## SCISPACE <br> formerly Typeset

〕 Open access • Journal Article • DOI:10.4007/ANNALS.2021.193.2.4
Integer multiplication in time $\mathbf{O}(\mathbf{n} \log \mathbf{n})$ - Source link
David Harvey, Joris van der Hoeven
Institutions: École Polytechnique
Published on: 01 Jan 2021 - Annals of Mathematics (Princeton University and the Institute for Advanced Study)
Topics: Strassen algorithm, Fast Fourier transform, Time complexity and Polynomial

Related papers:

- Modern Computer Algebra
- Asymptotically Fast Algorithms for the Numerical Multiplication and Division of Polynomials with Complex Coeficients
- Fast integer multiplication using generalized Fermat primes
- Fast on-line integer multiplication
- On the minimum computation time of functions


## Integer multiplication in time $\mathrm{O}(\mathrm{n} \log \mathrm{n})$

David Harvey, Joris van der Hoeven

## To cite this version:

David Harvey, Joris van der Hoeven. Integer multiplication in time $\mathrm{O}(\mathrm{n} \log \mathrm{n})$. Annals of Mathematics, Princeton University, Department of Mathematics, In press. hal-02070778v2

## HAL Id: hal-02070778 https://hal.archives-ouvertes.fr/hal-02070778v2

Submitted on 28 Nov 2020

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

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

# Integer multiplication in time $O(n \log n)$ 

David Harvey and Joris van der Hoeven


#### Abstract

We present an algorithm that computes the product of two $n$-bit integers in $O(n \log n)$ bit operations, thus confirming a conjecture of Schönhage and Strassen from 1971. Our complexity analysis takes place in the multitape Turing machine model, with integers encoded in the usual binary representation. Central to the new algorithm is a novel "Gaussian resampling" technique that enables us to reduce the integer multiplication problem to a collection of multidimensional discrete Fourier transforms over the complex numbers, whose dimensions are all powers of two. These transforms may then be evaluated rapidly by means of Nussbaumer's fast polynomial transforms.


## 1. Introduction

Let $\mathrm{M}(n)$ denote the time required to multiply two $n$-bit integers. We work in the multitape Turing model, in which the time complexity of an algorithm refers to the number of steps performed by a deterministic Turing machine with a fixed, finite number of linear tapes [35]. The main results of this paper also hold in the Boolean circuit model [41, Sec. 9.3], with essentially the same proofs.

For functions $f\left(n_{1}, \ldots, n_{k}\right)$ and $g\left(n_{1}, \ldots, n_{k}\right)$, we write $f(n)=O(g(n))$ to indicate that there exists a constant $C>0$ such that $f(n) \leqslant C g(n)$ for all tuples $n=\left(n_{1}, \ldots, n_{k}\right)$ in the domain of $f$. Similarly, we write $f(n)=\Omega(g(n))$ to mean that $f(n) \geqslant C g(n)$ for all $n$ in the domain of $f$, and $f(n)=\Theta(g(n))$ to indicate that both $f(n)=O(g(n))$ and $f(n)=\Omega(g(n))$ hold. ¿From Section 2 onwards we will always explicitly restrict the domain of $f$ to ensure that $g(n)>0$ throughout this domain. However, in this Introduction we will slightly abuse this notation: when writing for instance $f(n)=O(n \log n \log \log n)$, we tacitly assume that the domain of $f$ has been restricted to $\left[n_{0}, \infty\right)$ for some sufficiently large threshold $n_{0}$.

Schönhage and Strassen conjectured in 1971 that the true complexity of integer multiplication is given by $\mathrm{M}(n)=\Theta(n \log n)$ [40], and in the same paper established their famous upper bound $\mathrm{M}(n)=O(n \log n \log \log n)$. In 2007 their result was sharpened by Fürer to $\mathrm{M}(n)=O\left(n \log n K^{\log ^{*} n}\right)$ [12, 13] for some unspecified constant $K>1$, where $\log ^{*} n$ denotes the iterated logarithm, i.e., $\log ^{*} x:=\min \left\{k \geqslant 0: \log ^{\circ k} x \leqslant 1\right\}$. Prior to the present work, the record stood at $\mathrm{M}(n)=O\left(n \log n 4^{\log ^{*} n}\right)[22]$.

The main result of this paper is a verification of the upper bound in Schönhage and Strassen's conjecture, thus closing the remaining $4 \log ^{*} n$ gap:

Theorem 1.1. There is an integer multiplication algorithm achieving

$$
\mathrm{M}(n)=O(n \log n)
$$

[^0]If the Schönhage-Strassen conjecture is correct, then Theorem 1.1 is asymptotically optimal. Unfortunately, no super-linear lower bound for $\mathrm{M}(n)$ is known. Perhaps the best available evidence in favour of the conjecture is the $\Omega(n \log n)$ lower bound $[6,36]$ that has been proved for the "on-line" variant of the problem, in which the $k$-th bit of the product must be written before the $(k+1)$-th bits of the multiplicands are read. Again, the true complexity of on-line multiplication is not known: currently, the best known upper bound is $O(n \log n \exp (C \sqrt{\log \log n}))$ for $C=\sqrt{2 \log 2}+o(1)[29]$.

Theorem 1.1 has many immediate consequences, as many computational problems may be reduced to integer multiplication. For example, the theorem implies that quotients and $k$-th roots of real numbers may be computed to a precision of $n$ significant bits in time $O(n \log n)$, that transcendental functions and constants such as $e^{x}$ and $\pi$ may be computed to precision $n$ in time $O\left(n \log ^{2} n\right)$, and that the greatest common divisor of two $n$-bit integers may be found in time $O\left(n \log ^{2} n\right)$ [5].

Another interesting application is to the problem of computing DFTs (discrete Fourier transforms) over $\mathbb{C}$. Given a transform length $m \geqslant 2$ and a target accuracy of $p=\Omega(\log m)$ bits, it was pointed out in $[20,25]$ that one may use Bluestein's trick [2] followed by Kronecker substitution [14, Corollary 8.27] to reduce a given DFT of length $m$ to an integer multiplication problem of size $O(m p)$. Theorem 1.1 then implies that the DFT may be evaluated in time $O(m p \log (m p))$. This compares favourably with the traditional FFT (fast Fourier transform) approach, which requires $O(m \log m)$ operations in $\mathbb{C}$, and thus time $O(m \log m \mathrm{M}(p))=$ $O(m p \log m \log p)$ in the Turing model.

This faster method for computing DFTs over $\mathbb{C}$ leads to various further applications. One such application is the conversion of an $n$-digit integer from one base to another, for example from binary to decimal, in time $O\left(n \log ^{2} n / \log \log n\right)$ [30]. Alternatively, if one wishes to multiply two $n$-digit integers in a fixed base $\beta \neq 2$, then it is possible to adapt the new algorithm to obtain a direct $O(n \log n)$-time multiplication algorithm that works in base $\beta$ throughout. This is asymptotically faster than reducing the problem to binary via the above-mentioned base conversion algorithms.

All of the algorithms presented in this paper can be made completely explicit, and all implied big- $O$ constants are in principle effectively computable. On the other hand, we make no attempt to minimise these constants or to otherwise exhibit a practical multiplication algorithm. Our aim is to establish the theoretical $O(n \log n)$ bound as directly as possible.

We will actually describe two new multiplication algorithms. The first one depends on an unproved hypothesis concerning the least prime in an arithmetic progression. This hypothesis is weaker than standard conjectures in this area, but stronger than the best unconditional results currently available. We give only a brief sketch of this algorithm (see Section 1.2.1); a detailed treatment is given in the companion paper [24], which also presents an analogue of this algorithm for multiplication in $\mathbb{F}_{q}[x]$. The bulk of the present paper (Sections $2-5$ ) concentrates on working out the details of the second algorithm, which is technically more involved, but has the virtue of reaching the $O(n \log n)$ bound unconditionally.

In the remainder of Section 1, we review the literature on integer multiplication (Section 1.1), and give an overview of the new algorithms (Section 1.2).
1.1. Survey of integer multiplication algorithms. The first improvement on the classical $\mathrm{M}(n)=O\left(n^{2}\right)$ bound was found by Karatsuba in 1962. Significant progress was made during the 1960s by Toom, Cook, Schönhage and Knuth; see [25, Sec. 1.1] for further historical details and references for this period. FFTs were brought into the picture by Schönhage and Strassen [40] soon after the publication of the FFT by Cooley and Tukey [7]; see [28] for more on the history of the FFT. The multiplication algorithms published since [40] may be roughly classified into four families:
(1) Schönhage-Strassen's first algorithm [40] is, in retrospect, the most straightforward FFT-based integer multiplication algorithm imaginable. By splitting the $n$-bit multiplicands into chunks of $\operatorname{size} \Theta(\log n)$, they reduce to the problem of multiplying polynomials in $\mathbb{Z}[x]$ of degree $\Theta(n / \log n)$ and coefficient size $\Theta(\log n)$. The product in $\mathbb{Z}[x]$ is handled by means of FFTs over $\mathbb{C}$, i.e., evaluating the polynomials at suitable complex roots of unity, multiplying their values pointwise in $\mathbb{C}$, and then interpolating to obtain the product polynomial. Elements of $\mathbb{C}$ are represented approximately, with a precision of $\Theta(\log n)$ bits. Arithmetic operations in $\mathbb{C}$ (such as multiplication) are reduced to arithmetic in $\mathbb{Z}$ by scaling by a suitable power of two. This leads to the recursive estimate

$$
\mathrm{M}(n)=O\left(n \mathrm{M}\left(n^{\prime}\right)\right)+O(n \log n), \quad n^{\prime}=O(\log n)
$$

whose explicit solution is

$$
\mathrm{M}(n)=O\left(K^{\log ^{*} n} n \log n \log \log n \cdots \log ^{\circ\left(\left(\log ^{*} n\right)-1\right)} n\right)
$$

for some constant $K>0$. The algorithm achieves an exponential size reduction at each recursion level, from $n$ to $O(\log n)$, and the number of levels is $\log ^{*} n+O(1)$.

Pollard suggested a similar algorithm at around the same time [37], working over a finite field rather than $\mathbb{C}$. He did not analyse the bit complexity, but with some care one can prove essentially the same complexity bound as for the complex case (some technical difficulties arise due to the cost of finding suitable primes; these may be resolved by techniques similar to those discussed in [25, Sec. 8.2]).
(2) Schönhage-Strassen's second algorithm is the more famous and arguably the more ingenious of the two algorithms presented in [40]. It is probably the most widely used large-integer multiplication algorithm in the world today, due to the highly optimised implementation included in the free GNU Multiple Precision Arithmetic Library (GMP) [17, 15], which underlies the large-integer capabilities of all of the major contemporary computer algebra systems.

The basic recursive problem is taken to be multiplication in $\mathbb{Z} /\left(2^{n}+1\right) \mathbb{Z}$, where $n$ is a power of two. Let $n^{\prime}:=2^{\left\lceil\left(\log _{2} 2 n\right) / 2\right\rceil}=\Theta\left(n^{1 / 2}\right)$ and $T:=2 n / n^{\prime}=\Theta\left(n^{1 / 2}\right)$, so that $\left(n^{\prime}\right)^{2} \in\{2 n, 4 n\}$ and $T \mid n^{\prime}$; then by splitting the inputs into chunks of size $n^{\prime} / 2$, the problem is reduced to multiplication in $R[x] /\left(x^{T}+1\right)$ where $R:=\mathbb{Z} /\left(2^{n^{\prime}}+1\right) \mathbb{Z}$. The powers of 2 in $R$ are sometimes called "synthetic" roots of unity, as they have been synthesised algebraically, or "fast" roots of unity, as one can multiply an element of $R$ by an arbitrary power of 2 in linear time, i.e., in time $O\left(n^{\prime}\right)$. Consequently, for $\omega:=2^{n^{\prime} / T}$, one may evaluate a polynomial at $\omega, \omega^{3}, \ldots, \omega^{2 T-1}$ (the roots of $\left.x^{T}+1\right)$ via the FFT in time $O\left(\left(n^{\prime} \log n^{\prime}\right) n^{\prime}\right)=O(n \log n)$. The original multiplication problem is thus reduced to $T$ pointwise multiplications in $R$, which are handled recursively. Writing $\mathrm{M}_{1}(n)$ for the cost of a product in $\mathbb{Z} /\left(2^{n}+1\right) \mathbb{Z}$,
one obtains the recurrence

$$
\begin{equation*}
\mathrm{M}_{1}(n)<\frac{2 n}{n^{\prime}} \mathrm{M}_{1}\left(n^{\prime}\right)+O(n \log n), \quad n^{\prime}=O\left(n^{1 / 2}\right) . \tag{1.1}
\end{equation*}
$$

Unlike the first Schönhage-Strassen algorithm, this algorithm performs only a geometric size reduction, from $n$ to $O\left(n^{1 / 2}\right)$, at each recursion level, and the number of recursion levels is $\log _{2} \log n+O(1)=O(\log \log n)$.

The constant 2 in (1.1), which arises from zero-padding in the initial splitting stage, plays a crucial role in the complexity analysis: it ensures that at each recursion level, the total cost of the "fast" FFTs remains $O(n \log n)$, with the same implied constant at each level. The overall cost is thus $\mathrm{M}_{1}(n)=O(n \log n \log \log n)$.
(3) Fürer's algorithm $[12,13]$ combines the best features of the two SchönhageStrassen algorithms: the exponential size reduction from the first algorithm, and the fast roots of unity from the second one. The overall strategy is similar to the first algorithm, but instead of working over $\mathbb{C}$, one uses a bivariate splitting to reduce to a polynomial multiplication problem over $R:=\mathbb{C}[y] /\left(y^{r}+1\right)$, where $r=\Theta(\log n)$ is a power of two. This ring contains a synthetic root of unity $y$ of order $2 r$, but also inherits higher-order roots of unity from $\mathbb{C}$. Elements of $\mathbb{C}$ are represented approximately, with a precision of $O(\log n)$ bits; thus an element of $R$ occupies $O\left((\log n)^{2}\right)$ bits.

Fürer's key insight is to apply the Cooley-Tukey FFT decomposition in radix $2 r$ instead of radix two. He decomposes each "long" transform of length $\Theta\left(n /(\log n)^{2}\right)$ into many "short" transforms of length $2 r$, with one round of expensive "twiddle factor" multiplications interposed between each layer of short transforms. The short transforms take advantage of the synthetic roots of unity, and the twiddle factor multiplications are handled recursively (via Kronecker substitution). This leads to the recurrence

$$
\mathrm{M}(n)=O\left(\frac{n \log n}{n^{\prime} \log n^{\prime}} \mathrm{M}\left(n^{\prime}\right)\right)+O(n \log n), \quad n^{\prime}=O\left((\log n)^{2}\right),
$$

and then to the explicit bound $\mathrm{M}(n)=O\left(n \log n K^{\log ^{*} n}\right)$ for some constant $K>1$. Fürer did not give a specific value for $K$, but it is argued in [25, Sec. 7] that careful optimisation of his algorithm leads to the value $K=16$.

Several authors have given variants of Fürer's algorithm that also achieve $\mathrm{M}(n)=$ $O\left(n \log n K^{\log ^{*} n}\right)$, using essentially the same idea but working over different rings. De, Kurur, Saha and Saptharishi [10] replace $\mathbb{C}$ by a $p$-adic ring $\mathbb{Q}_{p}$; this has the benefit of avoiding numerical analysis over $\mathbb{C}$, but the value of $K$ becomes somewhat larger. Covanov and Thomé give another variant that achieves $K=4$, conditional on a conjecture on the distribution of generalised Fermat primes [8].
(4) The Harvey-van der Hoeven-Lecerf algorithm [25] follows Fürer in decomposing a "long" transform into many "short" transforms of exponentially smaller length. However, instead of working over a ring containing fast roots of unity, one works directly over $\mathbb{C}$ (as in the first Schönhage-Strassen algorithm), and converts the short transforms back to multiplication problems via Bluestein's trick [2]. These short products are then handled recursively.

The first version given in [25] achieved $\mathrm{M}(n)=O\left(n \log n K^{\log ^{*} n}\right)$ with $K=8$. The value of $K$ was improved gradually over a sequence of papers [18, 19, 21], reaching $K=4$ in [22]. All of these algorithms perform exponential size reduction, and the number of recursion levels is $\log ^{*} n+O(1)$.

An interesting feature of these algorithms - related to the fact that they dispense with the need for fast roots of unity - is that they can be adapted to prove bounds of the form $O\left(n \log n K^{\log ^{*} n}\right)$ for the cost of multiplying polynomials in $\mathbb{F}_{q}[x]$ of degree $n$ (for fixed $q$ ). This was first established with the constant $K=8$ in [26], and improved to $K=4$ in [23]. As mentioned previously, the first of the two new algorithms presented in this paper may be adapted to obtain an $O(n \log n)$ bound for the $\mathbb{F}_{q}[x]$ case [24], but unfortunately this result is still conditional and so does not yet supersede the unconditional $O\left(n \log n 4^{\log ^{*} n}\right)$ bound given in [23].
1.2. Overview of new algorithms. Our new algorithms are motivated by the observation that certain multivariate polynomial rings admit particularly efficient multiplication algorithms. Let $r$ be a power of two, and for $d \geqslant 2$ consider the ring

$$
\begin{equation*}
R\left[x_{1}, \ldots, x_{d-1}\right] /\left(x_{1}^{t_{1}}-1, \ldots, x_{d-1}^{t_{d-1}}-1\right), \quad R:=\mathbb{C}[y] /\left(y^{r}+1\right) \tag{1.2}
\end{equation*}
$$

where $t_{i} \mid 2 r$ for all $i$. One may multiply in this ring by first using FFTs to evaluate each $x_{i}$ at the synthetic $t_{i}$-th roots of unity (the powers of $y^{2 r / t_{i}}$ ), then multiplying pointwise in $R$, and finally performing inverse FFTs. Such transforms were studied extensively by Nussbaumer in the late 1970s (see for example [32]), and are sometimes known as fast polynomial transforms. They consist entirely of additions and subtractions in $\mathbb{C}$, and require no multiplications in $\mathbb{C}$ whatsoever.

In Sections 1.2 .1 and 1.2 .2 below, we outline two different ways of fashioning an integer multiplication algorithm from the polynomial multiplication algorithm just described. The key issue is to show how to transport an integer multiplication problem, which is intrinsically one-dimensional, to a ring of the type (1.2).

In both cases, we begin with the following setup. Suppose that we wish to multiply two $n$-bit integers. We choose a dimension parameter $d \geqslant 2$ and distinct primes $s_{1}, \ldots, s_{d} \approx(n / \log n)^{1 / d}$, subject to certain conditions that will be explained in Sections 1.2.1 and 1.2.2. Just as in the first Schönhage-Strassen algorithm, we split the inputs into around $n / \log n$ chunks of roughly $\log n$ bits, thereby reducing the problem to multiplication in $\mathbb{Z}[x] /\left(x^{s_{1} \cdots s_{d}}-1\right)$. Now, following a technique described by Agarwal and Cooley [1] (which is closely related to the Good-Thomas FFT algorithm $[16,42])$, we observe that the Chinese remainder theorem induces an isomorphism $\mathbb{Z}[x] /\left(x^{s_{1} \cdots s_{d}}-1\right) \cong \mathbb{Z}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{s_{1}}-1, \ldots, x_{d}^{s_{d}}-1\right)$, so the problem amounts to computing a product in the latter ring. For this, it suffices to show how to efficiently compute a multidimensional complex DFT of size $s_{1} \times \cdots \times s_{d}$, i.e., with respect to the complex $s_{i}$-th roots of unity, to an accuracy of $O(\log n)$ bits.
1.2.1. A conditional algorithm - Rader's trick. Suppose that we are able to choose the primes $s_{1}, \ldots, s_{d}$ so that $s_{i}=1(\bmod r)$, where $r \geqslant 2$ is a power of two, and where the $s_{i}$ are not much larger than $r$. We may then deploy a multidimensional generalisation of Rader's algorithm [38] to reduce the given DFT of size $s_{1} \times \cdots \times s_{d}$ to a multiplication problem in the ring $\mathbb{C}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{s_{1}-1}-1, \ldots, x_{d}^{s_{d}-1}-1\right)$ (together with some lower-dimensional multiplication problems of negligible cost). Crucially, the convolution lengths have been reduced from $s_{i}$ to $s_{i}-1$. Writing $s_{i}-1=q_{i} r$, where the $q_{i}$ are "small", we may further reduce this product to a collection of complex DFTs of size $q_{1} \times \cdots \times q_{d}$, plus a collection of multiplication problems in $\mathbb{C}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{r}-1, \ldots, x_{d}^{r}-1\right)$. After replacing $x_{d}$ with $e^{\pi \mathrm{i} / r} y$, we see that the latter products are exactly of the type (1.2). As discussed previously, we may use synthetic FFTs to reduce such a product to a collection of pointwise
products in $R=\mathbb{C}[y] /\left(y^{r}+1\right)$. These in turn are converted to integer multiplication problems via Kronecker substitution, and then handled recursively.

The main sticking point in the above algorithm is the cost of the auxiliary DFTs of size $q_{1} \times \cdots \times q_{d}$. There are various options available for evaluating these DFTs, but to ensure that this step does not dominate the complexity, the key issue is to keep the size of the $q_{i}$ under control. What we are able to prove is the following. For positive, relatively prime integers $r$ and $a$, define

$$
P(a, r):=\min \{s>0: s \text { prime and } s=a \bmod r\}
$$

and put $P(r):=\max _{a} P(a, r)$. Linnik's theorem states that there is an absolute constant $L>1$ such that $P(r)=O\left(r^{L}\right)$. (In our application, we are interested in bounding $P(1, r)$ when $r$ is a power of two.) The best published value for $L$ is currently $L=5.18$ [43], and under the Generalised Riemann Hypothesis one may take $L=2+\varepsilon$ for any $\varepsilon>0$ [27]. In the companion paper [24], we present an integer multiplication algorithm following the plan just described, but working over a finite field instead of $\mathbb{C}$. We prove that if Linnik's theorem holds for some $L<1+\frac{1}{303}$, and if we take $d$ near $10^{6}$, then the cost of the auxiliary DFTs can be controlled and one does in fact obtain an overall $\mathrm{M}(n)=O(n \log n)$ bound. We expect that the same argument works over $\mathbb{C}$, with a possibly different threshold for $L$, but we have not worked out the details.

On the other hand, it is widely expected that the bound $P(r)=O\left(r^{L}\right)$ should hold for any $L>1$. For this reason, we strongly suspect that the algorithm sketched above does run in time $O(n \log n)$, despite us being unable to supply a proof. For further discussion, and examples of even stronger bounds for $P(r)$ that are expected to hold, see [24].

Remark 1.2. The idea of evaluating a multidimensional transform via a combination of Rader's algorithm and polynomial transforms was previously suggested in a different context by Nussbaumer and Quandalle [33, p. 141].
1.2.2. An unconditional algorithm - Gaussian resampling. The rest of the paper is devoted to the second method. Here we choose the primes $s_{1}, \ldots, s_{d}$ in such a way that each $s_{i}$ is slightly smaller than a power of two $t_{i}$, and so that $t_{1} \cdots t_{d}=$ $O\left(s_{1} \cdots s_{d}\right)$. Finding such primes is easily accomplished using Eratosthenes' sieve and the prime number theorem with a suitable error term (see Lemma 5.1).

Assume as before that we wish to compute a complex multidimensional DFT of size $s_{1} \times \cdots \times s_{d}$, to an accuracy of $O(\log n)$ bits. Our key innovation is to show that this problem may be reduced directly to the problem of computing a complex multidimensional DFT of size $t_{1} \times \cdots \times t_{d}$.

The idea of the reduction is as follows. Suppose that we are given as input an $s_{1} \times \cdots \times s_{d}$ array of complex numbers $u=\left(u_{j_{1}, \ldots, j_{d}}\right)_{0 \leqslant j_{i}<s_{i}}$. We may regard this array as lying inside the $d$-dimensional unit torus $(\mathbb{R} / \mathbb{Z})^{d}$ : we imagine the coefficient $u_{j_{1}, \ldots, j_{d}}$ to be plotted at coordinates $\left(j_{1} / s_{1}, \ldots, j_{d} / s_{d}\right)$ in the torus (see Figure 1). We construct from $u$ an intermediate $t_{1} \times \cdots \times t_{d}$ array $v=\left(v_{k_{1}, \ldots, k_{d}}\right)_{0 \leqslant k_{i}<t_{i}}$. Again, we think of $v_{k_{1}, \ldots, k_{d}}$ as being plotted at coordinates $\left(k_{1} / t_{1}, \ldots, k_{d} / t_{d}\right)$ in the torus. The coefficients of $v$ are defined to be certain linear combinations of the coefficients of $u$. The weights are essentially $d$-dimensional Gaussians, so each coefficient of $v$ depends mainly on the "nearby" coefficients of $u$ within the torus.

This construction has two crucial properties. First, the rapid decay of the Gaussians allows us to compute (i.e., approximate) the coefficients of $v$ very quickly from


Figure 1. Torus $(\mathbb{R} / \mathbb{Z})^{2}$ with $13 \times 11$ source array (white circles) superimposed over $16 \times 16$ target array (black circles)
those of $u$; indeed, the cost of this step is asymptotically negligible compared to the cost of the DFTs themselves. Second, using the fact that the Fourier transform of a Gaussian is a Gaussian, we will show that $\hat{u}$ and $\hat{v}$ (the DFTs of $u$ and $v$ ) are related by a fairly simple system of linear equations. In fact, the matrix of this system is of the same type as the matrix relating $u$ and $v$. The system is somewhat overdetermined, because $t_{1} \cdots t_{d}>s_{1} \cdots s_{d}$. Provided that the ratios $t_{i} / s_{i}$ are not too close to 1 , we will show that this system may be solved in an efficient and numerically stable manner, and that we may therefore recover $\hat{u}$ from $\hat{v}$. This procedure forms the core of our "Gaussian resampling" method, and is developed in detail in Section 4. It is closely related to the Dutt-Rokhlin algorithm for non-equispaced FFTs [11]; see Section 4.4 .3 for a discussion of the similarities and differences.

We have therefore reduced to the problem of computing $\hat{v}$ from $v$, and we are free to do this by any convenient method. Note that this is a DFT of size $t_{1} \times \cdots \times t_{d}$ rather than $s_{1} \times \cdots \times s_{d}$. In Section 3 we will show how to use a multivariate generalisation of Bluestein's algorithm [2] to reduce this DFT to a multiplication problem in a ring of the form (1.2). As already pointed out, such a product may be handled efficiently via synthetic FFTs; the details of this step are also discussed in Section 3.

Analysis of this algorithm leads to a recurrence inequality of the form

$$
\begin{equation*}
\mathrm{M}(n)<\frac{K n}{n^{\prime}} \mathrm{M}\left(n^{\prime}\right)+O(n \log n), \quad n^{\prime}=n^{\frac{1}{d}+o(1)} \tag{1.3}
\end{equation*}
$$

where $K$ is an absolute constant, and in particular does not depend on $d$. (In Section 5 we establish (1.3) with the explicit constant $K=1728$, and in Section 5.4 we list some optimisations that improve it to $K=8$.) The first term arises from pointwise multiplications in a ring of the type $R=\mathbb{C}[y] /\left(y^{r}+1\right)$, and the second term from the fast FFTs and other auxiliary operations, including computing $v$ from $u$ and recovering $\hat{u}$ from $\hat{v}$.

We stress here the similarity with the corresponding bound (1.1) for the second Schönhage-Strassen algorithm; the difference is that we are now free to choose $d$. In Section 5 , we will simply take $d:=1729$ (any constant larger than $K$ would do), and then it is easy to see that (1.3) implies that $\mathrm{M}(n)=O(n \log n)$. (A similar
analysis holds for the conditional algorithm sketched in Section 1.2.1, for different values of $K$ and $d$.)

It is striking that for fixed $d$, the new algorithm performs only a geometric size reduction at each recursion level, just like the second Schönhage-Strassen algorithm, and unlike the first Schönhage-Strassen algorithm or any of the post-Fürer algorithms. In the new algorithm, the total cost of the FFTs actually decreases by the constant factor $d / K>1$ at each subsequent recursion level, unlike in the second Schönhage-Strassen algorithm, where it remains constant at each level, or any of the other algorithms mentioned, where it increases by a constant factor at each level.

Actually, it is possible to allow $d$ to grow with $n$, so as to achieve size reduction faster than geometric. With some care, this leads to a better constant in the main $O(n \log n)$ bound, by shifting more work from the pointwise multiplications into the fast FFTs. We will not carry out the details of this analysis.

Finally, we mention that our reduction from a DFT of size $s_{1} \times \cdots \times s_{d}$ to one of size $t_{1} \times \cdots \times t_{d}$ is highly non-algebraic, and depends heavily on the archimedean property of $\mathbb{R}$. Consequently, we do not know how to give an analogue of this algorithm for multiplication in $\mathbb{F}_{q}[x]$.

Acknowledgments. The authors would like to thank the anonymous referees, whose thoughtful comments helped to improve the presentation of these results.

## 2. DFTs, CONVOLUTIONS AND FIXED-POINT ARITHMETIC

In the Turing model we cannot compute with elements of $\mathbb{C}$ exactly. In this section we introduce a technical framework for systematic discussion of DFTs and convolutions in the setting of fixed-point arithmetic. This framework is loosely based on the presentation in $[25$, Sec. 3], and will be used throughout the rest of the paper. The impatient reader may skip most of the section and use it as a reference in case of doubt. To this effect, Table 1 contains a summary of the notation introduced in this section.
2.1. Integer arithmetic. Integers are assumed to be stored in the standard binary representation. We briefly recall several well-known results concerning integer arithmetic; see [5, Ch. 1] for further details and literature references.

