq 2m^2q^{1/m}$, we have

$$\|G\| \leq \frac{\|F\|}{r} + \frac{\|QP\|}{r} \leq \frac{m^3(q^{1/m})^2}{2m^2q^{1/m}} + \frac{mq^{1/m}}{2} \leq mq^{1/m}.$$

Finally, since $P(\theta) = 0 \pmod{q}$, and all arithmetic throughout the algorithm has been performed modulo $y^m + 1$, we see that $G(\theta) = F(\theta)/r \pmod{q}$. \square

Using the above reduction algorithm, we may give preliminary addition and multiplication algorithms for elements of $\mathbb{Z}/q\mathbb{Z}$ in θ -representation.

Lemma 2.10. *Assume that $P(y)$, r and $J(y)$ have been precomputed as in Lemma 2.3 and Lemma 2.7. Given as input θ -representations for $u, v \in \mathbb{Z}/q\mathbb{Z}$, we may compute θ -representations for uv/r and $(u \pm v)/r$ in $O(M_{\text{SS}}(\lg q))$ bit operations.*

Proof. Let the θ -representations be given by $U, V \in \mathbb{Z}[y]/(y^m + 1)$. We may compute $F_* := UV$ in $\mathbb{Z}[y]/(y^m + 1)$ using Kronecker substitution in $O(M_{\text{SS}}(\lg q))$ bit operations, and $F_{\pm} := U \pm V$ in $O(\lg q)$ bit operations. Note that $\|F_*\| \leq m\|U\|\|V\| \leq m^3(q^{1/m})^2$, and $\|F_{\pm}\| \leq \|U\| + \|V\| \leq 2mq^{1/m} \leq m^3(q^{1/m})^2$, so we may apply Lemma 2.9 to obtain the desired θ -representations. \square

Example 2.11. Continuing Example 2.2, we walk through an example of computing a product of elements in θ -representation. Let

$$\begin{aligned} u &= 1414213562373095048801689 \pmod{q}, \\ v &= 1732050807568877293527447 \pmod{q}. \end{aligned}$$

Suppose we are given as input the θ -representations

$$\begin{aligned} U(y) &= 3740635y^3 + 3692532y^2 - 3089740y + 4285386, \\ V(y) &= 4629959y^3 - 4018180y^2 - 2839272y - 3075767. \end{aligned}$$

We first compute the product of $U(y)$ and $V(y)$ modulo $y^m + 1$:

$$\begin{aligned} F(y) = U(y)V(y) &= 10266868543625y^3 - 37123194804209y^2 \\ &\quad - 4729783170300y + 26582459129078. \end{aligned}$$

We multiply $F(y)$ by $J(y)$ and reduce modulo r to obtain the quotient

$$Q(y) = 3932274y^3 - 14729381y^2 + 20464841y - 11934644.$$

Then the remainder

$$(F(y) - P(y)Q(y))/r = 995963y^3 - 1814782y^2 + 398819y + 777998$$

is a θ -representation for $uv/r \pmod{q}$.

The following precomputation will assist in eliminating the spurious $1/r$ factor appearing in Lemma 2.9 and Lemma 2.10.

Lemma 2.12. *Assume that $P(y)$, r and $J(y)$ have been precomputed as in Lemma 2.3 and Lemma 2.7. In $q^{1+o(1)}$ bit operations, we may compute a polynomial $D \in \mathbb{Z}[y]/(y^m + 1)$ such that $\|D\| \leq mq^{1/m}$ and $D(\theta) = r^2 \pmod{q}$.*

Proof. We may easily compute the totient function $\varphi(q)$ in $q^{1+o(1)}$ bit operations, by first factoring q . Since $(r, q) = 1$, we have $r^{-(\varphi(q)-2)} = r^2 \pmod{q}$. Repeatedly using the identity $r^{-i-1} = (r^{-i} \cdot 1)/r$, we may compute θ -representations for $r^{-1}, r^{-2}, \dots, r^{-(\varphi(q)-2)}$ by successively applying Lemma 2.10. \square

Remark 2.13. Assuming the factorisation of q is known (which will always be the case in the application in Section 3), the complexity of Lemma 2.12 may be improved to $O(M_{\text{SS}}(\lg q) \lg q)$ bit operations by using a modified “repeated squaring” algorithm.

Example 2.14. Continuing Example 2.2, we may take

$$D(y) = -1918607y^3 - 3680082y^2 + 2036309y - 270537.$$

Henceforth we write $\mathcal{P}(q, m, \theta)$ for the tuple $(P(y), r, J(y), D(y))$ of precomputed data generated by Lemma 2.3, Lemma 2.7, and Lemma 2.12. Given q , m and θ as input, the above results show that we may compute $\mathcal{P}(q, m, \theta)$ in $q^{1+o(1)}$ bit operations. With these precomputations out of the way, we may state complexity bounds for the main operations on θ -representations.

