# FASTER INTEGER MULTIPLICATION 

MARTIN FÜRER*


#### Abstract

For more than 35 years, the fastest known method for integer multiplication has been the Schönhage-Strassen algorithm running in time $O(n \log n \log \log n)$. Under certain restrictive conditions, there is a corresponding $\Omega(n \log n)$ lower bound. All this time, the prevailing conjecture has been that the complexity of an optimal integer multiplication algorithm is $\Theta(n \log n)$. We present a major step towards closing the gap from above by presenting an algorithm running in time $n \log n 2^{O\left(\log ^{*} n\right)}$. The running time bound holds for multitape Turing machines. The same bound is valid for the size of boolean circuits.


1. Introduction. All known methods for integer multiplication (except the trivial school method) are based on some version of the Chinese Remainder Theorem. Schönhage [Sch66] computes modulo numbers of the form $2^{k}+1$. Most methods can be interpreted as schemes for the evaluation of polynomials, multiplication of their values, followed by interpolation. The classical method of Karatsuba [KO62] can be viewed as selecting the values of homogeneous linear forms at $(0,1),(1,0)$, and $(1,1)$ to achieve time $T(n)=O\left(n^{\lg 3}\right)$ Toom [Too63] evaluates at small consecutive integer values to improve the circuit complexity to $T(n)=O\left(n^{1+\epsilon}\right)$ for every $\epsilon>0$. Cook [Coo66] presents a corresponding Turing machine implementation. Finally Schönhage and Strassen [SS71] use the usual fast Fourier transform (FFT) (i.e., evaluation and interpolation at $2^{m}$ th roots of unity) to compute integer products in time $O(n \log n \log \log n)$. They conjecture the optimal upper bound (for a yet unknown algorithm) to be $O(n \log n)$, but their result has remained unchallenged.

Schönhage and Strassen [SS71] really propose two distinct methods. The first one uses numerical approximation to complex arithmetic, and reduces multiplication of length $n$ to that of length $O(\log n)$. The complexity of this method is slightly higher. Even as a one level recursive approach, with the next level of multiplications done by a trivial algorithm, it is already very fast. The second method employs arithmetic in rings of integers modulo numbers of the form $F_{m}=2^{2^{m}}+1$ (Fermat numbers), and reduces the length of the factors from $n$ to $O(\sqrt{n})$. This second method is used recursively with $O(\log \log n)$ nested calls. In the ring $\mathbb{Z}_{F_{m}}$ of integers modulo $F_{m}$, the integer 2 is a particularly convenient root of unity for the FFT computation, because all multiplications with this root of unity are just modified cyclic shifts.

On the other hand, the first method has the advantage of a significant length reduction from $n$ to $O(\log n)$ in one level of recursive calls. If this method is applied recursively with $\lg ^{*} n-O(1)$ nested calls ${ }^{1}$ (i.e., until the factors are of constant length), then the running time is of order $n \log n \log \log n \cdots 2^{O\left(\lg ^{*} n\right)}$, because at level 0 , time $O(n \log n)$ is spent, and during the $k$ th of the $\lg ^{*} n-O(1)$ recursion levels, the amount of time increases by a factor of $O(\log \log \ldots \log n)$ (with the $\log$ iterated $k+1$ times) compared to the amount of time spent at the previous level $k$. The time spent at level $k$ refers to the time spent during $k$-fold nested recursive calls, excluding the time spent during the deeper nested recursive calls.

[^0]Note that for their second method, Schönhage and Strassen have succeeded with the difficult task of keeping the time at each level basically fixed. Even just a constant factor increase per level would have resulted in a factor of $2^{O(\log \log n)}=(\log n)^{O(1)}$ instead of $O(\log \log n)$.

Our version of the FFT allows us to combine the main advantages of both methods of Schönhage and Strassen. The reduction is from length $n$ to length $O\left(\log ^{2} n\right)$, and still most multiplications with roots of unity are just modified cyclic shifts. Unfortunately, we are not able to avoid the geometric increase over the $\lg ^{*} n$ levels.

Relative to the conjectured optimal time of $\Theta(n \log n)$, the first Schönhage and Strassen method had an overhead factor of $\log \log n \log \log \log n \ldots 2^{O\left(\lg ^{*} n\right)}$, representing a doubly exponential improvement compared to previous methods. Their second method with an overhead of $O(\log \log n)$ constitutes another poly-logarithmic decrease. Our new method reduces the overhead to $2^{O\left(\lg ^{*} n\right)}$, and thus represents a more than multiple exponential improvement of the overhead factor. Naturally, we have to admit that for practical values of $n$, say in the millions or billions, it is not immediately obvious how to benefit from this last improvement.

We use a divide-and-conquer approach to the $N$-point FFT. We are only interested in the case of $N$ being a power of 2 . Throughout this paper, integers denoted by capital letters are usually powers of 2 . It is well known and obvious that the $J K$-point FFT graph (butterfly graph, Figure 1.1) can be composed of two stages, one containing $K$


FIG. 1.1. The butterfly graph of a 16 point FFT
copies of a $J$-point FFT graph, and the other containing $J$ copies of a $K$-point FFT graph. Clearly $N=J K$ could be factored differently into $N=J^{\prime} K^{\prime}$ and the same $N$-point FFT graph could be viewed as being composed of $J^{\prime}$-point and $K^{\prime}$-point FFT graphs. The current author has been astonished that this is just true for the FFT graph and not for the FFT computation. Every way of (recursively) partitioning $N$ produces another FFT algorithm. Multiplications with other powers of $\omega$ (twiddle factors) appear when another recursive decomposition is used. Here $\omega$ is the principal $N$ th root of unity used in the FFT.

It seems that this fact has not been widely noticed within the algorithms community focusing on asymptotic complexity. On the other hand, there has been a long history of reducing the number of floating point operations (real additions and multiplications) by modifying the decomposition of the FFT over $\mathbb{C}$. Many papers have successfully decreased the number of real multiplications by increasing the use of the roots of unity $\pm i$ and to a lesser extent $\pm(1 \pm i) / \sqrt{2}$. The initial goal has been to speed up the implementations by constant factors. But minimizing this number of operations is also an interesting question in its own right.

Shortly after the publication of Cooley and Tukey [CT65], Gentleman and Sande [GS66] noticed that a radix 4 FFT (decomposition of $N$ into $4 \cdot N / 4$ instead of $2 \cdot N / 2$ ) can save complex multiplications. They have shown that such an implementation runs faster. Bergland [Ber68] has shown that more operations can be saved by radix 8 and radix 16 FFTs. Yavne [Yav68] followed by Duhamel and Hollman [DH84] have discovered the split-radix implementation. This decomposition of an $N$-point FFT into a $N / 2$-point FFT and two $N / 4$-point FFTs is particularly elegant and practical. It has been thought of having an optimal operations count until Johnson and Frigo [JF07] found an algorithm with an even lower number of floating point operations. At this point, the operations count is no longer the most practical concern on modern machines. Locality of data becomes more important than a complicated scheme of saving operations. Vitter and Shriver [VS94] provide a useful complexity theory for machines with a memory hierarchy (layers of memory with various speeds).

There seems to be only one previous attempt to decrease the asymptotic complexity of any algorithm by more than a constant factor based on a different decomposition of a $2^{m}$-point FFT. It was an earlier attempt to obtain a faster integer multiplication algorithm [Für89]. In that paper, the following result has been shown. If there is an integer $k>0$ such that for every $m$, there is a prime number in the sequence $F_{m+1}, F_{m+2}, \ldots, F_{2^{m+k}}$ of Fermat numbers $\left(F_{m}=2^{2^{m}}+1\right)$, then multiplication of binary integers of length $n$ can be done in time $n \log n 2^{O\left(\lg ^{*} n\right)}$. Hence, the Fermat primes could be extremely sparse and would still be sufficient for a faster integer multiplication algorithm. It turns out that the Fermat prime paper [Für89] provides some of the key ingredients of the current faster integer multiplication algorithm. Nevertheless, that paper by itself was not so exciting, because it is conjectured - based on probabilistic assumptions - that the number of Fermat primes is finite, and even that $F_{4}$ is the largest of them.

The FFT is often presented as an iterative process (see e.g., [SS71, AHU74]). A vector of coefficients at level 0 is transformed level by level, until the Fourier transformed vector at level $\lg N$ is reached. The operations at each level are additions, subtractions, and multiplications with powers of an $N$ th root of unity $\omega$. They are usually done as if the $N$-point FFT were recursively decomposed into two $N / 2$-point FFTs followed by $N / 2$ two-point FFTs or vice versa. The standard fast algorithm design principle, divide-and-conquer, calls for a balanced partition, but in this case it is not at all obvious that it will provide any significant benefit.

A balanced approach uses two stages. $K J$-point FFTs are followed by $J K$-point FFTs, where $J$ and $K$ are roughly $\sqrt{N}$. If $N=2^{m}$, then one can choose $J=2^{\lceil m / 2\rceil}$ and $K=2{ }^{\lfloor m / 2\rfloor}$. This can improve the FFT computation, because it turns out that "odd" powers of $\omega$ are then used very seldom. Most multiplications with powers of $\omega$ are then actually multiplications with powers of $\omega^{K}$. This key observation alone is not sufficiently powerful to obtain a better asymptotic running time, because usually $1,-1, i,-i$ and to a lesser extent $\pm(1 \pm i) / \sqrt{2}$ are the only powers of $\omega$ that are
easier to handle. We will achieve the desired speed-up by working over a ring with many "easy" powers of $\omega$. Hence, the new faster integer multiplication algorithm is based on two key ideas.

- An FFT version is used with the property that most occurring roots of unity are of low order.
- The computation is done over a ring where multiplications with many low order roots of unity are very simple and can be implemented as a kind of cyclic shifts. At the same time, this ring also contains high order roots of unity.
It is immediately clear that integers modulo a Fermat prime $F_{m}$ form such a ring. For $N=F_{m}-1=2^{2^{m}}$, the number 2 is a nice low order $(2 \lg N)$ th root of unity, while 3 is an $N$ th root of unity. Instead of 3 , it is computationally easy to find a number $\omega \in \mathbb{Z}_{F_{m}}$ such that $\omega^{N /(2 \lg N)}=2$. These properties form the basis of the Fermat prime result [Für89] mentioned above. The additional main accomplishment of the current paper is to provide a ring having similar properties as the field $\mathbb{Z}_{F_{m}}$.

The question remains whether the optimal running time for integer multiplication is indeed of the form $n \log n 2^{O\left(\lg ^{*} n\right)}$. Already, Schönhage and Strassen [SS71] have conjectured that the more elegant expression $O(n \log n)$ was optimal as we mentioned before. It would indeed be strange if such a natural operation as integer multiplication had such a complicated expression for its running time. But even for $O(n \log n)$ there is no unconditional corresponding lower bound. Still, long ago there have been some remarkable attempts. In the algebraic model, Morgenstern [Mor73] has shown that every $N$-point Fourier transform-done by just using linear combinations $a x+b y$ with $|a|+|b| \leq c$ for inputs or previously computed values $x$ and $y$-requires at least $(n \lg n) /(2(1+\lg c))$ operations. More recently, Bürgisser and Lotz [BL04] have extended this result to multiplying complex polynomials.

Under different assumptions on the computation graph, Papadimitriou [Pap79] and Pan [Pan86] have shown conditional lower bounds of $\Omega(n \log n)$ for the FFT. Both are for the interesting case of $n$ being a power of 2 . Cook and Anderaa [CA69] have developed a method for proving non-linear lower bounds for on-line computations of integer products and related functions. Based on this method, Paterson, Fischer and Meyer [PFM74] have improved the lower bound for on-line integer multiplication to $\Omega(n \log n)$. Naturally, one would like to see unconditional lower bounds, as the on-line requirement is a very severe restriction. On-line means that starting with the least significant bit, the $k$ th bit of the product is written before the $k+1$ st bits of the factors are read.

With the lack of tight lower bounds, it may be tempting to experiment with variations of our new algorithm with the goal of improving the running time or better understand the difficulties in trying to improve it. Indeed, after the publication of the conference version of our algorithm [Für07], a more discrete algorithm based on the same ideas has been obtained by De, Kurur, Saha, and Saptharishi [DKSS08]. Its running time is still $n \log n 2^{O\left(\lg ^{*} n\right)}$, but it is based on $p$-adic numbers instead of complex numbers. This might be a useful ingredient in the quest to decrease the constant factor blow-up from one recursion level to the next. The ambitious goal is to obtain a blow-up factor of $1+o(1)$.