Let $n \geqslant 1$, and assume that we are given as input $x, y \in \mathbb{Z}$ such that $|x|,|y| \leqslant 2^{n}$. We may compute $x+y$ and $x-y$ in time $O(n)$. For multiplication, we will often use the crude estimate $\mathrm{M}(n)=O\left(n^{1+\delta}\right)$, where for the rest of the paper $\delta$ denotes a small, fixed positive quantity; for definiteness, we assume that $\delta<\frac{1}{8}$. If $y>0$, then we may compute the quotients $\lfloor x / y\rfloor$ and $\lceil x / y\rceil$ in time $O\left(n^{1+\delta}\right)$. More generally, for a fixed positive rational number $a / b$, and assuming $x, y>0$, we may compute $\left\lfloor(x / y)^{a / b}\right\rfloor$ and $\left\lceil(x / y)^{a / b}\right\rceil$ in time $O\left(n^{1+\delta}\right)$.
2.2. Fixed-point coordinate vectors. Fix a precision parameter $p \geqslant 100$. Let $\mathbb{C}_{\circ}:=\{u \in \mathbb{C}:|u| \leqslant 1\}$ denote the complex unit disc, and set

$$
\tilde{\mathbb{C}}_{\circ}:=\left(2^{-p} \mathbb{Z}[\mathrm{i}]\right) \cap \mathbb{C}_{\circ}=\left\{2^{-p}(x+\mathrm{i} y): x, y \in \mathbb{Z} \text { and } x^{2}+y^{2} \leqslant 2^{2 p}\right\}
$$

In the Turing model, we represent an element $z=2^{-p}(x+i y) \in \tilde{\mathbb{C}}_{\circ}$ by the pair of integers $(x, y)$. It occupies $O(p)$ bits of storage, as $|x|,|y| \leqslant 2^{p}$. The precision $p$ is always known from context and does not need to be stored alongside $z$.

| $\delta \in\left(0, \frac{1}{8}\right)$ | a constant such that $\mathrm{M}(n)=O\left(n^{1+\delta}\right)$ | $(\S 2.1)$ |
| :--- | :--- | :--- |
| $p \geqslant 100$ | working precision in bits | $(\S 2.2)$ |
| $\\|\cdot\\|$ | supremum norm on $V$ (a finite-dimensional vector space | $(\S 2.2)$ |
|  | over $\mathbb{C}$ ) with respect to a specified basis $B_{V}$ | $(\S 2.2)$ |
| $V_{\circ}$ | closed unit ball in $V$ under the supremum norm | $(\S 2.2)$ |
| $\tilde{V}_{\circ}$ | set of vectors in $V_{\circ}$ whose coordinates are fixed-point |  |
|  | complex numbers with $p$ bits of accuracy | $(\S 2.2)$ |
| $\rho: V_{\circ} \rightarrow \tilde{V}_{\circ}$ | round-towards-zero function | $(\S 2.2)$ |
| $\tilde{v} \in \tilde{V}_{\circ}$ | a fixed-point approximation to a vector $v \in V_{\circ}$ | $(\S 2.2)$ |
| $\varepsilon(\tilde{v})$ | (scaled) error incurred in approximating $v$ by $\tilde{v}$ | $(\S 2.4)$ |
| $\mathcal{F}_{n}: \mathbb{C}^{n} \rightarrow \mathbb{C}^{n}$ | complex DFT of length $n$ | $(\S 2.4)$ |
| $\mathcal{F}_{n_{1}, \ldots, n_{d}}$ | multidimensional complex DFT of size $n_{1} \times \cdots \times n_{d}$ | $(\S 2.3)$ |
| $\mathscr{R}$ | the ring $\mathbb{C}[y] /\left(y^{r}+1\right)$, for a power of two $r \geqslant 2$ | $(\S 2.4)$ |
| $\mathcal{G}_{n}: \mathscr{R}^{n} \rightarrow \mathscr{R}^{n}$ | synthetic DFT of length $n$, where $n \mid 2 r$ | $(\S 2.4)$ |
| $\mathcal{G}_{n_{1}, \ldots, n_{d}}$ | multidimensional synthetic DFT of size $n_{1} \times \cdots \times n_{d}$ | $(\S 2.4)$ |
| $u \cdot v$ | pointwise product of vectors/tensors | $(\S 2.4)$ |
| $u * v$ | convolution product of vectors/tensors | $(\S 2.6)$ |
| $\mathcal{A}$ | a linear (or bilinear) map between finite-dimensional | $(\S 2.6)$ |
| $\\|\mathcal{A}\\|$ | complex vector spaces | operator norm of $\mathcal{A}$ |

Table 1. Glossary of notation introduced in Section 2

We define a round-towards-zero function $\rho: \mathbb{C} \rightarrow \mathbb{C}$ as follows. First, define $\rho_{0}: \mathbb{R} \rightarrow \mathbb{Z}$ by $\rho_{0}(x):=\lfloor x\rfloor$ for $x \geqslant 0$, and $\rho_{0}(x):=\lceil x\rceil$ for $x<0$. Then define $\rho_{0}: \mathbb{C} \rightarrow \mathbb{Z}[\mathrm{i}]$ by setting $\rho_{0}(x+\mathrm{i} y):=\rho_{0}(x)+\mathrm{i} \rho_{0}(y)$ for $x, y \in \mathbb{R}$. Observe that $\left|\rho_{0}(u)\right| \leqslant|u|$ and $\left|\rho_{0}(u)-u\right|<\sqrt{2}$ for any $u \in \mathbb{C}$. Finally, set

$$
\rho(u):=2^{-p} \rho_{0}\left(2^{p} u\right), \quad u \in \mathbb{C} .
$$

Thus $|\rho(u)| \leqslant|u|$ and $|\rho(u)-u|<\sqrt{2} \cdot 2^{-p}$ for any $u \in \mathbb{C}$. Clearly $\rho\left(\mathbb{C}_{\circ}\right) \subset \tilde{\mathbb{C}}_{\mathrm{o}}$.
Now let $V$ be a nonzero, finite-dimensional vector space over $\mathbb{C}$. In this paper, every such $V$ is understood to come equipped with a privileged choice of ordered basis $B_{V}=\left\{b_{0}, \ldots, b_{m-1}\right\}$, where $m=\operatorname{dim}_{\mathbb{C}} V \geqslant 1$. For the special case $V=\mathbb{C}^{m}$, we always take the standard basis; in particular, for $V=\mathbb{C}$ the basis is simply $\{1\}$.

We define a norm $\|\cdot\|: V \rightarrow[0, \infty)$ in terms of the basis $B_{V}$ by setting

$$
\left\|\lambda_{0} b_{0}+\cdots+\lambda_{m-1} b_{m-1}\right\|:=\max _{j}\left|\lambda_{j}\right|, \quad \lambda_{j} \in \mathbb{C}
$$

This norm satisfies $\|u+v\| \leqslant\|u\|+\|v\|$ and $\|\lambda u\|=|\lambda|\|u\|$ for any $u, v \in V, \lambda \in \mathbb{C}$. The unit ball in $V$ is defined to be

$$
V_{\circ}:=\{u \in V:\|u\| \leqslant 1\}=\left\{\lambda_{0} b_{0}+\cdots+\lambda_{m-1} b_{m-1}: \lambda_{j} \in \mathbb{C}_{\circ}\right\}
$$

and we also define

$$
\tilde{V}_{\circ}:=\left\{\lambda_{0} b_{0}+\cdots+\lambda_{m-1} b_{m-1}: \lambda_{j} \in \tilde{\mathbb{C}}_{o}\right\}
$$

We extend $\rho$ to a function $\rho: V \rightarrow V$ by acting componentwise, i.e., we put

$$
\rho\left(\lambda_{0} b_{0}+\cdots+\lambda_{m-1} b_{m-1}\right):=\sum_{j} \rho\left(\lambda_{j}\right) b_{j}, \quad \lambda_{j} \in \mathbb{C} .
$$

Then $\|\rho(u)\| \leqslant\|u\|$ and $\|\rho(u)-u\|<\sqrt{2} \cdot 2^{-p}$ for any $u \in V$. Clearly $\rho\left(V_{\circ}\right) \subset \tilde{V}_{\circ}$.
In the special case $V=\mathbb{C}$ we have simply $\|u\|=|u|$ for any $u \in \mathbb{C}$, and the notations $\mathbb{C}_{0}, \tilde{\mathbb{C}}_{\mathrm{o}}$ and $\rho: \mathbb{C} \rightarrow \mathbb{C}$ all agree with their previous definitions.

In the Turing model, an element $u \in \tilde{V}_{\circ}$ is represented by its coordinate vector with respect to $B_{V}$, i.e., as a list of $m$ elements of $\tilde{\mathbb{C}}_{\mathrm{o}}$, so $u$ occupies $O(m p)$ bits of storage.

For $u \in V_{\circ}$, we systematically use the notation $\tilde{u} \in \tilde{V}_{\circ}$ to indicate a fixed-point approximation for $u$ that has been computed by some algorithm. We write

$$
\varepsilon(\tilde{u}):=2^{p}\|\tilde{u}-u\|
$$

for the associated error, measured as a multiple of $2^{-p}$ (the "unit in the last place"). For example, we have the following result for addition and subtraction in $V$.

Lemma 2.1 (Addition/subtraction). Recall that $m=\operatorname{dim}_{\mathbb{C}} V \geqslant 1$. Given as input $u, v \in \tilde{V}_{0}$, in time $O(m p)$ we may compute an approximation $\tilde{w} \in \tilde{V}_{\circ}$ for $w:=\frac{1}{2}(u \pm v) \in V_{\circ}$ such that $\varepsilon(\tilde{w})<1$.
Proof. Consider first the case $m=1$, i.e., assume that $V=\mathbb{C}$. Let $u=2^{-p} a$ and $v=2^{-p} b$ where $a, b \in \mathbb{Z}[\mathrm{i}]$ and $|a|,|b| \leqslant 2^{p}$. Since the denominators of the real and imaginary parts of $2^{p} w=\frac{1}{2}(a \pm b)$ are at most 2 , we have $\left|\rho_{0}\left(2^{p} w\right)-2^{p} w\right| \leqslant$ $\left(\left(\frac{1}{2}\right)^{2}+\left(\frac{1}{2}\right)^{2}\right)^{1 / 2}=\frac{1}{\sqrt{2}}$. Define $\tilde{w}:=\rho(w)=2^{-p} \rho_{0}\left(2^{p} w\right)$. We may clearly compute $\tilde{w}$ in time $O(p)$, and $\varepsilon(\tilde{w})=2^{p}\|\rho(w)-w\| \leqslant \frac{1}{\sqrt{2}}<1$. The general case $(m \geqslant 1)$ follows by applying the same argument in each coordinate.

Occasionally we will encounter a situation in which we have computed an approximation $\tilde{u} \in \tilde{V}_{\circ}$ for some $u \in V$, and we wish to compute an approximation for $c u$, where $c \geqslant 1$ is a fixed integer scaling factor for which it is known that $c u \in V_{0}$. A typical example is the final scaling step in an inverse FFT. Unfortunately, the obvious approximation $c \tilde{u}$ might lie just outside $V_{0}$. To simplify subsequent estimates, it will be technically convenient to adjust $c \tilde{u}$ slightly to obtain a vector that is guaranteed to lie in $V_{0}$. This adjustment may be carried out as follows.

Lemma 2.2 (Scaling). Recall that $m=\operatorname{dim}_{\mathbb{C}} V \geqslant 1$. Let $u \in V$ and let $c$ be an integer such that $1 \leqslant c \leqslant 2^{p}$. Assume that $\|u\| \leqslant c^{-1}$, and let $v:=c u \in V_{\circ}$. Given as input $c$ and an approximation $\tilde{u} \in \tilde{V}_{0}$, in time $O\left(m p^{1+\delta}\right)$ we may compute an approximation $\tilde{v} \in \tilde{V}_{\circ}$ such that $\varepsilon(\tilde{v})<2 c \cdot \varepsilon(\tilde{u})+3$.

Proof. Again it suffices to handle the case $m=1, V=\mathbb{C}$.
We first compute $2^{-p}(x+\mathrm{i} y):=c \tilde{u}$ in time $O\left(p^{1+\delta}\right)$. Note that $c \tilde{u}$ might not lie in $\tilde{\mathbb{C}}_{\mathrm{o}}$, but $x$ and $y$ are certainly integers with $O(p)$ bits.

Next we compute $a:=x^{2}+y^{2}$ in time $O\left(p^{1+\delta}\right)$, so that $a^{1 / 2}=2^{p}|c \tilde{u}|$. If $a \leqslant 2^{2 p}$ then already $c \tilde{u} \in \tilde{\mathbb{C}}_{0}$, so we may simply take $\tilde{v}:=c \tilde{u}$, and then $\varepsilon(\tilde{v})=2^{p}|\tilde{v}-v|=$ $2^{p}|c \tilde{u}-c u|=c \cdot \varepsilon(\tilde{u})<2 c \cdot \varepsilon(\tilde{u})+3$.

Suppose instead that $a>2^{2 p}$ (i.e., $c \tilde{u} \notin \tilde{\mathbb{C}}_{\circ}$ ). We then compute $b:=\left\lceil a^{1 / 2}\right\rceil>2^{p}$, again in time $O\left(p^{1+\delta}\right)$. Let $z:=2^{p} c \tilde{u} / b=(x+\mathrm{i} y) / b$ and $\tilde{v}:=\rho(z)$. Note that $\tilde{v}=2^{-p}\left(x^{\prime}+\mathrm{i} y^{\prime}\right)$ where $x^{\prime}=\rho_{0}\left(2^{p} x / b\right)$ and $y^{\prime}=\rho_{0}\left(2^{p} y / b\right)$, so we may compute $\tilde{v}$ in time $O\left(p^{1+\delta}\right)$. We have $|\tilde{v}| \leqslant|z|=2^{p}|c \tilde{u}| / b \leqslant 2^{p}|c \tilde{u}| / a^{1 / 2}=1$, so indeed $\tilde{v} \in \tilde{\mathbb{C}}_{0}$. Moreover,

$$
|\tilde{v}-v| \leqslant|\tilde{v}-z|+|z-c \tilde{u}|+|c \tilde{u}-v|=|\rho(z)-z|+|z|\left|1-\frac{b}{2^{p}}\right|+c|\tilde{u}-u|,
$$

so $\varepsilon(\tilde{v})<\sqrt{2}+\left|2^{p}-b\right|+c \cdot \varepsilon(\tilde{u})$. We also have $2^{p}<b<a^{1 / 2}+1=2^{p}|c \tilde{u}|+1$, so

$$
0<b-2^{p}<2^{p}|c \tilde{u}|-2^{p}+1 \leqslant 2^{p}|c u|-2^{p}+2^{p}|c \tilde{u}-c u|+1 \leqslant c \cdot \varepsilon(\tilde{u})+1
$$

We conclude that $\left|2^{p}-b\right|<c \cdot \varepsilon(\tilde{u})+1$, and therefore $\varepsilon(\tilde{v})<2 c \cdot \varepsilon(\tilde{u})+(1+\sqrt{2})$.
2.3. Coefficient rings. By a coefficient ring we mean a finite-dimensional commutative $\mathbb{C}$-algebra $R$ with identity (together with a privileged basis $B_{R}$ ). We are chiefly interested in the following two examples:
(1) Complex case: $\mathbb{C}$ itself, with the basis $\{1\}$.
(2) Synthetic case: for a power of two $r \geqslant 2$, the $\operatorname{ring} \mathscr{R}:=\mathbb{C}[y] /\left(y^{r}+1\right)$, with the basis $\left\{1, y, \ldots, y^{r-1}\right\}$, so that $\left\|\lambda_{0}+\lambda_{1} y+\cdots+\lambda_{r-1} y^{r-1}\right\|=\max _{j}\left|\lambda_{j}\right|$.
Let $R$ be a coefficient ring of dimension $r \geqslant 1$ with basis $B_{R}$, and let $n \geqslant 1$. Then $R^{n}$ is a vector space of dimension $n r$ over $\mathbb{C}$. We associate to $R^{n}$ the "nested" basis formed by concatenating $n$ copies of $B_{R}$. In particular, we have $\|u\|=\max _{j}\left\|u_{j}\right\|$ for $u=\left(u_{0}, \ldots, u_{n-1}\right) \in R^{n}$. In place of the awkward expressions $\left(R^{n}\right)_{\circ}$ and $\left(\tilde{R}^{n}\right)_{\circ}$, we write more compactly $R_{\circ}^{n}$ and $\tilde{R}_{\circ}^{n}$. In the Turing model, an element of $\tilde{R}_{\circ}^{n}$ occupies $O(n r p)$ bits of storage.

Now let $d \geqslant 1$ and $n_{1}, \ldots, n_{d} \geqslant 1$. We write $\otimes_{i=1}^{d} R^{n_{i}}$, or just $\otimes_{i} R^{n_{i}}$ when $d$ is understood, for the tensor product $R^{n_{1}} \otimes_{R} \cdots \otimes_{R} R^{n_{d}}$. It is a free $R$-module of rank $n_{1} \cdots n_{d}$, and also a vector space over $\mathbb{C}$ of dimension $n_{1} \cdots n_{d} r$. An element $u \in \otimes_{i} R^{n_{i}}$ may be regarded as a $d$-dimensional array of elements of $R$ of size $n_{1} \times \cdots \times n_{d}$. For indices $j_{1}, \ldots, j_{d}$ where $0 \leqslant j_{i}<n_{i}$, we write $u_{j_{1}, \ldots, j_{d}} \in R$ for the $\left(j_{1}, \ldots, j_{d}\right)$-th component of $u$.

We associate to $\otimes_{i} R^{n_{i}}$ the nested basis consisting of $n_{1} \cdots n_{d}$ copies of $B_{R}$ arranged in lexicographical order, i.e., listing the coordinates of $u$ in the order $\left(u_{0, \ldots, 0}, u_{0, \ldots, 1}, \ldots, u_{n_{1}-1, \ldots, n_{d}-1}\right)$. Observe then that $\|u\|=\max _{j_{1}, \ldots, j_{d}}\left\|u_{j_{1}, \ldots, j_{d}}\right\|$. Instead of $\left(\otimes_{i} R^{n_{i}}\right)_{\circ}$ and $\left(\otimes_{i} \tilde{R}^{n_{i}}\right)_{0}$, we write $\otimes_{i} R_{\circ}^{n_{i}}$ and $\otimes_{i} \tilde{R}_{\circ}^{n_{i}}$. In the Turing model, an element of $\otimes_{i} \tilde{R}_{o}^{n_{i}}$ occupies $O\left(n_{1} \cdots n_{d} r p\right)$ bits of storage.

Let $u \in \otimes_{i} R^{n_{i}}$. By an $i$-slice of $u$ we mean a one-dimensional sub-array of $u$, consisting of the entries $u_{j_{1}, \ldots, j_{d}}$ where $j_{1}, \ldots, j_{i-1}, j_{i+1}, \ldots, j_{d}$ are held fixed and $j_{i}$ varies over $\left\{0, \ldots, n_{i}-1\right\}$. We will occasionally wish to apply a given algorithm separately to each of the $n_{1} \cdots n_{i-1} n_{i+1} \cdots n_{d}$ distinct $i$-slices of some $u \in \otimes_{i} \tilde{R}_{\circ}^{n_{i}}$. To accomplish this in the Turing model, we must first rearrange the data so that each $i$-slice is stored contiguously. In the lexicographical order specified above, this amounts to performing $n_{1} \cdots n_{i-1}$ matrix transpositions of size $n_{i} \times\left(n_{i+1} \cdots n_{d}\right)$. This data rearrangement may be performed in time $O\left(n_{1} \cdots n_{d} r p \log n_{i}\right)$ (assuming $n_{i} \geqslant 2$ ) using a fast matrix transposition algorithm [4, Appendix].
2.4. DFTs and convolutions. Let $R$ be a coefficient ring and let $n \geqslant 1$. Throughout the paper we adopt the convention that for a vector $u=\left(u_{0}, \ldots, u_{n-1}\right) \in R^{n}$
and an integer $j$, the expression $u_{j}$ always means $u_{j \bmod n}$. For $u, v \in R^{n}$, we define the pointwise product $u \cdot v \in R^{n}$ and the convolution product $u * v \in R^{n}$ by

$$
(u \cdot v)_{j}:=u_{j} v_{j}, \quad(u * v)_{j}:=\sum_{k=0}^{n-1} u_{k} v_{j-k}, \quad 0 \leqslant j<n
$$

Then $\left(R^{n}, \cdot\right)$ and $\left(R^{n}, *\right)$ are both commutative rings, isomorphic respectively to the direct sum of $n$ copies of $R$, and the polynomial ring $R[x] /\left(x^{n}-1\right)$.

A principal $n$-th root of unity in $R$ is an element $\omega \in R$ such that $\omega^{n}=1$, $\sum_{k=0}^{n-1}\left(\omega^{j}\right)^{k}=0$ for every integer $j \neq 0(\bmod n)$, and $\|\omega u\|=\|u\|$ for all $u \in R$. (The last condition is not part of the standard definition, but it is natural in our setting where $R$ carries a norm, and essential for error estimates.) We define an associated $R$-linear DFT map $F_{\omega}: R^{n} \rightarrow R^{n}$ by the formula

$$
\left(F_{\omega} u\right)_{j}:=\frac{1}{n} \sum_{k=0}^{n-1} \omega^{-j k} u_{k}, \quad u \in R^{n}, \quad 0 \leqslant j<n
$$

It is immediate that $\omega^{-1}\left(=\omega^{n-1}\right)$ is also a principal $n$-th root of unity in $R$, and that $\left\|F_{\omega} u\right\| \leqslant\|u\|$ for all $u \in R^{n}$.
Lemma 2.3 (Convolution formula). For any $u, v \in R^{n}$ we have

$$
\frac{1}{n} u * v=n F_{\omega^{-1}}\left(F_{\omega} u \cdot F_{\omega} v\right)
$$

Proof. For each $j$, the product $\left(F_{\omega} u\right)_{j}\left(F_{\omega} v\right)_{j}$ is equal to

$$
\frac{1}{n^{2}} \sum_{s=0}^{n-1} \sum_{t=0}^{n-1} \omega^{-j(s+t)} u_{s} v_{t}=\frac{1}{n^{2}} \sum_{k=0}^{n-1} \omega^{-j k} \sum_{s+t=k(\bmod n)} u_{s} v_{t}=\frac{1}{n} F_{\omega}(u * v)_{j}
$$

so $F_{\omega}(u * v)=n F_{\omega} u \cdot F_{\omega} v$. On the other hand, for any $w \in R^{n}$ we have

$$
\left(F_{\omega^{-1}}\left(F_{\omega} w\right)\right)_{j}=\frac{1}{n^{2}} \sum_{s=0}^{n-1} \omega^{s j} \sum_{t=0}^{n-1} \omega^{-s t} w_{t}=\frac{1}{n^{2}} \sum_{t=0}^{n-1}\left(\sum_{s=0}^{n-1} \omega^{s(j-t)}\right) w_{t}=\frac{1}{n} w_{j}
$$

so $F_{\omega^{-1}} F_{\omega} w=\frac{1}{n} w$. Taking $w:=u * v$, we obtain the desired result.
For the two coefficient rings mentioned earlier, we choose $\omega$ as follows:
(1) Complex case. For $R=\mathbb{C}$, let $n \geqslant 1$ be any positive integer, and put $\omega:=e^{2 \pi \mathrm{i} / n}$. We denote $F_{\omega}$ in this case by $\mathcal{F}_{n}: \mathbb{C}^{n} \rightarrow \mathbb{C}^{n}$. Explicitly,

$$
\left(\mathcal{F}_{n} u\right)_{j}=\frac{1}{n} \sum_{k=0}^{n-1} e^{-2 \pi \mathrm{i} j k / n} u_{k}, \quad u \in \mathbb{C}^{n}, \quad 0 \leqslant j<n
$$

We also write $\mathcal{F}_{n}^{*}: \mathbb{C}^{n} \rightarrow \mathbb{C}^{n}$ for $F_{\omega^{-1}}$.
(2) Synthetic case. For $R=\mathscr{R}=\mathbb{C}[y] /\left(y^{r}+1\right)$ where $r \geqslant 2$ is a power of two, let $n$ be any positive divisor of $2 r$. Then $\omega:=y^{2 r / n}$ is a principal $n$-th root of unity in $\mathscr{R}$. We denote $F_{\omega}$ in this case by $\mathcal{G}_{n}: \mathscr{R}^{n} \rightarrow \mathscr{R}^{n}$. Explicitly,

$$
\left(\mathcal{G}_{n} u\right)_{j}=\frac{1}{n} \sum_{k=0}^{n-1} y^{-2 r j k / n} u_{k}, \quad u \in \mathscr{R}^{n}, \quad 0 \leqslant j<n
$$

We also write $\mathcal{G}_{n}^{*}: \mathscr{R}^{n} \rightarrow \mathscr{R}^{n}$ for $F_{\omega^{-1}}$.

All of the concepts introduced above may be generalised to the multidimensional setting as follows. For $u, v \in \otimes_{i} R^{n_{i}}$, we define the pointwise product $u \cdot v \in \otimes_{i} R^{n_{i}}$ and the convolution product $u * v \in \otimes_{i} R^{n_{i}}$ by

$$
\begin{aligned}
(u \cdot v)_{j_{1}, \ldots, j_{d}} & :=u_{j_{1}, \ldots, j_{d}} v_{j_{1}, \ldots, j_{d}}, \\
(u * v)_{j_{1}, \ldots, j_{d}} & :=\sum_{k_{1}=0}^{n_{1}-1} \cdots \sum_{k_{d}=0}^{n_{d}-1} u_{k_{1}, \ldots, k_{d}} v_{j_{1}-k_{1}, \ldots, j_{d}-k_{d}} .
\end{aligned}
$$

Then $\left(\otimes_{i} R^{n_{i}}, \cdot\right)$ is isomorphic to the direct sum of $n_{1} \cdots n_{d}$ copies of $R$, and $\left(\otimes_{i} R^{n_{i}}, *\right)$ is isomorphic to $R\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{n_{1}}-1, \ldots, x_{d}^{n_{d}}-1\right)$.

Let $\omega_{1}, \ldots, \omega_{d} \in R$ be principal roots of unity of orders $n_{1}, \ldots, n_{d}$. We define an associated $R$-linear $d$-dimensional DFT map by taking the tensor product (over $R$ ) of the corresponding one-dimensional DFTs, that is,

$$
F_{\omega_{1}, \ldots, \omega_{d}}:=\otimes_{i} F_{\omega_{i}}: \otimes_{i} R^{n_{i}} \rightarrow \otimes_{i} R^{n_{i}}
$$

Explicitly, for $u \in \otimes_{i} R^{n_{i}}$ we have

$$
\left(F_{\omega_{1}, \ldots, \omega_{d}} u\right)_{j_{1}, \ldots, j_{d}}=\frac{1}{n_{1} \cdots n_{d}} \sum_{k_{1}=0}^{n_{1}-1} \cdots \sum_{k_{d}=0}^{n_{d}-1} \omega_{1}^{-j_{1} k_{1}} \cdots \omega_{d}^{-j_{d} k_{d}} u_{k_{1}, \ldots, k_{d}}
$$

The multidimensional analogue of Lemma 2.3 is

$$
\begin{equation*}
\frac{1}{n_{1} \cdots n_{d}} u * v=n_{1} \cdots n_{d} F_{\omega_{1}^{-1}, \ldots, \omega_{d}^{-1}}\left(F_{\omega_{1}, \ldots, \omega_{d}} u \cdot F_{\omega_{1}, \ldots, \omega_{d}} v\right) \tag{2.1}
\end{equation*}
$$

and is proved in exactly the same way.
In particular, in the "complex case" we obtain the $d$-dimensional transform

$$
\mathcal{F}_{n_{1}, \ldots, n_{d}}:=\otimes_{i} \mathcal{F}_{n_{i}}: \otimes_{i} \mathbb{C}^{n_{i}} \rightarrow \otimes_{i} \mathbb{C}^{n_{i}}
$$

(take $\omega_{i}:=e^{2 \pi \mathrm{i} / n_{i}}$ ), and in the "synthetic case" the $d$-dimensional transform

$$
\mathcal{G}_{n_{1}, \ldots, n_{d}}:=\otimes_{i} \mathcal{G}_{n_{i}}: \otimes_{i} \mathscr{R}^{n_{i}} \rightarrow \otimes_{i} \mathscr{R}^{n_{i}}
$$

(where each $n_{i}$ is a divisor of $2 r$, and $\omega_{i}:=y^{2 r / n_{i}}$ ). We define similarly $\mathcal{F}_{n_{1}, \ldots, n_{d}}^{*}:=$ $\otimes_{i} \mathcal{F}_{n_{i}}^{*}$ and $\mathcal{G}_{n_{1}, \ldots, n_{d}}^{*}:=\otimes_{i} \mathcal{G}_{n_{i}}^{*}$.

Any algorithm for computing $\mathcal{F}_{n}$ may easily be adapted to obtain an algorithm for computing $\mathcal{F}_{n}^{*}$, by adjusting signs appropriately. A similar remark applies to $\mathcal{G}_{n}$, and to the multidimensional generalisations of these maps. For the rest of the paper, we make use of these observations without further comment.
2.5. Fixed-point multiplication. We now consider the complexity of multiplication in the coefficient rings $R=\mathbb{C}$ and $R=\mathscr{R}$. In both cases we reduce the problem to integer multiplication. For the case $R=\mathscr{R}$ (Lemma 2.5) we will express the complexity in terms of $\mathrm{M}(\cdot)$ itself, as this eventually feeds into the main recurrence inequality for $\mathrm{M}(\cdot)$ that we prove in Section 5.3 . For the case $R=\mathbb{C}$ (Lemma 2.4) we do not need the best possible bound; to simplify the subsequent complexity analysis, we prefer to use the crude estimate $\mathrm{M}(p)=O\left(p^{1+\delta}\right)$.

Lemma 2.4 (Multiplication in $\mathbb{C}$ ). Given as input $u, v \in \tilde{\mathbb{C}}_{0}$, in time $O\left(p^{1+\delta}\right)$ we may compute an approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ for $w:=u v \in \mathbb{C}_{\circ}$ such that $\varepsilon(\tilde{w})<2$.

Proof. We take $\tilde{w}:=\rho(w)$, so that $\varepsilon(\tilde{w})=2^{p}\|\rho(w)-w\|<\sqrt{2}<2$. Writing $u=$ $2^{-p} a$ and $v=2^{-p} b$ where $a, b \in \mathbb{Z}[\mathrm{i}]$ and $|a|,|b| \leqslant 2^{p}$, we have $\tilde{w}=2^{-p} \rho_{0}\left(2^{-p} a b\right)$. Thus $\tilde{w}$ may be computed in time $O\left(p^{1+\delta}\right)$ by multiplying out the real and imaginary parts of $a$ and $b$, and then summing and rounding appropriately.

For the case $R=\mathscr{R}$, observe first that for any $u, v \in \mathscr{R}$ we have $\|u v\| \leqslant r\|u\|\|v\|$, as each coefficient of $u v=\left(u_{0}+\cdots+u_{r-1} y^{r-1}\right)\left(v_{0}+\cdots+v_{r-1} y^{r-1}\right) \bmod y^{r}+1$ is a sum of exactly $r$ terms of the form $\pm u_{i} v_{j}$. In particular, if $u, v \in \mathscr{R}_{\circ}$, then $u v / r \in \mathscr{R}_{\circ}$.