Proposition 2.15. *Assume that $\mathcal{P}(q, m, \theta)$ has been precomputed. Given as input θ -representations for $u, v \in \mathbb{Z}/q\mathbb{Z}$, we may compute θ -representations for uv and $u \pm v$ in $O(M_{\text{SS}}(\lg q))$ bit operations.*

Proof. For the product, we first use Lemma 2.10 to compute a θ -representation for $uv/r \pmod{q}$, and then we use Lemma 2.10 again to multiply by $D(y)$, to obtain a θ -representation for $(uv/r)(r^2)/r = uv \pmod{q}$. The sum and difference are handled similarly. \square

Remark 2.16. We suspect that the complexity bound for $u \pm v$ can be improved to $O(\lg q)$, but we do not currently know how to achieve this. This question seems closely related to Remark 2.22.

Proposition 2.17. *Assume that $\mathcal{P}(q, m, \theta)$ has been precomputed. Given as input a polynomial $F \in \mathbb{Z}[y]/(y^m + 1)$ (with no restriction on $\|F\|$), we may compute a θ -representation for $F(\theta) \pmod{q}$ in time $O(\lceil m \lg \|F\| / \lg q \rceil M_{\text{SS}}(\lg q))$.*

Proof. Let $b := \lg \lceil q^{1/m} \rceil$ and $n := \lceil 2m \lg \|F\| / \lg q \rceil$, so that

$$2^{nb} \geq (q^{1/m})^n \geq (q^{1/m})^{2m \lg \|F\| / \lg q} = 2^{\lg \|F\| (2 \log_2 q / \lg q)} \geq 2^{\lg \|F\|}.$$

We may therefore decompose the coefficients of F into n chunks of b bits, i.e., we may compute polynomials $F_0, \dots, F_{n-1} \in \mathbb{Z}[y]/(y^m + 1)$ such that $F = F_0 + 2^b F_1 + \dots + 2^{(n-1)b} F_{n-1}$ and $\|F_i\| \leq 2^b \leq 2q^{1/m}$. (This step implicitly requires an array transposition of cost $O(bmn \lg m) = O(n \lg q \lg \lg q)$.) Now we use Proposition 2.15 repeatedly to compute a θ -representation for F via Horner’s rule, i.e., first we compute a θ -representation for $2^b F_{n-1} + F_{n-2}$, then for $2^b(2^b F_{n-1} + F_{n-2}) + F_{n-3}$, and so on. \square

Proposition 2.18. *Assume that $\mathcal{P}(q, m, \theta)$ has been precomputed. Given as input an element $u \in \mathbb{Z}/q\mathbb{Z}$ in standard representation, we may compute a θ -representation for u in $O(m M_{\text{SS}}(\lg q))$ bit operations.*

Proof. Simply apply Proposition 2.17 to the constant polynomial $F(y) = u$, noting that $\|F\| \leq q$. \square

Remark 2.19. A corollary of Proposition 2.18 is that every $u \in \mathbb{Z}/q\mathbb{Z}$ admits a θ -representation. It would be interesting to have a direct proof of this fact that does not rely on the reduction algorithm. A related question is whether it is possible to tighten the bound in the definition of θ -representation from $mq^{1/m}$ to $q^{1/m}$, or even $\frac{1}{2}q^{1/m}$. We do not know whether such a representation exists for all $u \in \mathbb{Z}/q\mathbb{Z}$.

Proposition 2.20. *Given as input an element $u \in \mathbb{Z}/q\mathbb{Z}$ in θ -representation, we may compute the standard representation for u in $O(m M_{\text{SS}}(\lg q))$ bit operations.*

Proof. Let $U \in \mathbb{Z}[y]/(y^m + 1)$ be the input polynomial. The problem amounts to evaluating $U(\theta)$ in $\mathbb{Z}/q\mathbb{Z}$. Again we may simply use Horner's rule. \square

Remark 2.21. In both Proposition 2.18 and Proposition 2.20, the input and output have bit size $O(\lg q)$, but the complexity bounds given are not quasilinear in $\lg q$. It is possible to improve on the stated bounds, but we do not know a quasilinear time algorithm for the conversion in either direction.

Remark 2.22. In the reduction algorithm, the reader may wonder why we go to the trouble of introducing the auxiliary prime r . Why not simply precompute an approximation to a *real* inverse for $P(y)$, i.e., the inverse in $\mathbb{R}[y]/(y^m + 1)$, and use this to clear out the *high-order bits* of each coefficient of the dividend? In other words, why not replace the Montgomery-style division with the more natural Barrett-style division [1]?

The reason is that we cannot prove tight enough bounds on the size of the coefficients of this inverse: it is conceivable that $P(y)$ might accidentally take on a very small value near one of the complex roots of $y^m + 1$, or equivalently, that the resultant R in the proof of Lemma 2.7 might be unusually small. For the same reason, we cannot use a more traditional 2-adic Montgomery inverse to clear out the low-order bits of the dividend, because again $P(y)$ may take a 2-adically small value near one of the 2-adic roots of $y^m + 1$, or equivalently, the resultant R might be divisible by an unusually large power of 2.