In Section 2, we present the basics about roots of unity in rings, the Chinese Remainder Theorem for rings, and the discrete Fourier transform. In Section 3, we review the FFT in a way that shows which twiddle factors (powers of $\omega$ ) are used for any Cooley and Tukey type recursive decomposition of the algorithm. In

Section 4, we present a ring with many nice roots of unity allowing our faster FFT computation. In Section 5, we describe the new method of using this FFT for integer multiplication. It would be helpful for the reader to be familiar with the SchönhageStrassen integer multiplication algorithm, e.g., as described in the original paper, or in Aho, Hopcroft and Ullman [AHU74], but this is not a prerequisite. In Section 6, we study the precision requirements for the numerical approximations used in the Fourier transforms. Finally, in Section 7, we state the complexity results, followed by open problems in Section 8.
2. The Discrete Fourier Transform. Throughout this paper all rings are commutative rings with 1 . Primitive roots of unity are well known objects in algebra. An element $\omega$ in a ring $\mathcal{R}$ is a primitive $N$ th root of unity if it has the following properties.

1. $\omega^{N}=1$
2. $\omega^{k} \neq 1$ for $1 \leq k<N$

Closely related is the notion of a principal root of unity. An element $\omega$ in a ring $\mathcal{R}$ is a principal $N$ th root of unity if it has the following properties.

1. $\underset{N-1}{\omega^{N}}=1$
2. $\sum_{j=0}^{N-1} \omega^{j k}=0$ for $1 \leq k<N$

The two notions coincide in fields of characteristic 0 , but for the discrete Fourier transform over rings, we need the stronger principal roots of unity. A principal $N$ th root of unity is also a primitive $N$ th root of unity unless the characteristic of the ring $\mathcal{R}$ is a divisor of $N$. For integral domains (commutative rings with 1 and without zero divisors), every primitive root of unity is also a principal root of unity.

The definition of principal roots of unity immediately implies the following result. If $\omega$ is a principal $N$ th root of unity, then $\omega^{J}$ is a principal $(N / \operatorname{gcd}(N, J))$ th root of unity.

Example 1. In $\mathbb{C} \times \mathbb{C}$ the element $(1, i)$ is a primitive 4 th root of unity, but not a principal 4 th root of unity.

Lemma 2.1. If $N$ is a power of 2 , and $\omega^{N / 2}=-1$ in an arbitrary ring, then $\omega$ is a principal $N$ th root of unity.

Proof. Property 1 of principal roots of unity is trivial. Property 2 is shown as follows.

Let $0<k=(2 u+1) 2^{v}<N$, trivially implying that $k / 2^{v}=2 u+1$ is an odd integer and $2^{v+1} \leq N$.

$$
\begin{aligned}
\sum_{j=0}^{N-1} \omega^{j k} & =\sum_{i=0}^{2^{v+1}-1} \sum_{j=0}^{N / 2^{v+1}-1} \omega^{\left(i N / 2^{v+1}+j\right) k} \\
& =\sum_{j=0}^{N / 2^{v+1}-1} \omega^{j k} \underbrace{\sum_{i=0}^{2^{v+1}-1} \underbrace{\omega^{\left(i N / 2^{v+1}\right)(2 u+1) 2^{v}}}_{(-1)^{i}}}_{0} \\
& =0
\end{aligned}
$$

Definition 2.2. The $N$-point discrete Fourier transform (DFT) over a ring $\mathcal{R}$ is the linear function, mapping the vector $\boldsymbol{a}=\left(a_{0}, \ldots, a_{N-1}\right)^{\top}$ to $\boldsymbol{b}=\left(b_{0}, \ldots, b_{N-1}\right)^{\top}$
by

$$
\boldsymbol{b}=\Omega \boldsymbol{a}, \text { where } \Omega=\left(\omega^{j k}\right)_{0 \leq j, k \leq N-1}
$$

for a given principal $N$ th root of unity $\omega$.
In other words,

$$
\begin{equation*}
b_{j}=\sum_{k=0}^{N-1} \omega^{j k} a_{k} \tag{2.1}
\end{equation*}
$$

Hence, the discrete Fourier transform maps the vector of coefficients $\left(a_{0}, \ldots, a_{N-1}\right)^{\top}$ of a polynomial

$$
p(x)=\sum_{k=0}^{N-1} a_{k} x^{k}
$$

of degree $N-1$ to the vector

$$
\left(b_{0}, \ldots, b_{N-1}\right)^{\top}=\left(p(1), p(\omega), p\left(\omega^{2}\right), \ldots, p\left(\omega^{N-1}\right)\right)^{\top}
$$

of values at the $N$ powers of $\omega$.
Every $N$ th root of unity $\omega$ has an inverse $\omega^{-1}=\omega^{N-1}$. If $N$ also has an inverse in the ring $\mathcal{R}$, and $\omega$ is a principal root of unity, then $\omega^{-1}$ is also a principal root of unity, and the matrix $\Omega$ has an inverse $\Omega^{-1}$. The former is shown as follows

$$
\sum_{j=0}^{N-1} \omega^{-j k}=\sum_{j=0}^{N-1} \omega^{N k} \omega^{-j k}=\sum_{j=0}^{N-1} \omega^{(N-j) k}=\sum_{j=1}^{N} \omega^{j k}=0
$$

for $0<k<N$. The inverse of the discrete Fourier transform is $1 / N$ times the discrete Fourier transform with the principal $N$ th root of unity $\omega^{-1}$, because

$$
\sum_{j=0}^{N-1} \omega^{-i j} \omega^{j k}=\sum_{j=0}^{N-1} \omega^{(k-i) j}= \begin{cases}N & \text { if } i=k \\ 0 & \text { if } i \neq k\end{cases}
$$

If the Fourier transform operates over the field $\mathbb{C}$, then $\omega^{-1}=\bar{\omega}$, the complex conjugate of $\omega$. Therefore, the discrete Fourier transform scaled by $1 / \sqrt{N}$ is a unitary transformation, and $\frac{1}{\sqrt{N}} \Omega$ is a unitary matrix, i.e., $\bar{\Omega}^{\top} \Omega=N \mathbf{I}$, where $\mathbf{I}$ is the unit matrix.

Closely related to the $N$-point DFT, we also consider what we call the $N$-point half discrete Fourier transform (Half-DFT) in the case of $N$ being a power of 2. Here the evaluations are done at the $N$ odd powers of $\zeta$, where $\zeta$ is a principal $2 N$ th root of unity. The Half-DFT could be computed by extending $\boldsymbol{a}$ with $a_{N}=\cdots=a_{2 N-1}=0$ and doing a $2 N$-point DFT, but actually only about half the work is needed.

Definition 2.3. The $N$-point half discrete Fourier transform (Half-DFT) is the linear function, mapping the vector $\boldsymbol{a}=\left(a_{0}, \ldots, a_{N-1}\right)^{\top}$ to $\boldsymbol{b}=\left(b_{0}, \ldots, b_{N-1}\right)^{\top}$ by

$$
\boldsymbol{b}=Z \boldsymbol{a}, \text { where } Z=\left(\zeta^{j(2 k+1)}\right)_{0 \leq j, k \leq N-1}
$$

for a given principal $2 N$ th root of unity $\zeta$.

Hence, the Half-DFT maps the vector of coefficients $\left(a_{0}, \ldots, a_{N-1}\right)^{\top}$ of a polynomial

$$
p(x)=\sum_{j=0}^{N-1} a_{j} x^{j}
$$

of degree $N-1$ to the vector

$$
\boldsymbol{b}=\left(p(\zeta), p\left(\zeta^{3}\right), p\left(\zeta^{5}\right), \ldots, p\left(\zeta^{2 N-1}\right)\right)^{\top}
$$

of values at the $N$ odd powers of $\zeta$. Thus

$$
b_{j}=\sum_{k=0}^{N-1} \zeta^{(2 j+1) k} a_{k}=\sum_{k=0}^{N-1} \omega^{j k} \zeta^{k} a_{k} \quad(0 \leq j<N)
$$

for $\omega=\zeta^{2}$.
As usual, let $\Omega=\left(\omega^{j k}\right)_{0 \leq j, k \leq N-1}$. Define $\operatorname{diag}(\boldsymbol{z})$, for $\boldsymbol{z}=\left(\zeta^{0}, \zeta^{1}, \ldots, \zeta^{N-1}\right)^{\top}$ as the diagonal matrix with $\boldsymbol{z}$ in its diagonal. Then we have $Z=\Omega \operatorname{diag}(z)$ or

$$
\begin{aligned}
& Z=\left(\begin{array}{cccc}
\zeta^{1 \cdot 0} & \zeta^{1 \cdot 1} & \cdots & \zeta^{1(N-1)} \\
\zeta^{3 \cdot 0} & \zeta^{3 \cdot 1} & \cdots & \zeta^{3(N-1)} \\
\vdots & \vdots & & \vdots \\
\zeta^{(2 N-1) 0} & \zeta^{(2 N-1) 1} & \cdots & \zeta^{(2 N-1)(N-1)}
\end{array}\right) \\
&=\left(\begin{array}{cccc}
\omega^{0 \cdot 0} & \omega^{0 \cdot 1} & \cdots & \omega^{0(N-1)} \\
\omega^{1 \cdot 0} & \omega^{1 \cdot 1} & \cdots & \omega^{1(N-1)} \\
\vdots & \vdots & & \vdots \\
\omega^{(N-1) 0} & \omega^{(N-1) 1} & \cdots & \omega^{(N-1)(N-1)}
\end{array}\right)\left(\begin{array}{cccc}
\zeta^{0} & 0 & \cdots & 0 \\
0 & \zeta^{1} & \cdots & 0 \\
\vdots & & \ddots & \vdots \\
0 & 0 & \cdots & \zeta^{N-1}
\end{array}\right)
\end{aligned}
$$

Thus for $N$ a power of 2 , and $\zeta$ a principal $2 N$ th root of unity in a ring, an $N$-point half discrete Fourier transform (Half-DFT) is a scaling operation followed by a standard $N$-point DFT with $\omega=\zeta^{2}$. The Half-DFT is invertible if and only if the corresponding DFT is invertible, which is the case if and only if 2 has an inverse in the ring.

Over the field $\mathbb{C}$, the DFT as well as the Half-DFT are scaled unitary transformations, as the matrices $\frac{1}{\sqrt{N}} \Omega$ and $\operatorname{diag}(\boldsymbol{z})$ are both unitary.

The DFT and the Half-DFT perform a ring isomorphism as determined by the Chinese Remainder Theorem, which we copy from Bürgisser et al. [BCS97, p. 75].

Theorem 2.4. (Chinese Remainder Theorem) Let $\mathcal{R}$ be a commutative ring, $I_{1}, \ldots, I_{N}$ ideals in $\mathcal{R}$ which are pairwise coprime, i.e., $I_{i}+I_{j}=\mathcal{R}$ for all $i \neq j$. Then the ring homomorphism

$$
\mathcal{R} \ni z \mapsto\left(z+I_{1}, \ldots, z+I_{N}\right) \in \prod_{j} \mathcal{R} / I_{j}
$$

is surjective with kernel $I=\bigcap_{j=1}^{N} I_{j}$. This induces a ring isomorphism

$$
\mathcal{R} / I \rightarrow \mathcal{R} / I_{1} \times \cdots \times \mathcal{R} / I_{N}
$$

Consider a polynomial ring $\mathcal{R}=\mathcal{R}^{\prime}[x]$ over a $\operatorname{ring} \mathcal{R}^{\prime}$. We want to apply the Chinese Remainder Theorem to two cases. In the first case, the ideal $I_{j}$ is the ideal $\left(x-\omega^{j}\right)$ generated by $x-\omega^{j}$ for a principal $N$ th root of unity $\omega$. In this case, the induced ring isomorphism is the DFT. In the second case, the ideal $I_{j}$ is $\left(x-\zeta^{2 j+1}\right)$ for a principal $2 N$ th root of unity $\zeta$. In this case, the induced ring isomorphism is the half-DFT. In the former case, we want to prove the intersection of ideals $I=\bigcap_{j=1}^{N} I_{j}$ to be $\left(x^{N}-1\right)$ and in the latter case, we want to show $I=\left(x^{N}+1\right)$. These results are easily obtained if $\mathcal{R}$ is the ring of polynomials $F[x]$ over some field $F$. We will need much more effort to prove these results for some cases where $\mathcal{R}$ is just a ring.