Lemma 2.5 (Multiplication in $\mathscr{R}$ ). Assume that $r \geqslant 2$ is a power of two and that $r<2^{p-1}$. Given as input $u, v \in \tilde{\mathscr{R}}_{\circ}$, in time $4 \mathrm{M}(3 r p)+O(r p)$ we may compute an approximation $\tilde{w} \in \tilde{\mathscr{R}}_{\circ}$ for $w:=u v / r \in \mathscr{R}_{\circ}$ such that $\varepsilon(\tilde{w})<2$.

Proof. Write $2^{p} u=U_{0}(y)+\mathrm{i} U_{1}(y)$ and $2^{p} v=V_{0}(y)+\mathrm{i} V_{1}(y)$ where $U_{j}$ and $V_{j}$ are polynomials in $\mathbb{Z}[y]$ of degree less than $r$ and whose coefficients lie in the interval $\left[-2^{p}, 2^{p}\right]$. Then $2^{2 p} r w=W_{0}(y)+\mathrm{i} W_{1}(y)$ where

$$
W_{0}:=\left(U_{0} V_{0}-U_{1} V_{1}\right) \bmod y^{r}+1, \quad W_{1}:=\left(U_{0} V_{1}+U_{1} V_{0}\right) \bmod y^{r}+1
$$

We use the following algorithm, which is based on the well-known Kronecker substitution technique [14, Corollary 8.27].
(1) Pack coefficients. Evaluate $U_{j}\left(2^{3 p}\right), V_{j}\left(2^{3 p}\right) \in \mathbb{Z}$ for $j=0,1$. As the input coefficients have at most $p$ bits, this amounts to concatenating the coefficients with appropriate zero-padding (or one-padding in the case of negative coefficients), plus some carry and sign handling. The cost of this step is $O(r p)$.
(2) Multiply in $\mathbb{Z}$. Let $W_{j, k}:=U_{j} V_{k} \in \mathbb{Z}[y]$ for $j, k \in\{0,1\}$. Compute the four integer products $W_{j, k}\left(2^{3 p}\right)=U_{j}\left(2^{3 p}\right) V_{k}\left(2^{3 p}\right)$. The cost of this step is $4 \mathrm{M}(3 r p)$.
(3) Unpack coefficients. For each pair $(j, k)$, the coefficients of $W_{j, k} \in \mathbb{Z}[y]$ are bounded in absolute value by $r\left(2^{p}\right)^{2}<2^{3 p-1}$, so $W_{j, k}$ may be recovered from the integer $W_{j, k}\left(2^{3 p}\right)$ in time $O(r p)$. (In more detail: the constant term of $W_{j, k}$ lies in the interval $\left(-2^{3 p-1}, 2^{3 p-1}\right)$, so it is easily read off the last $3 p$ bits of $W_{j, k}\left(2^{3 p}\right)$. After stripping off this term, one proceeds to the linear term, and so on.) We then deduce the polynomials $W_{0}=\left(W_{0,0}-W_{1,1}\right) \bmod y^{r}+1$ and $W_{1}=\left(W_{0,1}+W_{1,0}\right) \bmod y^{r}+1$ in time $O(r p)$.
(4) Scale and round. Let $c_{\ell}:=\left(W_{0}\right)_{\ell}+\mathrm{i}\left(W_{1}\right)_{\ell} \in \mathbb{Z}[\mathrm{i}]$ for $\ell \in\{0, \ldots, r-1\}$. Then $w=\left(2^{2 p} r\right)^{-1}\left(c_{0}+\cdots+c_{r-1} y^{r-1}\right)$, so, recalling that $\|w\|=\|u v\| / r \leqslant 1$, we have $\left|c_{\ell}\right| \leqslant 2^{2 p} r$ for each $\ell$. In time $O(r p)$ we may compute $\tilde{w}:=\rho(w)=$ $2^{-p} \sum_{\ell=0}^{r-1} \rho_{0}\left(\left(2^{p} r\right)^{-1} c_{\ell}\right){\underset{\sim}{\mathscr{M}}}^{\ell}$ (each division by $2^{p} r$ amounts to a bit shift). Since $\|w\| \leqslant 1$, we have $\tilde{w} \in \tilde{\mathscr{R}}_{\circ}$, and as usual, $\varepsilon(\tilde{w})=2^{p}\|\rho(w)-w\|<\sqrt{2}<2$.
2.6. Linear and bilinear maps. Let $\mathcal{A}: V \rightarrow W$ be a $\mathbb{C}$-linear map between finite-dimensional vector spaces $V$ and $W$. We define the operator norm of $\mathcal{A}$ to be

$$
\|\mathcal{A}\|:=\sup _{v \in V_{0}}\|\mathcal{A} v\|
$$

Example 2.6. For the normalised DFT map $F_{\omega}$ defined in Section 2.4, we have $\left\|F_{\omega}\right\| \leqslant 1$. The same therefore holds for $\mathcal{F}_{n}, \mathcal{F}_{n}^{*}, \mathcal{G}_{n}, \mathcal{G}_{n}^{*}$, and for the multivariate generalisations of these maps. (In fact, all of these maps have norm exactly 1.)

Assume now that $\|\mathcal{A}\| \leqslant 1$. By a numerical approximation for $\mathcal{A}$ we mean a function $\tilde{\mathcal{A}}: \tilde{V}_{\circ} \rightarrow \tilde{W}_{\circ}$ that is computed by some algorithm, typically via fixedpoint arithmetic. The error of the approximation is defined to be

$$
\varepsilon(\tilde{\mathcal{A}}):=\max _{v \in \tilde{V}_{0}} 2^{p}\|\tilde{\mathcal{A}} v-\mathcal{A} v\|
$$

We write $\mathrm{C}(\tilde{\mathcal{A}})$ for the time required to compute $\tilde{\mathcal{A}} v$ from $v$ (taking the maximum over all possible inputs $v \in \tilde{V}_{\circ}$ ).
Lemma 2.7 (Error propagation). Let $\mathcal{A}: V \rightarrow W$ be a $\mathbb{C}$-linear map such that $\|\mathcal{A}\| \leqslant 1$, and let $v \in V_{\circ}$. Let $\tilde{\mathcal{A}}: \tilde{V}_{\circ} \rightarrow \tilde{W}_{\circ}$ be a numerical approximation for $\mathcal{A}$, and let $\tilde{v} \in \tilde{V}_{\circ}$ be an approximation for $v$. Then $\tilde{w}:=\tilde{\mathcal{A}} \tilde{v} \in \tilde{W}_{\circ}$ is an approximation for $w:=\mathcal{A} v \in W_{\circ}$ such that $\varepsilon(\tilde{w}) \leqslant \varepsilon(\tilde{\mathcal{A}})+\varepsilon(\tilde{v})$.

Proof. We have

$$
\begin{aligned}
\varepsilon(\tilde{w})=2^{p}\|\tilde{\mathcal{A}} \tilde{v}-\mathcal{A} v\| & \leqslant 2^{p}\|\tilde{\mathcal{A}} \tilde{v}-\mathcal{A} \tilde{v}\|+2^{p}\|\mathcal{A} \tilde{v}-\mathcal{A} v\| \\
& \leqslant \varepsilon(\tilde{\mathcal{A}})+2^{p}\|\mathcal{A}\|\|\tilde{v}-v\| \leqslant \varepsilon(\tilde{\mathcal{A}})+\varepsilon(\tilde{v})
\end{aligned}
$$

Lemma 2.7 yields the following estimate for compositions of linear maps.
Corollary 2.8 (Composition). Let $\mathcal{A}: U \rightarrow V$ and $\mathcal{B}: V \rightarrow W$ be $\mathbb{C}$-linear maps such that $\|\mathcal{A}\|,\|\mathcal{B}\| \leqslant 1$. Let $\tilde{\mathcal{A}}: \tilde{U}_{\circ} \rightarrow \tilde{V}_{\circ}$ and $\tilde{\mathcal{B}}: \tilde{V}_{\circ} \rightarrow \tilde{W}_{\circ}$ be numerical approximations. Then $\tilde{\mathcal{C}}:=\tilde{\mathcal{B}} \tilde{\mathcal{A}}: \tilde{U}_{\circ} \rightarrow \tilde{W}_{\circ}$ is a numerical approximation for $\mathcal{C}:=$ $\mathcal{B A}: U \rightarrow W$ such that $\varepsilon(\tilde{\mathcal{C}}) \leqslant \varepsilon(\tilde{\mathcal{B}})+\varepsilon(\tilde{\mathcal{A}})$.
Proof. For any $u \in \tilde{U}_{\circ}$, if we set $v:=\mathcal{A} u \in V_{\circ}$ and $\tilde{v}:=\tilde{\mathcal{A}} u \in \tilde{V}_{\circ}$, then

$$
2^{p}\|\tilde{\mathcal{B}} \tilde{\mathcal{A}} u-\mathcal{B} \mathcal{A} u\|=2^{p}\|\tilde{\mathcal{B}} \tilde{v}-\mathcal{B} v\| \leqslant \varepsilon(\tilde{\mathcal{B}})+\varepsilon(\tilde{v}) \leqslant \varepsilon(\tilde{\mathcal{B}})+\varepsilon(\tilde{\mathcal{A}})
$$

The above definitions and results may be adapted to the case of a $\mathbb{C}$-bilinear $\operatorname{map} \mathcal{A}: U \times V \rightarrow W$ as follows. We define

$$
\|\mathcal{A}\|:=\sup _{u \in U_{0}, v \in V_{0}}\|\mathcal{A}(u, v)\|
$$

If $\|\mathcal{A}\| \leqslant 1$, then a numerical approximation for $\mathcal{A}$ is a function $\tilde{\mathcal{A}}: \tilde{U}_{\circ} \times \tilde{V}_{\circ} \rightarrow \tilde{W}_{\circ}$ that is computed by some algorithm. The error of the approximation is

$$
\varepsilon(\tilde{\mathcal{A}}):=\max _{u \in \tilde{U}_{o}, v \in \tilde{V}_{0}} 2^{p}\|\tilde{\mathcal{A}}(u, v)-\mathcal{A}(u, v)\|
$$

and $\mathrm{C}(\tilde{\mathcal{A}})$ denotes the time required to compute $\tilde{\mathcal{A}}(u, v)$ from $u$ and $v$. Lemma 2.7 has the following analogue in the bilinear case.
Lemma 2.9 (Bilinear error propagation). Let $\mathcal{A}: U \times V \rightarrow W$ be a $\mathbb{C}$-bilinear map with $\|\mathcal{A}\| \leqslant 1$, and let $u \in U_{0}, v \in V_{0}$. Let $\tilde{\mathcal{A}}: \tilde{U}_{0} \times \tilde{V}_{\circ} \rightarrow \tilde{W}_{\circ}, \tilde{u} \in \tilde{U}_{0}$, $\tilde{v} \in \tilde{V}_{\circ}$ be approximations. Then $\tilde{w}:=\tilde{\mathcal{A}}(\tilde{u}, \tilde{v}) \in \tilde{W}_{\circ}$ is an approximation for $w:=\mathcal{A}(u, v) \in W_{\circ}$ such that $\varepsilon(\tilde{w}) \leqslant \varepsilon(\tilde{\mathcal{A}})+\varepsilon(\tilde{u})+\varepsilon(\tilde{v})$.
Proof. We have

$$
\begin{aligned}
\varepsilon(\tilde{w}) & \leqslant 2^{p}(\|\tilde{\mathcal{A}}(\tilde{u}, \tilde{v})-\mathcal{A}(\tilde{u}, \tilde{v})\|+\|\mathcal{A}(\tilde{u}, \tilde{v})-\mathcal{A}(u, \tilde{v})\|+\|\mathcal{A}(u, \tilde{v})-\mathcal{A}(u, v)\|) \\
& =2^{p}(\|\tilde{\mathcal{A}}(\tilde{u}, \tilde{v})-\mathcal{A}(\tilde{u}, \tilde{v})\|+\|\mathcal{A}(\tilde{u}-u, \tilde{v})\|+\|\mathcal{A}(u, \tilde{v}-v)\|) \\
& \leqslant \varepsilon(\tilde{\mathcal{A}})+2^{p}\|\mathcal{A}\|\|\tilde{u}-u\|\|\tilde{v}\|+2^{p}\|\mathcal{A}\|\|u\|\|\tilde{v}-v\| \\
& \leqslant \varepsilon(\tilde{\mathcal{A}})+\varepsilon(\tilde{u})+\varepsilon(\tilde{v})
\end{aligned}
$$

The following application of Lemma 2.9 will frequently be useful.
Corollary 2.10. Let $u, v \in \mathbb{C}_{0}$, and let $w:=u v \in \mathbb{C}_{0}$. Given as input approximations $\tilde{u}, \tilde{v} \in \tilde{\mathbb{C}}_{0}$, in time $O\left(p^{1+\delta}\right)$ we may compute an approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon(\tilde{w})<\varepsilon(\tilde{u})+\varepsilon(\tilde{v})+2$.

Proof. Define a bilinear map $\mathcal{A}: \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ by $\mathcal{A}(u, v):=u v$. Then $\|\mathcal{A}\| \leqslant 1$, and Lemma 2.4 yields an approximation $\tilde{\mathcal{A}}: \tilde{\mathbb{C}}_{\mathrm{o}} \times \tilde{\mathbb{C}}_{\mathrm{o}} \rightarrow \tilde{\mathbb{C}}_{\mathrm{o}}$ such that $\varepsilon(\tilde{\mathcal{A}})<2$ and $\mathrm{C}(\tilde{\mathcal{A}})=O\left(p^{1+\delta}\right)$. Applying Lemma 2.9 to $\mathcal{A}$ and $\tilde{\mathcal{A}}$ yields the desired result.
2.7. Tensor products. The following result shows how to construct numerical approximations for tensor products of linear maps over a coefficient ring.

Lemma 2.11 (Tensor products). Let $R$ be a coefficient ring of dimension $r \geqslant 1$. Let $d \geqslant 1$, let $m_{1}, \ldots, m_{d}, n_{1}, \ldots, n_{d} \geqslant 2$, and put $M:=\prod_{i} \max \left(m_{i}, n_{i}\right)$. For $i_{\tilde{\mathcal{A}}}^{i} \in\{1, \ldots, d\}$, let $\mathcal{A}_{i}: R^{m_{i}} \rightarrow R^{n_{i}}$ be an $R$-linear map with $\left\|\mathcal{A}_{i}\right\| \leqslant 1$, and let $\tilde{\mathcal{A}}_{i}: \tilde{R}_{\circ}^{m_{i}} \rightarrow \widetilde{R}_{\circ}^{n_{i}}$ be a numerical approximation. Let $\mathcal{A}:=\otimes_{i} \mathcal{A}_{i}: \otimes_{i} R^{m_{i}} \rightarrow \otimes_{i} R^{n_{i}}$ (note that automatically $\|\mathcal{A}\| \leqslant 1$ ).

Then we may construct a numerical approximation $\tilde{\mathcal{A}}: \otimes_{i} \tilde{R}_{\circ}^{m_{i}} \rightarrow \otimes_{i} \tilde{R}_{\circ}^{n_{i}}$ such that $\varepsilon(\tilde{\mathcal{A}}) \leqslant \sum_{i} \varepsilon\left(\tilde{\mathcal{A}}_{i}\right)$ and

$$
\mathrm{C}(\tilde{\mathcal{A}}) \leqslant M \sum_{i} \frac{\mathrm{C}\left(\tilde{\mathcal{A}}_{i}\right)}{\max \left(m_{i}, n_{i}\right)}+O(M r p \log M)
$$

Proof. For $i \in\{0,1, \ldots, d\}$, let

$$
U^{i}:=R^{n_{1}} \otimes \cdots \otimes R^{n_{i-1}} \otimes R^{n_{i}} \otimes R^{m_{i+1}} \otimes \cdots \otimes R^{m_{d}}
$$

In particular, $U^{0}=\otimes_{i} R^{m_{i}}$ and $U^{d}=\otimes_{i} R^{n_{i}}$. The map $\mathcal{A}: U^{0} \rightarrow U^{d}$ admits a decomposition $\mathcal{A}=\mathcal{B}_{d} \cdots \mathcal{B}_{1}$ where $\mathcal{B}_{i}: U^{i-1} \rightarrow U^{i}$ is given by

$$
\mathcal{B}_{i}:=\mathcal{I}_{n_{1}} \otimes \cdots \otimes \mathcal{I}_{n_{i-1}} \otimes \mathcal{A}_{i} \otimes \mathcal{I}_{m_{i+1}} \otimes \cdots \otimes \mathcal{I}_{m_{d}}
$$

(here $\mathcal{I}_{k}$ denotes the identity map on $R^{k}$ ). In other words, $\mathcal{B}_{i}$ acts by applying $\mathcal{A}_{i}$ separately on each $i$-slice. Explicitly, for any $u \in U^{i-1}$ we have $\left(\mathcal{B}_{i} u\right)_{j_{1}, \ldots, j_{d}}=$ $\left(\mathcal{A}_{i} v\right)_{j_{i}}$ where $v \in R^{m_{i}}$ is the vector defined by $v_{k}:=u_{j_{1}, \ldots, j_{i-1}, k, j_{i+1}, \ldots, j_{d}}$. In particular, $\left\|v_{k}\right\| \leqslant 1$ whenever $\|u\| \leqslant 1$, whence $\left\|\mathcal{B}_{i}\right\| \leqslant 1$.

We may define an approximation $\tilde{\mathcal{B}}_{i}: \tilde{U}_{\circ}^{i-1} \rightarrow \tilde{U}_{\circ}^{i}$ by mimicking the above formula for $\mathcal{B}_{i}$; i.e., for $u \in \tilde{U}_{\circ}^{i-1}$ we define $\left(\tilde{\mathcal{B}}_{i} u\right)_{j_{1}, \ldots, j_{d}}:=\left(\tilde{\mathcal{A}}_{i} v\right)_{j_{i}}$, where $v \in \tilde{R}_{\circ}^{m_{i}}$ is given by $v_{k}:=u_{j_{1}, \ldots, j_{i-1}, k, j_{i+1}, \ldots, j_{d}}$. We may evaluate $\tilde{\mathcal{B}}_{i}$ by first rearranging the data so that each $i$-slice is stored contiguously (see Section 2.3), then applying $\tilde{\mathcal{A}}_{i}$ to each $i$-slice, and finally rearranging the data back into the correct order. We then define $\tilde{\mathcal{A}}:=\tilde{\mathcal{B}}_{d} \cdots \tilde{\mathcal{B}}_{1}$.

We clearly have $\varepsilon\left(\tilde{\mathcal{B}}_{i}\right)=\varepsilon\left(\tilde{\mathcal{A}}_{i}\right)$ for all $i$, so by Corollary 2.8 we obtain $\varepsilon(\tilde{\mathcal{A}}) \leqslant$ $\sum_{i} \varepsilon\left(\tilde{\mathcal{B}}_{i}\right)=\sum_{i} \varepsilon\left(\tilde{\mathcal{A}}_{i}\right)$. The cost of the data rearrangement at stage $i$ is

$$
\begin{array}{r}
O\left(n_{1} \cdots n_{i-1} n_{i} m_{i+1} \cdots m_{d} r p \log n_{i}\right)+O\left(n_{1} \cdots n_{i-1} m_{i} m_{i+1} \cdots m_{d} r p \log m_{i}\right) \\
=O\left(M r p\left(\log n_{i}+\log m_{i}\right)\right)
\end{array}
$$

so the total over all $i$ is $O(M r p \log M)$. The total cost of the invocations of $\tilde{\mathcal{A}}_{i}$ is

$$
\sum_{i} n_{1} \cdots n_{i-1} m_{i+1} \cdots m_{d} \mathrm{C}\left(\tilde{\mathcal{A}}_{i}\right) \leqslant \sum_{i} \frac{M}{\max \left(m_{i}, n_{i}\right)} \mathrm{C}\left(\tilde{\mathcal{A}}_{i}\right)
$$

2.8. Exponential functions. The next three results concern the approximation of real and complex exponentials. We use the following facts:

- We may compute an $n$-bit approximation for $\pi$, i.e., an integer $u$ such that $\left|2^{-n} u-\pi\right| \leqslant 2^{-n}$, in time $O\left(n^{1+\delta}\right)$. Similarly for $\log 2$.
- For $z$ lying in a fixed bounded subset of $\mathbb{C}$, we may compute an $n$-bit approximation for $e^{z}$ in time $O\left(n^{1+\delta}\right)$. More precisely, for any constant $C>0$, given integers $x$ and $y$ such that $\left|2^{-n}(x+\mathrm{i} y)\right| \leqslant C$, we may compute integers $u$ and $v$ such that $\left|2^{-n}(u+\mathrm{i} v)-\exp \left(2^{-n}(x+\mathrm{i} y)\right)\right| \leqslant 2^{-n}$ in time $O\left(n^{1+\delta}\right)$.
In fact these tasks may be performed in time $O(\mathrm{M}(n) \log n)$; see [3, Ch. 6-7] or [5, Ch. 4]).

Lemma 2.12 (Complex exponentials). Let $k \geqslant 1$ and $j$ be integers such that $0 \leqslant j<k$, and let $w:=e^{2 \pi \mathrm{i} j / k} \in \mathbb{C}_{0}$. Given $j$ and $k$ as input, we may compute an approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon(\tilde{w})<2$ in time $O\left(\max (p, \log k)^{1+\delta}\right)$.
Proof. Let $p^{\prime}:=p+3$. We first compute a $p^{\prime}$-bit approximation $\tilde{r}$ for $r:=2 \pi j / k \in$ $[0,2 \pi]$, i.e., so that $|\tilde{r}-r| \leqslant 2^{-p^{\prime}}$, in time $O\left(\max (p, \log k)^{1+\delta}\right)$. We then compute a $p^{\prime}$-bit approximation $\tilde{u}$ for $u:=\exp \left(-2 \cdot 2^{-p^{\prime}}+\tilde{r} \mathrm{i}\right) \in \mathbb{C}_{0}$, i.e., so that $|\tilde{u}-u| \leqslant 2^{-p^{\prime}}$, in time $O\left(p^{1+\delta}\right)$ (the $2 \cdot 2^{-p^{\prime}}$ term ensures that $\tilde{u}$ lies within the unit circle). Let $\eta:=\left(-2 \cdot 2^{-p^{\prime}}+\tilde{r} \mathrm{i}\right)-r$; then $|\eta| \leqslant \sqrt{5} \cdot 2^{-p^{\prime}}$ and $|u-w|=$ $\left|\exp \left(-2 \cdot 2^{-p^{\prime}}+\tilde{r} \mathrm{i}\right)-\exp (r \mathrm{i})\right|=|\exp (\eta)-1| \leqslant 3 \cdot 2^{-p^{\prime}}$, so $|\tilde{u}-w| \leqslant 4 \cdot 2^{-p^{\prime}}$. We finally round $\tilde{u}$ towards zero to obtain the desired approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ for $w$ at the original precision $p$. This yields $\varepsilon(\tilde{w}) \leqslant 4 \cdot 2^{p-p^{\prime}}+\sqrt{2}<2$.

Lemma 2.13 (Real exponentials, negative case). Let $k \geqslant 1$ and $j \geqslant 0$ be integers, and let $w:=e^{-\pi j / k} \in \mathbb{C}_{\circ}$. Given $j$ and $k$ as input, we may compute an approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon(\tilde{w})<2$ in time $O\left(\max (p, \log (j+1), \log k)^{1+\delta}\right)$.
Proof. We first check whether $j>k p$ in time $O\left(\max (\log p, \log (j+1), \log k)^{1+\delta}\right)$. If so, then $e^{-\pi j / k}<e^{-\pi p}<2^{-p}$, so we may simply take $\tilde{w}:=0$.

Otherwise, we may assume that $0 \leqslant j / k \leqslant p$. In this case, we first compute an integer $\tau \geqslant 0$ such that $\tau \leqslant \frac{\pi}{\log 2} j / k \leqslant \tau+2$ in time $O\left(\max (\log p, \log k)^{1+\delta}\right)$ (note that $\tau=O(p)$ ). Let $p^{\prime}:=p+2$ and $z:=\tau \log 2-\pi j / k \in[-2 \log 2,0]$. We next compute a $p^{\prime}$-bit approximation $\tilde{z} \leqslant 0$ for $z$, i.e., so that $|\tilde{z}-z| \leqslant 2^{-p^{\prime}}$, in time $O\left(\max (p, \log k)^{1+\delta}\right)$. We then compute a $p^{\prime}$-bit approximation $\tilde{u} \leqslant 1$ for $u:=e^{\tilde{z}} \leqslant 1$, i.e., so that $|\tilde{u}-u| \leqslant 2^{-p^{\prime}}$, in time $O\left(p^{1+\delta}\right)$. Observe that $\left|e^{\tilde{z}}-e^{z}\right| \leqslant$ $2^{-p^{\prime}}$, so $\left|\tilde{u}-e^{z}\right| \leqslant 2 \cdot 2^{-p^{\prime}}$. We finally divide $\tilde{u}$ by $2^{\tau}$ (a bit shift) and round the result towards zero to obtain the desired approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ for $w$ at the original precision $p$. Since $\left|2^{-\tau} \tilde{u}-w\right|=\left|2^{-\tau} \tilde{u}-2^{-\tau} e^{z}\right| \leqslant 2 \cdot 2^{-\tau-p^{\prime}} \leqslant 2 \cdot 2^{-p^{\prime}}$, we obtain $\varepsilon(\tilde{w}) \leqslant 2 \cdot 2^{p-p^{\prime}}+1<2$.

Lemma 2.14 (Real exponentials, positive case). Let $k \geqslant 1, j \geqslant 0$ and $\sigma \geqslant 0$ be integers, and assume that $e^{\pi j / k} \leqslant 2^{\sigma}$ and $\sigma \leqslant 2 p$. Let $w:=2_{\tilde{C}^{-\sigma}} e^{\pi j / k} \in \mathbb{C}_{0}$. Given $j, k$ and $\sigma$ as input, we may compute an approximation $\tilde{w} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon(\tilde{w})<2$ in time $O\left(\max (p, \log k)^{1+\delta}\right)$.
Proof. The hypotheses automatically ensure that $j<k p$. We now proceed along similar lines to the proof of Lemma 2.13: we first compute an integer $\tau \geqslant 0$ near $\sigma-\frac{\pi}{\log 2} j / k$ (again with $\left.\tau=O(p)\right)$, and then at suitably increased precision we
approximate successively $z:=(\tau-\sigma) \log 2+\pi j / k$ and $e^{z}=2^{\tau-\sigma} e^{\pi j / k}$, and finally divide by $2^{\tau}$ and round towards zero to obtain an approximation for $2^{-\sigma} e^{\pi j / k}$ at the original precision $p$. We omit the details, which are similar to the proof of Lemma 2.13.

## 3. Complex transforms for power-of-TWO sizes

Let $p \geqslant 100$ be the working precision as defined in Section 2. The goal of this section is to construct an efficiently computable approximation for the $d$-dimensional complex transform $\mathcal{F}_{t_{1}, \ldots, t_{d}}: \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ (see Section 2.4) in the special case that the $t_{i}$ are powers of two. The following theorem is proved at the end of the section.

Theorem 3.1 (Power-of-two complex transforms). Let $d \geqslant 2$ and let $t_{1}, \ldots, t_{d}$ be powers of two such that $t_{d} \geqslant \cdots \geqslant t_{1} \geqslant 2$. Let $T:=t_{1} \cdots t_{d}$ and assume that $T<2^{p}$. Then we may construct a numerical approximation $\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ for $\mathcal{F}_{t_{1}, \ldots, t_{d}}$ such that $\varepsilon\left(\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}\right)<8 T \log _{2} T$ and

$$
\mathrm{C}\left(\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}\right)<\frac{4 T}{t_{d}} \mathrm{M}\left(3 t_{d} p\right)+O\left(T p \log T+T p^{1+\delta}\right)
$$

Throughout this section we set

$$
r:=t_{d}, \quad \mathscr{R}:=\mathbb{C}[y] /\left(y^{r}+1\right) .
$$

The basic idea of the proof of Theorem 3.1 is to use Bluestein's method [2] to reduce the DFT to the problem of computing a $(d-1)$-dimensional cyclic convolution of size $t_{1} \times \cdots \times t_{d-1}$ over $\mathscr{R}$, and then to perform that convolution by taking advantage of the synthetic roots of unity in $\mathscr{R}$. The $\mathrm{M}(\cdot)$ term in the complexity bound arises from the pointwise multiplications in $\mathscr{R}$. The $O(T p \log T)$ term covers the cost of the synthetic FFTs over $\mathscr{R}$, and the $O\left(T p^{1+\delta}\right)$ term covers various auxiliary operations.

For the rest of this section, $\otimes_{i} \mathscr{R}^{t_{i}}$ always means $\otimes_{i=1}^{d-1} \mathscr{R}^{t_{i}}$ (with $d-1$ factors), and $\otimes_{i} \mathbb{C}^{t_{i}}$ always means $\otimes_{i=1}^{d} \mathbb{C}^{t_{i}}$ (with $d$ factors). These are both vector spaces of dimension $T=t_{1} \cdots t_{d}$ over $\mathbb{C}$.
3.1. Transforms and convolutions over $\mathscr{R}$. We begin with the one-dimensional case. Recall that we have defined a synthetic transform $\mathcal{G}_{t}: \mathscr{R}^{t} \rightarrow \mathscr{R}^{t}$ (see Section 2.4) for each positive divisor $t$ of $2 r$, i.e., for $t \in\{1,2,4,8, \ldots, 2 r\}$.

Lemma 3.2 (FFT over $\mathscr{R})$. For $t \in\{2,4,8, \ldots, 2 r\}$, we may construct a numerical approximation $\tilde{\mathcal{G}}_{t}: \tilde{\mathscr{R}}_{\mathrm{o}}^{t} \rightarrow \tilde{\mathscr{R}}_{\mathrm{o}}^{t}$ for $\mathcal{G}_{t}$ such that $\varepsilon\left(\tilde{\mathcal{G}}_{t}\right) \leqslant \log _{2}$ t and $\mathrm{C}\left(\tilde{\mathcal{G}}_{t}\right)=$ $O(\operatorname{trp} \log t)$.
Proof. First observe that $\mathcal{G}_{1}: \mathscr{R} \rightarrow \mathscr{R}$ is the identity map, and admits the trivial approximation $\tilde{\mathcal{G}}_{1}: \tilde{\mathscr{R}}_{\circ} \rightarrow \tilde{\mathscr{R}}_{\circ}$ given by $\tilde{\mathcal{G}}_{1} u:=u$. This satisfies $\varepsilon\left(\tilde{\mathcal{G}}_{1}\right)=0$ and $\mathrm{C}\left(\tilde{\mathcal{G}}_{1}\right)=O(r p)$.