3. INTEGER MULTIPLICATION: THE RECURSIVE STEP

In this section we present a recursive routine TRANSFORM with the following interface. It takes as input a (sufficiently large) power-of-two transform length L , a prime $p = 1 \pmod{L}$, a prime power $q = p^\alpha$ such that

$$(3.1) \quad \lg L \leq \lg q \leq 3 \lg L \lg \lg L,$$

a principal L -th root of unity $\zeta \in \mathbb{Z}/q\mathbb{Z}$ (i.e., an L -th root of unity whose reduction modulo p is a primitive L -th root of unity in the field $\mathbb{Z}/p\mathbb{Z}$), certain precomputed data depending on L , q and ζ (see below), and a polynomial $F \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^L - 1)$. Its output is the DFT of F with respect to ζ , that is, the vector

$$\hat{F} := (F(1), F(\zeta), \dots, F(\zeta^{L-1})) \in (\mathbb{Z}/q\mathbb{Z})^L.$$

The coefficients of both F and \hat{F} are given in standard representation.

The precomputed data consists of the tuple $\mathcal{P}(q, m, \theta)$ defined in Section 2, where m and θ are defined as follows.

First, (3.1) implies that $\lg q \geq (\lg \lg L)^2 \lg \lg \lg L$ for large L , so we may take m to be the unique power of two lying in the interval

$$(3.2) \quad \frac{\lg q}{(\lg \lg L)^2 \lg \lg \lg L} \leq m < \frac{2 \lg q}{(\lg \lg L)^2 \lg \lg \lg L}.$$

Observe that (2.1) is certainly satisfied for this choice of m (for large enough L), as (3.1) implies that $\lg \lg L \sim \lg \lg q$.

Next, note that $2m \mid L$, because (3.1) and (3.2) imply that $m = o(\lg L) = o(L)$; therefore we may take $\theta := \zeta^{L/2m}$, so that $\theta^m = \zeta^{L/2} = -1$.

We remark that the role of the parameter α is to give us enough control over the bit size of q , to compensate for the fact that Linnik's theorem does not give us sufficiently fine control over the bit size of p (see Lemma 3.5).

Our implementation of TRANSFORM uses one of two algorithms, depending on the size of L . If L is below some threshold, say L_0 , then it uses any convenient base-case algorithm. Above this threshold, it reduces the given DFT problem to a collection of exponentially smaller DFTs of the same type, via a series of reductions that may be summarised as follows.

- (i) Use the conversion algorithms from Section 2 to reduce to a transform over $\mathbb{Z}/q\mathbb{Z}$ where the input and output coefficients are given in θ -representation. (During steps (ii) and (iii) below, all elements of $\mathbb{Z}/q\mathbb{Z}$ are stored and manipulated entirely in θ -representation.)
- (ii) Reduce the “long” transform of length L over $\mathbb{Z}/q\mathbb{Z}$ to many “short” transforms of exponentially small length $S := 2^{(\lg \lg L)^2}$ over $\mathbb{Z}/q\mathbb{Z}$, via the Cooley–Tukey decomposition.
- (iii) Reduce each short transform from step (ii) to a product in $(\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$, i.e., a cyclic convolution of length S , using Bluestein's algorithm.
- (iv) Use the definition of θ -representation to reinterpret each product from (iii) as a product in $\mathbb{Z}[x, y]/(x^S - 1, y^m + 1)$, where the coefficients in \mathbb{Z} are exponentially smaller than the original coefficients in $\mathbb{Z}/q\mathbb{Z}$.
- (v) Embed each product from (iv) into $(\mathbb{Z}/q'\mathbb{Z})[x, y]/(x^S - 1, y^m + 1)$ for a suitable prime power q' that is exponentially smaller than q , and large enough to resolve the coefficients of the products over \mathbb{Z} .
- (vi) Reduce each product from (v) to a collection of forward and inverse DFTs of length S over $\mathbb{Z}/q'\mathbb{Z}$, and recurse.

The structure of this algorithm is very similar to that of [10]. The main difference is that it is not necessary to explicitly split the coefficients into chunks in step (iv); this happens automatically as a consequence of storing the coefficients in θ -representation. In effect, the splitting (and reassembling) work has been shunted into the conversions in step (i).

We now consider each of the above steps in more detail. We write $T(L, q)$ for the running time of TRANSFORM. We always assume that L_0 is increased whenever necessary to accommodate statements that hold only for large L .

Step (i) — convert between representations. Let $T_{\text{long}}(L, q)$ denote the time required to compute a DFT of length L over $\mathbb{Z}/q\mathbb{Z}$ with respect to ζ , assuming that the coefficients of the input F and the output \hat{F} are given in θ -representation, and assuming that $\mathcal{P}(q, m, \theta)$ is known.

Lemma 3.1. *We have $\mathsf{T}(L, q) < \mathsf{T}_{\text{long}}(L, q) + O(L \lg L \lg q)$.*

Proof. We first convert F from standard to θ -representation using Proposition 2.18; we then compute \hat{F} from F (working entirely in θ -representation); at the end, we convert \hat{F} back to standard representation using Proposition 2.20. By (3.1) and (3.2), the total cost of the conversions is

$$\begin{aligned} O(Lm \mathsf{M}_{\text{SS}}(\lg q)) &= O\left(L \frac{\lg q}{(\lg \lg L)^2 \lg \lg L} \lg q \lg \lg q \lg \lg \lg q\right) \\ &= O\left(L \frac{\lg L \lg \lg L}{(\lg \lg L)^2 \lg \lg L} \lg q \lg \lg L \lg \lg \lg L\right) = O(L \lg L \lg q). \quad \square \end{aligned}$$

Henceforth all elements of $\mathbb{Z}/q\mathbb{Z}$ are assumed to be stored in θ -representation, and we will always use Proposition 2.15 to perform arithmetic operations on such elements in $O(\mathsf{M}_{\text{SS}}(\lg q))$ bit operations.