In order to apply the Chinese Remainder Theorem, we have to show that the ideals $\left(x-\zeta^{i}\right)$ and $\left(x-\zeta^{j}\right)$ are coprime for $i \neq j$, or equivalently that $\zeta^{i}-\zeta^{j}$ is a unit (i.e., an invertible element) in $\mathcal{R}$. For this purpose, we first show an auxiliary result.

Lemma 2.5. Let $N$ be a unit, and let $\zeta$ be a principal $2 N$ th root of unity in a ring. Then $\zeta^{N}=-1$.

Proof.

$$
0=\frac{1}{N} \sum_{j=0}^{2 N-1} \zeta^{j N}=\frac{1}{N} \sum_{j=0}^{2 N-1} \zeta^{2(j / 2-\lfloor j / 2\rfloor) N}=1+\zeta^{N}
$$

Here, the first term (1) is obtained as the sum over all even values of $j$, while the second term $\left(\zeta^{N}\right)$ is obtained as the sum over all odd values of $j$.

Lemma 2.6. For $N$ a power of 2, let $\omega$ be a principal $N$ th root of unity in a ring $\mathcal{R}$. Let 2 be a unit. Then $1-\omega^{k}$ is a unit for $1 \leq k<N$, and the ideals $\left(x-\omega^{i}\right)$ and $\left(x-\omega^{j}\right)$ are coprime for $i \not \equiv j(\bmod N)$.

Proof. Let $k=(2 u+1) 2^{v}$. Then

$$
\left(1-\omega^{k}\right) \sum_{j=0}^{N 2^{-v-1}-1} \omega^{j k}=1-\omega^{N 2^{-v-1} k}=1-\omega^{(N / 2)(2 u+1)}=1-(-1)^{2 u+1}=2
$$

Let $0 \leq i<j<N$. Then $\omega^{i}-\omega^{j}=\omega^{i}\left(1-\omega^{j-i}\right)$ is in the ideal $\left(x-\omega^{i}, x-\omega^{j}\right)$. As both, $\omega^{i}$ and $1-\omega^{j-i}$ are units, so is $\omega^{i}-\omega^{j}$, implying $\left(x-\omega^{i}, x-\omega^{j}\right)=\mathcal{R}[x]$, i.e., $\left(x-\omega^{i}\right)$ and $\left(x-\omega^{j}\right)$ are coprime.
3. The Fast Fourier Transform. Now let $N=J K$. Then $\omega^{J K}=1$. We want to represent the N-point DFT as a set of K parallel J-point DFTs (inner DFTs), followed by scalar multiplications and a set of J parallel K-point DFTs (outer DFTs). The inner DFTs employ the principal $J$ th root of unity $\omega^{K}$, while the outer DFTs work with the principal $K$ th root of unity $\omega^{J}$. Hence, most powers of $\omega$ (twiddle factors) used during the transformation are powers of $\omega^{J}$ or $\omega^{K}$. Only the scalar multiplications in the middle are by "odd" powers of $\omega$. This simple recursive decomposition of the DFT has in fact been presented in the original paper of Cooley and Tukey [CT65]. Any such recursive decomposition (even for $J=2$ or $K=2$ ) results in a fast algorithm for the DFT and is called Fast Fourier Transform (FFT). At one time, the FFT has been fully credited to Cooley and Tukey, but the FFT has appeared earlier. For the older history of the FFT back to Gauss, the reader is referred to [HJB84].

Here, we are only interested in the usual case of $N$ being a power of 2 . Instead of using $j$ and $k$ ranging from 0 to $N-1$, we use $j^{\prime} J+j$ and $k^{\prime} K+k$ with $0 \leq j, k^{\prime} \leq J-1$ and $0 \leq j^{\prime}, k \leq K-1$. Almost any textbook presenting the Fourier transformation recursively would use either $K=2$ or $J=2$.

For $0 \leq j \leq J-1$ and $0 \leq j^{\prime} \leq K-1$, Equation 2.1 transforms into the following equation, which after minor manipulations exhibits an arbitrary recursive decomposition of the Fourier transform.

$$
\begin{aligned}
b_{j^{\prime} J+j} & =\sum_{k=0}^{K-1} \sum_{k^{\prime}=0}^{J-1} \omega^{\left(j^{\prime} J+j\right)\left(k^{\prime} K+k\right)} a_{k^{\prime} K+k} \\
& =\underbrace{\sum_{k=0}^{K-1} \omega^{J j^{\prime} k} \omega^{\omega^{j k} \underbrace{\sum_{k^{\prime}=0}^{J-1} \omega^{K j k^{\prime}} a_{k^{\prime} K+k}}_{\text {inner }(\text { first }) \mathrm{DFTs}}}}
\end{aligned}
$$

outer (second) DFTs
For $N$ being a power of 2 , the fast Fourier transforms (FFTs) are obtained by recursive application of this method until $N=2$.

We could apply a balanced FFT with $J=K=\sqrt{N}$ or $J=2 K=\sqrt{2 N}$ depending on $N$ being an even or odd power of 2 . But actually, we just require the partition not to be extremely unbalanced.
4. The Ring $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$. We consider the ring of polynomials $\mathcal{R}[y]$ over the ring $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$. In all applications, we will assume $P$ to be a power of 2 . For a primitive $2 P$ th root of unity $\eta$ in $\mathbb{C}$, e.g., $\eta=e^{i \pi / P}$, we have

$$
\begin{aligned}
\mathcal{R} & =\mathbb{C}[x] /\left(x^{P}+1\right) \\
& =\mathbb{C}[x] / \prod_{j=0}^{P-1}\left(x-\eta^{2 j+1}\right) \\
& \cong \prod_{j=0}^{P-1} \mathbb{C}[x] /\left(x-\eta^{2 j+1}\right) \\
& \cong \mathbb{C}^{P}
\end{aligned}
$$

The first isomorphism is provided by an easy version of the Chinese Remainder Theorem for fields.

We want to do Half-DFTs over the ring $\mathcal{R}$. We notice that polynomials over $\mathcal{R}$ decompose into products of polynomials over $\mathbb{C}$.

$$
\mathcal{R}[y]=\mathbb{C}[x] /\left(x^{P}+1\right)[y] \cong \mathbb{C}^{P}[y] \cong \mathbb{C}[y]^{P}
$$

Each component $\mathbb{C}[y]$ is a principal ideal domain and therefore it is factorial, i.e., it has unique factorization. The isomorphic image $\left(\zeta_{0}, \zeta_{1}, \ldots, \zeta_{P-1}\right)^{\top}$ in $\mathbb{C}^{P}$ of a principal $2 N$ th root of unity $\zeta$ in $\mathcal{R}$ is a principal $2 N$ th root of unity. Thus each of its components $\zeta_{k}$ is a principal $2 N$ th root of unity in the field $\mathbb{C}$, implying that $\left(y-\zeta_{k}^{2 j+1}\right)$ divides $\left(y^{N}+1\right)$ and $\operatorname{gcd}\left(y-\zeta_{k}^{2 i+1}, y-\zeta_{k}^{2 j+1}\right)=\zeta_{k}^{2 i+1}-\zeta_{k}^{2 j+1}$ are units in $\mathbb{C}$ for all $i \neq j$ with $0 \leq i, j<N$. As a consequence of unique factorization in $\mathbb{C}[y]$, not only each $\left(y-\zeta_{k}^{2 j+1}\right)$, but also $\prod_{j=0}^{N-1}\left(y-\zeta_{k}^{2 j+1}\right)$ divides $y^{N}-1$. Just looking at the coefficient of $y^{N}$, we see that $\prod_{j=0}^{N-1}\left(y-\zeta_{k}^{2 j+1}\right)=y^{N}+1$.

We apply the Chinese Remainder Theorem for $I_{j}=\left(y-\zeta_{k}^{2 j+1}\right)$. We have seen that $I_{j}=\left(y-\zeta_{k}^{2 j+1}\right)$ divides $\left(y^{N}+1\right)$ for all $j$, implying $\left(y^{N}+1\right) \subseteq \bigcap_{j=0}^{N-1} I_{j}=I$.

The opposite containment is nontrivial. Every element of $I$ is not only a multiple of $y-\zeta_{k}^{2 j+1}$ for each $j$, but because of unique factorization also a multiple of their product $\prod_{j=0}^{N-1}\left(y-\zeta_{k}^{2 j+1}\right)$, i.e., $I \subseteq\left(y^{N}+1\right)$. Thus we obtain the following result.

LEMmA 4.1. $\mathcal{R}[y] /\left(y^{N}+1\right) \cong \prod_{j=0}^{N-1} \mathcal{R}[y] /\left(y-\zeta^{2 j+1}\right)$ and the Half-DFT produces this isomorphism.
$\mathcal{R}$ contains an interesting principal $2 P$ th root of unity, namely $x$. This follows from Lemma 2.1, because $x^{P}=-1$ in $\mathcal{R}$. Alternatively, because $\mathcal{R}$ is isomorphic to $\mathbb{C}^{P}$, a $\zeta \in \mathcal{R}$ is a principal $m$ th root of unity, if and only if it is a principal $m$ th root of unity in every factor $\mathbb{C}[x] /\left(x-\eta^{2 j+1}\right)$ of $\mathcal{R}$. But $x \bmod \left(x-\eta^{2 j+1}\right)$ is just $\eta^{2 j+1}$, which is a principal $2 P$ th root of unity in $\mathbb{C}$. Thus we obtain the following lemma.

Lemma 4.2. For $P$ a power of 2, the variable $x$ is a principal $2 P$ th root of unity in the ring $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$.

The variable $x$ is a very desirable root of unity, because multiplication by $x$ can be done very efficiently. As $x^{P}=-1$, multiplication by $x$ is just a cyclic shift of the polynomial coefficients with a sign change on wrap around.

There are many principal $2 N$ th roots of unity in

$$
\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)
$$

One can choose an arbitrary primitive $2 N$ th root of unity in every factor

$$
\mathbb{C}[x] /\left(x-\eta^{2 j+1}\right)
$$

independently. We want to pick one such $2 N$ th root of unity

$$
\rho \in \mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)
$$

with the convenient property

$$
\rho^{N / P}=x
$$

We write $\rho(x)$ for $\rho$ to emphasize that it is represented by a polynomial in $x$. Let $\sigma$ be a primitive $2 N$ th root of unity in $\mathbb{C}$, with

$$
\sigma^{N / P}=\eta
$$

e.g.,

$$
\sigma=e^{i \pi / N}
$$

Now we select the polynomial

$$
\rho(x)=\sum_{j=0}^{P-1} r_{j} x^{j}
$$

such that

$$
\rho(x) \equiv \sigma^{2 k+1} \quad\left(\bmod x-\eta^{2 k+1}\right) \quad \text { for } k=0,1, \ldots, P-1
$$

i.e., $\sigma^{2 k+1}$ is the value of the polynomial $\rho(x)$ at $\eta^{2 k+1}$. Then

$$
\rho(x)^{N / P} \equiv \sigma^{(2 k+1) N / P}=\eta^{2 k+1} \equiv x \quad\left(\bmod x-\eta^{2 k+1}\right)
$$

for $0 \leq k<P$, implying

$$
\rho(x)^{N / P} \equiv x \quad\left(\bmod x^{P}+1\right)
$$

For the FFT algorithm, the coefficients of $\rho(x)$ could be computed from Lagrange's interpolation formula

$$
\rho(x)=\sum_{k=0}^{P-1} \sigma^{2 k+1} \frac{\prod_{j \neq k}\left(x-\eta^{2 j+1}\right)}{\prod_{j \neq k}\left(\eta^{2 k+1}-\eta^{2 j+1}\right)}
$$

without affecting the asymptotic running time, because $P=O(\log N)$ is small.
In both products, $j$ ranges over $\{0, \ldots, P-1\} \backslash\{k\}$. The numerator in the previous expression is

$$
\frac{x^{P}+1}{x-\eta^{2 k+1}}=-\sum_{j=0}^{P-1} \eta^{-(j+1)(2 k+1)} x^{j}
$$

This implies that in our case, all coefficients of each of the additive terms in Lagrange's formula have the same absolute value. We also want to show that all coefficients of $\rho(x)$ have an absolute value of at most 1.

Definition 4.3. The $l_{2}$-norm of a polynomial $p(x)=\sum a_{k} x^{k}$ is $\|p(x)\|=$ $\sqrt{\sum\left|a_{k}\right|^{2}}$.

Our FFT will be done with the principal root of unity $\rho(x)$ defined above. In order to control the required numerical accuracy of our computations, we need a bound on the absolute value of the coefficients of $\rho(x)$. Such a bound is provided by the $l_{2}$-norm $\|\rho(x)\|$ of $\rho(x)$.

