

THE TRUNCATED FOURIER TRANSFORM AND APPLICATIONS

Joris van der Hoeven

Dépt. de Mathématiques (Bât. 425)
Université Paris-Sud
91405 Orsay Cedex
France
Email: joris@texmacs.org

January 9, 2004

In this paper, we present a truncated version of the classical Fast Fourier Transform. When applied to polynomial multiplication, this algorithm has the nice property of eliminating the “jumps” in the complexity at powers of two. When applied to the multiplication of multivariate polynomials or truncated multivariate power series, we gain a logarithmic factor with respect to the best previously known algorithms.

KEYWORDS: Fast Fourier Transform, jump phenomenon, truncated multiplication, FFT-multiplication, multivariate polynomials, multivariate power series.

A.M.S. SUBJECT CLASSIFICATION: 42-04, 68W25, 42B99, 30B10, 68W30.

1. INTRODUCTION

Let $\mathcal{R} \ni 1/2$ be an effective ring of constants (i.e. the usual arithmetic operations $+$, $-$ and \times can be carried out by algorithm). If \mathcal{R} has a primitive n -th root of unity with $n = 2^p$, then the product of two polynomials $P, Q \in \mathcal{R}[X]$ with $\deg PQ < n$ can be computed in time $O(n \log n)$ using the Fast Fourier Transform or FFT [CT65]. If \mathcal{R} does not admit a primitive n -th root of unity, then one needs an additional overhead of $O(\log \log n)$ in order to carry out the multiplication, by artificially adding new root of unity [SS71, CK91].

Besides the fact that the asymptotic complexity of the FFT involves a large constant factor, another classical drawback is that the complexity function admits important jumps at each power of two. These jumps can be reduced by using $(k 2^p)$ -th roots of unity for small k . They can also be smoothed by decomposing $(n + \delta) \times (n + \delta)$ -multiplications as $n \times n$ -, $n \times \delta$ - and $(n + \delta) \times \delta$ -multiplications. However, these tricks are not very elegant, cumbersome to implement, and they do not allow to completely eliminate the jump problem.

In section 3, we present a new kind of “Truncated Fourier Transform” or TFT, which allows for the fast evaluation of a polynomial $P \in \mathcal{R}[X]$ in any number n of well-chosen roots of unity. This algorithm coincides with the usual FFT if n is a power of two, but it behaves smoothly for intermediate values. In section 4, we also show that the inverse operation of interpolation can be carried out with the same complexity.

The TFT permits to speed up the multiplication of univariate polynomials with a constant factor between 1 and 2. In the case of multivariate polynomials, the repeated gain of such a constant factor leads to the gain of a non-trivial asymptotic factor. More precisely, assuming that \mathcal{R} admits sufficiently 2^p -th roots of unity, we will show in section 5 that the product of two multivariate polynomials $P, Q \in \mathcal{R}[z_1, \dots, z_d]$ can be computed in time $O(s \log s)$, where $s = \binom{r+d-1}{d}$ and $r = \deg P Q + 1$. The best previously known algorithm [CKL89], based on sparse polynomial multiplication, has time complexity $O(s \log^2 s)$.

In section 6 we finally give an algorithm for the multiplication of truncated multivariate power series. This algorithm, which has time complexity $O(s \log^2 s)$, again improves the best previously known algorithm [LS01] by a factor of $O(\log s)$. Moreover, both in the cases of multivariate polynomials and power series, we expect the corresponding constant factor to be much better.

2. THE FAST FOURIER TRANSFORM

Let \mathcal{R} be an effective ring of constants, $n = 2^p$ with $p \in \mathbb{N}$ and $\omega \in \mathcal{R}$ a primitive n -th root of unity (i.e. $\omega^{n/2} = -1$). The discrete Fast Fourier Transform (FFT) of an n -tuple $(a_0, \dots, a_{n-1}) \in \mathcal{R}^n$ (with respect to ω) is the n -tuple $(\hat{a}_0, \dots, \hat{a}_{n-1}) = \text{FFT}_\omega(a) \in \mathcal{R}^n$ with

$$\hat{a}_i = \sum_{j=0}^{n-1} a_j \omega^{ij}.$$

In other words, $\hat{a}_i = A(\omega^i)$, where $A \in \mathcal{R}[X]$ denotes the polynomial $A = a_0 + a_1 X + \dots + a_{n-1} X^{n-1}$.