Now let $t \in\{2,4,8, \ldots, 2 r\}$, and assume that we have already constructed $\tilde{\mathcal{G}}_{t / 2}: \tilde{\mathscr{R}}_{0}^{t / 2} \rightarrow \tilde{\mathscr{R}}_{0}^{t / 2}$ such that $\varepsilon\left(\tilde{\mathcal{G}}_{t / 2}\right) \leqslant \log _{2}(t / 2)$. Given as input $u \in \tilde{\mathscr{R}}_{\mathrm{o}}^{t}$, we will use the well-known Cooley-Tukey algorithm [7] to approximate $\mathcal{G}_{t} u \in \mathscr{R}_{\mathrm{o}}^{t}$.

For any $j \in\{0, \ldots, t-1\}$, observe that

$$
\left(\mathcal{G}_{t} u\right)_{j}=\frac{1}{t} \sum_{k=0}^{t-1} y^{-2 r j k / t} u_{k}=\frac{1}{t} \sum_{k=0}^{\frac{t}{2}-1} y^{-2 r j k / t} u_{k}+\frac{1}{t} \sum_{k=0}^{\frac{t}{2}-1} y^{-2 r j\left(k+\frac{t}{2}\right) / t} u_{k+\frac{t}{2}}
$$

$$
=\frac{1}{t / 2} \sum_{k=0}^{\frac{t}{2}-1} y^{-2 r j k / t} \frac{u_{k}+(-1)^{j} u_{k+\frac{t}{2}}}{2},
$$

where we have used the fact that $y^{r}=-1$. For $\ell \in\left\{0, \ldots, \frac{t}{2}-1\right\}$ this implies that $\left(\mathcal{G}_{t} u\right)_{2 \ell}=\left(\mathcal{G}_{t / 2} v\right)_{\ell}$ and $\left(\mathcal{G}_{t} u\right)_{2 \ell+1}=\left(\mathcal{G}_{t / 2} w\right)_{\ell}$, where $v, w \in \mathscr{R}_{\circ}^{t / 2}$ are given by

$$
v_{k}:=\frac{1}{2}\left(u_{k}+u_{k+\frac{t}{2}}\right), \quad w_{k}:=\frac{1}{2} y^{-2 r k / t}\left(u_{k}-u_{k+\frac{t}{2}}\right), \quad 0 \leqslant k<t / 2
$$

We may therefore use the following algorithm.
(1) Butterflies. For $k \in\left\{0, \ldots, \frac{t}{2}-1\right\}$, we use Lemma 2.1 to compute approximations $\tilde{v}_{k}, \tilde{w}_{k}^{\prime} \in \tilde{\mathscr{R}}_{\circ}$ for $v_{k}$ and $w_{k}^{\prime}:=\frac{1}{2}\left(u_{k}-u_{k+\frac{t}{2}}\right)$ such that $\varepsilon\left(\tilde{v}_{k}\right), \varepsilon\left(\tilde{w}_{k}^{\prime}\right)<1$. We then compute an approximation $\tilde{w}_{k} \in \tilde{\mathscr{R}}_{\circ}$ for $w_{k}=y^{-2 r k / t} w_{k}^{\prime}$; as $y^{r}=-1$, this amounts to cyclically permuting the coefficients of $\tilde{w}_{k}^{\prime}$ (and adjusting signs), and clearly $\varepsilon\left(\tilde{w}_{k}\right)=\varepsilon\left(\tilde{w}_{k}^{\prime}\right)<1$. The cost of this step is $O(\operatorname{trp})$.
(2) Recurse. We compute $\tilde{\mathcal{G}}_{t / 2} \tilde{v}$ and $\tilde{\mathcal{G}}_{t / 2} \tilde{w}$ using the previously constructed map $\tilde{\mathcal{G}}_{t / 2}$, and interleave the results (at a further cost of $O(\operatorname{trp})$ ) to obtain the output vector $\tilde{\mathcal{G}}_{t} u \in \tilde{\mathscr{R}}_{\circ}^{t}$ defined by $\left(\tilde{\mathcal{G}}_{t} u\right)_{2 \ell}:=\left(\tilde{\mathcal{G}}_{t / 2} \tilde{v}\right)_{\ell}$ and $\left(\tilde{\mathcal{G}}_{t} u\right)_{2 \ell+1}:=\left(\tilde{\mathcal{G}}_{t / 2} \tilde{w}\right)_{\ell}$ for $\ell \in\left\{0, \ldots, \frac{t}{2}-1\right\}$.

Recall from Example 2.6 that $\left\|\mathcal{G}_{n}\right\| \leqslant 1$ for all $n$. Applying Lemma 2.7 for $\mathcal{A}=\mathcal{G}_{t / 2}$ and using the induction hypothesis, we obtain

$$
\begin{aligned}
2^{p}\left\|\left(\tilde{\mathcal{G}}_{t} u\right)_{2 \ell}-\left(\mathcal{G}_{t} u\right)_{2 \ell}\right\| & =2^{p}\left\|\left(\tilde{\mathcal{G}}_{t / 2} \tilde{v}\right)_{\ell}-\left(\mathcal{G}_{t / 2} v\right)_{\ell}\right\| \\
& \leqslant \varepsilon\left(\tilde{\mathcal{G}}_{t / 2}\right)+\varepsilon(\tilde{v}) \leqslant \log _{2}(t / 2)+1=\log _{2} t
\end{aligned}
$$

A similar argument applies for $\left(\tilde{\mathcal{G}}_{t} u\right)_{2 \ell+1}$. Therefore $\varepsilon\left(\tilde{\mathcal{G}}_{t}\right) \leqslant \log _{2} t$.
As for the complexity, the above discussion shows that $\mathrm{C}\left(\tilde{\mathcal{G}}_{t}\right)<2 \mathrm{C}\left(\tilde{\mathcal{G}}_{t / 2}\right)+$ $O(\operatorname{trp})$. Together with the base case $\mathrm{C}\left(\tilde{\mathcal{G}}_{1}\right)=O(r p)$, this immediately yields the bound $\mathrm{C}\left(\tilde{\mathcal{G}}_{t}\right)=O(\operatorname{trp} \log t)$ for $t \geqslant 2$.

Combining Lemmas 3.2 and 2.11, we obtain the following approximation for the multidimensional transform $\mathcal{G}_{t_{1}, \ldots, t_{d-1}}: \otimes_{i} \mathscr{R}^{t_{i}} \rightarrow \otimes_{i} \mathscr{R}^{t_{i}}$ (defined in Section 2.4).
Proposition 3.3 (Multivariate FFT over $\mathscr{R}$ ). ${\text { Let } t_{1}, \ldots, t_{d} \text { and } T \text { be as in Theorem }}_{\text {G }}$ 3.1. We may construct a numerical approximation $\tilde{\mathcal{G}}_{t_{1}, \ldots, t_{d-1}}: \otimes_{i} \tilde{\mathscr{R}}_{0}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathscr{R}}_{0}^{t_{i}}$ for $\mathcal{G}_{t_{1}, \ldots, t_{d-1}}$ such that $\varepsilon\left(\tilde{\mathcal{G}}_{t_{1}, \ldots, t_{d-1}}\right)<\log _{2} T$ and $\mathrm{C}\left(\tilde{\mathcal{G}}_{t_{1}, \ldots, t_{d-1}}\right)=O(T p \log T)$.
Proof. Since $t_{i} \geqslant 2$, Lemma 3.2 yields approximations $\tilde{\mathcal{G}}_{t_{i}}$ for $i=1, \ldots, d-1$. We apply Lemma 2.11 (with $d$ replaced by $d-1$ ), taking $R:=\mathscr{R}, m_{i}:=t_{i}, n_{i}:=t_{i}$ and $\mathcal{A}_{i}:=\mathcal{G}_{t_{i}}$ for $i \in\{1, \ldots, d-1\}$. The quantity $M$ defined in Lemma 2.11 is given by $M=t_{1} \cdots t_{d-1}=T / r$ (recall that $r=t_{d}$ throughout Section 3). We obtain

$$
\varepsilon\left(\tilde{\mathcal{G}}_{t_{1}}, \ldots, t_{d-1}\right) \leqslant \sum_{i=1}^{d-1} \varepsilon\left(\tilde{\mathcal{G}}_{t_{i}}\right) \leqslant \sum_{i=1}^{d-1} \log _{2} t_{i}=\log _{2}(T / r)<\log _{2} T
$$

and

$$
\begin{aligned}
\mathrm{C}\left(\tilde{\mathcal{G}}_{t_{1}, \ldots, t_{d-1}}\right) & \leqslant \frac{T}{r} \sum_{i=1}^{d-1} \frac{\mathrm{C}\left(\tilde{\mathcal{G}}_{t_{i}}\right)}{t_{i}}+O\left(\frac{T}{r} r p \log \frac{T}{r}\right) \\
& \leqslant \frac{T}{r} \sum_{i=1}^{d-1} O\left(r p \log t_{i}\right)+O(T p \log T)
\end{aligned}
$$

$$
=O\left(\frac{T}{r} r p \log \frac{T}{r}\right)+O(T p \log T)=O(T p \log T)
$$

Next we will use the above result to approximate the normalised ( $d-1$ )-dimensional convolution map $\mathcal{M}_{\mathscr{R}}: \otimes_{i} \mathscr{R}^{t_{i}} \times \otimes_{i} \mathscr{R}^{t_{i}} \rightarrow \otimes_{i} \mathscr{R}^{t_{i}}$ defined by

$$
\mathcal{M}_{\mathscr{R}}(u, v):=\frac{1}{T} u * v, \quad u, v \in \otimes_{i} \mathscr{R}^{t_{i}}
$$

where $*$ is the convolution operator defined in Section 2.4. Note that $\left\|\mathcal{M}_{\mathscr{R}}\right\| \leqslant 1$; indeed, each component of $u * v$ is a sum of $t_{1} \cdots t_{d-1}=T / r$ terms of the form $u_{j_{1}, \ldots, j_{d-1}} v_{k_{1}, \ldots, k_{d-1}}$, and we saw just before Lemma 2.5 that $\|a b\| \leqslant r\|a\|\|b\|$ for all $a, b \in \mathscr{R}$.

Proposition 3.4 (Convolution over $\mathscr{R})$. Let $t_{1}, \ldots, t_{d}$ and $T$ be as in Theorem 3.1. We may construct a numerical approximation $\tilde{\mathcal{M}}_{\mathscr{R}}: \otimes_{i} \tilde{\mathscr{R}}_{o}^{t_{i}} \times \otimes_{i} \tilde{\mathscr{R}}_{o}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathscr{R}}_{o}^{t_{i}}$ for $\mathcal{M}_{\mathscr{R}}$ such that $\varepsilon\left(\tilde{\mathcal{M}}_{\mathscr{R}}\right)<3 T \log _{2} T+2 T+3$ and

$$
\mathrm{C}\left(\tilde{\mathcal{M}}_{\mathscr{R}}\right)<\frac{4 T}{r} \mathrm{M}(3 r p)+O\left(T p \log T+T p^{1+\delta}\right)
$$

Proof. We are given as input $u, v \in \otimes_{i} \tilde{\mathscr{R}}_{o}^{t_{i}}$. Let $w:=\mathcal{M}_{\mathscr{R}}(u, v)=\frac{1}{T} u * v \in \otimes_{i} \mathscr{R}_{o}^{t_{i}}$ be the exact (normalised) convolution. According to (2.1) we have

$$
r w=\frac{1}{t_{1} \cdots t_{d-1}} u * v=\left(t_{1} \cdots t_{d-1}\right) \mathcal{G}_{t_{1}, \ldots, t_{d-1}}^{*}\left(\left(\mathcal{G}_{t_{1}, \ldots, t_{d-1}} u\right) \cdot\left(\mathcal{G}_{t_{1}, \ldots, t_{d-1}} v\right)\right)
$$

Dividing both sides by $r=t_{d}$, we obtain $w=(T / r) w^{\prime}$ where

$$
w^{\prime}:=\mathcal{G}_{t_{1}, \ldots, t_{d-1}}^{*}\left(\frac{1}{r}\left(\mathcal{G}_{t_{1}, \ldots, t_{d-1}} u\right) \cdot\left(\mathcal{G}_{t_{1}, \ldots, t_{d-1}} v\right)\right) \in \otimes_{i} \mathscr{R}_{\circ}^{t_{i}} .
$$

We now use the following algorithm to approximate $w$.
(1) Forward transforms. We invoke Proposition 3.3 to compute approximations $\tilde{u}^{\prime}, \tilde{v}^{\prime} \in \otimes_{i} \tilde{\mathscr{R}}_{o}^{t_{i}}$ for $u^{\prime}:=\mathcal{G}_{t_{1}, \ldots, t_{d-1}} u \in \otimes_{i} \mathscr{R}_{o}^{t_{i}}$ and $v^{\prime}:=\mathcal{G}_{t_{1}, \ldots, t_{d-1}} v \in \otimes_{i} \mathscr{R}_{o}^{t_{i}}$, with $\varepsilon\left(\tilde{u}^{\prime}\right), \varepsilon\left(\tilde{v}^{\prime}\right)<\log _{2} T$. The cost of this step (and step (3) below) is $O(T p \log T)$.
(2) Pointwise multiplications. Let $\mathcal{A}: \mathscr{R} \times \mathscr{R} \rightarrow \mathscr{R}$ be the normalised multiplication map defined by $\mathcal{A}(a, b):=a b / r$; the bound $\|a b\| \leqslant r\|a\|\|b\|$ implies that $\|\mathcal{A}\| \leqslant 1$. Lemma 2.5 yields an approximation $\tilde{\mathcal{A}}: \tilde{\mathscr{R}}_{\circ} \times \tilde{\mathscr{R}}_{\circ} \rightarrow \tilde{\mathscr{R}}_{\circ}$ such that $\varepsilon(\tilde{\mathcal{A}})<2$ (note that $r=t_{d} \leqslant T / 2<2^{p-1}$ ). Applying $\tilde{\mathcal{A}}$ to each component of $\tilde{u}^{\prime}$ and $\tilde{v}^{\prime}$, we obtain an approximation $\tilde{z} \in \otimes_{i} \tilde{\mathscr{R}}_{0}^{t_{i}}$ for $z:=\frac{1}{r} u^{\prime} \cdot v^{\prime} \in \otimes_{i} \mathscr{R}_{o}^{t_{i}}$. This step requires time

$$
\frac{T}{r}(4 \mathrm{M}(3 r p)+O(r p))=\frac{4 T}{r} \mathrm{M}(3 r p)+O(T p)
$$

and by Lemma 2.9 we have

$$
\varepsilon(\tilde{z}) \leqslant \varepsilon(\tilde{\mathcal{A}})+\varepsilon\left(\tilde{u}^{\prime}\right)+\varepsilon\left(\tilde{v}^{\prime}\right)<2+\log _{2} T+\log _{2} T=2 \log _{2} T+2
$$

(3) Inverse transform. We use Proposition 3.3 again to compute an approximation $\tilde{w}^{\prime} \in \otimes_{i} \tilde{\mathscr{R}}_{0}^{t_{i}}$ for $w^{\prime}=\mathcal{G}_{t_{1}, \ldots, t_{d-1}}^{*} z \in \otimes_{i} \mathscr{R}_{o}^{t_{i}}$. Recall from Example 2.6 that $\left\|\mathcal{G}_{t_{1}, \ldots, t_{d-1}}^{*}\right\| \leqslant 1$. By Lemma 2.7, we obtain

$$
\varepsilon\left(\tilde{w}^{\prime}\right) \leqslant \varepsilon\left(\tilde{\mathcal{G}}_{t_{1}, \ldots, t_{d-1}}^{*}\right)+\varepsilon(\tilde{z})<\log _{2} T+\left(2 \log _{2} T+2\right)=3 \log _{2} T+2
$$

(4) Scaling. Recall that $w=(T / r) w^{\prime}$ and that $\|w\| \leqslant 1$. We may therefore apply Lemma 2.2 (with $c:=T / r \leqslant 2^{p}$ ) to compute an approximation $\tilde{w} \in \otimes_{i} \tilde{\mathscr{R}}_{0}^{t_{i}}$ such that

$$
\varepsilon(\tilde{w})<2(T / r) \varepsilon\left(\tilde{w}^{\prime}\right)+3 \leqslant T\left(3 \log _{2} T+2\right)+3
$$

(here we have used the hypothesis that $r=t_{d} \geqslant 2$ ). The cost of this scaling step is $O\left(T p^{1+\delta}\right)$. Finally we take $\tilde{\mathcal{M}}_{\mathscr{R}}(u, v):=\tilde{w}$.
3.2. Transforms and convolutions over $\mathbb{C}$. We now transfer the results of the previous section from $\mathscr{R}$ to $\mathbb{C}$. Consider the normalised $d$-dimensional convolution $\operatorname{map} \mathcal{M}_{\mathbb{C}}: \otimes_{i} \mathbb{C}^{t_{i}} \times \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ defined by

$$
\mathcal{M}_{\mathbb{C}}(u, v):=\frac{1}{T} u * v, \quad u, v \in \otimes_{i} \mathbb{C}^{t_{i}}
$$

As before we have $\left\|\mathcal{M}_{\mathbb{C}}\right\| \leqslant 1$.
Proposition 3.5 (Convolution over $\mathbb{C}$ ). Let $t_{1}, \ldots, t_{d}$ and $T$ be as in Theorem 3.1. We may construct a numerical approximation $\tilde{\mathcal{M}}_{\mathbb{C}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}} \times \otimes_{i} \tilde{\mathbb{C}}_{0}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ for $\mathcal{M}_{\mathbb{C}}$ such that $\varepsilon\left(\tilde{\mathcal{M}}_{\mathbb{C}}\right)<3 T \log _{2} T+2 T+15$ and

$$
\mathrm{C}\left(\tilde{\mathcal{M}}_{\mathbb{C}}\right)<\frac{4 T}{r} \mathrm{M}(3 r p)+O\left(T p \log T+T p^{1+\delta}\right)
$$

Proof. Let $\zeta:=e^{\pi \mathrm{i} / r}$ and consider the $\mathbb{C}$-linear map $\mathcal{S}: \mathbb{C}^{r} \rightarrow \mathscr{R}$ defined by

$$
\mathcal{S}\left(w_{0}, \ldots, w_{r-1}\right):=w_{0}+\zeta w_{1} y+\cdots+\zeta^{r-1} w_{r-1} y^{r-1}
$$

Then $\mathcal{S}$ is a isomorphism of rings between $\left(\mathbb{C}^{r}, *\right)$ and $\mathscr{R}=\mathbb{C}[y] /\left(y^{r}+1\right)$; in fact, recalling that $\left(\mathbb{C}^{r}, *\right) \cong \mathbb{C}[x] /\left(x^{r}-1\right)$, we may regard $\mathcal{S}$ as the map sending $x$ to $\zeta y$. Moreover, $\mathcal{S}$ induces an isomorphism of rings

$$
\mathcal{T}:\left(\otimes_{i=1}^{d} \mathbb{C}^{t_{i}}, *\right) \rightarrow\left(\otimes_{i=1}^{d-1} \mathscr{R}^{t_{i}}, *\right)
$$

Indeed, identifying these rings respectively as $\mathbb{C}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{t_{1}}-1, \ldots, x_{d}^{t_{d}}-1\right)$ and $\mathbb{C}\left[x_{1}, \ldots, x_{d-1}, y\right] /\left(x_{1}^{t_{1}}-1, \ldots, x_{d-1}^{t_{d-1}}-1, y^{r}+1\right)$, the isomorphism $\mathcal{T}$ sends $u\left(x_{1}, \ldots, x_{d-1}, x_{d}\right)$ to $u\left(x_{1}, \ldots, x_{d-1}, \zeta y\right)$. Writing $\mathcal{U}:=\mathcal{T}^{-1}$ for the inverse isomorphism, we obtain

$$
\mathcal{M}_{\mathbb{C}}(u, v)=\mathcal{U}\left(\mathcal{M}_{\mathscr{R}}(\mathcal{T} u, \mathcal{T} v)\right), \quad u, v \in \otimes_{i} \mathbb{C}^{t_{i}}
$$

Now we construct numerical approximations for the maps just introduced. We may construct an approximation $\tilde{\mathcal{S}}: \tilde{\mathbb{C}}^{r} \rightarrow \tilde{\mathscr{R}}_{\circ}$ by first using Lemma 2.12 to compute an approximation $\tilde{\zeta}_{j} \in \tilde{\mathbb{C}}_{\circ}$ for each $\zeta_{j}:=\zeta^{j}=e^{\pi \mathrm{i} j / r} \in \mathbb{C}_{0}$, and then using Corollary 2.10 to compute an approximation $\tilde{w}_{j}^{\prime} \in \tilde{\mathbb{C}}_{\circ}$ for each product $w_{j}^{\prime}:=\zeta_{j} w_{j} \in \mathbb{C}_{\circ}$. We obtain $\varepsilon\left(\tilde{\zeta}_{j}\right)<2$ and then $\varepsilon\left(\tilde{w}_{j}^{\prime}\right)<\varepsilon\left(\tilde{\zeta}_{j}\right)+2<4$, whence $\varepsilon(\tilde{\mathcal{S}})<4$. We also have $C(\tilde{\mathcal{S}})=O\left(r \max (p, \log r)^{1+\delta}\right)=O\left(r p^{1+\delta}\right)$, since $\log r=O(p)$ (recall from the proof of Proposition 3.4 that $r<2^{p-1}$ ). Then, applying $\tilde{\mathcal{S}}$ separately to the coefficient of each $x_{1}^{j_{1}} \cdots x_{d-1}^{j_{d-1}}$, we obtain an approximation $\tilde{\mathcal{T}}: \otimes_{i=1}^{d} \tilde{\mathbb{C}}_{o}^{t_{i}} \rightarrow \otimes_{i=1}^{d-1} \tilde{\mathscr{R}}_{0}^{t_{i}}$ such that $\varepsilon(\tilde{\mathcal{T}})<4$ and $\mathrm{C}(\tilde{\mathcal{T}})=O\left((T / r) r p^{1+\delta}\right)=O\left(T p^{1+\delta}\right)$. The inverse is handled similarly; we obtain an approximation $\tilde{\mathcal{U}}: \otimes_{i=1}^{d-1} \tilde{\mathscr{R}}_{o}^{t_{i}} \rightarrow \otimes_{i=1}^{d} \tilde{\mathbb{C}}_{o}^{t_{i}}$ such that $\varepsilon(\tilde{\mathcal{U}})<4$ and $\mathrm{C}(\tilde{\mathcal{U}})=O\left(T p^{1+\delta}\right)$.

Finally, given as input $u, v \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$, we define $\tilde{\mathcal{M}}_{\mathbb{C}}(u, v):=\tilde{\mathcal{U}}\left(\tilde{\mathcal{M}}_{\mathscr{R}}(\tilde{\mathcal{T}} u, \tilde{\mathcal{T}} v)\right)$. We already observed that $\left\|\mathcal{M}_{\mathscr{R}}\right\| \leqslant 1$, and clearly $\|\mathcal{T}\|,\|\mathcal{U}\| \leqslant 1$, so Lemma 2.7, Lemma 2.9 and Proposition 3.4 together imply that

$$
\varepsilon\left(\tilde{\mathcal{M}}_{\mathbb{C}}\right) \leqslant \varepsilon(\tilde{\mathcal{U}})+\varepsilon\left(\tilde{\mathcal{M}}_{\mathscr{R}}\right)+\varepsilon(\tilde{\mathcal{T}})+\varepsilon(\tilde{\mathcal{T}})<\left(3 T \log _{2} T+2 T+3\right)+4+4+4
$$

The estimate for $\mathrm{C}\left(\tilde{\mathcal{M}}_{\mathbb{C}}\right)$ follows immediately from Proposition 3.4.
Finally we use Bluestein's trick [2] to prove the main result of this section.

Proof of Theorem 3.1. We are given as input $u \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$. We wish to approximate its transform $v:=\mathcal{F}_{t_{1}, \ldots, t_{d}} u \in \otimes_{i} \mathbb{C}_{o}^{t_{i}}$, which is given explicitly by

$$
v_{j_{1}, \ldots, j_{d}}=\frac{1}{T} \sum_{k_{1}=0}^{t_{1}-1} \cdots \sum_{k_{d}=0}^{t_{d}-1} e^{-2 \pi \mathrm{i}\left(j_{1} k_{1} / t_{1}+\cdots+j_{d} k_{d} / t_{d}\right)} u_{k_{1}, \ldots, k_{d}}, \quad 0 \leqslant j_{i}<t_{i}
$$

For any $j_{1}, \ldots, j_{d} \in \mathbb{Z}$ set

$$
\begin{equation*}
a_{j_{1}, \ldots, j_{d}}:=e^{\pi \mathrm{i}\left(j_{1}^{2} / t_{1}+\cdots+j_{d}^{2} / t_{d}\right)} \in \mathbb{C}_{0} \tag{3.1}
\end{equation*}
$$

The identity $-2 j k=(j-k)^{2}-j^{2}-k^{2}$ implies that

$$
\begin{equation*}
v_{j_{1}, \ldots, j_{d}}=\bar{a}_{j_{1}, \ldots, j_{d}} \frac{1}{T} \sum_{k_{1}=0}^{t_{1}-1} \cdots \sum_{k_{d}=0}^{t_{d}-1} a_{j_{1}-k_{1}, \ldots, j_{d}-k_{d}}\left(\bar{a}_{k_{1}, \ldots, k_{d}} u_{k_{1}, \ldots, k_{d}}\right) \tag{3.2}
\end{equation*}
$$

where - denotes complex conjugation. Moreover, we observe that $a_{j_{1}, \ldots, j_{d}}$ is periodic in each $j_{i}$ with period $t_{i}$, as $e^{\pi \mathrm{i}\left(j_{i}+t_{i}\right)^{2} / t_{i}}=e^{\pi \mathrm{i} j_{i}^{2} / t_{i}}\left(e^{\pi \mathrm{i}}\right)^{2 j_{i}+t_{i}}=e^{\pi \mathrm{i} j_{i}^{2} / t_{i}}$ (using the fact that $t_{i}$ is even). Therefore, regarding (3.1) as defining a vector $a \in \otimes_{i} \mathbb{C}_{0}^{t_{i}}$, we may rewrite (3.2) in the form

$$
v=\bar{a} \cdot\left(\frac{1}{T} a *(\bar{a} \cdot u)\right) .
$$

We now use the following algorithm.
(1) Compute $a$. Recalling that each $t_{i}$ divides $r=t_{d}$, we may write

$$
a_{j_{1}, \ldots, j_{d}}=e^{2 \pi \mathrm{i} \eta_{j_{1}, \ldots, j_{d}} / 2 r}, \quad \eta_{j_{1}, \ldots, j_{d}}:=\frac{r}{t_{1}} j_{1}^{2}+\cdots+\frac{r}{t_{d}} j_{d}^{2} \quad(\bmod 2 r)
$$

Iterating over the tuples $\left(j_{1}, \ldots, j_{d}\right)$ in lexicographical order, we may compute $\eta_{j_{1}, \ldots, j_{d}}$ in amortised time $O(\log r)=O(p)$ per tuple (for example by repeatedly using the identity $(j+1)^{2}=j^{2}+(2 j+1)$, and the fact that each multiplication by $r / t_{i}$ is a bit shift), and then use Lemma 2.12 to compute an approximation $\tilde{a}_{j_{1}, \ldots, j_{d}} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon\left(\tilde{a}_{j_{1}, \ldots, j_{d}}\right)<2$ in time $O\left(p^{1+\delta}\right)$. We thus obtain $\tilde{a} \in \otimes_{i} \tilde{\mathbb{C}}_{0}^{t_{i}}$ with $\varepsilon(\tilde{a})<2$ in time $O\left(T p^{1+\delta}\right)$.
(2) Pre-multiply. We use Corollary 2.10 to compute an approximation $\tilde{b} \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ for $b:=\bar{a} \cdot u$ with $\varepsilon(\tilde{b})<\varepsilon(\tilde{a})+2<4$ in time $O\left(T p^{1+\delta}\right)$.
(3) Convolution. We use Proposition 3.5 to compute an approximation $\tilde{c} \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ for $c:=\frac{1}{T} a * b$. This requires time $(4 T / r) \mathrm{M}(3 r p)+O\left(T p \log T+T p^{1+\delta}\right)$, and by Lemma 2.9 we have

$$
\varepsilon(\tilde{c}) \leqslant \varepsilon\left(\tilde{\mathcal{M}}_{\mathbb{C}}\right)+\varepsilon(\tilde{a})+\varepsilon(\tilde{b})<3 T \log _{2} T+2 T+21
$$

(4) Post-multiply. We invoke Corollary 2.10 again to compute an approximation $\tilde{v} \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ for $v=\bar{a} \cdot c$ such that

$$
\varepsilon(\tilde{v}) \leqslant \varepsilon(\tilde{a})+\varepsilon(\tilde{c})+2<3 T \log _{2} T+2 T+25
$$

in time $O\left(T p^{1+\delta}\right)$. We have $2 T+25<5 T \log _{2} T$ (because $T \geqslant t_{1} t_{2} \geqslant 4$ ), so $\varepsilon(\tilde{v})<8 T \log _{2} T$. Finally we take $\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}} u:=\tilde{v}$.

Remark 3.6. In the algorithm developed in this section, it is essential that the multidimensional complex transform be performed "all at once". If instead we decompose it into one-dimensional transforms in the usual way, and then use Bluestein's method to convert each of these to a one-dimensional convolution, this would lead to an extraneous factor of $O(d)$ on the right hand side of (1.3).

Remark 3.7. An alternative method for reducing multidimensional complex transforms to synthetic transforms was described in [34]. Briefly, given as input $u \in$ $\mathbb{C}\left[x_{1}, \ldots, x_{d-1}, y\right] /\left(x_{1}^{t_{1}}-1, \ldots, x_{d-1}^{t_{d-1}}-1, y^{r}+1\right)$, assume that we wish to evaluate each $x_{i}$ at the complex $t_{i}$-th roots of unity, and $y$ at the "odd" complex $2 r$-th roots of unity (i.e., roots of $y^{r}+1$ ). We first use $(d-1)$-dimensional synthetic transforms to compute the polynomials $u_{i_{1}, \ldots, i_{d-1}}(y):=u\left(y^{2 r i_{1} / t_{1}}, \ldots, y^{2 r i_{d-1} / t_{d-1}}, y\right) \in$ $\mathbb{C}[y] /\left(y^{r}+1\right)$, for all $i_{k} \in\left\{0, \ldots, t_{k}-1\right\}$. It then suffices to compute the onedimensional complex DFT of each $u_{i_{1}, \ldots, i_{d-1}}(y)$, which could be done for example by Bluestein's method. This alternative method has the same complexity (up to a constant factor) as the method presented in this section.