Lemma 4.4. The $l_{2}$-norm of $\rho(x)$ is $\|\rho(x)\|=1$.
Proof. Note that the values of the polynomial $\rho(x)$ at all the primitive $2 P$ th roots of unity are also roots of unity, in particular complex numbers with absolute value 1. Thus the vector $\boldsymbol{b}$ of these values has $l_{2}$-norm $\sqrt{P}$. The coefficients of $\rho(x)$ are obtained by an inverse half discrete Fourier transform $Z^{\prime-1} \boldsymbol{b}$. As $\sqrt{P} Z^{\prime-1}$ is a unitary matrix, the vector of coefficients has norm 1.

Corollary 4.5. The absolute value of every coefficient of $\rho(x)$ is at most 1 .
5. The Algorithm. In order to multiply two non-negative integers modulo $2^{n}+$ 1 , we encode them as polynomials of $\mathcal{R}[y]$, where $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$, and we multiply these polynomials with the help of the Fourier transform as follows. Let $P=\Theta(\log n)$ be rounded to a power of 2 . The binary integers to be multiplied are decomposed into (large) pieces of length $P^{2} / 2$. Again, each such piece is decomposed into small pieces of length $P$. If $a_{i P / 2-1}, \ldots, a_{i 0}$ are the small pieces belonging to a common big piece $a_{i}$, then they are encoded as

$$
\tilde{a}_{i}=\sum_{j=0}^{P-1} a_{i j} x^{j} \in \mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)
$$

with

$$
a_{i P-1}=a_{i P-2}=\cdots=a_{i P / 2}=0
$$

Thus each large piece is encoded as an element of $\mathcal{R}$, which is a coefficient of a polynomial in $y$.

These elements of $\mathcal{R}$ are themselves polynomials in $x$. Their coefficients are integers at the beginning and at the end of the algorithm. The intermediate results, as well as the roots of unity are polynomials with complex coefficients, which themselves are represented by pairs of reals that have to be approximated numerically. In Section 6, we will show that it is sufficient to use fixed-point arithmetic with $O(P)=O(\log n)$ bits in the integer and fraction part.

Now every factor $\sum_{i=0}^{N-1} \tilde{a}_{i} 2^{i P^{2} / 2}$ is represented by a polynomial $\sum_{i=0}^{N-1} a_{i} y^{i} \in$ $\mathcal{R}[y]$. A Half-FFT computes the values of such a polynomial at those roots of unity which are odd powers of the $2 N$ th root of unity $\rho(x) \in \mathcal{R}$, defined in the previous section. The values are multiplied and an inverse Half-FFT produces another polynomial of $\mathcal{R}[y]$. From this polynomial the resulting integer product can be recovered by just doing some simple additions. The relevant parts of the coefficients have now grown to some length $O(P)$ from the initial length of $P$. (The constant factor growth could actually be decreased to a factor $2+o(1)$ by increasing the parameter $P$ from $\Theta(\log n)$ to $\Theta\left(\log ^{2} n\right)$, but this would only affect the constant factor in front of $\lg ^{*}$ in the exponent of the running time.)

Thus the algorithm runs pretty much like that of Schönhage and Strassen [SS71] except that the field $\mathbb{C}$ or the ring of integers modulo the $m$ th Fermat prime $F_{m}$ has been replaced by the ring $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$, and the FFT is decomposed more evenly. The standard decomposition of the $N$-point FFT into two $N / 2$-point FFTs and many 2-point FFTs would not allow such an improvement. Nevertheless, there is no need for balancing completely. Instead of recursively decomposing the $N=\Theta\left(\frac{n}{\log ^{2} n}\right)$-point FFT in the middle (in a divide-and-conquer fashion), we decompose into $2 P$-point FFTs and $N /(2 P)$-point FFTs. This is mainly done for simplicity. Both versions are efficient (even though with different constant factors in front of $\mathrm{lg}^{*}$ in the exponent), as only about every $\log P$ th level of the overall FFT requires complicated multiplications with difficult roots of unity (twiddle factors). At all the other levels, the twiddle factors are powers of $x$. Multiplications with these twiddle factors are just cyclic rotations of the $P$-tuple of coefficients of elements of $\mathcal{R}$, with a sign change on wrap around.

We use the auxiliary functions Decompose (Figure 5.1) and Compose (Figure 5.2).

## Procedure Decompose:

Input: Integer $a$ of length at most $n=N P^{2} / 2$ in binary; $N, P$ (powers of 2)
Output: $\boldsymbol{a} \in \mathcal{R}^{N}$ (or $\alpha \in \mathcal{R}[y]$ ) encoding the integer $a$
Comment: The integer $a$ is the concatenation of the $a_{i j}$ for $0 \leq i<N$ and $0 \leq$ $j<P / 2$ as binary integers of length $P$ defined by Equations (5.1), (5.2) and (5.4), and Inequalities (5.3). $a_{i 0}, a_{i 1}, \ldots, a_{i P-1}$ are the coefficients of $\alpha_{i} \in \mathcal{R}$. $\alpha_{0}, \alpha_{1}, \ldots, \alpha_{N-1}$ are the components of $\boldsymbol{a} \in \mathcal{R}^{N}$ as defined by Equation (5.5).
for $i=0$ to $N-1$ do
for $j=0$ to $P / 2-1$ do
$a_{i j}=a \bmod 2^{P}$
$a=\left\lfloor a / 2^{P}\right\rfloor$
for $j=P / 2$ to $P-1$ do
$a_{i j}=0$
$\alpha_{i}=a_{i 0}+a_{i 1} x+a_{i 2} x^{2}+\cdots+a_{i P-1} x^{P-1}$
Return $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)$
Fig. 5.1. The procedure Decompose

## Procedure Compose:

```
Input: \(\boldsymbol{a} \in \mathcal{R}^{N}, N, P\) (powers of 2)
Output: Integer \(a\) encoded by \(\boldsymbol{a}\)
Comment: \(\alpha_{0}, \alpha_{1}, \ldots, \alpha_{N-1}\) are the components of a vector \(\boldsymbol{a} \in \mathcal{R}^{N}\). For all \(i, j, a_{i j}\)
                is the coefficient of \(x^{j}\) in \(\alpha_{i}\). The integer \(a\) is obtained from the rounded \(a_{i j}\) as
                defined in Equation (5.2).
round all \(a_{i j}\) to the nearest integer
\(a=0\)
for \(j=P-1\) downto \(P / 2\) do
    \(a=a \cdot 2^{P}+a_{N-1 j}\)
for \(i=N-1\) downto 1 do
    for \(j=P / 2-1\) downto 0 do
        \(a=a \cdot 2^{P}+a_{i j}+a_{i-1 j+P / 2}\)
    for \(j=P / 2-1\) downto 0 do
    \(a=a \cdot 2^{P}+a_{0 j}\)
Return \(a \bmod \left(2^{n}+1\right)\)
```

Fig. 5.2. The procedure Compose
"Decompose" takes a binary number $a$ of length $n=N P^{2} / 2$. First $a$ is decomposed into $N$ pieces $a_{N-1}, a_{N-2}, \ldots, a_{0}$ of length $P^{2} / 2$ each. Then every $a_{i}$ is decomposed into $P / 2$ pieces $a_{i P / 2-1}, a_{i, P / 2-2}, \ldots, a_{i 0}$ of length $P$ each. The remaining $a_{i j}$ (for $0 \leq i<N$ and $P / 2 \leq j<P$ ) are defined to be 0 . This padding allows to properly recover the integer product from the product of the polynomials. In other words, we have

$$
\begin{equation*}
a=\sum_{i=0}^{N-1} a_{i} 2^{i P^{2} / 2} \quad \text { and } \quad a_{i}=\sum_{j=0}^{P-1} a_{i j} 2^{j P} \tag{5.1}
\end{equation*}
$$

implying

$$
\begin{equation*}
a=\sum_{i=0}^{N-1} \sum_{j=0}^{P-1} a_{i j} 2^{i\left(P^{2} / 2\right)+j P} \tag{5.2}
\end{equation*}
$$

with

$$
\begin{equation*}
0 \leq a_{i j}<2^{P} \text { for all } i, j \tag{5.3}
\end{equation*}
$$

and

$$
\begin{equation*}
a_{i j}=0 \text { for } 0 \leq i<N \text { and } P / 2 \leq j<P \tag{5.4}
\end{equation*}
$$

We use Greek letters to denote elements of the ring $\mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)$. The number $a_{i}$ is encoded as an element $\alpha_{i} \in \mathcal{R}$ and $a$ is encoded as a polynomial

$$
\alpha=\sum_{i=0}^{N-1} \alpha_{i} y^{i} \in \mathcal{R}[y]
$$

represented by its vector of coefficients $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)^{\top} \in \mathcal{R}^{N}$, with

$$
\begin{equation*}
\alpha_{i}=\sum_{j=0}^{P-1} a_{i j} x^{j}=a_{i 0}+a_{i 1} x+a_{i 2} x^{2}+\cdots+a_{i P-1} x^{P-1} \tag{5.5}
\end{equation*}
$$

We say that $\boldsymbol{a}$ represents the integer $a$, when (5.1) and (5.2) hold. Typically an integer a ( $0 \leq a<2^{N P^{2} / 2}$ ) has many representations. "Decompose" selects the unique representation in normal form, defined by (5.3) and (5.4).

In normal form, the padding (defined by (5.4)) is designed to avoid any wrap around modulo $x^{P}+1$ when doing multiplication in $\mathcal{R}$. "Compose" not only reverses the effect of "Decompose", but it works just as well for arbitrary representations not in normal form, which are produced during the computation.

To be precise, no wrap around modulo $x^{P}+1$ would occur, if the product of polynomials were computed with the standard school multiplication. During the actual computation of the product using FFT, wrap around happens frequently. But naturally the final result is just the product of the polynomials, i.e., the same as if it were computed with school multiplication.

## Procedure Select:

Input: $N \geq 4$ (a power of 2 ), $P$ (a power of 2 )
Output: $J \geq 2$ (a power of 2 dividing $N / 2$ )
Comment: The procedure selects $J$ such that the $N$-point FFT is decomposed into $J$ point FFTs followed by $K=N / J$-point FFTs.
if $N \leq 2 P$ then Return 2 else Return $2 P$
Fig. 5.3. The procedure Select determining the recursive decomposition
The procedure Select ( $N$ ) (Figure 5.3) determines how the FFT is broken down recursively, corresponding to a factorization $N=J K$. Schönhage has used Select $(N)=$ 2, Aho, Hopcroft and Ullman use $\operatorname{Select}(N)=N / 2$, a balanced approach corresponds to Select $(N)=2^{\lfloor(\lg N) / 2\rfloor}$. We choose $\operatorname{Select}(N)=2 P$, which is slightly better than a balanced solution (only by a constant factor in front of $\lg ^{*} n$ ), because with our choice of Select $(N)$ only every $2 P$ th level (instead of every $i$ th for some $P<i \leq 2 P$ ) requires expensive multiplications with twiddle factors.

With the help of these auxiliary procedures, we can now give an overview of the whole integer multiplication algorithm in Figure 5.4. A more detailed description of the various procedures follows. Integer-Multiplication and Modular-IntegerMultiplication also use the auxiliary functions of Figure 5.5.

The previously presented procedures Decompose and Compose are simple format conversions. The three major parts of Modular-Integer-Multiplication are HalfFFT (Figure 5.8) for both factors, Componentwise-Multiplication (Figure 5.10), and Inverse-Half-FFT (Figure 5.11). The crucial part, FFT (Figure 5.9), is presented as a recursive algorithm for simplicity and clarity. It uses the auxiliary procedure Select (Figure 5.3). The algorithms FFT and Componentwise-Multiplication, use the operation $*$ (Figure 5.12), which is the multiplication in the ring $\mathcal{R}$. This operation is implemented by modular integer multiplication, which is executed by recursive calls to the procedure Modular-Integer-Multiplication (Figure 5.7).