The F.F.T can be computed efficiently using binary splitting: writing

$$(a_0, \dots, a_{n-1}) = (b_0, c_0, \dots, b_{n/2-1}, c_{n/2-1}),$$

we recursively compute the Fourier transforms of $(b_0, \dots, b_{n/2-1})$ and $(c_0, \dots, c_{n/2-1})$

$$\begin{aligned} \text{FFT}_{\omega^2}(b_0, \dots, b_{n/2-1}) &= (\hat{b}_0, \dots, \hat{b}_{n/2-1}); \\ \text{FFT}_{\omega^2}(c_0, \dots, c_{n/2-1}) &= (\hat{c}_0, \dots, \hat{c}_{n/2-1}). \end{aligned}$$

Then we have

$$\begin{aligned} \text{FFT}_\omega(a_0, \dots, a_{n-1}) &= (\hat{b}_0 + \hat{c}_0, \dots, \hat{b}_{n/2-1} + \hat{c}_{n/2-1} \omega^{n/2-1} \\ &\quad \hat{b}_0 - \hat{c}_0, \dots, \hat{b}_{n/2-1} - \hat{c}_{n/2-1} \omega^{n/2-1}). \end{aligned}$$

This algorithm requires $n p = n \log_2 n$ multiplications with powers of ω and $2 n p$ additions (or subtractions).

In practice, it is most efficient to implement an in-place variant of the above algorithm. We will denote by $[i]_p$ the bitwise mirror of i at length p (for instance, $[11]_5 = 26$). At step 0, we start with the vector

$$x_0 = (x_{0,0}, \dots, x_{0,n-1}) = (a_0, \dots, a_{n-1}).$$

At step $s \in \{1, \dots, p\}$, we set

$$\begin{pmatrix} x_{s,im+j} \\ x_{s,(i+1)m+j} \end{pmatrix} = \begin{pmatrix} 1 & \omega^{[i]_s m} \\ 1 & -\omega^{[i]_s m} \end{pmatrix} \begin{pmatrix} x_{s-1,im+j} \\ x_{s-1,(i+1)m+j} \end{pmatrix}. \quad (1)$$

for all $i \in \{0, 2, \dots, n/m - 2\}$ and $j \in \{0, \dots, m - 1\}$, where $m = 2^{p-s}$. Using induction over s , it can easily be seen that

$$x_{s,im+j} = (\text{FFT}_{\omega^m}(a_j, a_{m+j}, \dots, a_{n-m+j}))_{[i]_s},$$

for all $i \in \{0, \dots, n/m - 1\}$ and $j \in \{0, \dots, m - 1\}$. In particular,

$$\begin{aligned} x_{p,i} &= \hat{a}_{[i]_p} \\ \hat{a}_i &= x_{p,[i]_p} \end{aligned}$$

for all $i \in \{0, \dots, n - 1\}$. This algorithm of “repeated crossings” is illustrated in figure 1.




Figure 1. Schematic representation of a Fast Fourier Transform for $n = 16$. The black dots correspond to the $x_{s,i}$, the upper row being $(x_{0,0}, \dots, x_{0,15}) = (a_0, \dots, a_{15})$ and the lower row $(x_{4,0}, \dots, x_{4,15}) = (\hat{a}_0, \hat{a}_8, \hat{a}_4, \hat{a}_{12}, \dots, \hat{a}_{15})$.

A classical application of the FFT is the multiplication of polynomials $A = a_0 + \dots + a_{n-1} X^{n-1}$ and $B = b_0 + \dots + b_{n-1} X^{n-1}$. Assuming that $\deg AB < n$, we first evaluate A and B in $1, \omega, \dots, \omega^{n-1}$ using the FFT:

$$\begin{aligned} (A(1), \dots, A(\omega^{n-1})) &= \text{FFT}_{\omega}(a_0, \dots, a_{n-1}) \\ (B(1), \dots, B(\omega^{n-1})) &= \text{FFT}_{\omega}(b_0, \dots, b_{n-1}) \end{aligned}$$