## 4. Gaussian Resampling

Let $p \geqslant 100$ be the working precision as defined in Section 2. The aim of this section is to show how to reduce the problem of approximating a multidimensional complex transform $\mathcal{F}_{s_{1}, \ldots, s_{d}}: \otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$ (defined in Section 2.4), for a given "source" size $s_{1} \times \cdots \times s_{d}$, to the problem of approximating another transform $\mathcal{F}_{t_{1}, \ldots, t_{d}}: \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ for a somewhat larger "target" size $t_{1} \times \cdots \times t_{d}$. (In Section 5 we will specialise to the case that the $s_{i}$ are primes and the $t_{i}$ are powers of two.) The following theorem is proved at the end of Section 4.3. It may be strengthened in various ways; see the discussion in Section 4.4.

Theorem 4.1 (Gaussian resampling). Let $d \geqslant 1$, let $s_{1}, \ldots, s_{d}$ and $t_{1}, \ldots, t_{d}$ be integers such that $2 \leqslant s_{i}<t_{i}<2^{p}$ and $\operatorname{gcd}\left(s_{i}, t_{i}\right)=1$, and let $T:=t_{1} \cdots t_{d}$. Let $\alpha$ be an integer in the interval $2 \leqslant \alpha<p^{1 / 2}$. For each $i$, let $\theta_{i}:=t_{i} / s_{i}-1$, and assume that $\theta_{i} \geqslant p / \alpha^{4}$.

Then there exist linear maps $\mathcal{A}: \otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ and $\mathcal{B}: \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$, with $\|\mathcal{A}\|,\|\mathcal{B}\| \leqslant 1$, such that

$$
\mathcal{F}_{s_{1}, \ldots, s_{d}}=2^{\gamma} \mathcal{B} \mathcal{F}_{t_{1}, \ldots, t_{d}} \mathcal{A}, \quad \gamma:=2 d \alpha^{2}
$$

Moreover, we may construct numerical approximations $\tilde{\mathcal{A}}: \otimes_{i} \widetilde{\mathbb{C}}_{0}^{s_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{0}^{t_{i}}$ and $\tilde{\mathcal{B}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ such that $\varepsilon(\tilde{\mathcal{A}}), \varepsilon(\tilde{\mathcal{B}})<d p^{2}$ and

$$
\mathrm{C}(\tilde{\mathcal{A}}), \mathrm{C}(\tilde{\mathcal{B}})=O\left(d T p^{3 / 2+\delta} \alpha+T p \log T\right)
$$

This theorem shows that to approximate a transform of size $s_{1} \times \cdots \times s_{d}$, one may first apply $\tilde{\mathcal{A}}$, then compute a transform of size $t_{1} \times \cdots \times t_{d}$, then apply $\tilde{\mathcal{B}}$, and finally multiply by $2^{\gamma}$. The $d T p^{3 / 2+\delta} \alpha$ term in the complexity bound arises from the numerical computations at the heart of the "Gaussian resampling" method. The $T p \log T$ term covers the cost of various data rearrangements in the Turing model (this term would not appear if we worked in the Boolean circuit model). In the application in Section 5, the parameters will be chosen so that the first term is negligible compared to the $O(T p \log T)$ cost of evaluating $\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}$.

Throughout this section we use the notation

$$
[x]:=\left\lfloor x+\frac{1}{2}\right\rfloor, \quad\langle x\rangle:=x-[x], \quad x \in \mathbb{R}
$$

Thus $[x]$ is the nearest integer to $x$, rounding upwards in case of a tie, and $\langle x\rangle$ is the corresponding fractional part with $-\frac{1}{2} \leqslant\langle x\rangle<\frac{1}{2}$. For convenience of the reader, we provide in Table 2 a list of the linear maps appearing in this section and where they are defined.

| $\mathcal{F}_{n}$ | complex DFT of length $n$ | $(\S 2.4)$ |
| :--- | :--- | ---: |
| $\mathcal{F}_{n_{1}, \ldots, n_{d}}$ | complex multidimensional DFT of size $n_{1} \times \cdots \times n_{d}$ | $(\S 2.4)$ |
| $\mathcal{A}, \mathcal{B}$ | the main maps we want to construct in Theorem 4.1 |  |
| $\mathcal{S}, \mathcal{T}$ | resampling maps | $(\S 4.1)$ |
| $\mathcal{P}_{s}, \mathcal{P}_{t}$ | permutation maps | $(\S 4.1)$ |
| $\mathcal{C}$ | row-deleting map | $(\S 4.2)$ |
| $\mathcal{T}^{\prime}$ | $\mathcal{C} \mathcal{T}$ (a certain square submatrix of $\mathcal{T})$ | $(\S 4.2)$ |
| $\mathcal{D}$ | a diagonal map | $(\S 4.2)$ |
| $\mathcal{N}$ | normalised version of $\mathcal{T}$ | $(\S 4.2)$ |
| $\mathcal{E}$ | $\mathcal{N}-\mathcal{I}($ a map with small norm $)$ | $(\S 4.2)$ |
| $\mathcal{J}$ | inverse of $\mathcal{N}$ | $(\S 4.3)$ |
| $\mathcal{S}^{\prime}, \mathcal{J}^{\prime}, \mathcal{D}^{\prime}$ | normalised versions of $\mathcal{S}, \mathcal{J}, \mathcal{D}$ | $(\S 4.3)$ |

Table 2. Glossary of linear maps appearing in Section 4
4.1. The resampling identity. Throughout Sections 4.1 and 4.2 , let $s$ and $t>s$ be positive integers such that $\operatorname{gcd}(s, t)=1$, and let $\alpha \in(0, \infty)$. Recall from Section 2.4 that $u_{j}$ always means $u_{j \bmod s}$ if $u \in \mathbb{C}^{s}$, whereas $u_{k}$ always means $u_{k \bmod t}$ if $u \in \mathbb{C}^{t}$.

Define "resampling maps" $\mathcal{S}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{t}$ and $\mathcal{T}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{t}$ by

$$
\begin{array}{ll}
(\mathcal{S} u)_{k}:=\alpha^{-1} \sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{-2} s^{2}\left(\frac{k}{t}-\frac{j}{s}\right)^{2}} u_{j}, & u \in \mathbb{C}^{s}, \\
(\mathcal{T} u)_{k}:=\quad \sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2} t^{2}\left(\frac{k}{t}-\frac{j}{s}\right)^{2}} u_{j}, & u \in \mathbb{C}^{s}, \quad 0 \leqslant k<t
\end{array}
$$

These sums certainly converge due to the rapid decay of the function $e^{-x^{2}}$. Each entry $(\mathcal{S} u)_{k}$ and $(\mathcal{T} u)_{k}$ is a weighted linear combination of $u_{0}, \ldots, u_{s-1}$, with the largest weightings given to those $u_{j}$ for which $j / s$ is closest to $k / t$ modulo 1. Figure 2 shows examples of the matrices of $\mathcal{S}$ and $\mathcal{T}$. They have relatively large entries near the "diagonal" of slope $t / s$, and the entries decay rapidly away from the diagonal according to a Gaussian law. The parameter $\alpha$ controls the rate of decay.

We also define permutation maps $\mathcal{P}_{s}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s}$ and $\mathcal{P}_{t}: \mathbb{C}^{t} \rightarrow \mathbb{C}^{t}$ by

$$
\begin{array}{lll}
\left(\mathcal{P}_{s} u\right)_{j}:=u_{t j}, & u \in \mathbb{C}^{s}, & 0 \leqslant j<s, \\
\left(\mathcal{P}_{t} u\right)_{k}:=u_{-s k}, & u \in \mathbb{C}^{t}, & 0 \leqslant k<t .
\end{array}
$$

Then we have the following fundamental identity, which uses $\mathcal{S}$ and $\mathcal{T}$ to transform $\mathcal{F}_{s}$ into $\mathcal{F}_{t}$.

Theorem 4.2 (Resampling identity). We have $\mathcal{T} \mathcal{P}_{s} \mathcal{F}_{s}=\mathcal{P}_{t} \mathcal{F}_{t} \mathcal{S}$. In other words, the following diagram commutes:


|  | 0.5000 | 0.2280 | 0.0216 | $4.3 \mathrm{e}-4$ | $1.7 \mathrm{e}-6$ | 3.0e-9 | $1.7 \mathrm{e}-6$ | $4.3 \mathrm{e}-4$ | 0.0216 | 0.2280 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.3142 | 0.4795 | 0.1522 | 0.0100 | $1.4 \mathrm{e}-4$ | $3.9 \mathrm{e}-7$ | $9.0 \mathrm{e}-9$ | $7.1 \mathrm{e}-6$ | 0.0012 | 0.0428 |
|  | 0.0779 | 0.3982 | 0.4230 | 0.0934 | 0.0043 | $4.1 \mathrm{e}-5$ | $8.1 \mathrm{e}-8$ | $4.7 \mathrm{e}-8$ | $2.7 \mathrm{e}-5$ | 0.0032 |
|  | 0.0076 | - 0.1305 | 0.4642 | 0.3432 | 0.0527 | 0.0017 | $1.1 \mathrm{e}-5$ | $1.6 \mathrm{e}-8$ | $2.3 \mathrm{e}-7$ | $9.3 \mathrm{e}-5$ |
|  | $2.9 \mathrm{e}-4$ | 0.0169 | 0.2011 | 0.4977 | 0.2561 | 0.0274 | $6.1 \mathrm{e}-4$ | $2.8 \mathrm{e}-6$ | $3.5 \mathrm{e}-9$ | $1.1 \mathrm{e}-6$ |
|  | $4.5 \mathrm{e}-6$ | $8.6 \mathrm{e}-4$ | 0.0344 | 0.2849 | 0.4908 | 0.1757 | 0.0131 | $2.0 \mathrm{e}-4$ | $6.5 \mathrm{e}-7$ | $5.3 \mathrm{e}-9$ |
|  | $2.7 \mathrm{e}-8$ | $1.7 \mathrm{e}-5$ | 0.0023 | 0.0644 | 0.3714 | 0.4452 | 0.1109 | 0.0057 | $6.2 \mathrm{e}-5$ | $1.4 \mathrm{e}-7$ |
|  | $2.7 \mathrm{e}-8$ | $1.4 \mathrm{e}-7$ | $6.2 \mathrm{e}-5$ | 0.0057 | 0.1109 | 0.4452 | 0.3714 | 0.0644 | 0.0023 | $1.7 \mathrm{e}-5$ |
|  | $4.5 \mathrm{e}-6$ | $5.3 \mathrm{e}-9$ | $6.5 \mathrm{e}-7$ | $2.0 \mathrm{e}-4$ | 0.0131 | 0.1757 | 0.4908 | 0.2849 | 0.0344 | $8.6 \mathrm{e}-4$ |
|  | $2.9 \mathrm{e}-4$ | $1.1 \mathrm{e}-6$ | $3.5 \mathrm{e}-9$ | $2.8 \mathrm{e}-6$ | $6.1 \mathrm{e}-4$ | 0.0274 | 0.2561 | 0.4977 | 0.2011 | 0.0169 |
|  | 0.0076 | 9.3e-5 | $2.3 \mathrm{e}-7$ | $1.6 \mathrm{e}-8$ | $1.1 \mathrm{e}-5$ | 0.0017 | 0.0527 | 0.3432 | 0.4642 | 0.1305 |
|  | 0.0779 | 0.0032 | $2.7 \mathrm{e}-5$ | $4.7 \mathrm{e}-8$ | $8.1 \mathrm{e}-8$ | $4.1 \mathrm{e}-5$ | 0.0043 | 0.0934 | 0.4230 | 0.3982 |
|  | [ 0.3142 | - 0.0428 | 0.0012 | $7.1 \mathrm{e}-6$ | $9.0 \mathrm{e}-9$ | $3.9 \mathrm{e}-7$ | $1.4 \mathrm{e}-4$ | 0.0100 | 0.1522 O | 0.4795 |
|  | 1.0000 | $6.0 \mathrm{e}-10$ | $1.3 \mathrm{e}-37$ | $9.8 \mathrm{e}-84$ | $2.7 \mathrm{e}-148$ | 5.3e-231 | $2.7 \mathrm{e}-148$ | $9.8 \mathrm{e}-84$ | $1.3 \mathrm{e}-37$ | 7 6.0e-10 |
|  | $3.5 \mathrm{e}-6$ | 0.3227 | $1.1 \mathrm{e}-14$ | $1.3 \mathrm{e}-46$ | $5.4 \mathrm{e}-97$ | $8.1 \mathrm{e}-166$ | $1.6 \mathrm{e}-210$ | $9.2 \mathrm{e}-132$ | $1.9 \mathrm{e}-71$ | $1.3 \mathrm{e}-29$ |
|  | $1.5 \mathrm{e}-22$ | 0.0021 | 0.0108 | $2.0 \mathrm{e}-20$ | $1.3 \mathrm{e}-56$ | $3.1 \mathrm{e}-111$ | $2.6 \mathrm{e}-184$ | $1.1 \mathrm{e}-190$ | $3.3 \mathrm{e}-116$ | $6 \quad 3.7 \mathrm{e}-60$ |
|  | $7.6 \mathrm{e}-50$ | $1.7 \mathrm{e}-16$ | 0.1339 | $3.8 \mathrm{e}-5$ | $3.9 \mathrm{e}-27$ | $1.4 \mathrm{e}-67$ | $1.8 \mathrm{e}-126$ | $8.4 \mathrm{e}-204$ | $7.1 \mathrm{e}-172$ | $2 \quad 1.2 \mathrm{e}-101$ |
|  | $4.8 \mathrm{e}-88$ | $1.6 \mathrm{e}-40$ | $2.0 \mathrm{e}-11$ | 0.8819 | $1.4 \mathrm{e}-8$ | $7.8 \mathrm{e}-35$ | $1.6 \mathrm{e}-79$ | $1.1 \mathrm{e}-142$ | $2.9 \mathrm{e}-224$ | $4 \quad 5.0 \mathrm{e}-154$ |
|  | $3.7 \mathrm{e}-137$ | $1.9 \mathrm{e}-75$ | $3.7 \mathrm{e}-32$ | $2.5 \mathrm{e}-7$ | 0.6049 | $5.3 \mathrm{e}-13$ | $1.6 \mathrm{e}-43$ | $1.8 \mathrm{e}-92$ | $7.2 \mathrm{e}-160$ | $0 \quad 2.5 \mathrm{e}-217$ |
|  | $3.4 \mathrm{e}-197 \quad 2$ | $2.8 \mathrm{e}-121$ | $8.2 \mathrm{e}-64$ | $8.6 \mathrm{e}-25$ | $3.2 \mathrm{e}-4$ | 0.0432 | $2.1 \mathrm{e}-18$ | $3.6 \mathrm{e}-53$ | $2.2 \mathrm{e}-106$ | $6 \quad 4.8 \mathrm{e}-178$ |
|  | $3.4 \mathrm{e}-197$ 4 | $4.8 \mathrm{e}-178$ | $2.2 \mathrm{e}-106$ | 3.6e-53 | $2.1 \mathrm{e}-18$ | 0.0432 | $3.2 \mathrm{e}-4$ | $8.6 \mathrm{e}-25$ | $8.2 \mathrm{e}-64$ | $4 \quad 2.8 \mathrm{e}-121$ |
|  | $3.7 \mathrm{e}-137 \quad 2$ | $2.5 \mathrm{e}-217$ | $7.2 \mathrm{e}-160$ | $1.8 \mathrm{e}-92$ | $1.6 \mathrm{e}-43$ | $5.3 \mathrm{e}-13$ | 0.6049 | $2.5 \mathrm{e}-7$ | $3.7 \mathrm{e}-32$ | 1.9e-75 |
|  | $4.8 \mathrm{e}-88$ 5 | $5.0 \mathrm{e}-154$ | $2.9 \mathrm{e}-224$ | 1.1e-142 | $1.6 \mathrm{e}-79$ | $7.8 \mathrm{e}-35$ | $1.4 \mathrm{e}-8$ | 0.8819 | $2.0 \mathrm{e}-11$ | $1.6 \mathrm{e}-40$ |
|  | $7.6 \mathrm{e}-50 \quad 1$ | $1.2 \mathrm{e}-101$ | $7.1 \mathrm{e}-172$ | $8.4 \mathrm{e}-204$ | $1.8 \mathrm{e}-126$ | $1.4 \mathrm{e}-67$ | $3.9 \mathrm{e}-27$ | $3.8 \mathrm{e}-5$ | 0.1339 | -1.7e-16 |
|  | $1.5 \mathrm{e}-22$ | $3.7 \mathrm{e}-60$ | $3.3 \mathrm{e}-116$ | $1.1 \mathrm{e}-190$ | $2.6 \mathrm{e}-184$ | $3.1 \mathrm{e}-111$ | $1.3 \mathrm{e}-56$ | $2.0 \mathrm{e}-20$ | 0.0108 | 0.0021 |
| L | $3.5 \mathrm{e}-6$ | $1.3 \mathrm{e}-29$ | $1.9 \mathrm{e}-71$ | $9.2 \mathrm{e}-132$ | $1.6 \mathrm{e}-210$ | $8.1 \mathrm{e}-166$ | $5.4 \mathrm{e}-97$ | $1.3 \mathrm{e}-46$ | $1.1 \mathrm{e}-14$ | 4 |

Figure 2. Matrices of $\mathcal{S}$ and $\mathcal{T}$ for $s=10, t=13, \alpha=2$. Maximal entries in each column are shown in bold. All entries are rounded to the number of significant figures shown.

Proof. Given $u \in \mathbb{C}^{s}$, define a smooth, 1-periodic function $f_{u}: \mathbb{R} \rightarrow \mathbb{C}$ by

$$
f_{u}(x):=\sum_{m \in \mathbb{Z}} u_{m} g\left(x-\frac{m}{s}\right), \quad g(x):=e^{-\pi \alpha^{-2} s^{2} x^{2}}
$$

It has an absolutely and uniformly convergent Fourier expansion

$$
f_{u}(x)=\sum_{r \in \mathbb{Z}} \hat{f}_{u}(r) e^{2 \pi \mathrm{i} r x}
$$

where the Fourier coefficients are given by

$$
\begin{aligned}
\hat{f}_{u}(r) & =\int_{0}^{1} e^{-2 \pi \mathrm{i} r x} f_{u}(x) d x \\
& =\int_{0}^{1} \sum_{m \in \mathbb{Z}} u_{m} e^{-2 \pi \mathrm{i} r x} g\left(x-\frac{m}{s}\right) d x \\
& =\int_{0}^{1} \sum_{j=0}^{s-1} \sum_{q \in \mathbb{Z}} u_{j} e^{-2 \pi \mathrm{i} r x} g\left(x-q-\frac{j}{s}\right) d x \\
& =\sum_{j=0}^{s-1} u_{j} \int_{-\infty}^{\infty} e^{-2 \pi \mathrm{i} r x} g\left(x-\frac{j}{s}\right) d x \\
& =\sum_{j=0}^{s-1} u_{j} e^{-2 \pi \mathrm{i} r j / s} \int_{-\infty}^{\infty} e^{-2 \pi \mathrm{i} r x} g(x) d x
\end{aligned}
$$

Using the well-known fact that the Fourier transform of $g(x)$ on $\mathbb{R}$ is given by

$$
\int_{-\infty}^{\infty} e^{-2 \pi \mathrm{i} y x} g(x) d x=\alpha s^{-1} e^{-\pi \alpha^{2} s^{-2} y^{2}}, \quad y \in \mathbb{R}
$$

we obtain

$$
\hat{f}_{u}(r)=\alpha e^{-\pi \alpha^{2} s^{-2} r^{2}}\left(\mathcal{F}_{s} u\right)_{r}, \quad r \in \mathbb{Z}
$$

By definition $(\mathcal{S} u)_{\ell}=\alpha^{-1} f_{u}(\ell / t)$ for any $\ell \in \mathbb{Z}$, so for any $k \in\{0, \ldots, t-1\}$ we have

$$
\begin{aligned}
\left(\mathcal{P}_{t} \mathcal{F}_{t} \mathcal{S} u\right)_{k} & =\left(\mathcal{F}_{t} \mathcal{S} u\right)_{-s k} \\
& =\alpha^{-1} t^{-1} \sum_{\ell=0}^{t-1} e^{2 \pi \mathrm{i} s k \ell / t} f_{u}(\ell / t) \\
& =\alpha^{-1} t^{-1} \sum_{\ell=0}^{t-1} e^{2 \pi \mathrm{i} s k \ell / t} \sum_{r \in \mathbb{Z}} \hat{f}_{u}(r) e^{2 \pi \mathrm{i} r \ell / t} \\
& =\alpha^{-1} \sum_{r=-s k \bmod t} \hat{f}_{u}(r) \\
& =\sum_{r=-s k \bmod t} e^{-\pi \alpha^{2} s^{-2} r^{2}}\left(\mathcal{F}_{s} u\right)_{r} \\
& =\sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2} s^{-2}(t j-s k)^{2}}\left(\mathcal{F}_{s} u\right)_{t j-s k} \\
& =\sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2} t^{2}\left(\frac{j}{s}-\frac{k}{t}\right)^{2}}\left(\mathcal{P}_{s} \mathcal{F}_{s} u\right)_{j} \\
& =\left(\mathcal{T} \mathcal{P}_{s} \mathcal{F}_{s} u\right)_{k} .
\end{aligned}
$$

Remark 4.3. Another interpretation of the above proof is that the measure $f_{u}(x) d x$ is the convolution of the measures $\sum_{j=0}^{s-1} u_{j} \delta_{j / s}$ and $\sum_{j \in \mathbb{Z}} g(x-j) d x$ on $\mathbb{R} / \mathbb{Z}$, where $\delta_{x}$ means a unit mass concentrated at $x$. The key point is that the Fourier transform maps convolution of measures to pointwise multiplication of their Fourier coefficients.

Remark 4.4. If the hypothesis $\operatorname{gcd}(s, t)=1$ is not satisfied, then the results of Sections 4.1-4.2 still hold as stated. However, in this situation Theorem 4.2 is of no use, as the $\operatorname{map} \mathcal{P}_{s}$ fails to be invertible. For further discussion, see Section 4.4.1.

We conclude this section with a straightforward bound for $\|\mathcal{S}\|$.
Lemma 4.5. We have $\|\mathcal{S}\|<1+\alpha^{-1}$.
Proof. For any $u \in \mathbb{C}_{\circ}^{s}$ and $k \in\{0, \ldots, t-1\}$, we have

$$
\left|(\mathcal{S} u)_{k}\right| \leqslant \alpha^{-1} \sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{-2} s^{2}\left(\frac{k}{t}-\frac{j}{s}\right)^{2}}=\alpha^{-1} \sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{-2}\left(j-\frac{s k}{t}\right)^{2}}=\alpha^{-1} \sum_{j \in \mathbb{Z}} G(\eta+j),
$$

where $\eta:=\left\langle\frac{-s k}{t}\right\rangle \in\left[-\frac{1}{2}, \frac{1}{2}\right]$ and $G(x):=e^{-\pi \alpha^{-2} x^{2}}$.
First suppose that $\eta \in\left[-\frac{1}{2}, 0\right]$. Then $G(x)$ is increasing on $(-\infty, \eta)$ and decreasing on $(\eta+1, \infty)$, so

$$
\int_{-\infty}^{\eta} G(x) d x>\sum_{j=-\infty}^{-1} G(\eta+j), \quad \int_{\eta+1}^{\infty} G(x) d x>\sum_{j=2}^{\infty} G(\eta+j)
$$

For the remaining interval $(\eta, \eta+1)$, we observe that $G(x) \geqslant G(\eta+1)$ for $x \in$ $(0, \eta+1)$ and $G(x) \geqslant G(\eta)$ for $x \in(\eta, 0)$; but we have additionally $G(\eta) \geqslant G(\eta+1)$
because $|\eta| \leqslant \frac{1}{2} \leqslant|\eta+1|$, so in fact $G(x) \geqslant G(\eta+1)$ on the whole interval $(\eta, \eta+1)$. This implies that $\int_{\eta}^{\eta+1} G(x) d x \geqslant G(\eta+1)$, and adding the three integrals yields

$$
\sum_{j \in \mathbb{Z}} G(\eta+j)=G(\eta)+\sum_{j \neq 0} G(\eta+j)<1+\int_{-\infty}^{\infty} G(x) d x=1+\alpha
$$

A symmetrical argument yields the same bound for the case $\eta \in\left[0, \frac{1}{2}\right]$. We conclude that $\left|(\mathcal{S} u)_{k}\right|<\alpha^{-1}(1+\alpha)=1+\alpha^{-1}$, and hence $\|\mathcal{S}\|<1+\alpha^{-1}$.
4.2. Solving the system. We wish to use Theorem 4.2 to express $\mathcal{F}_{s}$ in terms of $\mathcal{F}_{t}$. To do this, we must show how to solve a system of the form $\mathcal{T} x=y$. This system is overdetermined, as $t>s$. For fixed $\alpha$, it turns out that the system is numerically unstable if $t / s$ is too close to 1 , or in other words, if the quantity $\theta:=t / s-1$ is too close to zero. On the other hand, we will show that imposing the condition $\theta \geqslant 1 / \alpha^{2}$ is enough to ensure that the system becomes numerically tractable, and in this case we may even construct an explicit left inverse for $\mathcal{T}$.

We begin by reducing from a rectangular to a square system. Consider the function $\ell \mapsto[t \ell / s]$, which maps $\{0, \ldots, s-1\}$ (injectively) into $\{0, \ldots, t-1\}$. We use this to define a map $\mathcal{C}: \mathbb{C}^{t} \rightarrow \mathbb{C}^{s}$ by the formula

$$
(\mathcal{C} u)_{\ell}:=u_{[t \ell / s]}, \quad u \in \mathbb{C}^{t}, \quad 0 \leqslant \ell<s
$$

We then set

$$
\mathcal{T}^{\prime}:=\mathcal{C T}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s}
$$

Note that the matrix of $\mathcal{T}^{\prime}$ is obtained by deleting $t-s$ rows from the matrix of $\mathcal{T}$. If we can show that $\mathcal{T}^{\prime}$ is invertible, then a left inverse for $\mathcal{T}$ is given by $\left(\mathcal{T}^{\prime}\right)^{-1} \mathcal{C}$.

The entries of $\mathcal{T}^{\prime} u$ are given explicitly by

$$
\begin{aligned}
\left(\mathcal{T}^{\prime} u\right)_{\ell}=(\mathcal{T} u)_{[t \ell / s]}=\sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2} t^{2}\left(\frac{1}{t}\left[\frac{t \ell}{s}\right]-\frac{j}{s}\right)^{2}} u_{j} & =\sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2}\left(\frac{t j}{s}-\left[\frac{t \ell}{s}\right]\right)^{2}} u_{j} \\
& =\sum_{h \in \mathbb{Z}} e^{-\pi \alpha^{2}\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}} u_{\ell+h}
\end{aligned}
$$

where

$$
\beta_{\ell}:=\frac{t \ell}{s}-\left[\frac{t \ell}{s}\right]=\left\langle\frac{t \ell}{s}\right\rangle, \quad \ell \in \mathbb{Z} .
$$

Observe that $\beta_{\ell}$ is periodic in $\ell$, i.e., $\beta_{\ell}=\beta_{\ell^{\prime}}$ if $\ell \equiv \ell^{\prime}(\bmod s)$, and that $\left|\beta_{\ell}\right| \leqslant \frac{1}{2}$ for all $\ell$.

We normalise $\mathcal{T}^{\prime}$ as follows. Let $\mathcal{D}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s}$ be the diagonal map defined by $(\mathcal{D} u)_{\ell}:=d_{\ell} u_{\ell}$, where $d_{\ell}:=e^{\pi \alpha^{2} \beta_{\ell}^{2}}$ for $\ell \in \mathbb{Z}$. Since $\beta_{\ell}^{2} \leqslant \frac{1}{4}$ we have $1 \leqslant d_{\ell} \leqslant e^{\pi \alpha^{2} / 4}$, and in particular

$$
\begin{equation*}
\|\mathcal{D}\| \leqslant e^{\pi \alpha^{2} / 4} \tag{4.1}
\end{equation*}
$$

Define

$$
\mathcal{N}:=\mathcal{T}^{\prime} \mathcal{D}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s}
$$

In other words, the matrix of $\mathcal{N}$ is obtained by multiplying the $\ell$-th column of the matrix of $\mathcal{T}^{\prime}$ by $d_{\ell}$. Explicitly,

$$
(\mathcal{N} u)_{\ell}=\sum_{h \in \mathbb{Z}} e^{-\pi \alpha^{2}\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}} d_{\ell+h} u_{\ell+h}=\sum_{h \in \mathbb{Z}} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)} u_{\ell+h}
$$

$\left[\begin{array}{cccccccccc}9.6 \mathrm{e}-923 & 1.9 \mathrm{e}-9 & 9.6 \mathrm{e}-37 & 1.1 \mathrm{e}-83 & 4.4 \mathrm{e}-148 & 1.2 \mathrm{e}-229 & 4.4 \mathrm{e}-148 & 1.1 \mathrm{e}-83 & 9.6 \mathrm{e}-37 & 1.9 \mathrm{e}-9 \\ 3.5 \mathrm{e}-6 & 1.8 \mathrm{e}-880 & 8.0 \mathrm{e}-14 & 1.4 \mathrm{e}-46 & 8.9 \mathrm{e}-97 & 1.9 \mathrm{e}-164 & 2.7 \mathrm{e}-210 & 1.0 \mathrm{e}-131 & 1.4 \mathrm{e}-70 & 4.2 \mathrm{e}-29 \\ 7.6 \mathrm{e}-50 & 5.2 \mathrm{e}-16 & 2.8 \mathrm{e}-866 & 4.3 \mathrm{e}-5 & 6.4 \mathrm{e}-27 & 3.2 \mathrm{e}-66 & 3.0 \mathrm{e}-126 & 9.6 \mathrm{e}-204 & 5.3 \mathrm{e}-171 & 3.8 \mathrm{e}-101 \\ 4.8 \mathrm{e}-88 & 5.1 \mathrm{e}-40 & 1.5 \mathrm{e}-10 & 7.4 \mathrm{e}-909 & 2.3 \mathrm{e}-8 & 1.8 \mathrm{e}-33 & 2.6 \mathrm{e}-79 & 1.3 \mathrm{e}-142 & 2.2 \mathrm{e}-223 & 1.5 \mathrm{e}-153 \\ 3.7 \mathrm{e}-137 & 6.0 \mathrm{e}-75 & 2.7 \mathrm{e}-31 & 2.8 \mathrm{e}-7 & 1.1 \mathrm{e}-894 & 1.2 \mathrm{e}-11 & 2.7 \mathrm{e}-43 & 2.1 \mathrm{e}-92 & 5.4 \mathrm{e}-159 & 7.6 \mathrm{e}-217 \\ 3.4 \mathrm{e}-197 & 1.5 \mathrm{e}-177 & 1.6 \mathrm{e}-105 & 4.1 \mathrm{e}-53 & 3.4 \mathrm{e}-18 & 4.3 \mathrm{e}-852 & 5.3 \mathrm{e}-4 & 9.7 \mathrm{e}-25 & 6.1 \mathrm{e}-63 & 8.6 \mathrm{e}-121 \\ 3.7 \mathrm{e}-137 & 7.6 \mathrm{e}-217 & 5.4 \mathrm{e}-159 & 2.1 \mathrm{e}-92 & 2.7 \mathrm{e}-43 & 1.2 \mathrm{e}-11 & 1.1 \mathrm{e}-894 & 2.8 \mathrm{e}-7 & 2.7 \mathrm{e}-31 & 6.0 \mathrm{e}-75 \\ 4.8 \mathrm{e}-88 & 1.5 \mathrm{e}-153 & 2.2 \mathrm{e}-223 & 1.3 \mathrm{e}-142 & 2.6 \mathrm{e}-79 & 1.8 \mathrm{e}-33 & 2.3 \mathrm{e}-8 & 7.4 \mathrm{e}-909 & 1.5 \mathrm{e}-10 & 5.1 \mathrm{e}-40 \\ 7.6 \mathrm{e}-50 & 3.8 \mathrm{e}-101 & 5.3 \mathrm{e}-171 & 9.6 \mathrm{e}-204 & 3.0 \mathrm{e}-126 & 3.2 \mathrm{e}-66 & 6.4 \mathrm{e}-27 & 4.3 \mathrm{e}-5 & 2.8 \mathrm{e}-866 & 5.2 \mathrm{e}-16 \\ 3.5 \mathrm{e}-6 & 4.2 \mathrm{e}-29 & 1.4 \mathrm{e}-70 & 1.0 \mathrm{e}-131 & 2.7 \mathrm{e}-210 & 1.9 \mathrm{e}-164 & 8.9 \mathrm{e}-97 & 1.4 \mathrm{e}-46 & 8.0 \mathrm{e}-14 & 1.8 \mathrm{e}-880\end{array}\right]$

Figure 3. Matrix of $\mathcal{E}$ for $s=10, t=13, \alpha=2$.