Step (ii) — reduce to short DFTs. Let $S := 2^{(\lg \lg L)^2}$. Given as input polynomials $F_1, \dots, F_{L/S} \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$ (presented sequentially on tape), let $\mathsf{T}_{\text{short}}(L, q)$ denote the time required to compute the transforms $\hat{F}_1, \dots, \hat{F}_{L/S} \in (\mathbb{Z}/q\mathbb{Z})^S$ with respect to the principal S -th root of unity $\omega := \zeta^{L/S}$. (Here and below, we continue to assume that $\mathcal{P}(q, m, \theta)$ is known.)

Lemma 3.2. *We have $\mathsf{T}_{\text{long}}(L, q) < \frac{\lg L}{(\lg \lg L)^2} \mathsf{T}_{\text{short}}(L, q) + O(L \lg L \lg q)$.*

Proof. Let $d := \lceil \lg L / \lg S \rceil$, so that $\lg L = d \lg S + d'$ where $0 \leq d' < \lg S$. Applying the Cooley–Tukey method [4] to the factorisation $L = S^d 2^{d'}$, the given transform of length L may be decomposed into d layers, each consisting of L/S transforms of length S (with respect to ω), followed by d' layers, each consisting of $L/2$ transforms of length 2. Between each of these layers, we must perform $O(L)$ multiplications by “twiddle factors” in $\mathbb{Z}/q\mathbb{Z}$, which are given by certain powers of ζ . (For further details of the Cooley–Tukey decomposition, see for example [12, §2.3].)

The total cost of the twiddle factor multiplications, including the cost of computing the twiddle factors themselves, is

$$\begin{aligned} O((d + d')L \mathsf{M}_{\text{SS}}(\lg q)) &= O\left(\left(\frac{\lg L}{(\lg \lg L)^2} + (\lg \lg L)^2\right) L \lg q \lg \lg q \lg \lg \lg q\right) \\ &= O\left(\frac{\lg L}{(\lg \lg L)^2} L \lg q \lg \lg L \lg \lg \lg L\right) = O(L \lg L \lg q). \end{aligned}$$

This bound also covers the cost of the length 2 transforms (‘butterflies’), each of which requires one addition and one subtraction in $\mathbb{Z}/q\mathbb{Z}$.

In the Turing model, we must also account for the cost of rearranging data so that the inputs for each layer of short DFTs are stored sequentially on tape. The cost per layer is $O(L \lg S \lg q)$ bit operations, so $O(L \lg L \lg q)$ altogether (see [12, §2.3] for further details). \square

Step (iii) — reduce to short convolutions. Given polynomials $G_1, \dots, G_{L/S}, H \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$ as input, let $\mathsf{M}_{\text{short}}(L, q)$ denote the time required to compute the products $G_1 H, \dots, G_{L/S} H$.

Lemma 3.3. *We have $\mathsf{T}_{\text{short}}(L, q) < \mathsf{M}_{\text{short}}(L, q) + O(L(\lg \lg L)^2 \lg q)$.*

Proof. We use Bluestein's method [2], which reduces the the problem of computing the DFT of $F \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$ to the problem of computing the product of certain polynomials $G, H \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$, plus $O(S)$ auxiliary multiplications in $\mathbb{Z}/q\mathbb{Z}$ (for further details see [12, §2.5]). Here G depends on F and ζ , but H depends only on ζ . The total cost of the auxiliary multiplications is

$$O((L/S)S M_{\text{SS}}(\lg q)) = O(L \lg q \lg \lg q \lg \lg q) = O(L(\lg \lg L)^2 \lg q). \quad \square$$

Step (iv) — reduce to bivariate products over \mathbb{Z} . Given as input polynomials $\tilde{G}_1, \dots, \tilde{G}_{L/S}, \tilde{H} \in \mathbb{Z}[x, y]/(x^S - 1, y^m + 1)$, all whose of coefficients are bounded in absolute value by $mq^{1/m}$, let $M_{\text{bivariate}}(L, q)$ denote the cost of computing the products $\tilde{G}_1 \tilde{H}, \dots, \tilde{G}_{L/S} \tilde{H}$.

Lemma 3.4. *We have $M_{\text{short}}(L, q) < M_{\text{bivariate}}(L, q) + O(L(\lg \lg L)^2 \lg q)$.*

Proof. We are given as input polynomials $G_1, \dots, G_{L/S}, H \in (\mathbb{Z}/q\mathbb{Z})[x]/(x^S - 1)$. Since their coefficients are given in θ -representation, we may immediately reinterpret them as polynomials $\tilde{G}_1, \dots, \tilde{G}_{L/S}, \tilde{H} \in \mathbb{Z}[x, y]/(x^S - 1, y^m + 1)$, with coefficients bounded by $mq^{1/m}$. By definition of θ -representation, we have $\tilde{H}(x, \theta) = H(x) \pmod{q}$, and similarly for the G_i .