We next compute the evaluations $(A(1)B(1), \dots, A(\omega^{n-1})B(\omega^{n-1}))$ of AB at $1, \dots, \omega^{n-1}$. We finally have to recover AB from these values using the inverse FFT. But the inverse FFT with respect to ω is nothing else as $1/n$ times the direct FFT with respect to ω^{-1} . Indeed, for all $(a_0, \dots, a_{n-1}) \in \mathcal{R}^n$ and all $i \in \{0, \dots, n - 1\}$, we have

$$\text{FFT}_{\omega^{-1}}(\text{FFT}_{\omega}(a))_i = \sum_{k=0}^{n-1} \sum_{j=0}^{n-1} a_j \omega^{(i-k)j} = n a_i, \quad (2)$$

since

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

whenever $i \neq k$. This yields a multiplication algorithm of time complexity $O(n \log n)$ in $\mathcal{R}[X]$, when assuming that \mathcal{R} admits enough primitive 2^p -th roots of unity. In the case that \mathcal{R} does not, then new roots of unity can be added artificially [SS71, CK91, vdH02] so as to yield an algorithm of time complexity $O(n \log n \log \log n)$.

3. THE TRUNCATED FOURIER TRANSFORM

The algorithm from the previous section has the disadvantage that n needs to be a power of two. If we want to multiply two polynomials $A, B \in \mathcal{R}[X]$ such that $\deg AB + 1 = n + \delta$, then we need to carry out the FFT at precision $2n$, thereby losing a factor of 2. This factor can be reduced using several tricks. For instance, one may decompose the $(n + \delta) \times (n + \delta)$ -product into an $n \times n$ product, an $n \times \delta$ -product and an $(n + \delta) \times \delta$ -product. This is efficient for small δ , but not very good if $\delta \approx n/2$. In the latter case, one may also use an FFT at precision $3n/2$, by using 3×3 -matrices at one step of the FFT computation. However, all these tricks of the trade require a large amount of hacking and one always continues to lose a non-trivial factor between 1 and 2.

The idea behind the Truncated Fourier Transform is to provide an efficient algorithm for the evaluation of polynomials in any number of distinct points. Moreover, the inverse operation of interpolation can be carried out with the same complexity. This technique will eliminate the “jumps” in the complexity of FFT multiplication.

So let $n = 2^p$, $l < n$ (usually, $l \geq n/2$) and let ω be a primitive n -th root of unity. Given an l -tuple (a_0, \dots, a_{l-1}) , we will evaluate the corresponding polynomial $A = a_0 + \dots + a_{l-1} X^{l-1}$ in $\omega^{[0]_p}, \omega^{[1]_p}, \dots, \omega^{[l-1]_p}$. We call $(A(\omega^{[0]_p}), \dots, A(\omega^{[l-1]_p}))$ the *Truncated Fourier Transform* (TFT) of (a_0, \dots, a_{l-1}) . Now consider the completion of the l -tuple (a_0, \dots, a_{l-1}) into an n -tuple $(a_0, \dots, a_{l-1}, 0, \dots, 0)$. When using the in-place algorithm from the previous section in order to compute $(A(\omega^{[0]_p}), \dots, A(\omega^{[l-1]_p}))$, we claim that many of the computations of the $x_{s,i}$ can actually be skipped (see figure 2). Indeed, at stage s , it suffices to compute the vector $(x_{s,0}, \dots, x_{s, \lfloor l/m \rfloor (m+1) - 1})$. Besides $x_{s,0}, \dots, x_{s,l-1}$, we therefore compute at most $m = 2^{p-s}$ additional values. In total, we therefore compute at most $pl + 2^{p-1} + 2^{p-2} + \dots + 1 < pl + n$ values $x_{s,i}$. This proves the following result:

THEOREM 1. *Let $n = 2^p$, $l < n$ and let $\omega \in \mathcal{R}$ be a primitive n -th root of unity in \mathcal{R} . Then the Truncated Fourier Transform of an l -tuple (a_0, \dots, a_{l-1}) w.r.t. ω can be computed using at most $lp + n$ additions (or subtractions) and $\lceil (lp + n)/2 \rceil$ multiplications with powers of ω .*




Figure 2. Schematic representation of a TFT for $n = 16$ and $l = 11$.