We compute the product of two complex polynomials (elements of $\mathcal{R}$ ) by first writing each as a sum of a real and an imaginary polynomial, and then computing four products of real polynomials. Alternatively, we could achieve the same result with three real polynomial multiplications based on the basic idea of [KO62]. Real polynomials (with coefficients given in fixed point representation) are multiplied by multiplying their values at a good power of 2 as proposed by Schönhage [Sch82]. A good power of 2 makes the values not too big, but still allows the coefficients of the

## The Structure of the Complete Algorithm:

Comment: Lower level (indented) algorithms are called from the higher level algorithms. In addition, the algorithm FFT calls itself recursively, and it calls the auxiliary procedure Select. Furthermore, in FFT and Componentwise-Multiplication, the Operation $*$ calls Modular-Integer-Multiplication recursively.
Integer-Multiplication
Modular-Integer-Multiplication
Decompose
Half-FFT
FFT
Componentwise-Multiplication
Half-Inverse-FFT
FFT
Compose
Fig. 5.4. Overall structure of the multiplication algorithm

## Various functions:

$\lg$ : the log to the base 2
length: the length in binary
round: rounded up to the next power of 2
Fig. 5.5. Various functions
product polynomial to be easily recovered from the binary representation of the integer product. The case of positive integer coefficients is particularly intuitive. A binary integer is formed by concatenating the coefficients after padding them with leading zeros such that all have equal length and the non-zero parts are well separated.

Schönhage [Sch82] has shown that such an encoding can easily be achieved with a constant factor blow-up. Actually, he proposes an even better method for handling complex polynomials. He does integer multiplication modulo $2^{N}+1$ and notices that $2^{N / 2}$ can serve as the imaginary unit $i$. We don't further elaborate on this method, as it only affects the constant factor in front of $\lg ^{*} n$ in the exponent of the running time. An even better method has been suggested by Bernstein [Ber]. In the definition of the ring $\mathcal{R}$, the field $\mathbb{C}$ could be replaced by $\mathbb{R}$, as $x^{P / 2}$ could be viewed as the imaginary unit $i$.

## Algorithm Integer-Multiplication:

Input: Integers $a$ and $b$ in binary
Output: Product $d=a b$
Comment: The product $d=a b$ is computed with Modular-Integer-Multiplication. The length $n$ is chosen to be a power of 2 sufficiently large to avoid any warp around.
$n=$ round $($ length $(a)+$ length $(b))$
Return Modular-Integer-Multiplication $(n, a, b)$
Fig. 5.6. The Algorithm Integer-Multiplication

```
Algorithm Modular-Integer-Multiplication:
Input: Integer n, Integers a and b modulo 2}\mp@subsup{2}{}{n}+1\mathrm{ in binary
Output: Product d=ab mod 2n}+
Comment: n is a power of 2. The product d=ab mod 2n}+1\mathrm{ is computed with
half Fourier transforms over \mathcal{R}}\mathrm{ . Let }\zeta\mathrm{ be the 2Nth root of unity in }\mathcal{R}\mathrm{ with
value }\mp@subsup{e}{}{i\pi(2k+1)/N}\mathrm{ at }\mp@subsup{e}{}{i\pi(2k+1)/P}\mathrm{ (for }k=0,\ldots,P-1). Let no \geq16 be some
constant. For n\leqnn, a trivial multiplication algorithm is used.
if n\leqno then Return ab mod 2n}+
P= round(lgn)
N=2n/P}\mp@subsup{P}{}{2
a=\operatorname{Half-FFT(Decompose(a), \zeta,N,P)}
b}=\operatorname{Half-FFT(Decompose(b),\zeta,N,P)
c=Componentwise-Multiplication(a,b,N,P)
d= Inverse-Half-FFT(c,\zeta,N,P)
Return Compose(d)
```

Fig. 5.7. The Algorithm Modular-Integer-Multiplication

## Algorithm Half-FFT:

Input: $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)^{\top} \in \mathcal{R}^{N}, \zeta \in \mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)(\zeta$ is a principal $2 N$ th root of unity in $\mathcal{R}$ with $\left.\zeta^{N / P}=x\right), N, P$ (powers of 2),
Output: $\boldsymbol{b} \in \mathcal{R}^{N}$ the $N$-point half DFT of the input
Comment: The $N$-point half DFT is the evaluation of the polynomial with coefficient vector $\boldsymbol{a}$ at the odd powers of $\zeta$, i.e., in those powers of $\zeta$ that are principal $2 N$ th roots of unity in $\mathcal{R}$.

```
    for \(k=0\) to \(N-1\) do
        \(\alpha_{k}=\alpha_{k} * \zeta^{k}\)
    \(\omega=\zeta^{2}\)
Return \(\operatorname{FFT}(\boldsymbol{a}, \omega, N, P)\)
```

Fig. 5.8. The algorithm half-FFT
6. Precision for the FFT over $\mathcal{R}$. We compute Fourier transforms over the ring $\mathcal{R}$. Elements of $\mathcal{R}$ are polynomials over $\mathbb{C}$ modulo $x^{P}+1$. The coefficients are represented by pairs of reals with fixed precision for the real and imaginary part. We want to know the numerical precision needed for the coefficients of these polynomials. We start with integer coefficients. After doing two Half-FFTs in parallel, and multiplying the corresponding values followed by an inverse Half-FFT, we know that the result has again integer coefficients. Therefore, the precision has to be such that at the end the absolute errors are less than $\frac{1}{2}$. Hence, a set of final rounding operations provably produces the correct result.

We do all computations with at least $S$ bits of precision, where $S$ is fixed (as a function of $n$ ) and will be determined later. We use at least $S$ bits for the real as well as the imaginary part of each complex number occurring in the FFT algorithms. In addition, there is a sign to be stored with each number. Of the $S$ bits, we use at least $V$ bits before the binary point and at least $S-V$ bits after the binary point. $V$ varies throughout the algorithm. We are very generous with $V$ and $S$, meaning that the bits in the integer part might include many leading zeros and the number

## Algorithm FFT:

Input: $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)^{\top} \in \mathcal{R}^{N}, \omega \in \mathcal{R}=\mathbb{C}[x] /\left(x^{P}+1\right)(\omega$ is an $N$ th root of unity in $\mathcal{R}$ with $\omega^{N / 2 P}=x ; \omega=x^{2 P / N}$ for $\left.N<2 P\right), N, P$ (powers of 2),
Output: $\boldsymbol{b} \in \mathcal{R}^{N}$ the $N$-point DFT of the input
Comment: The $N$-point FFT is the composition of $J$-point inner FFTs and $K$-point outer FFTs. We use the vectors $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{R}^{N}, \boldsymbol{c}^{k}=\left(\gamma_{0}^{k}, \ldots, \gamma_{J-1}^{k}\right)^{\top} \in \mathcal{R}^{J}$ $(k=0, \ldots, K-1)$, and $\boldsymbol{d}^{j}=\left(\delta_{0}^{j}, \ldots, \delta_{K-1}^{j}\right)^{\top} \in \mathcal{R}^{K}(j=0, \ldots, J-1)$.
if $N=1$ then Return $\boldsymbol{a}$
if $N=2$ then $\left\{\beta_{0}=\alpha_{0}+\alpha_{1} ; \beta_{1}=\alpha_{0}-\alpha_{1} ;\right.$ Return $\left.\boldsymbol{b}=\left(\beta_{0}, \ldots, \beta_{N-1}\right)^{\top}\right\}$
$J=\operatorname{Select}(N, P) ; K=N / J$
for $k=0$ to $K-1$ do
for $k^{\prime}=0$ to $J-1$ do
$\gamma_{k^{\prime}}^{k}=\alpha_{k^{\prime} K+k}$
$c^{k}=\operatorname{FFT}\left(c^{k}, \omega^{K}, J\right) / /$ inner FFTs
for $j=0$ to $J-1$ do
for $k=0$ to $K-1$ do
$\delta_{k}^{j}=\gamma_{j}^{k} * \omega^{j k}$
$\boldsymbol{d}^{j}=\mathrm{FFT}\left(\boldsymbol{d}^{j}, \omega^{J}, K\right) / /$ outer FFTs
for $j^{\prime}=0$ to $K-1$ do
$\beta_{j^{\prime} J+j}=\delta_{j^{\prime}}^{j}$
Return $\boldsymbol{b}=\left(\beta_{0}, \ldots, \beta_{N-1}\right)^{\top}$
Fig. 5.9. The algorithm FFT

## Algorithm Componentwise-Multiplication:

Input: $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)^{\top}, \boldsymbol{b}=\left(\beta_{0}, \ldots, \beta_{N-1}\right)^{\top} \in \mathcal{R}^{N}, N, P$ (powers of 2)
Output: $\boldsymbol{c} \in \mathcal{R}^{N}$ (the componentwise product of $\boldsymbol{a}$ and $\boldsymbol{b}$ )
for $j=0$ to $N-1$ do
$\gamma_{j}=\alpha_{j} * \beta_{j}$
Return $\boldsymbol{c}=\left(\gamma_{0} \ldots, \gamma_{N-1}\right)^{\top}$
Fig. 5.10. The algorithm Componentwise-Multiplication

## Algorithm Inverse-Half-FFT:

Input: $\boldsymbol{a}=\left(\alpha_{0}, \ldots, \alpha_{N-1}\right)^{\top} \in \mathcal{R}^{N}, \zeta \in \mathcal{R}$ (a principal $2 N$ th root of unity in $\mathcal{R}$ ), $N$, $P$ (powers of 2)
Output: $\boldsymbol{b} \in \mathcal{R}^{N}$ (the inverse of the half $N$-point DFT applied to the input)

```
    \(\omega=\zeta^{2}\)
    \(\boldsymbol{b}=\frac{1}{N} \mathrm{FFT}\left(\boldsymbol{a}, \omega^{-1}, N, P\right)\)
    for \(k=0\) to \(N-1\) do
        \(\beta_{k}=\beta_{k} * \zeta^{-k}\)
    Return \(\boldsymbol{b}=\left(\beta_{0}, \ldots, \beta_{N-1}\right)^{\top}\)
```

FIG. 5.11. The algorithm Inverse-FFT
of bits in the fractional part might be unnecessarily high. The main purpose is to prove correctness. Tighter bounds would improve the constant factor hidden by the $O$-notation in the running time.

Operation ${ }^{*}(=$ Multiplication in $\mathcal{R})$ :
Input: $\alpha, \beta \in \mathcal{R}, P$ (a power of 2 )
Output: $\gamma$ (the product $\alpha \cdot \beta \in \mathcal{R}$ )
Comment: If the second factor $\beta$ is a power of $x$, then obtain $\gamma$ as a cyclic shift of $\alpha$ with sign change on wrap around. Otherwise, start by writing each of the two polynomials as a sum of a real and an imaginary polynomial. Then compute the 4 products of real polynomials by multiplying their values at a good power of 2 , which pads the space between the coefficients nicely such that the coefficients of the product polynomial can easily be recovered from the binary representation of the integer product.
The details can easily be filled in, as the coefficients are presented in fixed-point arithmetic. Precise bounds on the lengths of the integer and fractional parts are given in Section 6.

FIG. 5.12. The multiplication in $\mathcal{R}$ (= operation *)

In a practical implementation, one can either use floating point arithmetic with at least $S$ bits in the mantissa, or one could always scale with the known appropriate power of 2 and use integer arithmetic.

For the initial call to Half-FFT, the fractional part after the binary point is 0 , and we use $V=P$ bits in the integer part in front of the binary point. In each level of the FFT, we do everywhere an addition or subtraction and a multiplication with a twiddle factor (which might be 1). We generously increase $V$ by 1 for each addition or subtraction. Most multiplications with twiddle factors are handled by cyclic shifts producing no errors. At every level divisible by $\lg P+1$, the multiplications by twiddle factors are general multiplications in $\mathcal{R}$.

We first investigate the growth of the value and error bounds during these multiplications. Elements of $\mathcal{R}$ are represented by polynomials in $x$ of degree $P-1$.

Notation 1. We refer to the real or imaginary part of any coefficient of an element of $\mathcal{R}$ simply as a part.

Let $r$ be a part of any element of $\mathcal{R}$ occurring in an idealized infinite precision algorithm. In reality, a finite precision algorithm uses an approximation $r+\varepsilon_{r}$ instead of $r$.