In this last expression, the $h=0$ term is simply $u_{\ell}$. Therefore, setting $\mathcal{E}:=\mathcal{N}-\mathcal{I}$, where $\mathcal{I}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s}$ is the identity map, we have

$$
(\mathcal{E} u)_{\ell}=\sum_{h \in \mathbb{Z} \backslash\{0\}} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)} u_{\ell+h}, \quad u \in \mathbb{C}^{s}, \quad 0 \leqslant \ell<s
$$

An example of the matrix of $\mathcal{E}$ is shown in Figure 3.
The following estimate is crucial for establishing left-invertibility of $\mathcal{T}$ and for obtaining a fast algorithm for solving the system $\mathcal{T} x=y$.

Lemma 4.6. Assume that $\alpha^{2} \theta \geqslant 1$. Then

$$
\|\mathcal{E}\|<2.01 \cdot e^{-\pi \alpha^{2} \theta / 2}<2^{-\alpha^{2} \theta}
$$

Proof. For any $u \in \mathbb{C}_{0}^{s}$, the above formula for $(\mathcal{E} u)_{\ell}$ implies that

$$
\left|(\mathcal{E} u)_{\ell}\right| \leqslant \sum_{h \in \mathbb{Z} \backslash\{0\}} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)}, \quad 0 \leqslant \ell<s
$$

Since $\left|\beta_{\ell}\right| \leqslant \frac{1}{2}<\frac{t}{2 s}$, we have $\left|\frac{t h}{s}+\beta_{\ell}\right| \geqslant\left|\frac{t h}{s}\right|-\left|\beta_{\ell}\right|>\frac{t}{s}\left(|h|-\frac{1}{2}\right)$. For $h \neq 0$ we have $|h|-\frac{1}{2} \geqslant \frac{1}{2}>0$, so

$$
\begin{aligned}
\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2} & >(t / s)^{2}\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4} \\
& =(1+\theta)^{2}\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4} \\
& \geqslant(1+2 \theta)\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4} \\
& =\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4}+2 \theta\left(|h|-\frac{1}{2}\right)^{2} \\
& \geqslant 2 \theta\left(|h|-\frac{1}{2}\right)^{2} .
\end{aligned}
$$

Therefore

$$
\left|(\mathcal{E} u)_{\ell}\right|<\sum_{h \in \mathbb{Z} \backslash\{0\}} e^{-2 \pi \alpha^{2} \theta\left(|h|-\frac{1}{2}\right)^{2}}=2\left(w^{1 / 4}+w^{9 / 4}+w^{25 / 4}+\cdots\right)
$$

where $w:=e^{-2 \pi \alpha^{2} \theta}$. Since $\alpha^{2} \theta \geqslant 1$ we have $w \leqslant e^{-2 \pi}<0.002$, so

$$
\left|(\mathcal{E} u)_{\ell}\right|<2 w^{1 / 4}\left(1+w^{2}+w^{6}+\cdots\right)<2.01 \cdot w^{1 / 4}=2.01 \cdot e^{-\pi \alpha^{2} \theta / 2}<2^{-\alpha^{2} \theta}
$$

where we have used the fact that $2.01 \cdot e^{-\pi x / 2}<2^{-x}$ for all $x \geqslant 1$.
Under the hypothesis of Lemma 4.6 , we see that $\|\mathcal{E}\|<\frac{1}{2}$, so $\mathcal{N}=\mathcal{I}+\mathcal{E}$ is invertible, with inverse given by $\mathcal{N}^{-1}=\mathcal{I}-\mathcal{E}+\mathcal{E}^{2}-\cdots$. Moreover, $\mathcal{D N}{ }^{-1} \mathcal{C}$ is the promised left inverse for $\mathcal{T}$, as

$$
\begin{equation*}
\left(\mathcal{D} \mathcal{N}^{-1} \mathcal{C}\right) \mathcal{T}=\mathcal{D} \mathcal{N}^{-1} \mathcal{T}^{\prime}=\mathcal{D}\left(\mathcal{T}^{\prime} \mathcal{D}\right)^{-1} \mathcal{T}^{\prime}=\mathcal{I} \tag{4.2}
\end{equation*}
$$

4.3. Proof of the theorem. We may now prove the following special case of Theorem 4.1.

Proposition 4.7 (Gaussian resampling in one dimension). Let $s$ and $t$ be integers such that $2 \leqslant s<t<2^{p}$ and $\operatorname{gcd}(s, t)=1$. Let $\alpha$ be an integer in the interval $2 \leqslant \alpha<p^{1 / 2}$. Let $\theta:=t / s-1>0$, and assume that $\theta \geqslant p / \alpha^{4}$. Then:
(i) There exist linear maps $\mathcal{A}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{t}$ and $\mathcal{B}: \mathbb{C}^{t} \rightarrow \mathbb{C}^{s}$ with $\|\mathcal{A}\|,\|\mathcal{B}\| \leqslant 1$ such that $\mathcal{F}_{s}=2^{2 \alpha^{2}} \mathcal{B} \mathcal{F}_{t} \mathcal{A}$.
(ii) We may construct numerical approximations $\tilde{\mathcal{A}}: \tilde{\mathbb{C}}^{s} \rightarrow \tilde{\mathbb{C}}^{t}$ and $\tilde{\mathcal{B}}: \tilde{\mathbb{C}}_{\mathrm{o}}^{t} \rightarrow \tilde{\mathbb{C}}^{s}$ such that $\varepsilon(\tilde{\mathcal{A}}), \varepsilon(\tilde{\mathcal{B}})<p^{2}$ and $\mathrm{C}(\tilde{\mathcal{A}}), \mathrm{C}(\tilde{\mathcal{B}})=O\left(t p^{3 / 2+\delta} \alpha+t p \log t\right)$.

Proof of (i). We apply the results of Sections 4.1-4.2 with the given $s, t$ and $\alpha$. Lemma 4.5 implies that $\|\mathcal{S}\|<\frac{3}{2}$. The hypotheses $\alpha<p^{1 / 2}$ and $\theta \geqslant p / \alpha^{4}$ imply that $\alpha^{2} \theta \geqslant 1$, so Lemma 4.6 yields $\|\mathcal{E}\|<2.01 \cdot e^{-\pi / 2}<0.42$. In particular, $\mathcal{N}=\mathcal{I}+\mathcal{E}$ is invertible and $\left\|\mathcal{N}^{-1}\right\|<1+0.42+(0.42)^{2}+\cdots<\frac{7}{4}$.

Let $\mathcal{J}:=\mathcal{N}^{-1}$, and define normalised maps

$$
\mathcal{S}^{\prime}:=\mathcal{S} / 2, \quad \mathcal{J}^{\prime}:=\mathcal{J} / 2, \quad \mathcal{D}^{\prime}:=\mathcal{D} / 2^{2 \alpha^{2}-2}
$$

Then $\left\|\mathcal{S}^{\prime}\right\|<\frac{3}{4}<1$ and $\left\|\mathcal{J}^{\prime}\right\|<\frac{7}{8}<1$. By (4.1) we have $\|\mathcal{D}\| \leqslant e^{\pi \alpha^{2} / 4}<2^{1.14 \alpha^{2}}<$ $2^{2 \alpha^{2}-2}$, as $1.14 x<2 x-2$ for all $x \geqslant 4$; hence also $\left\|\mathcal{D}^{\prime}\right\|<1$.

Now define

$$
\mathcal{A}:=\mathcal{S}^{\prime}, \quad \mathcal{B}:=\mathcal{P}_{s}^{-1} \mathcal{D}^{\prime} \mathcal{J}^{\prime} \mathcal{C} \mathcal{P}_{t}
$$

where $\mathcal{P}_{s}$ and $\mathcal{P}_{t}$ are as in Theorem 4.2, and $\mathcal{C}$ is as in Section 4.2. Note that $\mathcal{P}_{s}$ is invertible thanks to the hypothesis $\operatorname{gcd}(s, t)=1$. It is clear that $\left\|\mathcal{P}_{t}\right\|=\left\|\mathcal{P}_{s}^{-1}\right\|=$ $\|\mathcal{C}\|=1$, so $\|\mathcal{A}\|<1$ and $\|\mathcal{B}\|<1$. Moreover, by (4.2) and Theorem 4.2 we have

$$
\begin{aligned}
& 2^{2 \alpha^{2}} \mathcal{B} \mathcal{F}_{t} \mathcal{A}=\mathcal{P}_{s}^{-1}\left(2^{2 \alpha^{2}-2} \mathcal{D}^{\prime}\right)\left(2 \mathcal{J}^{\prime}\right) \mathcal{C} \mathcal{P}_{t} \mathcal{F}_{t}\left(2 \mathcal{S}^{\prime}\right) \\
&=\mathcal{P}_{s}^{-1}(\mathcal{D N} \\
&=\mathcal{P}_{s}^{-1}(\mathcal{C})\left(\mathcal{P}_{t} \mathcal{F}_{t} \mathcal{S}\right) \\
&\mathcal{C})\left(\mathcal{T} \mathcal{P}_{s} \mathcal{F}_{s}\right)=\mathcal{P}_{s}^{-1} \mathcal{P}_{s} \mathcal{F}_{s}=\mathcal{F}_{s}
\end{aligned}
$$

We break up the proof of (ii) into several lemmas. We begin with a straightforward algorithm for approximating $\mathcal{D}^{\prime}$ (Lemma 4.8). Next we give algorithms for approximating $\mathcal{S}^{\prime}=\mathcal{S} / 2$ and $\mathcal{E}$ (Lemmas 4.9 and 4.11 ); these amount to merely evaluating sufficiently many terms of the defining series, which converge quickly thanks to the rapid decay of the Gaussian weights. We then give an algorithm for approximating $\mathcal{J}^{\prime}=\mathcal{N}^{-1} / 2$, using the series $\mathcal{N}^{-1}=\mathcal{I}-\mathcal{E}+\mathcal{E}^{2}-\cdots$ (Lemma 4.12); here the fast convergence is guaranteed by the bound on $\|\mathcal{E}\|$ given in Lemma 4.6.
Lemma 4.8. Assume the hypotheses of Proposition 4.7. We may construct a numerical approximation $\tilde{\mathcal{D}}^{\prime}: \tilde{\mathbb{C}}_{\circ}^{s} \rightarrow \tilde{\mathbb{C}}_{\circ}^{s}$ for $\mathcal{D}^{\prime}$ such that $\varepsilon\left(\tilde{\mathcal{D}}^{\prime}\right)<4$ and $\mathrm{C}\left(\tilde{\mathcal{D}}^{\prime}\right)=$ $O\left(t p^{1+\delta}\right)$.
Proof. We are given as input $u \in \tilde{\mathbb{C}}_{0}^{s}$. For each $\ell \in\{0, \ldots, s-1\}$, by definition $\left(\mathcal{D}^{\prime} u\right)_{\ell}=d_{\ell}^{\prime} u_{\ell}$ where

$$
d_{\ell}^{\prime}:=d_{\ell} / 2^{2 \alpha^{2}-2}=e^{\pi \alpha^{2} \beta_{\ell}^{2}} / 2^{2 \alpha^{2}-2}<1
$$

(the last inequality follows from the estimate $\left\|\mathcal{D}^{\prime}\right\|<1$ in the proof of part (i) of Proposition 4.7). We may rewrite the rational part of the exponent of $d_{\ell}^{\prime}$ as $\alpha^{2} \beta_{\ell}^{2}=\alpha^{2}\left\langle\frac{t \ell}{s}\right\rangle^{2}=\alpha^{2} k_{\ell} / s^{2}$ for some non-negative integer $k_{\ell} \leqslant s^{2} / 4$. As $\alpha, s, t$ and $\ell$ are all integers with $O(p)$ bits (here we have used the hypotheses $s, t<2^{p}$
and $\alpha<p^{1 / 2}$ ), we may compute $\alpha^{2} k_{\ell}$ and $s^{2}$ in $O\left(p^{1+\delta}\right)$ bit operations. Feeding this as input to Lemma 2.14 (with $\sigma:=2 \alpha^{2}-2<2 p$ ), we obtain an approximation $\tilde{d}_{\ell}^{\prime} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon\left(\tilde{d}_{\ell}^{\prime}\right)<2$ in time $O\left(p^{1+\delta}\right)$. We then use Corollary 2.10 to compute an approximation $\tilde{z}_{\ell} \in \tilde{\mathbb{C}}_{\circ}$ for $z_{\ell}:=d_{\ell}^{\prime} u_{\ell}$ such that $\varepsilon\left(\tilde{z}_{\ell}\right)<\varepsilon\left(\tilde{d}_{\ell}^{\prime}\right)+2<4$ in time $O\left(p^{1+\delta}\right)$. Finally we set $\left(\tilde{\mathcal{D}}^{\prime} u\right)_{\ell}:=\tilde{z}_{\ell}$. The total cost over all $\ell$ is $O\left(s p^{1+\delta}\right)=$ $O\left(t p^{1+\delta}\right)$.

Lemma 4.9. Assume the hypotheses of Proposition 4.7. We may construct a numerical approximation $\tilde{\mathcal{S}}^{\prime}: \widetilde{\mathbb{C}}_{\circ}^{s} \rightarrow \tilde{\mathbb{C}}_{\circ}^{t}$ for $\mathcal{S}^{\prime}$ such that $\varepsilon\left(\tilde{\mathcal{S}}^{\prime}\right)<16$ p and $\mathrm{C}\left(\tilde{\mathcal{S}}^{\prime}\right)=$ $O\left(t p^{3 / 2+\delta} \alpha\right)$.
Proof. We are given as input $u \in \tilde{\mathbb{C}}_{0}^{s}$. For each $k=0, \ldots, t-1$ in turn, we approximate $\left(\mathcal{S}^{\prime} u\right)_{k}$ as follows. By definition

$$
\left(\mathcal{S}^{\prime} u\right)_{k}=\left(\frac{1}{2} \mathcal{S} u\right)_{k}=\sum_{j \in \mathbb{Z}} \frac{1}{2} \alpha^{-1} e^{-\pi \alpha^{-2}\left(j-\frac{s k}{t}\right)^{2}} u_{j}
$$

Let $m:=\left\lceil p^{1 / 2}\right\rceil \alpha$ and consider the truncated sum

$$
\begin{equation*}
T_{k}:=\sum_{\left|j-\frac{s k}{t}\right|<m} \frac{1}{2} \alpha^{-1} e^{-\pi \alpha^{-2}\left(j-\frac{s k}{t}\right)^{2}} u_{j} . \tag{4.3}
\end{equation*}
$$

Since $\|u\| \leqslant 1$ and $\alpha \geqslant 2$ we have

$$
\left|\left(\mathcal{S}^{\prime} u\right)_{k}-T_{k}\right| \leqslant \sum_{\left|j-\frac{s k}{t}\right| \geqslant m} \frac{1}{2} \alpha^{-1} e^{-\pi \alpha^{-2}\left(j-\frac{s k}{t}\right)^{2}}\left|u_{j}\right| \leqslant \frac{1}{4} \sum_{\left|j-\frac{s k}{t}\right| \geqslant m} e^{-\pi \alpha^{-2}\left(j-\frac{s k}{t}\right)^{2}} .
$$

Let $w:=e^{-\pi \alpha^{-2}}<1$; then

$$
\begin{aligned}
\left|\left(\mathcal{S}^{\prime} u\right)_{k}-T_{k}\right| & \leqslant \frac{1}{2}\left(w^{m^{2}}+w^{(m+1)^{2}}+w^{(m+2)^{2}}+\cdots\right) \\
& =\frac{1}{2} w^{m^{2}}\left(1+w^{2 m+1}+w^{4 m+4}+\cdots\right)
\end{aligned}
$$

Since $\alpha<p^{1 / 2}$ we have

$$
w^{m}=e^{-\pi\left\lceil p^{1 / 2}\right\rceil / \alpha} \leqslant e^{-\pi p^{1 / 2} / \alpha}<e^{-\pi}<0.05
$$

so certainly $1+w^{2 m+1}+w^{4 m+4}+\cdots<2$. We conclude that

$$
\left|\left(\mathcal{S}^{\prime} u\right)_{k}-T_{k}\right|<w^{m^{2}} \leqslant e^{-\pi\left\lceil p^{1 / 2}\right\rceil^{2}} \leqslant e^{-\pi p}<2^{-p}
$$

Now we explain how to compute a suitable fixed-point approximation for $T_{k}$. There are at most $2 m$ terms in the sum (4.3). Let $\beta:=\frac{1}{2} \alpha^{-1}$, and for each $j$ appearing in (4.3), let $x_{j}:=e^{-\pi \alpha^{-2}(j-s k / t)^{2}}, y_{j}:=\beta x_{j}, z_{j}:=y_{j} u_{j}$, so that $T_{k}=\sum_{j} z_{j}$. We first compute $\tilde{\beta}:=\rho(\beta)=2^{-p} \rho_{0}\left(2^{p-1} / \alpha\right) \in \widetilde{\mathbb{C}}_{\circ}$ in time $O\left(p^{1+\delta}\right)$; clearly $\varepsilon(\tilde{\beta})=2^{p}|\rho(\beta)-\beta|<1$ (as $\beta$ is real). Then for each $j$ we perform the following steps. As $s, t, j, k$ and $\alpha$ are all integers with $O(p)$ bits, the same holds for the numerator and the denominator of the rational number $\alpha^{-2}(j-s k / t)^{2}$, so we may use Lemma 2.13 to compute an approximation $\tilde{x}_{j} \in \widetilde{\mathbb{C}}_{\circ}$ with $\varepsilon\left(\tilde{x}_{j}\right)<2$ in time $O\left(p^{1+\delta}\right)$. We then use Corollary 2.10 to compute an approximation $\tilde{y}_{j} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon\left(\tilde{y}_{j}\right)<\varepsilon(\tilde{\beta})+\varepsilon\left(\tilde{x}_{j}\right)+2<5$, and again to obtain $\tilde{z}_{j} \in \tilde{\mathbb{C}}_{\mathrm{o}}$ such that $\varepsilon\left(\tilde{z}_{j}\right)<\varepsilon\left(\tilde{y}_{j}\right)+2<7$, in time $O\left(p^{1+\delta}\right)$. Finally, we form the sum $\tilde{T}_{k}:=\sum_{j} \tilde{z}_{j}$; that is, writing $\tilde{z}_{j}=2^{-p} a_{j}$ for integers $a_{j}$, we compute $\sum_{j} a_{j}$ and set $\tilde{T}_{k}:=2^{-p} \sum_{j} a_{j}$.

Defining $\left(\tilde{\mathcal{S}}^{\prime} u\right)_{k}:=\tilde{T}_{k}$, we must check that $\left|\tilde{T}_{k}\right| \leqslant 1$ (so that $\tilde{T}_{k} \in \tilde{\mathbb{C}}_{\mathrm{o}}$ ) and that $2^{p}\left|\tilde{T}_{k}-\left(\mathcal{S}^{\prime} u\right)_{k}\right|<16 p$. For the latter, observe that

$$
\begin{aligned}
2^{p}\left|\tilde{T}_{k}-\left(\mathcal{S}^{\prime} u\right)_{k}\right| & \leqslant 2^{p}\left|\tilde{T}_{k}-T_{k}\right|+2^{p}\left|T_{k}-\left(\mathcal{S}^{\prime} u\right)_{k}\right| \\
& <\left(\sum_{j} 2^{p}\left|\tilde{z}_{j}-z_{j}\right|\right)+1<(2 m) \cdot 7+1=14 m+1
\end{aligned}
$$

As $m<\left\lceil p^{1 / 2}\right\rceil p^{1 / 2} \leqslant p+p^{1 / 2}$ and $p \geqslant 100$, we find that

$$
2^{p}\left|\tilde{T}_{k}-\left(\mathcal{S}^{\prime} u\right)_{k}\right| \leqslant 14 p+14 p^{1 / 2}+1<16 p
$$

as desired. Recalling from the proof of Proposition $4.7(\mathrm{i})$ that $\left\|\mathcal{S}^{\prime}\right\|<\frac{3}{4}$, we also have

$$
\left|\tilde{T}_{k}\right| \leqslant\left|\tilde{T}_{k}-\left(\mathcal{S}^{\prime} u\right)_{k}\right|+\left|\left(\mathcal{S}^{\prime} u\right)_{k}\right|<16 p \cdot 2^{-p}+\left\|\mathcal{S}^{\prime}\right\|\|u\| \leqslant 10^{-26}+\frac{3}{4} \cdot 1<1
$$

The cost of the above procedure for each $k$ is $O\left(m p^{1+\delta}\right)=O\left(p^{3 / 2+\delta} \alpha\right)$, so the total over all $k$ is $O\left(t p^{3 / 2+\delta} \alpha\right)$.

Remark 4.10. The algorithm in the proof of Lemma 4.9 amounts to multiplying a vector by a matrix of the type shown in Figure 2, including only those entries that are numerically significant, which form a strip of width roughly $2 m$ around the diagonal. In the Turing model we must account for the cost of moving the tape head to access the input data needed to process each row, i.e., for row $k$ we must access those $u_{j \bmod s}$ such that $\left|j-\frac{s k}{t}\right|<m$. For most rows this is straightforward, as the relevant $u_{j}$ 's lie in a contiguous interval, and the $(k+1)$-th interval is obtained by shifting the $k$-th interval $O(1)$ cells to the right. However, for the rows near the top and bottom of the matrix, namely for $k<(t / s)(m-1)$ and $k>(t / s)(s-m)$, the relevant $u_{j}$ actually lie in two intervals separated by a gap of about $s-2 m$ cells. For example, when $k=0$, the relevant values are $u_{0}, \ldots, u_{m-1}$ and $u_{s-m+1}, \ldots, u_{s-1}$. As there are $O(m t / s)$ exceptional rows, the extra cost of jumping over these gaps is $O(m t p)=O\left(t p^{3 / 2} \alpha\right)$, which still fits within the target time bound. Similar remarks apply to the proof of Lemma 4.11 below.

Lemma 4.11. Assume the hypotheses of Proposition 4.7. We may construct a numerical approximation $\tilde{\mathcal{E}}: \widetilde{\mathbb{C}}_{\circ}^{s} \rightarrow \tilde{\mathbb{C}}_{\circ}^{s}$ for $\mathcal{E}$ such that $\varepsilon(\tilde{\mathcal{E}})<\frac{1}{3} p$ and $\mathrm{C}(\tilde{\mathcal{E}})=$ $O\left(t p^{3 / 2+\delta} \alpha^{-1}\right)$.

Proof. The argument is similar to the proof of Lemma 4.9. Given as input $u \in \tilde{\mathbb{C}}^{s}$, for each $\ell=0, \ldots, s-1$ we approximate $(\mathcal{E} u)_{\ell}$ as follows. By definition

$$
(\mathcal{E} u)_{\ell}=\sum_{h \in \mathbb{Z} \backslash\{0\}} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)} u_{\ell+h}
$$

As in the proof of Lemma 4.6, for $h \neq 0$ we have

$$
\begin{aligned}
\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2} & >(t / s)^{2}\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4} \\
& >\left(|h|-\frac{1}{2}\right)^{2}-\frac{1}{4}=|h|(|h|-1) \geqslant(|h|-1)^{2}
\end{aligned}
$$

Let $m:=\left\lceil\left(p / 4 \alpha^{2}\right)^{1 / 2}\right\rceil=\left\lceil p^{1 / 2} / 2 \alpha\right\rceil \geqslant 1$, and consider the truncated sum

$$
\begin{equation*}
T_{\ell}:=\sum_{\substack{h \in \mathbb{Z} \backslash\{0\} \\|h| \leqslant m}} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)} u_{\ell+h} . \tag{4.4}
\end{equation*}
$$

As $\|u\| \leqslant 1$ we have

$$
\left|(\mathcal{E} u)_{\ell}-T_{\ell}\right| \leqslant \sum_{|h|>m} e^{-\pi \alpha^{2}\left(\left(\frac{t h}{s}+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)}\left|u_{\ell+h}\right|<\sum_{|h|>m} e^{-\pi \alpha^{2}(|h|-1)^{2}}
$$

Let $w:=e^{-\pi \alpha^{2}} \leqslant e^{-4 \pi}<10^{-5}$; then $w^{m}<10^{-5}$ and

$$
\begin{aligned}
\left|(\mathcal{E} u)_{\ell}-T_{\ell}\right| & \leqslant 2\left(w^{m^{2}}+w^{(m+1)^{2}}+w^{(m+2)^{2}}+\cdots\right) \\
& =2 w^{m^{2}}\left(1+w^{2 m+1}+w^{4 m+4}+\cdots\right) \\
& <3 w^{m^{2}} \leqslant 3 e^{-\pi \alpha^{2}\left(p / 4 \alpha^{2}\right)}=3 e^{-\pi p / 4}<3 \cdot 2^{-p} .
\end{aligned}
$$

Now we explain how to approximate $T_{\ell}$. The sum (4.4) has exactly $2 m$ terms. For each $h$ appearing in (4.4), let $x_{h}:=e^{-\pi \alpha^{2}\left(\left(t h / s+\beta_{\ell}\right)^{2}-\beta_{\ell+h}^{2}\right)}$ and $z_{h}:=x_{h} u_{\ell+h}$, so that $T_{\ell}=\sum_{h} z_{h}$. As in the proof of Lemma 4.9, we may use Lemma 2.13 to compute an approximation $\tilde{x}_{h} \in \tilde{\mathbb{C}}_{\circ}$ such that $\varepsilon\left(\tilde{x}_{h}\right)<2$ in time $O\left(p^{1+\delta}\right)$. Using Corollary 2.10 , we then compute $\tilde{z}_{h} \in \tilde{\mathbb{C}}_{0}$ such that $\varepsilon\left(\tilde{z}_{h}\right)<\varepsilon\left(\tilde{x}_{h}\right)+2<4$. Finally we set $\tilde{T}_{\ell}:=\sum_{h} \tilde{z}_{h}$. We have

$$
\begin{aligned}
2^{p}\left|\tilde{T}_{\ell}-(\mathcal{E} u)_{\ell}\right| & \leqslant 2^{p}\left|\tilde{T}_{\ell}-T_{\ell}\right|+2^{p}\left|T_{\ell}-(\mathcal{E} u)_{\ell}\right| \\
& <\left(\sum_{h} 2^{p}\left|\tilde{z}_{h}-z_{h}\right|\right)+3<(2 m) \cdot 4+3=8 m+3
\end{aligned}
$$

As $m \leqslant\left(p^{1 / 2} / 2 \alpha\right)+1 \leqslant \frac{1}{4} p^{1 / 2}+1$ and $p \geqslant 100$, we find that

$$
2^{p}\left|\tilde{T}_{\ell}-(\mathcal{E} u)_{\ell}\right|<2 p^{1 / 2}+11<\frac{1}{3} p
$$

Recalling that $\|\mathcal{E}\|<0.42$ (see the proof of Proposition 4.7(i)), it follows that

$$
\left|\tilde{T}_{\ell}\right| \leqslant\left|\tilde{T}_{\ell}-(\mathcal{E} u)_{\ell}\right|+\left|(\mathcal{E} u)_{\ell}\right|<\frac{1}{3} p \cdot 2^{-p}+\|\mathcal{E}\|\|u\|<10^{-28}+0.42 \cdot 1<1
$$

so we may define $(\tilde{\mathcal{E}} u)_{\ell}:=\tilde{T}_{\ell} \in \tilde{\mathbb{C}}_{0}$. The cost of the above procedure for each $\ell$ is $O\left(m p^{1+\delta}\right)$. The hypothesis $\alpha<p^{1 / 2}$ implies that $m \leqslant \frac{1}{2} p^{1 / 2} \alpha^{-1}+1=O\left(p^{1 / 2} \alpha^{-1}\right)$, so the total cost over all $\ell$ is $O\left(t p^{3 / 2+\delta} \alpha^{-1}\right)$.

Lemma 4.12. Assume the hypotheses of Proposition 4.7. We may construct $a$ numerical approximation $\tilde{\mathcal{J}}^{\prime}: \widetilde{\mathbb{C}}_{\circ}^{s} \rightarrow \tilde{\mathbb{C}}_{\circ}^{s}$ for $\mathcal{J}^{\prime}$ such that $\varepsilon\left(\tilde{\mathcal{J}}^{\prime}\right)<\frac{3}{4} p^{2}$ and $\mathrm{C}\left(\tilde{\mathcal{J}}^{\prime}\right)=$ $O\left(t p^{3 / 2+\delta} \alpha\right)$.