REMARK 2. Assume that \mathcal{R} admits a privileged primitive n -th root of unity ω_n for every $n \in 2^{\mathbb{N}}$, such that $\omega_{2^n}^2 = \omega_n$ for all n . Then the TFT $(\hat{a}_0, \dots, \hat{a}_{l-1})$ of an l -tuple (a_0, \dots, a_{l-1}) w.r.t. ω_n with $n \geq l$ does not depend on the choice of n . We call $(\hat{a}_0, \dots, \hat{a}_{l-1})$ the TFT of (a_0, \dots, a_{l-1}) w.r.t. the privileged sequence $(\omega_1, \omega_2, \omega_4, \dots)$ of roots of unity.

REMARK 3. Since the only operations we need for computing the TFT are additions, subtractions and multiplications by powers of ω , the algorithm naturally combines with Schönhage-Strassen's algorithm when ω is a symbolically added root of unity.

REMARK 4. If $f_0 = \dots = f_{l'-1}$, then the TFT of (f_0, \dots, f_{l-1}) can be computed using $O((l-l')p + 2n)$ operations using a similar algorithm as above. More generally, this allows the rapid transformation of “unions of segments”.

4. INVERTING THE TRUNCATED FOURIER TRANSFORM

Unfortunately, the inverse TFT cannot be computed using a similar formula as (2). Indeed, starting with the $x_{l,i}$, we need to compute an increasing number of $x_{s,i}$ when s decreases. Therefore we will rather invert the algorithm which computes the TFT, but with this difference that we will sometimes need $x_{s',i'}$ with $s' < s$ in order to compute $x_{s,i}$. We will use the fact that whenever one value among $x_{s,im+j}$, $x_{s-1,im+j}$ and one value among $x_{s,(i+1)m+j}$, $x_{s-1,(i+1)m+j}$ are known in the cross relation (1), then we can deduce the others from them using one multiplication by a power of ω and two “shifted” additions or subtractions (i.e. the results may have to be divided by 2).

More precisely, let us denote $k_s = \lfloor l/m \rfloor m$ and $l_s = k_s + m$ at each stage s . We use a recursive algorithm which takes the values $x_{p,k_s}, \dots, x_{p,l-1}$ and $x_{s,l}, \dots, x_{s,l_s}$ on input, and which computes $x_{s,k_s}, \dots, x_{s,l-1}$. If $s = p$, then we have nothing to do. Otherwise, we distinguish two cases:

- If $l_s = l_{s+1}$, then we first compute $x_{s,k_s}, \dots, x_{s,k_{s+1}-1}$ from $x_{p,k_s}, \dots, x_{p,k_{s+1}-1}$ using repeated crossings. We next deduce $x_{s,i}$ and $x_{s+1,i+m/2}$ from $x_{s+1,i}$ and $x_{s,i+m/2}$ for all $i \in \{l - m/2, \dots, k_{s+1} - 1\}$. Invoking our algorithm recursively, we now obtain $x_{s+1,k_{s+1}}, \dots, x_{s+1,l-1}$. We finally compute $x_{s,i}$ and $x_{s,i+m/2}$ from $x_{s+1,i}$ and $x_{s+1,i+m/2}$ for $i \in \{k_s, \dots, l - m/2 - 1\}$.
- If $l_s > l_{s+1}$, then we first compute $x_{s+1,i}$ from $x_{s,i}$ and $x_{s,i+m/2}$ for $i \in \{l, \dots, l_{s+1} - 1\}$. Invoking our algorithm recursively, we next compute $x_{s+1,k_{s+1}}, \dots, x_{s+1,l-1}$. We finally deduce $x_{s,i}$ from $x_{s+1,i}$ and $x_{s,i+m/2}$ for all $i \in \{k_s, \dots, l - 1\}$.