We say that at some stage of an algorithm, we have a value bound $v$ and an error bound $e$, if we have the following bounds for all parts $r$.

$$
|r| \leq v, \quad\left|\varepsilon_{r}\right| \leq e
$$

Our bounds are always powers of 2 . Whenever we have a value bound $v$, we do all computations with $V=\lg v$ bits before the binary point. For twiddle factors (which are also elements of $\mathcal{R}$ ), we have the following stricter requirements for all its parts $t$.

$$
\left|t+\varepsilon_{t}\right| \leq 1, \quad\left|\varepsilon_{t}\right| \leq 2^{-S}
$$

Lemma 6.1. Let $v_{\boldsymbol{c}}$ be a value bound and $e_{\boldsymbol{c}}$ an error bound on the parts of an element of $\mathcal{R}$ before a multiplication with a twiddle factor. Then

$$
v_{\boldsymbol{d}}=2 P v_{\boldsymbol{c}}
$$

is a value bound and

$$
\begin{aligned}
e_{\boldsymbol{d}} & =2 P e_{\boldsymbol{c}}+v_{\boldsymbol{d}} 2^{-S+1} \\
& =2 P\left(e_{\boldsymbol{c}}+v_{\boldsymbol{c}} 2^{-S+1}\right)
\end{aligned}
$$

is an error bound after the multiplication.
Proof. All parts (real and imaginary parts of coefficients) $t$ of twiddle factors have an absolute value of at most 1. Therefore, the value bound on $|r t|$ is the same as the bound on $|r|$. All multiplications of parts are of the form $\left(r+\varepsilon_{r}\right)\left(t+\varepsilon_{t}\right)$, where $r+\varepsilon_{r}$ is the current approximation to a part $r$, and $t+\varepsilon_{t}$ is the approximation to a part $t$ of a twiddle factor. Using the bounds $|r| \leq v_{\boldsymbol{c}},\left|\varepsilon_{r}\right| \leq e_{\boldsymbol{c}},\left|t+\varepsilon_{t}\right| \leq 1$, and $\left|\varepsilon_{t}\right| \leq 2^{-S}$, where $v_{\boldsymbol{c}}$ and $e_{\boldsymbol{c}}$ are the current value and error bounds respectively, associated with parts $r$, we see that the error after an exact multiplication of the approximated parts is

$$
\left|\varepsilon_{r t}\right|=\left|\left(t+\varepsilon_{t}\right) \varepsilon_{r}+r \varepsilon_{t}\right| \leq e_{\boldsymbol{c}}+v_{\boldsymbol{c}} 2^{-S}
$$

Thus the absolute value of the old error $\left|\varepsilon_{r}\right| \leq e_{\boldsymbol{c}}$ does not increase during this multiplication of parts, but due to the error $\varepsilon_{t}$ in a part of the twiddle factor, a new error of at most $v_{\boldsymbol{c}} 2^{-S}$ is created.

Every coefficient of the product in $\mathcal{R}$ is the sum of $P$ products of coefficients of the two factors. As these coefficients are complex numbers, each part of the product of two coefficients involves two products of real numbers. Thus, we obtain a trivial upper bound $v_{\boldsymbol{d}}$ on the absolute values of the parts of the product if we multiply the upper bound on products of parts $v_{\boldsymbol{c}}$ by $2 P$.

Similarly, the error bound of $e_{\boldsymbol{c}}+v_{\boldsymbol{c}} 2^{-S}$ for the product of two parts is multiplied by $2 P$ to obtain the error bound $e_{\boldsymbol{d}}$ for the parts of the product. Note that non-trivial multiplication in $\mathcal{R}$ is done by reduction to integer multiplication. Thus, starting from the representations with $S$ bits per part, initially all multiplications and additions are done exactly, i.e., without any rounding in between. But finally all parts are rounded, creating an additional new error of at most $v_{\boldsymbol{d}} 2^{-S}$. Thus the total new error for multiplication with roots of unity in $\mathcal{R}$ is at most $4 P v_{c} 2^{-S}$, resulting in a total error bound as claimed by the lemma. $\square$

The proof of the following lemma shows value and error bounds by induction based on the recursive structure of the FFT algorithm. For any vector $\boldsymbol{a}$ over $\mathcal{R}$, let $v_{\boldsymbol{a}}$ be a bound on the absolute values of the parts (real and imaginary parts of any coefficient) of any component of $\boldsymbol{a}$, and let $e_{\boldsymbol{a}}$ be a bound on the absolute values of the error in the parts of any component of $\boldsymbol{a}$. Note that the following lemma considers an arbitrary error bound at the start, because FFT as a recursive procedure is also called in the middle of other FFT computations.

Lemma 6.2. Let $N, P$ be powers of 2 with $N \geq 2$ and $P \geq 1$. Let $L=$ $\lceil(\lg N) / \lg (2 P)\rceil-1$ be the number of levels with computationally intensive twiddle factors. If the input $\boldsymbol{a}$ of an $N$-point FFT has a value bound $v_{\boldsymbol{a}}$ and an error bound $e_{\boldsymbol{a}}$, then the output $\boldsymbol{b}$ has a value bound

$$
v_{\boldsymbol{b}}=N(2 P)^{L} v_{\boldsymbol{a}} \leq N^{2} v_{\boldsymbol{a}}
$$

and an error bound

$$
\begin{aligned}
e_{\boldsymbol{b}} & =N(2 P)^{L} e_{\boldsymbol{a}}+v_{\boldsymbol{b}} 2^{-S}(\lg N+2 L) \\
& =N(2 P)^{L}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S}(\lg N+2 L)\right) \\
& \leq N^{2}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S+1} \lg N\right)
\end{aligned}
$$

Proof. $(2 P)^{L} \leq N$ immediately follows from the definition of $L$ implying both inequalities. We show the other bounds by induction on $N$. We note that $L=0$ for $N \leq 2 P$.

For $N=2$, the algorithm does one addition or subtraction at most doubling the previous values (bounded by $v_{\boldsymbol{a}}$ ) and the previous errors (bounded by $e_{\boldsymbol{a}}$ ). Furthermore, the result is rounded in the last position, creating a new error of at most $v_{b} 2^{-S}$. Thus $v_{\boldsymbol{b}}=2 v_{\boldsymbol{a}}$ is a value bound and $e_{\boldsymbol{b}}=2 e_{\boldsymbol{a}}+v_{\boldsymbol{b}} 2^{-S}$ is an error bound.

For $N>2$, the FFT is a composition of inner FFTs (computing $\boldsymbol{c}$ from $\boldsymbol{a}$ ), multiplications with twiddle factors (computing $\boldsymbol{d}$ from $\boldsymbol{c}$ ), and outer FFTs (computing $\boldsymbol{b}$ from $\boldsymbol{d})$. We use the inductive hypotheses for the inner and outer FFTs.

For $2<N \leq 2 P$, after the inner 2-point FFTs, we have a value bound of $v_{\boldsymbol{c}}=2 v_{\boldsymbol{a}}$ and an error bound $e_{\boldsymbol{c}}=2 e_{\boldsymbol{a}}+v_{\boldsymbol{c}} 2^{-S}$. The multiplications with twiddle factors are just cyclic shifts without any new errors or bound increases, i.e., $v_{\boldsymbol{d}}=v_{\boldsymbol{c}}$ and $e_{\boldsymbol{d}}=e_{\boldsymbol{c}}$ are value and error bounds respectively. After the outer $N / 2$-point FFTs, the inductive hypothesis implies a value bound of

$$
v_{\boldsymbol{b}}=(N / 2) v_{\boldsymbol{d}}=N v_{\boldsymbol{a}}
$$

and the error bound of

$$
\begin{aligned}
e_{\boldsymbol{b}} & =(N / 2) e_{\boldsymbol{d}}+v_{\boldsymbol{b}} 2^{-S} \lg (N / 2) \\
& =(N / 2)\left(2 e_{\boldsymbol{a}}+v_{\boldsymbol{d}} 2^{-S}\right)+v_{\boldsymbol{b}} 2^{-S} \lg (N / 2) \\
& =N e_{\boldsymbol{a}}+v_{\boldsymbol{b}} 2^{-S} \lg N \\
& =N\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S} \lg N\right)
\end{aligned}
$$

For $N>2 P$, after the inner $2 P$-point FFT,

$$
v_{c}=2 P v_{\boldsymbol{a}}
$$

is a value bound, and

$$
e_{\boldsymbol{c}}=2 P e_{\boldsymbol{a}}+v_{\boldsymbol{c}} 2^{-S} \lg (2 P)=2 P\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S} \lg (2 P)\right)
$$

is an error bound.
Multiplication with the twiddle factors increases the value and the previous error by at most a factor of $2 P$ and introduces a new rounding error of at most $v_{\boldsymbol{d}} 2^{-S}$ as well as an error with the same bound due to the inaccuracy in the twiddle factor as shown in Lemma 6.1. Thus

$$
v_{\boldsymbol{d}}=2 P v_{\boldsymbol{c}}=(2 P)^{2} v_{\boldsymbol{a}}
$$

is a value bound and

$$
\begin{aligned}
e_{\boldsymbol{d}} & =2 P e_{\boldsymbol{c}}+v_{\boldsymbol{d}} 2^{-S+1} \\
& =(2 P)^{2}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S} \lg (2 P)\right)+(2 P)^{2} v_{\boldsymbol{a}} 2^{-S+1} \\
& =(2 P)^{2}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S}(\lg (2 P)+2)\right)
\end{aligned}
$$

is an error bound.

Finally, after the outer $(\lg N-\lg (2 P))$-point FFT, the inductive hypothesis provides a value bound of

$$
v_{\boldsymbol{b}}=\frac{N}{2 P}(2 P)^{L-1} v_{\boldsymbol{d}}=N(2 P)^{L} v_{\boldsymbol{a}}
$$

and an error bound of

$$
\begin{aligned}
e_{\boldsymbol{b}} & =\frac{N}{2 P}(2 P)^{L-1} e_{\boldsymbol{d}}+v_{\boldsymbol{b}} 2^{-S}\left(\lg \frac{N}{2 P}+2(L-1)\right) \\
& =N(2 P)^{L-2}(2 P)^{2}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S}(\lg (2 P)+2)\right)+v_{\boldsymbol{b}} 2^{-S}\left(\lg \frac{N}{2 P}+2(L-1)\right) \\
& =N(2 P)^{L} e_{\boldsymbol{a}}+v_{\boldsymbol{b}} 2^{-S}(\lg (2 P)+2)+v_{\boldsymbol{b}} 2^{-S}\left(\lg \frac{N}{2 P}+2(L-1)\right) \\
& =N(2 P)^{L} e_{\boldsymbol{a}}+v_{\boldsymbol{b}} 2^{-S}(\lg N+2 L) \\
& =N(2 P)^{L}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S}(\lg N+2 L)\right)
\end{aligned}
$$

After having computed the FFTs of both factors, we compute products of corresponding values. These are elements of $\mathcal{R}$. The following lemma controls the value and error bound for these multiplications. Let $v_{\boldsymbol{c}}$ and $e_{\boldsymbol{c}}$ refer to the value and error bounds of parts, i.e., real or imaginary parts of coefficients of the vectors $\boldsymbol{c}$ and $\boldsymbol{c}^{\prime}$ representing the two factors.

LEMMA 6.3. If $v_{\boldsymbol{c}}$ is a value bound and $e_{\boldsymbol{c}}$ with $2^{-S} \leq e_{\boldsymbol{c}} \leq v_{\boldsymbol{c}}$ is an error bound for $\boldsymbol{c}$ and $\boldsymbol{c}^{\prime}$ before the multiplications of the values, then $v_{\boldsymbol{d}}=2 P v_{\boldsymbol{c}}^{2}$ is a value bound and $e_{\boldsymbol{d}}=8 P v_{\boldsymbol{c}} e_{\boldsymbol{c}}$ is an error bound afterwards.

Proof. As $v_{\boldsymbol{c}}$ is a value bound on the parts of the factors in $\boldsymbol{c}$ and $\boldsymbol{c}^{\prime}$, obviously $v_{\boldsymbol{c}}^{2}$ is a value bound for their products. Because every part of a product in $\mathcal{R}$ is the sum of $2 P$ products of parts, $v_{\boldsymbol{d}}=2 P v_{\boldsymbol{c}}^{2}$ is a value bound for the result. The multiplications are of the form $\left(r+\varepsilon_{r}\right)\left(r^{\prime}+\varepsilon_{r^{\prime}}\right)$, where $r+\varepsilon_{r}$ and $r^{\prime}+\varepsilon_{r^{\prime}}$ are the current approximations to parts of $\boldsymbol{c}$ and $\boldsymbol{c}^{\prime}$. Using the fact that $|r|,\left|r^{\prime}\right| \leq v_{\boldsymbol{c}}$ and $\left|\varepsilon_{r}\right|,\left|\varepsilon_{r}^{\prime}\right| \leq e_{\boldsymbol{c}}$, where $e_{\boldsymbol{c}}$ is the current error bound, we see that the error after an exact multiplication of the approximated parts has a bound of $\left|r \varepsilon_{r^{\prime}}+r^{\prime} \varepsilon_{r}+\varepsilon_{r} \varepsilon_{r^{\prime}}\right|$, which is generously bounded by $3 v_{\boldsymbol{c}} e_{\boldsymbol{c}}$. During the additions, the error bound increases by a factor $2 P$, and an additional error of at most $v_{\boldsymbol{d}} 2^{-S}=2 P v_{c}^{2} 2^{-S}$ occurs due to rounding. Therefore, using the condition that $v_{\boldsymbol{c}} 2^{-S} \leq e_{\boldsymbol{c}}$ the rounded product has an error bound of

$$
e_{\boldsymbol{d}} \leq 2 P 3 v_{\boldsymbol{c}} e_{\boldsymbol{c}}+v_{\boldsymbol{d}} 2^{-S} \leq 6 P v_{\boldsymbol{c}} e_{\boldsymbol{c}}+2 P v_{\boldsymbol{c}}^{2} 2^{-S} \leq 8 P v_{\boldsymbol{c}} e_{\boldsymbol{c}}
$$

$\square$
Lemma 6.4. For $P=\operatorname{round}(\lg n) \geq 2$ and $N=\operatorname{round}\left(2 n / P^{2}\right)$, where round is rounding to the next power of 2, precision $S \geq 5 \lg N+\lg \lg (2 N)+2 P+4 \lg P+9$ is sufficient for the multiplication of integers of length $n$.