After computing the products $\tilde{G}_i \tilde{H}$ for $i = 1, \dots, L/S$, suppose that

$$(\tilde{G}_i \tilde{H})(x, y) = \sum_{j=0}^{S-1} A_{ij}(y) x^j, \quad A_{ij} \in \mathbb{Z}[y]/(y^m + 1).$$

Then we have $(G_i H)(x) = (\tilde{G}_i \tilde{H})(x, \theta) = \sum_j A_{ij}(\theta) x^j \pmod{q}$ for each i . Therefore, to compute the desired products $G_i H$ with coefficients in θ -representation, it suffices to apply Proposition 2.17 to each A_{ij} , to compute θ -representations for all of the $A_{ij}(\theta)$.

Let us estimate the cost of the invocations of Proposition 2.17. We have $\|A_{ij}\| \leq Sm(mq^{1/m})^2 = Sm^3(q^{1/m})^2$, so

$$\lg \|A_{ij}\| \leq \frac{2 \lg q}{m} + \lg S + 3 \lg m < \frac{2 \lg q}{m} + (\lg \lg L)^2 + O(\lg \lg L).$$

From (3.2) we have $\frac{\lg q}{m} > \frac{1}{2}(\lg \lg L)^2 \lg \lg L$, so for large L ,

$$(3.3) \quad \lg \|A_{ij}\| < \left(2 + \frac{3}{\lg \lg \lg L}\right) \frac{\lg q}{m}.$$

The cost of applying Proposition 2.17 for all A_{ij} is thus

$$O\left((L/S)S \left\lceil \frac{m \lg \|A_{ij}\|}{\lg q} \right\rceil M_{\text{SS}}(\lg q)\right) = O(L M_{\text{SS}}(\lg q)) = O(L(\lg \lg L)^2 \lg q). \quad \square$$

Step (v) — Reduce to bivariate products over $\mathbb{Z}/q'\mathbb{Z}$. Let p' be the smallest prime such that $p' \equiv 1 \pmod{S}$; by Linnik's theorem we have $p' = O(S^{5.2})$. Put $q' := (p')^{\alpha'}$ where

$$\alpha' := \left\lceil \left(2 + \frac{4}{\lg \lg \lg L}\right) \frac{\lg q}{m} / \lg \lfloor p'/2 \rfloor \right\rceil.$$

We have the following bounds for q' .

Lemma 3.5. *Let A_{ij} be as in the proof of Lemma 3.4, for $i = 1, \dots, L/S$ and $j = 0, \dots, S-1$. Then $q' \geq 4\|A_{ij}\|$ and*

$$\lg q' < \left(2 + \frac{O(1)}{\lg \lg \lg L}\right) \frac{\lg q}{m}.$$

Proof. In what follows, we frequently use the fact that $\frac{\lg q}{m} \asymp (\lg \lg L)^2 \lg \lg \lg L$ (see (3.2)). Now, observe that $\log_2 q' = \alpha' \log_2 p' \geq \alpha' \lg \lfloor p'/2 \rfloor$, so by (3.3),

$$\log_2 q' \geq \left(2 + \frac{4}{\lg \lg \lg L}\right) \frac{\lg q}{m} \geq \left(2 + \frac{3}{\lg \lg \lg L}\right) \frac{\lg q}{m} + 2 \geq \lg \|A_{ij}\| + 2.$$

Thus $q' \geq 4\|A_{ij}\|$. For the other direction, since $\lg p' \asymp \lg S = (\lg \lg L)^2$, we have

$$\lg q' \leq \alpha' \lg p' \leq \left[\frac{\left(2 + \frac{4}{\lg \lg \lg L}\right) \frac{\lg q}{m}}{\lg \lfloor p'/2 \rfloor} + 1 \right] \lg p' < \left(2 + \frac{O(1)}{\lg \lg \lg L}\right) \frac{\lg q}{m} \cdot \frac{\lg p'}{\lg \lfloor p'/2 \rfloor},$$

and $\lg p' / \lg \lfloor p'/2 \rfloor < 1 + O(1)/\lg p' < 1 + O(1)/(\lg \lg L)^2$. \square

Now, given as input polynomials $g_1, \dots, g_{L/S}, h \in (\mathbb{Z}/q'\mathbb{Z})[x, y]/(x^S - 1, y^m + 1)$, let $M'_{\text{bivariate}}(L, q)$ denote the cost of computing the products $g_1 h, \dots, g_{L/S} h$, where all input and output coefficients in $\mathbb{Z}/q'\mathbb{Z}$ are in standard representation.