The two cases are illustrated in figures 3 resp. 4. Since $x_{0,l} = \dots = x_{0,n-1} = 0$, the application of our algorithm for $s = 0$ computes the inverse TFT. We notice that the values $x_{s,i}$ with $i < l$ are computed in decreasing order (for s) and the values $x_{s,i}$ with $i \geq l$ in increasing order. In other words, the algorithm may be designed in such a way to remain in place. We have proved:

THEOREM 5. *Let $n = 2^p$, $l < n$ and let $\omega \in \mathcal{R}$ be a primitive n -th root of unity in \mathcal{R} . Then the l -tuple (a_0, \dots, a_{l-1}) can be recovered from its Truncated Fourier Transform w.r.t. ω using at most $lp + n$ shifted additions (or subtractions) and $\lceil (lp + n)/2 \rceil$ multiplications with powers of ω .*

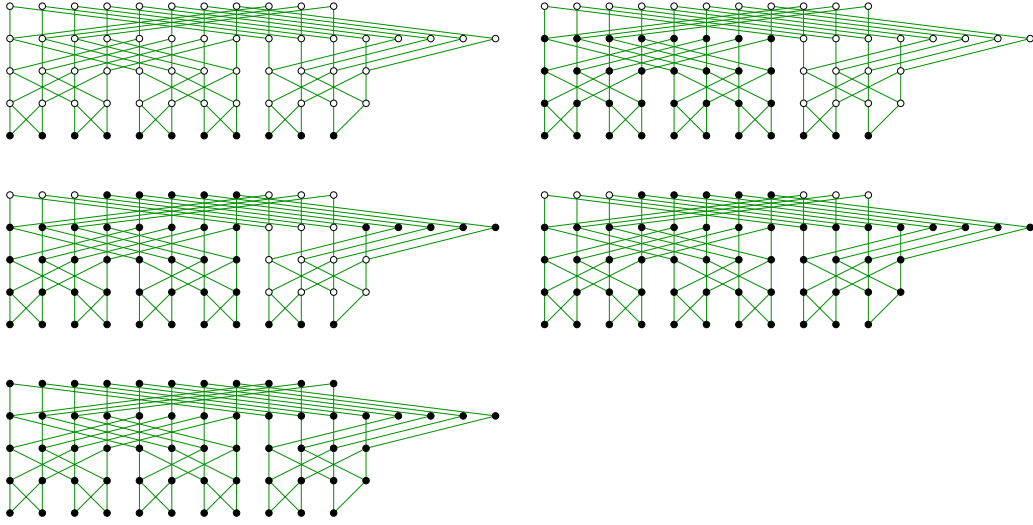


Figure 3. Schematic representation of the recursive computation of the inverse TFT for $n = 16$, $l = 11$. The different images show the progression of the known values $x_{i,j}$ (the black dots) during the different computations at stage $s = 0$. Since $l_0 = l_1 = 16$, we fall into the first case of our algorithm and the recursive invocation of the algorithm is done between the third and the fourth image.




Figure 4. Schematic representation of the recursive computation of the inverse TFT for $n = 16$, $l = 11$ at stage $s = 1$. Since $l_1 = 16$ and $l_2 = 12$, we now fall into the second case of our algorithm and the recursive invocation of the algorithm is done between the third and the fourth image.

5. MULTIPLYING MULTIVARIATE POLYNOMIALS

Let \mathcal{R} be a ring with a privileged sequence $(\omega_1, \omega_2, \omega_4, \dots)$ of roots of unity (see remark 2). Given a non-zero multivariate polynomial

$$f = \sum_{i_1, \dots, i_d} f_{i_1, \dots, i_d} z_1^{i_1} \cdots z_d^{i_d} \in \mathcal{R}[z_1, \dots, z_d],$$

in $d > 1$ variables, we define the *total degree* of f by

$$\deg f = \max \{i_1 + \cdots + i_d : f_{i_1, \dots, i_d} \neq 0\} \in \mathbb{N}$$

We let $\deg 0 = -1$. Now let $f, g \in \mathcal{R}[z_1, \dots, z_d]$ be such that $\deg fg < r$. In this section we present an algorithm to compute fg , which has a good complexity in terms of the number

$$s = \binom{r + d - 1}{d}$$

of expected coefficients of fg . When computing fg using the classical FFT with respect to each of the variables z_1, \dots, z_d , we need a time $O(d(2r)^d \log r)$ which is much bigger than s , in general. When using multiplication of sparse polynomials [CKL89], we need a time $O(s \log^2 s)$ with a very large constant factor. Our algorithm is based on the TFT w.r.t. all variables and we will show that it has a complexity $O(s \log s)$.