Proof. We start with a value bound of $2^{P}$ and an error bound of 0 . The initial multiplication of the Half-FFT with twiddle factors is analyzed in Lemma 6.1. Thus for the subsequent FFT, the value bound is $v_{\boldsymbol{a}}=2 P 2^{P}$ and the error bound is $e_{\boldsymbol{a}}=v_{\boldsymbol{a}} 2^{-S+1}=2 P 2^{P} 2^{-S+1}$. By Lemma 6.2, the bounds after the $N$-point FFT are

$$
v_{\boldsymbol{c}} \leq N^{2} v_{\boldsymbol{a}}=2 P 2^{P} N^{2}
$$

and

$$
e_{\boldsymbol{c}} \leq N^{2}\left(e_{\boldsymbol{a}}+v_{\boldsymbol{a}} 2^{-S+1} \lg N\right)=N^{2} v_{\boldsymbol{a}} 2^{-S+1} \lg (2 N)=2 P 2^{P} N^{2} \lg (2 N) 2^{-S+1}
$$

By Lemma 6.3, the bounds after the multiplication of values stage are

$$
v_{\boldsymbol{d}}=2 P v_{\boldsymbol{c}}^{2} \leq(2 P)^{3} 2^{2 P} N^{4}
$$

and

$$
e_{\boldsymbol{d}}=8 P v_{\boldsymbol{c}} e_{\boldsymbol{c}} \leq 8 P(2 P)^{2} 2^{2 P} N^{4} \lg (2 N) 2^{-S+1}=32 P^{3} 2^{2 P} N^{4} \lg (2 N) 2^{-S+1}
$$

The inverse $N$-point FFT obeys the bounds of Lemma 6.2 followed by a scaling by $1 / N$ of the value and error bounds. Thus after the inverse FFT, we have the following bounds.

$$
\begin{aligned}
v_{\boldsymbol{b}} & \leq \frac{1}{N} N^{2} v_{\boldsymbol{d}} \leq(2 P)^{3} 2^{2 P} N^{5} \\
e_{\boldsymbol{b}} & \leq \frac{1}{N} N^{2}\left(e_{\boldsymbol{d}}+v_{\boldsymbol{d}} 2^{-S+1} \lg N\right) \\
& =N\left(32 P^{3} 2^{2 P} N^{4} \lg (2 N) 2^{-S+1}+8 P^{3} 2^{2 P} N^{4} \lg N 2^{-S+1}\right) \\
& \leq 40 P^{3} 2^{2 P} N^{5} \lg (2 N) 2^{-S+1}
\end{aligned}
$$

The final part of Inverse-Half-FFT consists of multiplications with twiddle factors. It results in a value bound of

$$
2 P v_{b} \leq(2 P)^{4} 2^{2 P} N^{5}
$$

and an error bound of

$$
2 P\left(e_{\boldsymbol{b}}+v_{\boldsymbol{b}} 2^{-S+1}\right)<96 P^{4} 2^{2 P} N^{5} \lg (2 N) 2^{-S+1}
$$

which is less than $1 / 2$ if

$$
7+4 \lg P+5 \lg N+2 P+\lg \lg (2 N)-S+1 \leq-1
$$

proving the claim.
The bounds of the previous proof are summarized in Table 6.1. As an immediate

| Position in Algorithm | Value bound | Absolute error bound |  |
| :--- | :--- | :--- | :---: |
| Start | $2^{P}$ | 0 |  |
| After first level of Half-FFT | $2 P 2^{P}$ | $2 P 2^{P} 2^{-S+1}$ |  |
| After $N$-point FFT | $2 P 2^{P} N^{2}$ | $2 P 2^{P} N^{2} \lg (2 N) 2^{-S+1}$ |  |
| After multiplication of values | $(2 P)^{3} 2^{2 P} N^{4}$ | $32 P^{3} 2^{2 P} N^{4} \lg (2 N) 2^{-S+1}$ |  |
| After inverse FFT | $(2 P)^{3} 2^{2 P} N^{5}$ | $40 P^{3} 2^{2 P} N^{5} \lg (2 N) 2^{-S+1}$ |  |
| After last level of inverse Half-FFT | $(2 P)^{4} 2^{2 P} N^{5}$ | $96 P^{4} 2^{2 P} N^{5} \lg (2 N) 2^{-S+1}$ |  |
| TABLE 6.1 |  |  |  |
| Bounds on absolute values and errors |  |  |  |

implication of Lemma 6.4, we obtain the following result.
Theorem 6.5. For some $S=\Theta(\lg n)$, doing Half-FFT with precision $S$ is sufficient for the Algorithm Modular-Integer-Multiplication.
7. Complexity. Independently of how an $N$-point Fourier transform is recursively decomposed, the computation can always be visualized by the well known butterfly graph with $\lg N+1$ rows. Every row represents $N$ elements of the ring $\mathcal{R}$. Row 0 represents the input, row $N$ represents the output, and every entry of row $j+1$ is obtained from row $j(0 \leq j<N)$ by an addition or subtraction and possibly a multiplication with a power of $\omega$. When investigating the complexity of performing the multiplications in $\mathcal{R}$ recursively, it is best to still think in terms of the same $\lg N+1$ rows. At the next level of recursion, non-trivial multiplications are done for every $\lg P+1$ st row. It is important to observe that the sum of the lengths of the representations of all entries in such a row grows just by a constant factor from each level of recursion to the next. The blow-up by a constant factor is due to the padding with 0's, and due to the precision needed to represent numerical approximations of complex roots of unity. Padding with 0's occurs when reducing multiplication in $\mathcal{R}$ to modular integer multiplication and during the procedure Decompose.

We do $O\left(\lg ^{*} n\right)$ levels of recursive calls to Modular-Integer-Multiplication. As the total length of a row grows by a constant factor from level to level, we obtain the factor $2^{O\left(\lg ^{*} n\right)}$ in the running time. From a practical point of view, one should not worry too much about this factor. The function $\lg ^{*} n$ in the exponent of the running time actually represents $\lg ^{*} n-4$ or $\lg ^{*} n-3$, which for all practical purposes could be thought as being 1 or 2 , because at a low recursion level, one would switch to a more traditional multiplication method.

The crucial advantage of our new FFT algorithm is the fact that most multiplications with twiddle factors can be done in linear time, as each of them only involves a cyclic shift (with sign change on wrap around) of a vector of coefficients representing an element of $\mathcal{R}$. Indeed, only every $O(\log \log N)$ th row of the FFT requires recursive calls for non-trivial multiplications with roots of unity. We recall that our Fourier transform is over the ring $\mathcal{R}$, whose elements are represented by polynomials of degree $P-1$ with coefficients of length $O(P)=O(\log N)$.

Based on these arguments, one obtains the following recurrence equations for the boolean circuit complexity $T(n)$ of Modular-Integer-Multiplication and $T^{\prime}(N)$ of FFT.

$$
\begin{aligned}
T(n) & =O\left(T^{\prime}\left(n / \log ^{2} n\right)\right) \\
T^{\prime}(N) & =O\left(N \log ^{3} N+\frac{N \log N}{\log \log N} T\left(O\left(\log ^{2} N\right)\right)\right)
\end{aligned}
$$

These recurrence equations have the following solutions.

$$
\begin{aligned}
T(n) & =n \log n 2^{O\left(\lg ^{*} n\right)} \\
T^{\prime}(N) & =N \log ^{3} N 2^{O\left(\lg ^{*} N\right)}
\end{aligned}
$$

A reader convinced by these intuitive arguments may jump directly to Theorem 7.5. We are more formal here, producing the recurrence equations step by step based on the recursive structure of the algorithms.

First we count the number of additions $\operatorname{Add}(N)$ and the number of multiplications $\operatorname{Mult}(N)$ of the $N$-point FFT. The counts refer to operations in $\mathcal{R}$. As always, we
assume $J, K$ and $N$ to be powers of 2 with $J K=N$.

$$
\operatorname{Add}(N)= \begin{cases}0 & \text { if } N=1 \\ 2 & \text { if } N=2 \\ K \operatorname{Add}(J)+J \operatorname{Add}(K) & \text { otherwise }\end{cases}
$$

The solution $\operatorname{Add}(N)=N \lg N$ is immediate.

$$
\operatorname{Mult}(N)= \begin{cases}0 & \text { if } N \leq 2 \\ K \operatorname{Mult}(J)+J \operatorname{Mult}(K)+K J & \text { otherwise }\end{cases}
$$

Induction on $N$ verifies the solution

$$
\operatorname{Mult}(N)= \begin{cases}0 & \text { if } N=1 \\ N(\lg N-1) & \text { if } N \geq 2\end{cases}
$$

One should note that more than half of the multiplications counted by Mult $(N)$ are actually multiplications by 1 . For the sake of simplicity of the presentation, we did not do the corresponding obvious optimization (for $j=0$ and $k=0$ ) in the algorithm FFT.

More interesting than the total number of multiplications $\operatorname{Mult}(N)$, is the number of expensive multiplications EMult $(N)$. Multiplications with (low order) $2 P$ th roots of unity are inexpensive, as they are done by cyclic shifts (with sign changes on wrap around). The recurrence equations for Mult and EMult differ in the start conditions.

$$
\operatorname{EMult}(N)= \begin{cases}0 & \text { if } N \leq 2 P \\ K \operatorname{EMult}(J)+J \operatorname{EMult}(K)+K J & \text { otherwise }\end{cases}
$$

Note that for $N>2 P$, the procedure Select chooses $J=2 P$ and $K=N /(2 P)$, implying $\operatorname{EMult}(J)=0$ simplifying the recurrence equation.

$$
\operatorname{EMult}(N)= \begin{cases}0 & \text { if } N \leq 2 P \\ 2 P \operatorname{EMult}(N /(2 P))+N & \text { otherwise }\end{cases}
$$

Lemma 7.1. This recurrence equation has the solution

$$
E M u l t(N)=N\left(\left\lceil\log _{2 P} N\right\rceil-1\right) \leq N \frac{\lg N}{\lg (2 P)}
$$

Proof. Only the case $N>2 P$ is non-trivial.

$$
\begin{aligned}
\operatorname{EMult}(N) & =2 P E M u l t(N /(2 P))+N \\
& =2 P \frac{N}{2 P}\left(\left\lceil\log _{2 P}(N /(2 P))\right\rceil-1\right)+N \\
& =N\left(\left\lceil\log _{2 P} N\right\rceil-2\right)+N \\
& =N\left(\left\lceil\log _{2 P} N\right\rceil-1\right)
\end{aligned}
$$

LEMMA 7.2. The number of expensive multiplications is
(a) $N\left(\left\lceil\log _{2 P} N\right\rceil-1\right)$ for $F F T$
(b) $N\left\lceil\log _{2 P} N\right\rceil$ for Half-FFT
(c) $N\left(3\left\lceil\log _{2 P} N\right\rceil+1\right)$ for Modular-Integer-Multiplication

Proof. The number of expensive multiplications for FFT is EMult $(N)$. The other results immediately follow from the definitions of the algorithms.

Let $T(n)$ be the boolean circuit complexity of Modular-Integer-Multiplication. $T(n)+O(n)$ is then also the circuit complexity of Integer Multiplication with $0 \leq$ product $\leq 2^{n}$.