Lemma 3.6. *We have $M_{\text{bivariate}}(L, q) < M'_{\text{bivariate}}(L, q) + O(L \lg q)$.*

Proof. We may locate p' by testing $S+1, 2S+1, \dots$, in $S^{O(1)} = 2^{O((\lg \lg L)^2)} = O(L)$ bit operations, and we may easily compute α' and q' within the same time bound. Now, given as input $\tilde{G}_1, \dots, \tilde{G}_{L/S}, \tilde{H} \in \mathbb{Z}[x, y]/(x^S - 1, y^m + 1)$, we first convert them (in linear time) to polynomials $g_1, \dots, g_{L/S}, h \in (\mathbb{Z}/q'\mathbb{Z})[x, y]/(x^S - 1, y^m + 1)$, and then multiply them in the latter ring. The bound $q' \geq 4\|A_{ij}\|$ in Lemma 3.5 shows that the products over \mathbb{Z} may be unambiguously recovered from those over $\mathbb{Z}/q'\mathbb{Z}$; again, this lifting can be done in linear time. \square

Step (vi) — reduce to DFTs over $\mathbb{Z}/q'\mathbb{Z}$. In this step we will call TRANSFORM recursively to handle certain transforms of length S over $\mathbb{Z}/q'\mathbb{Z}$. To check that these calls are permissible, we must verify the precondition corresponding to (3.1), namely $\lg S \leq \lg q' \leq 3 \lg S \lg \lg S$. The first inequality is clear since $q' \geq p' > S$. The second inequality follows from (3.2), Lemma 3.5, and the observation that $\lg S \lg \lg S \geq (\lg \lg L)^2 \lg \lg L$.

Lemma 3.7. *We have $M'_{\text{bivariate}}(L, q) < \left(\frac{2L}{S} + 1\right) m \mathsf{T}(S, q') + O(L(\lg \lg L)^2 \lg q)$.*

Proof. We start by computing various data needed for the recursive calls. We may compute a primitive S -th root of unity in $\mathbb{Z}/p'\mathbb{Z}$ in $(p')^{O(1)} = O(L)$ bit operations, and then Hensel lift it to a principal S -th root of unity $\zeta' \in \mathbb{Z}/q'\mathbb{Z}$ in $(\lg p' \lg q')^{O(1)} = O(L)$ bit operations. Just as before, we define m' to be the unique power of two in the interval

$$(3.4) \quad \frac{\lg q'}{(\lg \lg S)^2 \lg \lg \lg S} \leq m' < \frac{2 \lg q'}{(\lg \lg S)^2 \lg \lg \lg S},$$

and set $\theta' := (\zeta')^{S/2m'}$. Using Lemma 2.3, Lemma 2.7, and Lemma 2.12, we may compute $\mathcal{P}(q', m', \theta')$ in $(q')^{1+o(1)} = 2^{O((\lg \lg L)^2 \lg \lg \lg L)} = O(L)$ bit operations.

Now suppose we wish to compute the products $g_1 h, \dots, g_{L/S} h$, for polynomials $g_1, \dots, g_{L/S}, h \in (\mathbb{Z}/q'\mathbb{Z})[x, y]/(x^S - 1, y^m + 1)$. We use the following algorithm.

First we use `TRANSFORM` to transform all $L/S+1$ polynomials with respect to x , that is, we compute $g_i((\zeta')^j, y)$ and $h((\zeta')^j, y)$ as elements of $(\mathbb{Z}/q'\mathbb{Z})[y]/(y^m+1)$, for $i = 1, \dots, L/S$ and $j = 0, \dots, S-1$. Since `TRANSFORM` must be applied separately to every coefficient $1, y, \dots, y^{m-1}$, the total number of calls is $(L/S+1)m$. Accessing the coefficient of each y^k also implies a number of array transpositions whose total cost is $O((L/S)Sm \lg m \lg q') = O(L \lg L \lg q)$.

Next we compute the $(L/S)S = L$ pointwise products $g_i((\zeta')^j, y)h((\zeta')^j, y)$. Using Kronecker substitution, each such product in $(\mathbb{Z}/q'\mathbb{Z})[y]/(y^m+1)$ costs $O(M_{SS}(\lg q))$ bit operations, as $m(\lg q' + \lg m) = O(\lg q)$.

Finally, we perform $(L/S)m$ inverse transforms with respect to x . It is well known that these may be computed by the same algorithm as the forward transform, with ζ' replaced by $(\zeta')^{-1}$, followed by a division by S . The division may be accomplished by simply multiplying through by $S^{-1} \pmod{q'}$; this certainly costs no more than the pointwise multiplication step. \square

Corollary 3.8. *We have $\mathsf{T}(L, q) < \frac{\lg L}{(\lg \lg L)^2} \left(\frac{2L}{S} + 1\right) m \mathsf{T}(S, q') + O(L \lg L \lg q)$.*

Proof. This follows immediately by chaining together Lemma 3.1, Lemma 3.2, Lemma 3.3, Lemma 3.4, Lemma 3.6, and Lemma 3.7. \square

Define

$$\mathsf{T}(L) := \max_q \frac{\mathsf{T}(L, q)}{L \lg L \lg q},$$

where the maximum is taken over all prime powers q satisfying (3.1). (For large L , at least one such q always exists. For example, take $\alpha := 1$ and take $q = p$ to be the smallest prime satisfying $p = 1 \pmod{L}$; then Linnik's theorem implies that (3.1) holds for this q .)