Proof. We are given as input $u \in \tilde{\mathbb{C}}_{\mathrm{o}}^{s}$. Let $v:=u / 2 \in \mathbb{C}_{\mathrm{o}}^{s}$, and define $v^{(j)}:=\mathcal{E}^{j} v \in$ $\mathbb{C}_{\circ}^{s}$ for $j \geqslant 0$ (recall that $\|\mathcal{E}\|<0.42$ ). We wish to approximate

$$
\mathcal{J}^{\prime} u=\left(\mathcal{N}^{-1} / 2\right) u=\mathcal{N}^{-1} v=v-\mathcal{E} v+\mathcal{E}^{2} v-\cdots=v^{(0)}-v^{(1)}+v^{(2)}-\cdots .
$$

Let $n:=\left\lceil p / \alpha^{2} \theta\right\rceil=\left\lceil p s / \alpha^{2}(t-s)\right\rceil \geqslant 1$. We compute a sequence of approximations $\tilde{v}^{(0)}, \ldots, \tilde{v}^{(n-1)} \in \tilde{\mathbb{C}}_{\text {o }}^{s}$ as follows. First set $\tilde{v}^{(0)}:=\rho\left(v^{(0)}\right)=2^{-p} \rho_{0}\left(2^{p-1} u\right) \in$ $\tilde{\mathbb{C}}_{0}^{s}$, so that $\varepsilon\left(\tilde{v}^{(0)}\right)<2$. Then compute in sequence $\tilde{v}^{(j)}:=\tilde{\mathcal{E}} \tilde{v}^{(j-1)} \in \tilde{\mathbb{C}}_{0}^{s}$ for $j=1, \ldots, n-1$, using Lemma 4.11. We claim that $\varepsilon\left(\tilde{v}^{(j)}\right)<\frac{2}{3} p$ for each $j$. This is clear for $j=0$. For $j \geqslant 1$ it follows by induction, as

$$
\begin{aligned}
\varepsilon\left(\tilde{v}^{(j)}\right)=2^{p}\left\|\tilde{v}^{(j)}-v^{(j)}\right\| & =2^{p}\left\|\tilde{\mathcal{E}} \tilde{v}^{(j-1)}-\mathcal{E} v^{(j-1)}\right\| \\
& \leqslant 2^{p}\left\|\tilde{\mathcal{E}} \tilde{v}^{(j-1)}-\mathcal{E} \tilde{v}^{(j-1)}\right\|+2^{p}\left\|\mathcal{E} \tilde{v}^{(j-1)}-\mathcal{E} v^{(j-1)}\right\| \\
& \leqslant \varepsilon(\tilde{\mathcal{E}})+\|\mathcal{E}\| \varepsilon\left(\tilde{v}^{(j-1)}\right)<\frac{1}{3} p+\frac{1}{2} \cdot \frac{2}{3} p=\frac{2}{3} p
\end{aligned}
$$

Finally, define $\tilde{\mathcal{J}}^{\prime} u:=\tilde{v}^{(0)}-\tilde{v}^{(1)}+\cdots \pm \tilde{v}^{(n-1)}$. Then, as $\left\|v^{(j)}\right\| \leqslant\left\|\mathcal{E}^{j}\right\|\|v\| \leqslant \frac{1}{2}\|\mathcal{E}\|^{j}$ for all $j$, we have

$$
\begin{aligned}
2^{p}\left\|\tilde{\mathcal{J}}^{\prime} u-\mathcal{J}^{\prime} u\right\| \leqslant \sum_{j=0}^{n-1} 2^{p}\left\|\tilde{v}^{(j)}-v^{(j)}\right\|+2^{p} \sum_{j=n}^{\infty}\left\|v^{(j)}\right\| & <\frac{2}{3} n p+\frac{2^{p}\|\mathcal{E}\|^{n}}{2(1-\|\mathcal{E}\|)} \\
& <\frac{2}{3} n p+2^{p}\|\mathcal{E}\|^{n}
\end{aligned}
$$

We already saw in the proof of Proposition $4.7(\mathrm{i})$ that $\alpha^{2} \theta \geqslant 1$; this implies that $n \leqslant p$, and also (via Lemma 4.6) that $\|\mathcal{E}\|^{n}<2^{-\alpha^{2} \theta n} \leqslant 2^{-p}$. Thus

$$
2^{p}\left\|\tilde{\mathcal{J}}^{\prime} u-\mathcal{J}^{\prime} u\right\|<\frac{2}{3} p^{2}+1<\frac{3}{4} p^{2}
$$

(since $p \geqslant 100$ ). This also shows that $\tilde{\mathcal{J}}^{\prime} u \in \tilde{\mathbb{C}}^{s}$, as

$$
\left\|\tilde{\mathcal{J}}^{\prime} u\right\| \leqslant\left\|\mathcal{J}^{\prime} u\right\|+\left\|\tilde{\mathcal{J}}^{\prime} u-\mathcal{J}^{\prime} u\right\|<\left\|\mathcal{J}^{\prime}\right\|\|u\|+\frac{3}{4} p^{2} \cdot 2^{-p}<\frac{7}{8} \cdot 1+10^{-26}<1
$$

where the estimate $\left\|\mathcal{J}^{\prime}\right\|<\frac{7}{8}$ again comes from the proof of Proposition 4.7(i).
In the above algorithm, the bulk of the work consists of $n-1$ invocations of $\tilde{\mathcal{E}}$. The hypothesis $\theta \geqslant p / \alpha^{4}$ implies that

$$
n \leqslant \frac{p}{\alpha^{2} \theta}+1 \leqslant \alpha^{2}+1
$$

so the total cost is $O\left(t n p^{3 / 2+\delta} \alpha^{-1}\right)=O\left(t p^{3 / 2+\delta} \alpha\right)$.
Now we may complete the proof of Proposition 4.7.
Proof of Proposition 4.7 (ii). For $\tilde{\mathcal{A}}$ we simply take $\tilde{\mathcal{A}}:=\tilde{\mathcal{S}}^{\prime}$ where $\tilde{\mathcal{S}}^{\prime}$ is as described in Lemma 4.9; then $\varepsilon(\tilde{\mathcal{A}})<16 p<p^{2}$ (as $p \geqslant 100$ ), and $\mathrm{C}(\tilde{\mathcal{A}})=O\left(t p^{3 / 2+\delta} \alpha\right)$.

For $\tilde{\mathcal{B}}$ we take $\tilde{\mathcal{B}}:=\tilde{\mathcal{P}}_{s}^{-1} \tilde{\mathcal{D}}^{\prime} \tilde{\mathcal{J}}^{\prime} \tilde{\mathcal{C}} \tilde{\mathcal{P}}_{t}$, where $\tilde{\mathcal{D}}^{\prime}$ and $\tilde{\mathcal{J}}^{\prime}$ are as described in Lemmas 4.8 and 4.12 , and where $\tilde{\mathcal{C}}: \tilde{\mathbb{C}}_{o}^{t} \rightarrow \tilde{\mathbb{C}}_{0}^{s}, \tilde{\mathcal{P}}_{s}^{-1}: \tilde{\mathbb{C}}_{o}^{s} \rightarrow \tilde{\mathbb{C}}_{0}^{s}$ and $\tilde{\mathcal{P}}_{t}: \tilde{\mathbb{C}}_{\mathrm{o}}^{t} \rightarrow \tilde{\mathbb{C}}_{0}^{t}$ are the maps performing the obvious data rearrangements corresponding to $\mathcal{C}, \mathcal{P}_{s}^{-1}$ and $\mathcal{P}_{t}$, namely

$$
\begin{aligned}
(\tilde{\mathcal{C}} u)_{\ell}:=u_{[t \ell / s]}, & & u \in \tilde{\mathbb{C}}_{\circ}^{t}, & 0 \leqslant \ell<s \\
\left(\tilde{\mathcal{P}}_{s}^{-1} u\right)_{j}:=u_{\left(t^{-1} \bmod s\right) j}, & & u \in \tilde{\mathbb{C}}_{\circ}^{s}, & 0 \leqslant j<s \\
\left(\tilde{\mathcal{P}}_{t} u\right)_{k} & :=u_{-s k}, & & u \in \tilde{\mathbb{C}}_{\circ}^{t},
\end{aligned}
$$

These do not perform any arithmetic in $\tilde{\mathbb{C}}_{\circ}$ so $\varepsilon(\tilde{\mathcal{C}})=\varepsilon\left(\tilde{\mathcal{P}}_{s}^{-1}\right)=\varepsilon\left(\tilde{\mathcal{P}}_{t}\right)=0$. By Corollary 2.8 we obtain $\varepsilon(\tilde{\mathcal{B}}) \leqslant \varepsilon\left(\tilde{\mathcal{D}}^{\prime}\right)+\varepsilon\left(\tilde{\mathcal{J}}^{\prime}\right)<4+\frac{3}{4} p^{2}<p^{2}$.

As for the complexity, first observe that $\tilde{\mathcal{C}}$ simply copies its input in order, skipping $t-s$ unwanted entries, so $\mathrm{C}(\tilde{\mathcal{C}})=O(t p)$. To compute $\tilde{\mathcal{P}}_{s}^{-1} u$ for $u \in \tilde{\mathbb{C}}_{0}^{s}$, we use a "label-and-sort" algorithm: we first construct the list of ordered pairs $\left(t j \bmod s, u_{j}\right)$ for $j=0, \ldots, s-1$ in time $O(t p)$ (each label occupies $O(p)$ bits as $\left.s<2^{p}\right)$, then sort the list by the first entry using merge sort in time $O(t p \log t)$ [31], and finally extract the second entries to obtain the desired output in time $O(t p)$. Thus $\mathrm{C}\left(\tilde{\mathcal{P}}_{s}^{-1}\right)=O(t p \log t)$, and similarly $\mathrm{C}\left(\tilde{\mathcal{P}}_{t}\right)=O(t p \log t)$. Altogether we have

$$
\mathrm{C}(\tilde{\mathcal{B}})=\mathrm{C}\left(\tilde{\mathcal{D}}^{\prime}\right)+\mathrm{C}\left(\tilde{\mathcal{J}}^{\prime}\right)+O(t p \log t)=O\left(t p^{1+\delta}+t p^{3 / 2+\delta} \alpha+t p \log t\right)
$$

Finally we show how to deduce the general case from the one-dimensional case.

Proof of Theorem 4.1. Let $s_{1}, \ldots, s_{d}$ and $t_{1}, \ldots, t_{d}$ be as in the statement of the theorem. Applying Proposition 4.7 for each $i$, we obtain maps $\mathcal{A}_{i}: \mathbb{C}^{s_{i}} \rightarrow \mathbb{C}^{t_{i}}$ and $\mathcal{B}_{i}: \mathbb{C}^{t_{i}} \rightarrow \mathbb{C}^{s_{i}}$ with $\left\|\mathcal{A}_{i}\right\|,\left\|\mathcal{B}_{\tilde{\mathcal{B}}}\right\| \leqslant 1$ such that $\mathcal{F}_{s_{i}}=2^{2 \alpha^{2}} \mathcal{B}_{i} \mathcal{F}_{t_{i}} \mathcal{A}_{i}$, and approximations $\tilde{\mathcal{A}}_{i}: \tilde{\mathbb{C}}_{0}^{s_{i}} \rightarrow \tilde{\mathbb{C}}_{0}^{t_{i}}$ and $\tilde{\mathcal{B}}_{i}: \tilde{\mathbb{C}}_{0}^{t_{i}} \rightarrow \tilde{\mathbb{C}}_{0}^{s_{i}}$ such that $\varepsilon\left(\tilde{\mathcal{A}}_{i}\right), \varepsilon\left(\tilde{\mathcal{B}}_{i}\right)<p^{2}$ and $\mathrm{C}\left(\tilde{\mathcal{A}}_{i}\right), \mathrm{C}\left(\tilde{\mathcal{B}}_{i}\right)=O\left(t_{i} p^{3 / 2+\delta} \alpha+t_{i} p \log t_{i}\right)$.

Now observe that

$$
\mathcal{F}_{s_{1}, \ldots, s_{d}}=\otimes_{i} \mathcal{F}_{s_{i}}=2^{2 d \alpha^{2}}\left(\otimes_{i} \mathcal{B}_{i}\right)\left(\otimes_{i} \mathcal{F}_{t_{i}}\right)\left(\otimes_{i} \mathcal{A}_{i}\right)=2^{2 d \alpha^{2}} \mathcal{B} \mathcal{F}_{t_{1}, \ldots, t_{d}} \mathcal{A}
$$

where $\mathcal{A}:=\otimes_{i} \mathcal{A}_{i}: \otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ and $\mathcal{B}:=\otimes_{i} \mathcal{B}_{i}: \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$. Applying Lemma 2.11 (with $R:=\mathbb{C}, r:=1, m_{i}:=s_{i}, n_{i}:=t_{i}$ ), we may construct an approximation $\tilde{\mathcal{A}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ such that $\varepsilon(\tilde{\mathcal{A}}) \leqslant \sum_{i} \varepsilon\left(\tilde{\mathcal{A}}_{i}\right)<d p^{2}$. Moreover, let $M$ be as in Lemma 2.11; then $M=\prod_{i} \max \left(s_{i}, t_{i}\right)=t_{1} \cdots t_{d}=T$, so

$$
\begin{aligned}
\mathrm{C}(\tilde{\mathcal{A}}) & <T \sum_{i} \frac{\mathrm{C}\left(\tilde{\mathcal{A}}_{i}\right)}{t_{i}}+O(T p \log T) \\
& <T \sum_{i} O\left(p^{3 / 2+\delta} \alpha+p \log t_{i}\right)+O(T p \log T) \\
& =O\left(d T p^{3 / 2+\delta} \alpha+T p \log T\right)
\end{aligned}
$$

We may similarly construct an approximation $\tilde{\mathcal{B}}$ satisfying exactly the same error and cost bounds, and this completes the proof.
4.4. Further remarks. Our presentation of the Gaussian resampling technique has been optimised in favour of giving the simplest possible proof of the main $\mathrm{M}(n)=O(n \log n)$ bound. In this section we outline several ways in which these results may be improved and generalised, with an eye towards practical applications.
4.4.1. Minor technical issues. In our presentation we insisted that $\alpha$ be an integer and that $\alpha \geqslant 2$. Neither of these restrictions are essential; they were made for technical reasons to simplify certain proofs.

Similarly, the assumption $\operatorname{gcd}(s, t)=1$ is not necessary. We briefly outline what modifications must be made to handle the case $g:=\operatorname{gcd}(s, t)>1$. For $0 \leqslant h<g$, define maps $\mathcal{P}_{s, h}: \mathbb{C}^{s} \rightarrow \mathbb{C}^{s / g}, \mathcal{P}_{t, h}: \mathbb{C}^{t} \rightarrow \mathbb{C}^{t / g}$ and $\mathcal{T}_{h}: \mathbb{C}^{s / g} \rightarrow \mathbb{C}^{t / g}$ by

$$
\begin{aligned}
\left(\mathcal{P}_{s, h} u\right)_{j} & :=u_{t j+h}, & u \in \mathbb{C}^{s}, & 0 \leqslant j<s / g \\
\left(\mathcal{P}_{t, h} u\right)_{k} & :=u_{-s k+h}, & u \in \mathbb{C}^{t}, & 0 \leqslant k<t / g \\
\left(\mathcal{T}_{h} u\right)_{k} & :=\sum_{j \in \mathbb{Z}} e^{-\pi \alpha^{2} t^{2}\left(\frac{k}{t}-\frac{j}{s}-\frac{h}{s t}\right)^{2}} u_{j}, & & u \in \mathbb{C}^{s / g},
\end{aligned} \quad 0 \leqslant k<t / g .
$$

Then one may prove (analogously to Theorem 4.2) that $\mathcal{T}_{h} \mathcal{P}_{s, h} \mathcal{F}_{s}=\mathcal{P}_{t, h} \mathcal{F}_{t} \mathcal{S}$ for each $h$. In other words, for $u \in \mathbb{C}^{s}$, the matrix $\mathcal{T}_{h}$ gives a system of linear equations that relate the coefficients $\left(\mathcal{F}_{t} \mathcal{S} u\right)_{j}$ and $\left(\mathcal{F}_{s} u\right)_{j}$ for those $j$ congruent to $h$ modulo $g$. One may use this to prove an analogue of Theorem 4.1, by first constructing a left inverse for each $\mathcal{T}_{h}$ along the lines of Section 4.2.
4.4.2. Faster system solving. The iterative method used in Lemma 4.12 to approximate $\mathcal{J}=\mathcal{N}^{-1}$ (i.e., to solve the system $\mathcal{T} x=y$ ) has complexity $O\left(t p^{5 / 2+\delta} / \alpha^{3} \theta\right)$. To ensure that this step does not dominate the $O\left(p^{3 / 2+\delta} \alpha\right)$ complexity of approximating $\mathcal{S}$ (Lemma 4.9), we were compelled to introduce the hypothesis $\theta \geqslant p / \alpha^{4}$. On the other hand, to make the target DFT of length $t$ as cheap as possible, it
is desirable for $\theta$ to be as close to zero as possible. Together, these considerations imply that we cannot take $\alpha$ smaller than about $p^{1 / 4}$. (Indeed, in Section 5 , for fixed $d$, we do take $\alpha=\Theta\left(p^{1 / 4}\right)$ for this very reason.) For the choice $\alpha=\Theta\left(p^{1 / 4}\right)$, the overall complexity in Proposition 4.7 is $O\left(t p^{7 / 4+\delta}\right)$.

A better complexity bound may be obtained by precomputing an LU decomposition for $\mathcal{N}$, and then solving the system directly. The cost of the precomputation is $O\left(t p^{2+\delta} / \alpha^{2}\right)$ (assuming classical matrix arithmetic, while exploiting the circular banded structure), and then the cost of each application of $\mathcal{J}$ becomes $O\left(t p^{3 / 2+\delta} / \alpha\right)$. This allows us to relax the condition $\theta \geqslant p / \alpha^{4}$ to merely $\theta \geqslant 1 / \alpha^{2}$. Taking $\alpha=\Theta(1)$, the overall complexity in Proposition 4.7 (discounting the precomputation) falls to $O\left(t p^{3 / 2+\delta}\right)$. We did not use this method in our presentation because the error analysis is considerably more intricate than for the iterative method.

After making this modification, it would be interesting to investigate whether this method is competitive for practical computations of complex DFTs of length $s$ when $s$ is a large prime. One would choose a smooth transform length $t$ somewhat larger than $s$, say $1.25 s<t<1.5 s$, and use the algorithm to reduce the desired DFT of length $s$ to a DFT of length $t$; the latter could be handled via existing software libraries implementing the Cooley-Tukey algorithm. For large enough $s$, perhaps around $2^{20}$ or $2^{30}$, we expect that the invocations of $\tilde{\mathcal{S}}$ and $\tilde{\mathcal{J}}$ would be quite cheap compared to the FFT of length $t$. Indeed, $\tilde{\mathcal{S}}$ can be computed in a single pass over the input vector, and $\tilde{\mathcal{J}}$ in two passes (one for each of the L and U matrices), so they have excellent locality. It is conceivable that a highly optimised implementation could outperform existing software libraries, which handle transforms of prime length by techniques such as Rader's algorithm [38]. Such techniques introduce a large constant factor overhead that does not arise in the method just sketched.
4.4.3. Comparison with the Dutt-Rokhlin method. There is an enormous literature on "non-uniform FFTs" (sometimes called "non-equispaced FFTs"), going back to the seminal paper of Dutt and Rokhlin [11]. They consider transforms of the type

$$
\begin{equation*}
v_{j}:=\sum_{k=0}^{t-1} e^{2 \pi \mathrm{i} \omega_{k} y_{j}} u_{k}, \quad 0 \leqslant j<t \tag{4.5}
\end{equation*}
$$

The ordinary "uniform" DFT may be regarded as the special case where $\omega_{k}:=k$ and $y_{j}:=j / t$, but the Dutt-Rokhlin algorithms may be applied in cases where the frequencies $\omega_{k}$ are not necessarily integers, and/or the sample points $y_{j}$ are not necessarily integer multiples of $1 / t$. In these cases the algorithms reduce the problem to an ordinary FFT of length $t$ (or in some variants, a small multiple of $t$ ). The complexity, counting floating-point operations, is $O(t \log t+t p)$, where $p$ is the desired precision in bits.

If we now take instead $\omega_{k}:=k$ and $y_{j}:=j / s$, where $s$ is the "source" transform length, we see that (4.5) is exactly a DFT of length $s$ (apart from some inconsequential zero-padding), so the Dutt-Rokhlin algorithms may be used to compute a DFT of length $s$ by means of an FFT of length $t$. Inspection of their algorithms in this case reveals them to be essentially equivalent to our method in the special case that $\alpha=\Theta\left(p^{1 / 2}\right)$.

For example, consider [11, Algorithm 2], which corresponds roughly to a "transposed" version of our algorithm. Step 3 of that algorithm is analogous to approximating $\mathcal{S}$ (see Lemma 4.9). For the choice $\alpha=\Theta\left(p^{1 / 2}\right)$, the complexity for this step is $O\left(t p^{2+\delta}\right)$ bit operations, corresponding to the $O(t p)$ term in their complexity bound. Step 2 corresponds to our $\mathcal{F}_{t}$, and yields the $O(t \log t)$ term. The most interesting point of comparison is Step 1 , which corresponds roughly to solving the system $\mathcal{T} x=y$. The choice $\alpha=\Theta\left(p^{1 / 2}\right)$ implies that this system is essentially diagonal, i.e., the off-diagonal entries of $\mathcal{T}$ decay so rapidly that for numerical purposes they may be discarded. Solving the system is therefore trivial: their Step 1 consists of simply dividing each coefficient by the corresponding diagonal entry of $\mathcal{T}$ (in the literature these are often called "scale factors"). This step contributes only $O(t)$ floating-point operations.

The reason that Dutt and Rokhlin are (in effect) unable to take $\alpha$ smaller than about $p^{1 / 2}$ is essentially due to the approximation error committed when they truncate the Gaussian, for example in [11, Theorem 2.7]. Our Theorem 4.1 may be viewed as an "exact" replacement for that theorem. Rather than truncate the Gaussian, we take into account the effect of the Gaussian tail, which manifests as the off-diagonal entries of our $\mathcal{T}$ matrix. For $\alpha$ considerably smaller than $p^{1 / 2}$, these entries are numerically significant and cannot be ignored.

In our algorithm, assuming that we use the LU decomposition method mentioned in Section 4.4.2, as $\alpha$ decreases from $\Theta\left(p^{1 / 2}\right)$ to $\Theta(1)$ we see that the complexity of approximating $\mathcal{S}$ decreases from $O(t p)$ to $O\left(t p^{1 / 2}\right)$ floating-point operations, and the complexity of approximating $\mathcal{J}$ (i.e., solving the $\mathcal{T}$ system) increases from $O(t)$ to $O\left(t p^{1 / 2}\right)$ floating-point operations. When $\alpha=\Theta(1)$ they are balanced, and the overall complexity drops to $O\left(t \log t+t p^{1 / 2}\right)$; the last term improves on the DuttRokhlin bound by a factor of $p^{1 / 2}$. Note that the Dutt-Rokhlin bound is not strong enough for our application to integer multiplication; using their bound, the error term in Proposition 5.2 would grow to $O\left(n(\log n)^{1+\delta}\right)$ which is unacceptably large.

Of course, our discussion has only considered the case corresponding to the DFT of length $s$, i.e., the choice $y_{j}:=j / s$. An interesting question is whether the bound $O\left(t \log t+t p^{1 / 2}\right)$ can be proved for the general non-equispaced case, and if so, whether this method outperforms the Dutt-Rokhlin algorithm in practice.

## 5. The main algorithm

In this section we present the main integer multiplication algorithm. We actually give a family of algorithms, parameterised by a dimension parameter $d \geqslant 2$. Let

$$
n_{0}:=2^{d^{12}} \geqslant 2^{4096}
$$

and suppose that we wish to multiply integers with $n$ bits. For $n<n_{0}$, we may use any convenient base-case multiplication algorithm, such as the classical $O\left(n^{2}\right)$ algorithm. For $n \geqslant n_{0}$ we will describe a recursive algorithm that reduces the problem to a collection of multiplication problems of size roughly $n^{1 / d}$. We will show that this algorithm achieves $\mathrm{M}(n)=O(n \log n)$, provided that $d \geqslant 1729$.
5.1. Parameter selection. Henceforth we assume that $n \geqslant n_{0}$. We first discuss the computation of several parameters depending on $n$ that will be needed later.

Let

$$
\begin{equation*}
b:=\left\lceil\log _{2} n\right\rceil \geqslant d^{12} \geqslant 4096 \tag{5.1}
\end{equation*}
$$

be the "chunk size", and let the working precision be

$$
\begin{equation*}
p:=6 b=6\left\lceil\log _{2} n\right\rceil \geqslant 6 d^{12} \geqslant 24576>100 \tag{5.2}
\end{equation*}
$$

Define

$$
\begin{equation*}
\alpha:=\left\lceil\left(12 d^{2} b\right)^{1 / 4}\right\rceil \tag{5.3}
\end{equation*}
$$

Clearly $\alpha \geqslant 2$, and as $d \leqslant b^{1 / 12}$ and $b \geqslant 4096$, we also have

$$
\begin{equation*}
\alpha \leqslant\left\lceil 12^{1 / 4} b^{7 / 24}\right\rceil \leqslant 1.87 \cdot b^{7 / 24}+1<2 b^{7 / 24}<p^{1 / 2} \tag{5.4}
\end{equation*}
$$

As in Theorem 4.1, set

$$
\begin{equation*}
\gamma:=2 d \alpha^{2}<2 b^{1 / 12} \cdot 4 b^{7 / 12}=8 b^{2 / 3} \tag{5.5}
\end{equation*}
$$

Let $T$ be the unique power of two lying in the interval

$$
\begin{equation*}
4 n / b \leqslant T<8 n / b \tag{5.6}
\end{equation*}
$$

and let $r$ be the unique power of two in the interval

$$
\begin{equation*}
T^{1 / d} \leqslant r<2 T^{1 / d} \tag{5.7}
\end{equation*}
$$

We certainly have $b \leqslant 4 n^{1 / 2}$, so

$$
\begin{equation*}
r \geqslant T^{1 / d} \geqslant(4 n / b)^{1 / d} \geqslant\left(n^{1 / 2}\right)^{1 / d} \geqslant n^{1 / d^{2}} \geqslant n_{0}^{1 / d^{2}} \geqslant 2^{d^{10}} \tag{5.8}
\end{equation*}
$$

We now construct a factorisation $T=t_{1} \cdots t_{d}$ satisfying the hypotheses of Theorem 3.1. Let $d^{\prime}:=\log _{2}\left(r^{d} / T\right)$. As $T \leqslant r^{d}<2^{d} T$ we have $1 \leqslant r^{d} / T<2^{d}$ and hence $0 \leqslant d^{\prime}<d$. Define

$$
t_{1}, \ldots, t_{d^{\prime}}:=\frac{r}{2}, \quad t_{d^{\prime}+1}, \ldots, t_{d}:=r
$$

Then $t_{d} \geqslant \cdots \geqslant t_{1} \geqslant 2$ and

$$
t_{1} \cdots t_{d}=(r / 2)^{d^{\prime}} r^{d-d^{\prime}}=r^{d} / 2^{d^{\prime}}=T
$$

Also

$$
\begin{equation*}
T<8 n / b<n \leqslant 2^{b}<2^{p} \tag{5.9}
\end{equation*}
$$

so the hypotheses of Theorem 3.1 are indeed satisfied. The parameters $b, p, \alpha, \gamma$, $T, r$ and $t_{1}, \ldots, t_{d}$ may all be computed in time $(\log n)^{O(1)}$.

Our next task is to choose distinct primes $s_{1}, \ldots, s_{d}$ that are slightly smaller than the corresponding $t_{1}, \ldots, t_{d}$. In a moment we will use Theorem 4.1 to reduce a transform of size $s_{1} \times \cdots \times s_{d}$ to a transform of size $t_{1} \times \cdots \times t_{d}$; to avoid excessive data expansion in this reduction, we must ensure that the ratio $t_{1} \cdots t_{d} / s_{1} \cdots s_{d}$ is not too large. On the other hand, to satisfy the requirements of the theorem, we must also ensure that the individual ratios $t_{i} / s_{i}$ are not too close to 1 . We will achieve this by means of the following result.