Lemma 7.3. Let $n_{0}$ and $N$ be the positive integers from the algorithm Modular-Integer-Multiplication. $n_{0} \geq 16$ is a constant, and $\frac{1}{2} n / \lg ^{2} n \leq N \leq 2 n / \lg ^{2} n \leq n$ for all $n \geq n_{0}$. For some real constants $c, c^{\prime}>1, T(n)$ satisfies the following recurrence

$$
T(n) \leq N\left(3\left\lceil\log _{2 P} N\right\rceil+1\right) \cdot T\left(c \lg ^{2} N\right)+c^{\prime} N \lg N \cdot \lg ^{2} N \quad \text { for } n \geq n_{0}
$$

Proof. This recurrence is based on the counts of additions, multiplications and expensive multiplications, and on the following facts. Binary integers are chopped into $N=O\left(n / \log ^{2} n\right)$ pieces, which are represented by elements of $\mathcal{R}$ encoded by strings of length $O\left(\log ^{2} N\right)$. In this encoding, additions, easy multiplications and all bookkeeping operations are done in linear time. Expensive multiplications in $\mathcal{R}$ are done recursively after encoding the elements of $\mathcal{R}$ as modular integers, which causes a constant factor blow-up .

Now we can claim

$$
T(n) \leq n \lg n 2^{O\left(\lg ^{*} n\right)}
$$

but such a claim resists a direct induction proof, because $\lg ^{*} n$ is not continuous. Even though there are only $O\left(\lg ^{*} n\right)$ recursion levels, $\lg ^{*} n$ does not decrease at each level due to the reduction from $n$ to $O\left(\log ^{2} n\right)$ not $\lg n$. As a trick, we use the fact that $\lg ^{*} \sqrt[4]{n}$ decreases at each level.

Lemma 7.4.

$$
T(n) \leq n \lg n\left(2^{d \lg ^{*} \sqrt[4]{n}}-d^{\prime}\right)
$$

for some $d, d^{\prime}>0$ and all $n \geq 2$.
Proof. From the algorithm Modular-Integer-Multiplication, recall the definitions, $P=\operatorname{round}(\lg n)$ and $N=2 n / P^{2}$. The implications $N \leq \min \left(n, 2 n / \lg ^{2} n\right)$ for $n \geq 2$, and $\left\lceil\log _{2 P} N\right\rceil<\log _{2 P} n$ for $n \geq 16$ are used in Ineq. 7.3 below.

First we do the inductive step. $d, d^{\prime}$ and a constant $n_{0}^{\prime}$ will be determined later. Let $n \geq n_{0}^{\prime} \geq n_{0} \geq 16$. Assume the claim of the lemma holds for all $n^{\prime}$ with $2 \leq n^{\prime}<n$.

$$
\begin{align*}
T(n) & \leq N\left(3\left\lceil\log _{2 P} N\right\rceil+1\right) T\left(c \lg ^{2} N\right)+c^{\prime} N \lg ^{3} N  \tag{7.1}\\
& \leq 4 N\left\lceil\log _{2 P} N\right\rceil c \lg ^{2} N \lg \left(c \lg ^{2} N\right)\left(2^{d \lg ^{*} \sqrt[4]{c \lg ^{2} N}}-d^{\prime}\right)+c^{\prime} N \lg ^{3} N  \tag{7.2}\\
& \leq 8 \frac{n}{\lg ^{2} n} \log _{2 P} n c \lg ^{2} n 2 \lg \lg n\left(2^{d \lg ^{*}\left(\frac{1}{4} \lg n\right)}-d^{\prime}\right)+2 c^{\prime} n \lg n  \tag{7.3}\\
& =16 c n \lg n \frac{\lg \lg n}{\lg 2 P}\left(2^{d\left(\lg ^{*} \sqrt[4]{n}-1\right)}-d^{\prime}\right)+2 c^{\prime} n \lg n  \tag{7.4}\\
& \leq n \lg n\left(2^{d \lg ^{*} \sqrt[4]{n}}-d^{\prime}\right) \tag{7.5}
\end{align*}
$$

Ineq. 7.1 is the recurrence from Lemma 7.3. The inductive hypothesis is used in Ineq. 7.2. For Ineq. 7.3, we use the definitions of $P$ and $N$, and we select $n_{0}^{\prime}$ sufficiently
big such that $\sqrt[4]{c \lg ^{2} N} \leq \frac{1}{4} \lg n$ for all $n \geq n_{0}^{\prime}$. Finally, for Ineq. 7.5, we use $\lg \lg n \leq$ $\lg 2 P \leq \lg \lg n+2 \leq 2 \lg \lg n$, and we just choose $d$ and $d^{\prime}$ sufficiently big such that $16 c \leq 2^{d}$ and $-8 c d^{\prime}+2 c^{\prime} \leq-d^{\prime}$. Furthermore, we make sure $d$ is big enough that the claim of the lemma holds for all $n$ with $2 \leq n<n_{0}^{\prime}$.

Lemma 7.4 implies our main results for circuit complexity and (except for the organizational details) for multitape Turing machines.

THEOREM 7.5. Multiplication of binary integers of length $n$ can be done by a boolean circuit of size $n \log n 2^{O\left(\lg ^{*} n\right)}$.

Theorem 7.6. Multiplication of binary integers of length $n$ can be done in time $n \log n 2^{O\left(\lg ^{*} n\right)}$ on a 2-tape Turing machine.

A detailed proof of Theorem 7.6 would be quite tedious. Nevertheless, it should be obvious that due to the relatively simple structure of the algorithms, there is no principle problem to implement them on Turing machines.

As an important application of integer multiplication, we obtain corresponding bounds for the multiplication of polynomials by boolean circuits or Turing machines. We are looking at bit complexity, not assuming that products of coefficients can be obtained in one step.

Corollary 7.7. Products of polynomials of degree less than $n$, with a $2^{O(m)}$ upper bound on the absolute values of their real or complex coefficients, can be approximated in time $m n \log m n 2^{O\left(\log ^{*} m n\right)}$ with an absolute error bound of $2^{-m}$, for a given $m=\Omega(\log n)$.

Proof. Schönhage [Sch82] has shown how to reduce the multiplication of polynomials with complex coefficients to integer multiplication with only a constant factor in time increase.

Indeed, multiplying polynomials with real or complex coefficients is a major area where long integer multiplication is very useful. Long integer multiplication is used extensively for finding large prime numbers. Another application is the computation of billions of digits of $\pi$ to study patterns. A very practical application is the testing computational hardware.
8. Open Problem. Besides the obvious question whether integer multiplication is in $O(n \log n)$, a multiplication algorithm running in time $O\left(n \log n \lg ^{*} n\right)$ would also be very desirable. It could be achieved, if one could avoid the constant factor cost increase from one recursion level to the next. Furthermore, it would be nice to have an implementation that compares favorably with current implementations of the algorithm of Schönhage and Strassen. The asymptotic improvement from $O(n \log n \log \log n)$ to $n \log n 2^{O\left(\log ^{*} n\right)}$ might suggest that an actual speed-up only shows up for astronomically large numbers. Indeed, the expressions are not very helpful to judge the performance for reasonable values of $n$. But one should notice that $\lg ^{*} n$ in the exponent really just represents an upper bound on the nesting of recursive calls to integer multiplication. For any practical purposes, one would nest these calls at most twofold.

## REFERENCES

[AHU74] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman, The design and analysis of computer algorithms, Addison-Wesley, Reading, Massachusetts, 1974.
[BCS97] Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi, Algebraic complexity theory, Springer-Verlag, Berlin, 1997.
[Ber] Daniel J. Bernstein, Personal communication.
[Ber68] Glenn D. Bergland, A Fast Fourier Transform algorithm using base 8 iterations, Mathematics of Computation 22 (1968), no. 102, 275-279.
[BL04] Peter Bürgisser and Martin Lotz, Lower bounds on the bounded coefficient complexity of bilinear maps, Journal of the ACM 51 (2004), no. 3, 464-482.
[CA69] Stephen A. Cook and Stål O. Aanderaa, On the minimum computation time of functions, Transactions of the American Mathematical Society 142 (1969), 291-314.
[Coo66] Stephen A. Cook, On the minimum computation time of functions, Ph.D. thesis, Harvard University, 1966.
[CT65] James W. Cooley and John W. Tukey, An algorithm for the machine calculation of complex Fourier series, Math. of Comput. 19 (1965), 297-301.
[DH84] P. Duhamel and H. Hollman, Split-radix FFT algorithms, Electronics Letters 20 (1984), 14-16.
[DKSS08] Anindya De, Piyush Kurur, Chandan Saha, and Ramprasad Saptharishi, Fast integer multiplication using modular arithmetic, STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing (New York, NY, USA), ACM, 2008, pp. 499-506.
[Für89] Martin Fürer, On the complexity of integer multiplication, (Extended Abstract), Tech. Report Technical report CS-89-17, Department of Computer Science, The Pennsylvania State University, 1989.
[Für07] , Faster integer multiplication, STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing (New York, NY, USA), ACM, 2007, pp. 57-66.
[GS66] W. Morven Gentleman and Gordon Sande, Fast Fourier transforms: For fun and profit, AFIPS '66 (Fall) Joint Computer Conferences: Proc. Fall Joint Computer Conference (NY), vol. 29, ACM, 1966, pp. 563-578.
[HJB84] Michael T. Heideman, D. H. Johnson, and C. S. Burrus, Gauss and the history of the FFT, IEEE Acoustics, Speech, and Signal Processing 1 (1984), 14-21.
[JF07] Steven G. Johnson and Matteo Frigo, A modified split-radix FFT with fewer arithmetic operations, IEEE Transactions on Signal Processing 55 (2007), no. 1, 111-119.
[KO62] Anatolii Karatsuba and Yu Ofman, Multiplication of multidigit numbers on automata, Doklady Akademii Nauk SSSR 145 (1962), no. 2, 293-294, (in Russian). English translation in Soviet Physics-Doklady 7, 595-596,1963.
[Mor73] Jacques Morgenstern, Note on a lower bound on the linear complexity of the fast Fourier transform, Journal of the ACM 20 (1973), no. 2, 305-306.
[Pan86] Victor Ya. Pan, The trade-off between the additive complexity and the asynchronicity of linear and bilinear algorithms, Information Processing Letters 22 (1986), no. 1, 11-14.
[Pap79] Christos H. Papadimitriou, Optimality of the fast Fourier transform, Journal of the ACM 26 (1979), no. 1, 95-102.
[PFM74] Michael S. Paterson, Michael J. Fischer, and Albert R. Meyer, An improved overlap argument for on-line multiplication, Tech. Report 40, Project MAC, MIT, January 1974.
[Sch66] Arnold Schönhage, Multiplikation großer Zahlen, Computing 1 (1966), no. 3, 182-196 (German).
[Sch82] , Asymptotically fast algorithms for the numerical multiplication and division of polynomials with complex coeficients, Computer Algebra, EUROCAM '82, European Computer Algebra Conference, Marseille, France, 5-7 April, 1982, Proceedings, Lecture Notes in Computer Science, vol. 144, Springer, 1982, pp. 3-15.
[SS71] Arnold Schönhage and Volker Strassen, Schnelle Multiplikation grosser Zahlen, Computing 7 (1971), 281-292.
[Too63] Andre L. Toom, The complexity of a scheme of functional elements simulating the multiplication of integers, Dokl. Akad. Nauk SSSR 150 (1963), 496-498, (in Russian). English translation in Soviet Mathematics 3, 714-716, 1963.
[VS94] Jeffrey Scott Vitter and Elizabeth A. M. Shriver, Algorithms for parallel memory II: Hierarchical multilevel memories, Algorithmica 12 (1994), no. 2/3, 148-169.
[Yav68] R. Yavne, An economical method for calculating the discrete Fourier transform, Proc. AFIPS Fall Joint Computer Conf. 33 (1968), 115-125.


[^0]:    *Research supported in part by the Penn State Grove Award and the National Science Foundation under Grant CCF-0728921, Department of Computer Science and Engineering, Pennsylvania State University, Visiting: Laboratoire d'Algorithmique, EPFL, Lausanne, and Institut für Mathematik, Universität Zürich, furer@cse.psu.edu
    ${ }^{1} \lg$ refers to the logarithm to the base 2 , and $\lg ^{*} n=\min \left\{i \geq 0: \lg ^{(i)} n \leq 1\right\}$ with $\lg ^{(0)} n=n$ and $\lg ^{(i+1)} n=\lg \lg ^{(i)} n$