Proposition 3.9. *We have $\mathsf{T}(L) < \left(4 + \frac{O(1)}{\lg \lg L}\right) \mathsf{T}(2^{(\lg \lg L)^2}) + O(1)$.*

Proof. Dividing the bound in Corollary 3.8 by $L \lg L \lg q$ yields

$$\frac{\mathsf{T}(L, q)}{L \lg L \lg q} < \left(2 + \frac{S}{L}\right) \frac{m \lg q'}{\lg q} \cdot \frac{\mathsf{T}(S, q')}{S \lg S \lg q'} + O(1).$$

Applying Lemma 3.5 and the estimate $S/L < O(1)/\lg \lg L$ yields

$$\frac{\mathsf{T}(L, q)}{L \lg L \lg q} < \left(4 + \frac{O(1)}{\lg \lg L}\right) \mathsf{T}(S) + O(1).$$

Taking the maximum over allowable q yields the desired bound. \square

Corollary 3.10. *We have $\mathsf{T}(L) = O(4^{\log^* L})$.*

Proof. This follows by applying the “master theorem” [12, Prop. 8] to the recurrence in Proposition 3.9. Alternatively, it follows by the same method used to deduce [10, Cor. 3] from [10, Prop. 2]. The key point is that $2^{(\lg \lg L)^2}$ is dominated by a “logarithmically slow” function of L , such as $\Phi(x) := 2^{(\log \log x)^3}$ (see [12, §5]). \square

Remark 3.11. When working with θ -representations, it is possible to multiply an element of $\mathbb{Z}/q\mathbb{Z}$ by any power of θ in linear time, by simply permuting the coefficients. In other words, we have available “fast roots of unity” in the sense of Fürer. Notice however that the algorithm presented in this section makes no use of this fact!

This raises the question of whether one can design an integer multiplication algorithm that uses these fast roots in the same way as in Fürer's original algorithm, instead of our appeal to Bluestein's trick. This is indeed possible, and one does obtain a bound of the form $O(n \lg n K^{\lg^* n})$. In this algorithm, instead of the running time being dominated by the short transforms, it is dominated by the twiddle factor multiplications, just as in Fürer's algorithm. Unfortunately, this leads to a worse value of K , because of the implied constant in Proposition 2.15.

4. INTEGER MULTIPLICATION: THE TOP LEVEL

The only complication in building an integer multiplication algorithm on top of the TRANSFORM routine is ensuring that the precomputations do not dominate the complexity. We achieve this by means of a multivariate splitting, as follows.

Proof of Theorem 1.1. Suppose that we wish to compute the product of two n -bit integers u and v , for some sufficiently large n . Let $b := \lg n$ and $t := \lceil n/b \rceil^{1/6}$, so that $t^6 b \geq n$. Decompose u into t^6 chunks of b bits, say $u = u_0 + u_1 2^b + \dots + u_{t^6-1} 2^{(t^6-1)b}$ where $0 \leq u_i < 2^b$ for each i , and similarly for v . Let

$$U(x_0, \dots, x_5) := \sum_{i_0=0}^{t-1} \dots \sum_{i_5=0}^{t-1} u_{i_0+t i_1+\dots+t^5 i_5} x_0^{i_0} \dots x_5^{i_5} \in \mathbb{Z}[x_0, \dots, x_5],$$

so that $u = U(2^b, 2^{2b}, \dots, 2^{t^5 b})$, and define $V(x_0, \dots, x_5)$ similarly. The product UV has degree less than $2t$ in each variable, so at most $64t^6$ terms altogether, and its coefficients are bounded by $2^{2b} n \leq 4n^3$. We may therefore reconstruct uv from UV using a straightforward overlap-add procedure (essentially, evaluating at $(2^b, 2^{2b}, \dots, 2^{t^5 b})$) in $O(t^6 \lg n) = O(n)$ bit operations.

Now we consider the computation of UV . Let L be the unique power of two in the interval $2t \leq L < 4t$; then it suffices to compute the product UV in the ring $\mathbb{Z}[x_0, \dots, x_5]/(x_0^L - 1, \dots, x_5^L - 1)$.

For $i = 0, \dots, 18$, let q_i be the least prime such that $q_i \equiv 1 \pmod{L}$ and $q_i \equiv i \pmod{19}$. Then the q_i are distinct, and by Linnik's theorem they satisfy $q_i = O(L^{5.2}) = O(t^{5.2}) = O(n^{0.9})$, so we may locate the q_i in $n^{0.9+o(1)}$ bit operations, and they certainly satisfy (3.1). Moreover, for large n we have $q_0 \dots q_{18} > L^{19} > 2^{19} t^{19} \geq 2^{19} (n/\lg n)^{19/6} > 4n^3$, so to compute UV it suffices to compute $UV \pmod{q_i}$ for each i and then reconstruct UV by the Chinese remainder theorem. The cost of this reconstruction is $(\lg n)^{1+o(1)}$ bit operations per coefficient, so $(n/\lg n)(\lg n)^{1+o(1)} = n(\lg n)^{o(1)}$ altogether.