Given $f \in \mathcal{R}[z_1, \dots, z_d]$ with $\deg f < r$, the TFT of f with respect to one variable z_v at order r is defined by

$$\text{TFT}_{v;r}(f) = \left(f_{i_1, \dots, i_{v-1}, i_{v+1}, \dots, i_d} (z_v = \omega_v^{[i_v]n}) \right)_{i_1 + \dots + i_d < r},$$

where $n = 2^p \geq r$. We recall that the result does not depend on the choice of n . The TFT with respect to all variables z_1, \dots, z_d at order r is defined by

$$\text{TFT}_{;r}(f) = \left(f(\omega_n^{[i_1]n}, \dots, \omega_n^{[i_d]n}) \right)_{i_1 + \dots + i_d < r},$$

where $n = 2^p \geq r$ (see figure 5). We have

$$\text{TFT}_{;r}(f) = \text{TFT}_{d;r}(\dots \text{TFT}_{1;r}(f) \dots).$$

Given $f, g \in \mathcal{C}[z_1, \dots, z_d]$ with $\deg fg < r$, we will use the formula

$$fg = \text{TFT}_{;r}^{-1}(\text{TFT}_{;r}(f) \text{TFT}_{;r}(g))$$

in order to compute the product fg .

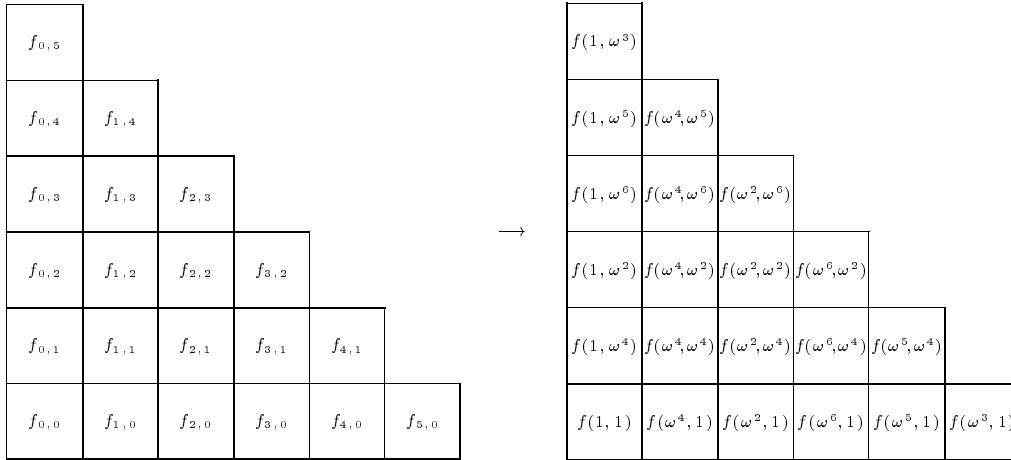


Figure 5. Illustration of the TFT in two variables ($\omega = \omega_8$).

In order to compute $\text{TFT}_{v;r}(f)$, say for $v = 1$, we compute the TFT of $(f_{0, i_2, \dots, i_d}, \dots, f_{l-1, i_2, \dots, i_d})$ with $l = r - i_2 - \dots - i_r$ for all i_2, \dots, i_d with $i_2 + \dots + i_n < r - 1$ (if $i_2 + \dots + i_n = r - 1$, then the TFT of (f_{0, i_2, \dots, i_d}) is given by itself, so we have nothing to do). One such computation takes a time $\leq Cl \log l$ for some universal constant C , by using the TFT w.r.t. ω_n with minimal $n = 2^p \geq l$ (so n may vary as a function of i_2, \dots, i_d , but not C). The computation of $\text{TFT}_{;r}(f)$ therefore takes a time $T_{d,r}$ with

$$T_{d,r} \leq C d \sum_{l=2}^r \binom{r-l+d-2}{d-2} l \log l.$$

Dividing by s , we obtain

$$\begin{aligned} \frac{T_{d,r}}{s} &\leq C d^2 (d-1) \sum_{l=2}^r \frac{(r+d-l-2)!}{(r+d-1)!} \frac{(r-1)!}{(r-l)!} l \log l \\ &\leq C d^3 \sum_{l=2}^r \frac{(r+d-l)!}{(r+d-1)!} \frac{(r-1)!}{(r-l)!} \frac{l \log l}{(r-l+d)(r-l+d-1)} \end{aligned} \tag{3}$$