Lemma 5.1. Let $\eta \in\left(0, \frac{1}{4}\right)$ and let $x \geqslant e^{2 / \eta}$. Then there are at least $\frac{1}{2} \eta x / \log x$ primes $q$ in the interval $(1-2 \eta) x<q \leqslant(1-\eta) x$.
Proof. Let $\vartheta(y):=\sum_{q \leqslant y} \log q$ (sum taken over primes) denote the usual Chebyshev function. According to [39, Thm. 4], for all $y \geqslant 563$ we have

$$
y-\frac{y}{2 \log y}<\vartheta(y)<y+\frac{y}{2 \log y}
$$

As the function $y / \log y$ is increasing for $y \geqslant 563$, we see that

$$
y-\frac{x}{2 \log x}<\vartheta(y)<y+\frac{x}{2 \log x}, \quad 563 \leqslant y \leqslant x
$$

Applying this result for $y_{0}:=(1-2 \eta) x>x / 2>e^{8} / 2>563$, and then again for $y_{1}:=(1-\eta) x>3 x / 4>563$, we obtain

$$
\vartheta\left(y_{1}\right)-\vartheta\left(y_{0}\right)>y_{1}-\frac{x}{2 \log x}-y_{0}-\frac{x}{2 \log x}=\eta x-\frac{x}{\log x} \geqslant \eta x-\frac{x}{2 / \eta}=\frac{\eta x}{2}
$$

and hence

$$
\sum_{y_{0}<q \leqslant y_{1}} 1 \geqslant \frac{1}{\log x} \sum_{y_{0}<q \leqslant y_{1}} \log q=\frac{\vartheta\left(y_{1}\right)-\vartheta\left(y_{0}\right)}{\log x}>\frac{\eta x}{2 \log x}
$$

Let us apply Lemma 5.1 with

$$
\begin{equation*}
\eta:=\frac{1}{4 d} \leqslant \frac{1}{8} \tag{5.10}
\end{equation*}
$$

and $x:=r / 2$, noting that (5.8) implies that $r / 2 \geqslant 2^{d^{10}-1} \geqslant e^{8 d}=e^{2 / \eta}$. We find that there are at least

$$
\frac{1}{8 d} \cdot \frac{r / 2}{\log (r / 2)} \geqslant \frac{1}{16 d} \cdot \frac{r}{\log r} \geqslant \frac{1}{16 d} \cdot \frac{2^{d^{10}}}{\log \left(2^{d^{10}}\right)}=\frac{2^{d^{10}}}{(16 \log 2) d^{11}} \geqslant d \geqslant d^{\prime}
$$

primes $q$ in the interval

$$
(1-2 \eta) \frac{r}{2}<q \leqslant(1-\eta) \frac{r}{2} .
$$

Using Eratosthenes' sieve, we may find $d^{\prime}$ such primes $s_{1}, \ldots, s_{d^{\prime}}$ in time $r^{1+o(1)}=$ $o(n)$. Applying Lemma 5.1 again, with the same $\eta$ but now with $x:=r \geqslant e^{2 / \eta}$, we find that there are at least $d \geqslant d-d^{\prime}$ primes $q$ in the interval

$$
(1-2 \eta) r<q \leqslant(1-\eta) r
$$

Again, we may find $d-d^{\prime}$ such primes $s_{d^{\prime}+1}, \ldots, s_{d}$ in time $o(n)$. These two collections of primes are disjoint, as

$$
(1-\eta) \frac{r}{2}<\frac{r}{2}<\frac{3 r}{4} \leqslant(1-2 \eta) r .
$$

In summary, we have found $d$ distinct primes $s_{1}, \ldots, s_{d}$ such that

$$
\begin{equation*}
(1-2 \eta) t_{i}<s_{i} \leqslant(1-\eta) t_{i}, \quad i \in\{1, \ldots, d\} \tag{5.11}
\end{equation*}
$$

Setting $S:=s_{1} \cdots s_{d}<T$, we see that

$$
\begin{equation*}
\frac{S}{T}>(1-2 \eta)^{d}=\left(1-\frac{1}{2 d}\right)^{d}>\frac{1}{2} \tag{5.12}
\end{equation*}
$$

as $(1-x)^{d}>1-d x$ for all $x \in(0,1)$.
5.2. DFTs and convolutions for coprime sizes. The following result combines Theorems 3.1 and 4.1 to obtain an approximation for the complex transform $\mathcal{F}_{s_{1}, \ldots, s_{d}}: \otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$ (defined in Section 2.4). Recall that the working precision was chosen to be $p:=6 b=6\left\lceil\log _{2} n\right\rceil$.
Proposition 5.2. We may construct a numerical approximation $\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}} \rightarrow$ $\otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ for $\mathcal{F}_{s_{1}, \ldots, s_{d}}$ such that $\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)<2^{\gamma+5} T \log _{2} T$ and

$$
\mathrm{C}\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)<\frac{4 T}{r} \mathrm{M}(3 r p)+O(n \log n)
$$

Proof. Let us verify the hypotheses of Theorem 4.1. We have already shown that $2 \leqslant s_{i}<t_{i}<2^{p}$ for all $i$ (see (5.9) and (5.11)) and that $2 \leqslant \alpha<p^{1 / 2}$ (see (5.4)). We certainly have $\operatorname{gcd}\left(s_{i}, t_{i}\right)=1$, as $s_{i}$ is an odd prime and $t_{i}$ is a power of two. Let $\theta_{i}:=t_{i} / s_{i}-1$; then by (5.10) and (5.11) we have

$$
\theta_{i} \geqslant \frac{1}{1-\eta}-1=\frac{\eta}{1-\eta}>\eta=\frac{1}{4 d}
$$

¿From (5.3) and (5.2) we see that

$$
\alpha^{4} \theta_{i} \geqslant \frac{12 d^{2} b}{4 d}=3 d b=\frac{d}{2} \cdot p \geqslant p
$$

so $\theta_{i} \geqslant p / \alpha^{4}$ as required.
Theorem 4.1 thus produces maps $\mathcal{A}: \otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{t_{i}}$ and $\mathcal{B}: \otimes_{i} \mathbb{C}^{t_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$ such that $\mathcal{F}_{s_{1}, \ldots, s_{d}}=2^{\gamma} \mathcal{B} \mathcal{F}_{t_{1}, \ldots, t_{d}} \mathcal{A}$ (where $\gamma$ is given by (5.5)), and approximations $\tilde{\mathcal{A}}: \otimes_{i} \widetilde{\mathbb{C}}_{o}^{s_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}}$ and $\tilde{\mathcal{B}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{t_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$. Applying Theorem 3.1 (whose hypotheses were verified in Section 5.1), we obtain furthermore an approximation $\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}$ for $\mathcal{F}_{t_{1}, \ldots, t_{d}}$. Now consider the scaled transform

$$
\mathcal{F}_{s_{1}, \ldots, s_{d}}^{\prime}:=2^{-\gamma} \mathcal{F}_{s_{1}, \ldots, s_{d}}=\mathcal{B} \mathcal{F}_{t_{1}, \ldots, t_{d}} \mathcal{A}
$$

and the approximation $\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}^{\prime}:=\tilde{\mathcal{B}} \tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}} \tilde{\mathcal{A}}$. The maps $\mathcal{A}, \mathcal{B}$ and $\mathcal{F}_{t_{1}, \ldots, t_{d}}$ all have norm at most 1 (by Theorem 4.1 and Example 2.6), so Corollary 2.8 implies that

$$
\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}^{\prime}\right) \leqslant \varepsilon(\tilde{\mathcal{B}})+\varepsilon\left(\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}\right)+\varepsilon(\tilde{\mathcal{A}})<2 d p^{2}+8 T \log _{2} T
$$

As $\left\|\mathcal{F}_{s_{1}, \ldots, s_{d}}\right\| \leqslant 1$, we obtain the desired approximation $\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}$ by applying Lemma 2.2 with $c:=2^{\gamma}$ to the output of $\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}^{\prime}$ (the condition $c \leqslant 2^{p}$ holds as $\gamma<8 b^{2 / 3}<p$; see (5.5)). We therefore obtain

$$
\begin{aligned}
\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right) & <2^{\gamma+1} \varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}^{\prime}\right)+3 \\
& <2^{\gamma+1}\left(2 d p^{2}+8 T \log _{2} T\right)+3 \\
& <2^{\gamma+1}\left(3 d p^{2}+8 T \log _{2} T\right)
\end{aligned}
$$

Moreover, by (5.1), (5.2) and (5.6) we have

$$
3 d p^{2} \leqslant 108 b^{1 / 12} b^{2}=108\left\lceil\log _{2} n\right\rceil^{25 / 12}<n^{1 / 2}<T<8 T \log _{2} T
$$

so we conclude that $\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)<2^{\gamma+5} T \log _{2} T$.
The cost of the scaling step is $O\left(S p^{1+\delta}\right)=O\left(T p^{1+\delta}\right)$, so by Theorem 3.1 and Theorem 4.1 the overall complexity is

$$
\begin{aligned}
\mathrm{C}\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right) & =\mathrm{C}(\tilde{\mathcal{A}})+\mathrm{C}\left(\tilde{\mathcal{F}}_{t_{1}, \ldots, t_{d}}\right)+\mathrm{C}(\tilde{\mathcal{B}})+O\left(T p^{1+\delta}\right) \\
& =\frac{4 T}{r} \mathrm{M}(3 r p)+O\left(d T p^{3 / 2+\delta} \alpha+T p \log T+T p^{1+\delta}\right)
\end{aligned}
$$

Recalling (5.4) and the assumption $\delta<\frac{1}{8}$ from Section 2.1, we see that

$$
d p^{3 / 2+\delta} \alpha=O\left(p^{1 / 12} p^{3 / 2} p^{1 / 8} p^{7 / 24}\right)=O\left(p^{2}\right)
$$

By definition $p=O(\log n)$, and (5.6) yields $T=O(n / \log n)$. The bound for $\mathrm{C}\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)$ thus simplifies to $(4 T / r) \mathrm{M}(3 r p)+O(n \log n)$.

Next we construct an approximation for the scaled convolution map $\mathcal{M}: \otimes_{i} \mathbb{C}^{s_{i}} \times$ $\otimes_{i} \mathbb{C}^{s_{i}} \rightarrow \otimes_{i} \mathbb{C}^{s_{i}}$ given by $\mathcal{M}(u, v):=\frac{1}{S} u * v$, where $*$ is the convolution operator defined in Section 2.4. Note that $\|\mathcal{M}\| \leqslant 1$.
Proposition 5.3. We may construct a numerical approximation $\tilde{\mathcal{M}}: \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}} \times$ $\otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}} \rightarrow \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ for $\mathcal{M}$ such that $\varepsilon(\tilde{\mathcal{M}})<2^{\gamma+8} T^{2} \log _{2} T$ and

$$
\begin{equation*}
\mathrm{C}(\tilde{\mathcal{M}})<\frac{12 T}{r} \mathrm{M}(3 r p)+O(n \log n) \tag{5.13}
\end{equation*}
$$

Proof. We are given as input $u, v \in \otimes_{i} \widetilde{\mathbb{C}}_{o}^{s_{i}}$. Let $w:=\mathcal{M}(u, v)=\frac{1}{S} u * v \in \otimes_{i} \mathbb{C}^{s_{i}}$. According to (2.1) we have $w=S w^{\prime}$ where $w^{\prime}:=\mathcal{F}_{s_{1}, \ldots, s_{d}}^{*}\left(\mathcal{F}_{s_{1}, \ldots, s_{d}} u \cdot \mathcal{F}_{s_{1}, \ldots, s_{d}} v\right)$. We use the following algorithm (essentially the same as in the proof of Proposition 3.4, but working over $\mathbb{C}$ instead of $\mathscr{R})$. We first compute an approximation $\tilde{w}^{\prime} \in$ $\otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ by using Proposition 5.2 to handle the forward and inverse transforms, and Corollary 2.10 to handle the pointwise multiplications. Applying Lemma 2.7 in the usual way, we obtain

$$
\begin{aligned}
\varepsilon\left(\tilde{w}^{\prime}\right) & \leqslant \varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)+\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}\right)+\varepsilon\left(\tilde{\mathcal{F}}_{s_{1}, \ldots, s_{d}}^{*}\right)+2 \\
& <3 \cdot 2^{\gamma+5} T \log _{2} T+2<\frac{7}{2} \cdot 2^{\gamma+5} T \log _{2} T .
\end{aligned}
$$

Then we apply Lemma 2.2 (with $c:=S \leqslant T \leqslant 2^{p}$, thanks to (5.9)) to obtain an approximation $\tilde{w} \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ such that

$$
\varepsilon(\tilde{w})<2 S \varepsilon\left(\tilde{w}^{\prime}\right)+3<7 S \cdot 2^{\gamma+5} T \log _{2} T+3<2^{\gamma+8} T^{2} \log _{2} T
$$

The cost of the pointwise multiplications and scalings is $O\left(S p^{1+\delta}\right)=O(n \log n)$, and the constant 12 accounts for the three invocations of Proposition 5.2.
5.3. Integer multiplication. We are now in a position to describe the recursive step of the main integer multiplication algorithm.

Proposition 5.4. For $n \geqslant n_{0}$ we have

$$
\begin{equation*}
\mathrm{M}(n)<\frac{12 T}{r} \mathrm{M}(3 r p)+O(n \log n) \tag{5.14}
\end{equation*}
$$

Proof. We are given as input integers $0 \leqslant f, g<2^{n}$. The algorithm consists of the following series of reductions.
(1) Reduce to one-dimensional convolution over $\mathbb{Z}$. Let $N:=\lceil n / b\rceil$ where $b:=$ $\left\lceil\log _{2} n\right\rceil$ as in (5.1). We split $f$ and $g$ into $N$ chunks of $b$ bits, i.e., we write $f=F\left(2^{b}\right)$ and $g=G\left(2^{b}\right)$ where

$$
F(x)=\sum_{j=0}^{N-1} F_{j} x^{j} \in \mathbb{Z}[x], \quad G(x)=\sum_{j=0}^{N-1} G_{j} x^{j} \in \mathbb{Z}[x], \quad 0 \leqslant F_{j}, G_{j}<2^{b}
$$

We have $f g=(F G)\left(2^{b}\right)$, so it suffices to compute the polynomial product $H(x):=$ $F(x) G(x)$ and then evaluate at $x=2^{b}$. The coefficients of $H(x)$ lie in the interval $0 \leqslant H_{j}<2^{2 b} N<2^{3 b}$; in particular, they have at most $3 b=O(\log n)$ bits, so the
evaluation may be achieved via a straightforward overlap-add algorithm in time $O(N \log n)=O(n)$. By (5.6) and (5.12) we have

$$
\operatorname{deg} H \leqslant 2 N-2 \leqslant \frac{2 n}{b} \leqslant \frac{T}{2}<S
$$

so it suffices to compute $F(x) G(x)\left(\bmod x^{S}-1\right)$.
(2) Reduce to d-dimensional convolution over $\mathbb{Z}$. We now use the AgarwalCooley method [1] to reduce to a multidimensional convolution. Consider the ring

$$
\mathscr{A}:=\mathbb{Z}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{s_{1}}-1, \ldots, x_{d}^{s_{d}}-1\right)
$$

As the integers $s_{1}, \ldots, s_{d}$ are pairwise relatively prime, there is an isomorphism of rings $\mathbb{Z}[x] /\left(x^{S}-1\right) \cong \mathscr{A}$ induced by the Chinese remainder theorem, namely by sending $x$ to $x_{1} \cdots x_{d}$. Let $F^{\prime}, G^{\prime}, H^{\prime} \in \mathscr{A}$ be the images of $F, G$ and $H$, so that $H^{\prime}=F^{\prime} G^{\prime}$. In the Turing model, $F^{\prime}, G^{\prime}$ and $H^{\prime}$ are represented as $d$-dimensional arrays of integers of $3 b=O(\log n)$ bits, using a similar layout to the tensor products in Section 2.3. The isomorphism amounts to a data rearrangement, and may be computed by attaching labels and sorting, as in the proof of Proposition 4.7(ii). The label $\left(i_{1} \bmod s_{1}, \ldots, i_{d} \bmod s_{d}\right)$ occupies $O\left(\sum_{i} \log s_{i}\right)=O(\log n)$ bits, and may be incremented to $\left(i_{1}+1 \bmod s_{1}, \ldots, i_{d}+1 \bmod s_{d}\right)$ in time $O(\log n)$. Therefore the isomorphism may be computed in either direction in time

$$
O(S \log n \log S)=O\left(T \log ^{2} n\right)=O(n \log n)
$$

(For an alternative algorithm that does not rely on sorting, see [23, Sec. 2.3].) We have thus reduced to the problem of computing $H^{\prime}=F^{\prime} G^{\prime}$ in $\mathscr{A}$.
(3) Reduce to approximate d-dimensional convolution over $\mathbb{C}$. Regarding $\mathscr{A}$ as a subring of

$$
\mathbb{C}\left[x_{1}, \ldots, x_{d}\right] /\left(x_{1}^{s_{1}}-1, \ldots, x_{d}^{s_{d}}-1\right) \cong\left(\otimes_{i} \mathbb{C}^{s_{i}}, *\right)
$$

let $F^{\prime \prime}, G^{\prime \prime}, H^{\prime \prime} \in \otimes_{i} \mathbb{C}^{s_{i}}$ be the elements corresponding to $F^{\prime}, G^{\prime}$ and $H^{\prime}$, so that $H^{\prime \prime}=F^{\prime \prime} * G^{\prime \prime}$. Let $u:=2^{-b} F^{\prime \prime}, v:=2^{-b} G^{\prime \prime}$ and $w:=\mathcal{M}(u, v)=\frac{1}{S} u * v$. Then $\|u\|,\|v\|,\|w\| \leqslant 1$, and $H^{\prime \prime}=2^{2 b} S w$. Recalling our choice of working precision $p:=6 b=6\left\lceil\log _{2} n\right\rceil$, we may use Proposition 5.3 to compute an approximation $\tilde{w}:=\tilde{\mathcal{M}}(u, v) \in \otimes_{i} \tilde{\mathbb{C}}_{o}^{s_{i}}$ such that $\varepsilon(\tilde{w})<2^{\gamma+8} T^{2} \log _{2} T$ in time $(12 T / r) \mathrm{M}(3 r p)+$ $O(n \log n)$.

Now observe that

$$
\left\|H^{\prime \prime}-2^{2 b} S \tilde{w}\right\|=2^{2 b} S\|w-\tilde{w}\| \leqslant 2^{2 b} T \cdot 2^{-p} \varepsilon(\tilde{w})<2^{2 b+\gamma+8-p} T^{3} \log _{2} T .
$$

Since $T<n \leqslant 2^{b}$ and $T \log _{2} T \leqslant T \log _{2} n \leqslant T b<8 n \leqslant 2^{b+3}$ (by (5.6)), this yields

$$
\left\|H^{\prime \prime}-2^{2 b} S \tilde{w}\right\|<2^{2 b+\gamma+8-p} \cdot 2^{2 b} \cdot 2^{b+3}=2^{5 b+\gamma+11-p}=2^{-b+\gamma+11}
$$

But (5.5) yields $\gamma<8 b^{2 / 3}<b-13$ (as $b \geqslant 4096$ ), so

$$
\left\|H^{\prime \prime}-2^{2 b} S \tilde{w}\right\|<\frac{1}{4}
$$

In particular, we may recover $H^{\prime \prime}$ in time $O\left(S p^{1+\delta}\right)=O(n \log n)$ by multiplying each coefficient of $\tilde{w}$ by $2^{2 b} S$ and then rounding to the nearest integer.

Corollary 5.5. Define $\mathrm{T}(n):=\mathrm{M}(n) /(n \log n)$ for $n \geqslant 2$. For $n \geqslant n_{0}$ we have

$$
\mathrm{T}(n)<\frac{1728}{d-\frac{1}{2}} \mathrm{~T}(3 r p)+O(1)
$$

Proof. Dividing (5.14) by $n \log n$ yields

$$
\mathrm{T}(n)<\frac{36 T p}{n} \cdot \frac{\log (3 r p)}{\log n} \cdot \mathrm{~T}(3 r p)+O(1)
$$

By (5.2) and (5.6) we have $36 T p / n=216 T b / n<1728$. Moreover, (5.7) implies that $r<2 T^{1 / d}<2 n^{1 / d}$, so
$\frac{\log (3 r p)}{\log n}=\frac{\log (r / 2)+\log (36 b)}{\log n}<\frac{1}{d}+\frac{\log (36 b)}{\log n}=\frac{1}{d}+\frac{\log _{2}(36 b)}{\log _{2} n} \leqslant \frac{1}{d}+\frac{\log _{2}(36 b)}{b-1}$.
Since $b \geqslant 4096$ and $d \leqslant b^{1 / 12}$ (see (5.1)) we have

$$
\frac{\log _{2}(36 b)}{b-1}<\frac{1}{2 b^{1 / 6}} \leqslant \frac{1}{2 d^{2}}
$$

and the result follows from the observation that

$$
\frac{1}{d}+\frac{1}{2 d^{2}}<\frac{1}{d}\left(1-\frac{1}{2 d}\right)^{-1}=\frac{1}{d-\frac{1}{2}}
$$

Finally we may prove the main result of the paper.
Proof of Theorem 1.1. According to Corollary 5.5, there is an absolute constant $A>0$ such that

$$
\mathrm{T}(n)<\frac{1728}{d-\frac{1}{2}} \cdot \mathrm{~T}(3 r p)+A
$$

for all $n \geqslant n_{0}$. We now take $d:=1729$. Then for all $n \geqslant n_{0}=2^{1729^{12}}$ we have

$$
\mathrm{T}(n)<0.9998 \cdot \mathrm{~T}(3 r p)+A
$$

Define

$$
B:=\max _{2 \leqslant n<n_{0}} \mathrm{~T}(n), \quad C:=\max (B, 5000 A)
$$

(Recall that for $n<n_{0}$, we use any convenient base-case multiplication algorithm to define $\mathrm{M}(n)$, and hence $\mathrm{T}(n)$.)

We prove by induction that $\mathrm{T}(n) \leqslant C$ for all $n \geqslant 2$. The choice of $B$ ensures that the statement holds for $n<n_{0}$. Now assume that $n \geqslant n_{0}$. By (5.7), (5.1) and (5.6) we have

$$
3 r p<6 T^{1 / d} p<36 n^{1 / d} b=36 n^{1 / 1729}\left\lceil\log _{2} n\right\rceil<n
$$

By induction,

$$
\mathrm{T}(n)<0.9998 C+A \leqslant 0.9998 C+0.0002 C=C
$$

Hence $\mathrm{T}(n)=O(1)$ and $\mathrm{M}(n)=O(n \log n)$.
5.4. Optimising the dimension threshold. It is possible to improve the factor $K=1728$ appearing in Corollary 5.5, at the expense of introducing various technical complications into the algorithm. In this section we outline a number of such modifications that together reduce the constant to $K=8+\epsilon$, so that the modified algorithm achieves $\mathrm{M}(n)=O(n \log n)$ for any $d \geqslant 9$ (rather than $d \geqslant 1729$ ). The techniques described here are similar to those used in [25] to optimise the value of $K$ in the Fürer-type bound $\mathrm{M}(n)=O\left(n \log n K^{\log ^{*} n}\right)$.
(1) We may increase the chunk size from $b=\Theta(\log n)$ to $b=\Theta(\log n \log \log n)$, and then take the working precision to be $p=2 b+O(\log n)=(2+o(1)) b$ rather than $6 b$. This improves $K$ by a factor of 3 . Note that the $O(\log n)$ term must be chosen large enough to ensure correct rounding at the end of the proof of Proposition 5.4. (We cannot take $b$ as large as $(\log n)^{2}$, as is done in [25], because then the Gaussian resampling becomes too expensive.)
(2) The choice of $b$ in the previous item allows us to improve the term $3 r p$ in Lemma 2.5 to $(2+o(1)) r p$, by packing the coefficients together more tightly in the Kronecker substitution step (the hypothesis $r<2^{p-1}$ must also be tightened somewhat). This improves $K$ by a factor of $3 / 2$.
(3) In Bluestein's algorithm (see proof of Proposition 3.1), the multiplicand $a$ is invariant, i.e., does not depend on the input vector $u$. To take advantage of this, we change the basic problem from multiplication of two arbitrary integers to multiplication of an arbitrary integer by a fixed integer. Consequently we save one forward transform in the proof of Proposition 5.3, reducing the factor 12 in (5.13) to 8 . This improves $K$ by a factor of $3 / 2$.
(4) We may choose the primes $s_{i}$ closer to $t_{i}$, so that instead of $T<2 S$ (see (5.12)) we have $T<(1+o(1)) S$. This improves $K$ by a factor of 2 . Some care is needed to avoid excessive precision loss in the Gaussian resampling step, due to the larger values of $\alpha$ and $\gamma$.
(5) By allowing $T$ to contain small odd prime factors, or alternatively by introducing flexibility into the choice of $b$, we may improve the choice of $T$ to $(4+o(1)) n / b$ (see (5.6)). This improves $K$ by a factor of 2 .
(6) We may change the basic problem from multiplication in $\mathbb{Z}$ to multiplication in $\mathbb{Z}[\mathrm{i}]$. In step (1) of the proof of Proposition 5.4 , the chunks $F_{j}$ and $G_{j}$ are taken in $\mathbb{Z}[\mathrm{i}]$ instead of $\mathbb{Z}$, and in step (1) of the proof of Lemma 2.5, the evaluations lie in $\mathbb{Z}[\mathrm{i}]$ instead of $\mathbb{Z}$. This improves $K$ by a factor of 4 , essentially by eliminating the factor 4 appearing in Lemma 2.5.
(7) We may change the basic problem from multiplication in $\mathbb{Z}[i]$ to multiplication in $\mathbb{Z}[\mathrm{i}] /\left(2^{n}+1\right) \mathbb{Z}[\mathrm{i}]$. Note that the Kronecker substitution in Lemma 2.5 maps the multiplication modulo $y^{r}+1$ naturally onto this problem. This improves $K$ by a factor of 2 , because it avoids the degree growth in step (1) in the proof of Proposition 5.4. It also introduces a technical difficulty into that step: to reduce a multiplication modulo $2^{n}+1$ to a polynomial multiplication modulo $x^{S}-1$ (or $x^{S}+1$ ), we must split an $n$-bit integer into $S$ chunks, even though $n$ will not in general be divisible by $S$. This may be addressed by means of the Crandall-Fagin algorithm [9].

## References

1. R. Agarwal and J. Cooley, New algorithms for digital convolution, IEEE Transactions on Acoustics, Speech, and Signal Processing 25 (1977), no. 5, 392-410.
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. J. M. Borwein and P. B. Borwein, Pi and the AGM, Canadian Mathematical Society Series of Monographs and Advanced Texts, John Wiley \& Sons, Inc., New York, 1987, A study in analytic number theory and computational complexity, A Wiley-Interscience Publication. MR 877728
4. 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)
5. R. P. Brent and P. Zimmermann, Modern Computer Arithmetic, Cambridge Monographs on Applied and Computational Mathematics, vol. 18, Cambridge University Press, Cambridge, 2011. MR 2760886
6. S. A. Cook and S. O. Aanderaa, On the minimum computation time of functions, Trans. Amer. Math. Soc. 142 (1969), 291-314. MR 0249212
7. 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
8. S. Covanov and E. Thomé, Fast integer multiplication using generalized Fermat primes, Math. Comp. 88 (2019), no. 317, 1449-1477. MR 3904152
9. R. Crandall and B. Fagin, Discrete weighted transforms and large-integer arithmetic, Math. Comp. 62 (1994), no. 205, 305-324. MR 1185244
10. A. De, P. Kurur, C. Saha, and R. Saptharishi, Fast integer multiplication using modular arithmetic, SIAM Journal on Computing 42 (2013), no. 2, 685-699.
11. A. Dutt and V. Rokhlin, Fast Fourier transforms for nonequispaced data, SIAM J. Sci. Comput. 14 (1993), no. 6, 1368-1393. MR 1241591
12. 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)
13. , Faster integer multiplication, SIAM J. Comput. 39 (2009), no. 3, 979-1005. MR 2538847 (2011b:68296)
14. J. von zur Gathen and J. Gerhard, Modern Computer Algebra, 3rd ed., Cambridge University Press, Cambridge, 2013. MR 3087522
15. P. Gaudry, A. Kruppa, and P. Zimmermann, A GMP-based implementation of SchönhageStrassen's large integer multiplication algorithm, ISSAC 2007, ACM, New York, 2007, pp. 167174. MR 2396199
16. I. J. Good, The interaction algorithm and practical Fourier analysis, J. Roy. Statist. Soc. Ser. B 20 (1958), 361-372. MR 0102888 (21 \#1674)
17. T. Granlund, The GNU Multiple Precision Arithmetic Library (Version 6.1.2), http:// gmplib.org/.
18. D. Harvey, Faster truncated integer multiplication, https://arxiv.org/abs/1703.00640, 2017.
19. D. Harvey and J. van der Hoeven, Faster integer and polynomial multiplication using cyclotomic coefficient rings, https://arxiv.org/abs/1712.03693, 2017.
20. ._ On the complexity of integer matrix multiplication, J. Symbolic Comput. 89 (2018), 1-8. MR 3804803
21. _ Faster integer multiplication using plain vanilla FFT primes, Math. Comp. 88 (2019), no. $315,501-514$. MR 3854069
22. , Faster integer multiplication using short lattice vectors, Proceedings of the Thirteenth Algorithmic Number Theory Symposium, Open Book Series 2 (R. Scheidler and J. Sorenson, eds.), Mathematical Sciences Publishers, Berkeley, 2019, pp. 293-310.
23. , Faster polynomial multiplication over finite fields using cyclotomic coefficient rings, J. Complexity 54 (2019), 101404, 18. MR 3983215
24. , Polynomial multiplication over finite fields in time $O(n \log n)$, HAL preprint http: //hal.archives-ouvertes.fr/hal-02070816, 2019.
25. D. Harvey, J. van der Hoeven, and G. Lecerf, Even faster integer multiplication, J. Complexity 36 (2016), 1-30. MR 3530637
26. _, Faster polynomial multiplication over finite fields, J. ACM 63 (2017), no. 6, 52:152:23.
27. D. R. Heath-Brown, Zero-free regions for Dirichlet L-functions, and the least prime in an arithmetic progression, Proc. London Math. Soc. (3) 64 (1992), no. 2, 265-338. MR 1143227 (93a:11075)
28. M. T. Heideman, D. H. Johnson, and C. S. Burrus, Gauss and the history of the fast Fourier transform, Arch. Hist. Exact Sci. 34 (1985), no. 3, 265-277. MR 815154 (87f:01018)
29. J. van der Hoeven, Faster relaxed multiplication, ISSAC 2014-Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation, ACM, New York, 2014, pp. 405-412. MR 3239953
30. , Faster Chinese remaindering, HAL preprint http://hal.archives-ouvertes.fr/ hal-01403810, 2016.
31. D. E. Knuth, The Art of Computer Programming. Vol. 3, Addison-Wesley, Reading, MA, 1998, Sorting and searching, Second edition [of MR0445948]. MR 3077154
32. H. J. Nussbaumer, Fast polynomial transform algorithms for digital convolution, IEEE Trans. Acoust. Speech Signal Process. 28 (1980), no. 2, 205-215. MR MR563380 (80m:94004)
33. H. J. Nussbaumer and P. Quandalle, Computation of convolutions and discrete Fourier transforms by polynomial transforms, IBM J. Res. Develop. 22 (1978), no. 2, 134-144. MR 0471260
34._, Fast computation of discrete Fourier transforms using polynomial transforms, IEEE Trans. Acoust. Speech Signal Process. 27 (1979), no. 2, 169-181. MR 523618
34. C. H. Papadimitriou, Computational Complexity, Addison-Wesley Publishing Company, Reading, MA, 1994. MR 1251285 (95f:68082)
35. M. S. Paterson, M. J. Fischer, and A. R. Meyer, An improved overlap argument for on-line multiplication, (1974), 97-111. SIAM-AMS Proc., Vol. VII. MR 0423875
36. J. M. Pollard, The fast Fourier transform in a finite field, Math. Comp. 25 (1971), 365-374. MR 0301966
37. C. M. Rader, Discrete Fourier transforms when the number of data samples is prime, Proc. IEEE 56 (1968), no. 6, 1107-1108.
38. J. B. Rosser and L. Schoenfeld, Approximate formulas for some functions of prime numbers, Illinois J. Math. 6 (1962), 64-94. MR 0137689
39. A. Schönhage and V. Strassen, Schnelle Multiplikation grosser Zahlen, Computing (Arch. Elektron. Rechnen) 7 (1971), 281-292. MR 0292344 (45 \#1431)
40. M. Sipser, Introduction to the Theory of Computation, 2nd ed., Thomson Course Technology, 2006.
41. L. H. Thomas, Using computers to solve problems in physics, Applications of digital computers 458 (1963), 42-57.
42. 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

Email address: d.harvey@unsw.edu.au
CNRS, Laboratoire D'informatique, École polytechnique, 91128 Palaiseau, France
Email address: vdhoeven@lix.polytechnique.fr


[^0]:    Harvey was supported by the Australian Research Council (grant FT160100219).