We have therefore reduced to the problem of computing a product in the ring $(\mathbb{Z}/q_i\mathbb{Z})[x_0, \dots, x_5]/(x_0^L - 1, \dots, x_5^L - 1)$ for each $i = 0, \dots, 19$. To do this, we use TRANSFORM to perform forward DFTs of length L with respect to a suitable primitive L -th root of unity in $\mathbb{Z}/q_i\mathbb{Z}$, for each variable x_0, \dots, x_5 successively; then we multiply pointwise in $\mathbb{Z}/q_i\mathbb{Z}$; finally we perform inverse DFTs and scale the results. The necessary precomputations for each prime (finding a suitable root of unity and computing the appropriate tuple $\mathcal{P}(q_i, m_i, \theta_i)$) require only $q_i^{1+o(1)} = n^{0.9+o(1)}$ bit operation per prime. The total cost of the pointwise multiplications is $n(\lg n)^{o(1)}$. The total number of calls to TRANSFORM for each prime is $12L^5$, so by

Corollary 3.10 we obtain

$$\begin{aligned}
 M(n) &= O(L^5 \sum_{i=0}^{18} T(L, q_i)) + n(\lg n)^{o(1)} \\
 &= O(L^6 \sum_{i=0}^{18} T(L) \lg L \lg q_i) + n(\lg n)^{o(1)} \\
 &= O((n/\lg n)4^{\log^* L} \lg n \lg n) + n(\lg n)^{o(1)} \\
 &= O(n \lg n 4^{\log^* n}).
 \end{aligned}$$

□

REFERENCES

1. P. Barrett, *Implementing the Rivest Shamir and Adleman public key encryption algorithm on a standard digital signal processor*, Advances in cryptology—CRYPTO '86 (Santa Barbara, Calif., 1986), Lecture Notes in Comput. Sci., vol. 263, Springer, Berlin, 1987, pp. 311–323. MR 907099 (88i:94015)
2. L. I. Bluestein, *A linear filtering approach to the computation of discrete Fourier transform*, IEEE Transactions on Audio and Electroacoustics **18** (1970), no. 4, 451–455.
3. A. Bostan, P. Gaudry, and É. Schost, *Linear recurrences with polynomial coefficients and application to integer factorization and Cartier-Manin operator*, SIAM J. Comput. **36** (2007), no. 6, 1777–1806. MR 2299425 (2008a:11156)
4. J. W. Cooley and J. W. Tukey, *An algorithm for the machine calculation of complex Fourier series*, Math. Comp. **19** (1965), 297–301. MR 0178586
5. S. Covanov and E. Thomé, *Fast integer multiplication using generalized Fermat primes*, <http://arxiv.org/abs/1502.02800>, 2016.
6. R. Crandall and B. Fagin, *Discrete weighted transforms and large-integer arithmetic*, Math. Comp. **62** (1994), no. 205, 305–324. MR 1185244
7. M. Fürer, *Faster integer multiplication*, STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing, ACM, New York, 2007, pp. 57–66. MR 2402428 (2009e:68124)
8. ———, *Faster integer multiplication*, SIAM J. Comput. **39** (2009), no. 3, 979–1005. MR 2538847 (2011b:68296)
9. D. Harvey, *Faster truncated integer multiplication*, <https://arxiv.org/abs/1703.00640>, 2017.
10. D. Harvey and J. van der Hoeven, *Faster integer multiplication using plain vanilla FFT primes*, <https://arxiv.org/abs/1611.07144>, to appear in Mathematics of Computation, 2016.
11. ———, *Faster integer and polynomial multiplication using cyclotomic coefficient rings*, <https://arxiv.org/abs/1712.03693>, 2017.
12. D. Harvey, J. van der Hoeven, and G. Lecerf, *Even faster integer multiplication*, J. Complexity **36** (2016), 1–30. MR 3530637
13. S. Lang, *Algebraic number theory*, second ed., Graduate Texts in Mathematics, vol. 110, Springer-Verlag, New York, 1994. MR 1282723
14. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann. **261** (1982), no. 4, 515–534. MR 682664
15. D. Micciancio and P. Voulgaris, *A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations*, SIAM J. Comput. **42** (2013), no. 3, 1364–1391. MR 3504632
16. P. L. Montgomery, *Modular multiplication without trial division*, Math. Comp. **44** (1985), no. 170, 519–521. MR 777282 (86e:11121)
17. C. H. Papadimitriou, *Computational complexity*, Addison-Wesley Publishing Company, Reading, MA, 1994. MR 1251285 (95f:68082)
18. A. Schönhage and V. Strassen, *Schnelle Multiplikation grosser Zahlen*, Computing (Arch. Elektron. Rechnen) **7** (1971), 281–292. MR 0292344 (45 #1431)
19. J. von zur Gathen and J. Gerhard, *Modern computer algebra*, third ed., Cambridge University Press, Cambridge, 2013. MR 3087522
20. T. Xylouris, *On the least prime in an arithmetic progression and estimates for the zeros of Dirichlet L-functions*, Acta Arith. **150** (2011), no. 1, 65–91. MR 2825574 (2012m:11129)

SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NEW SOUTH WALES, SYDNEY NSW
2052, AUSTRALIA

E-mail address: `d.harvey@unsw.edu.au`

LABORATOIRE D'INFORMATIQUE, UMR 7161 CNRS, ÉCOLE POLYTECHNIQUE, 91128 PALAISEAU,
FRANCE

E-mail address: `vdhoeven@lix.polytechnique.fr`