If $r \leq d$, then the summand rapidly decreases when $l > 2$, so that

$$\frac{T_{d,r}}{s} = O(d^3 \frac{r}{(r+d)^3}) = O(r) = O(\log s).$$

Consequently, $T_{d,r} = O(s \log s)$ and even $T_{d,r} = O(s)$ for fixed r . If $r > d$, then for $d = \varepsilon r$ and $l = \delta r$, Stirling's formula yields

$$\log \frac{(r+d-l)! (r-1)!}{(r+d-1)! (r-l)!} = -\varepsilon \delta r + \dots$$

It follows that only the first $O(r/d) = O(1/\varepsilon)$ terms in (3) contribute to the asymptotic behaviour of $\frac{T_{d,r}}{s}$, so that

$$T_{d,r} = O(d^3 \frac{1}{\varepsilon} \frac{\log(1/\varepsilon)}{\varepsilon r^2}) = O(d \log(r/d)) = O(\log s).$$

Again, we find that $T_{d,r} = O(s \log s)$. We have proved:

THEOREM 6. *Let \mathcal{R} be a ring with a privileged sequence $(\omega_1, \omega_2, \omega_4, \dots)$ of roots of unity. Let $f, g \in \mathcal{R}[z_1, \dots, z_d]$ be polynomials with $\deg f + \deg g < r$ and let $s = \binom{r+d-1}{r}$. Then the product fg can be computed using $O(s \log s)$ operations in \mathcal{R} .*

6. MULTIPLYING MULTIVARIATE POWER SERIES

Since power series have infinitely many terms, implementing an operation on power series really corresponds to implementing the operation for polynomial approximations at all degrees. As usual, multiplication is a particularly important operation. Given $f, g \in C[z_1, \dots, z_d]$ with $\deg f < r$ and $\deg g < r$, we will show how to compute the truncated product $h = \sum_{i_1, \dots, i_d < d} (fg)_{i_1, \dots, i_d} z_1^{i_1} \dots z_d^{i_d} \in \mathcal{R}[z_1, \dots, z_d]$ of f and g .

The first idea [LS01] is to use homogeneous coordinates instead of the usual ones:

$$\begin{aligned} \tilde{f}(z_1, z_2, \dots, z_d) &= f(z_1, z_1 z_2, \dots, z_1 z_d) \\ \tilde{g}(z_1, z_2, \dots, z_d) &= g(z_1, z_1 z_2, \dots, z_1 z_d). \end{aligned}$$

This transformation takes no time since it corresponds to some re-indexing. We next compute the TFTs \hat{f} and \hat{g} in z_2, \dots, z_d at order r :

$$\begin{aligned} \hat{f} &= \text{TFT}_{d;r}(\dots \text{TFT}_{2;r}(\tilde{f}) \dots) \\ \hat{g} &= \text{TFT}_{d;r}(\dots \text{TFT}_{2;r}(\tilde{g}) \dots). \end{aligned}$$

We next compute the $s' = \binom{r+d-2}{d-1}$ truncated products $\hat{h}_{\cdot, i_2, \dots, i_d}(z_1)$ of the obtained polynomials $\hat{f}_{\cdot, i_2, \dots, i_d}(z_1)$ and $\hat{g}_{\cdot, i_2, \dots, i_d}(z_1)$. After transforming the results of these multiplication back using

$$\tilde{h} = \text{TFT}_{2;r}^{-1}(\dots \text{TFT}_{d;r}^{-1}(\hat{h}) \dots),$$

we obtain the truncated product h of f and g by

$$h(z_1, z_2, \dots, z_d) = \tilde{h}(z_1, z_2/z_1, \dots, z_d/z_1).$$

The total computation time is bounded by $O(r s' \log s' + r s' \log r)$. Using the fact that $r s' = O(s \log s)$, we have proved the following theorem:

THEOREM 7. *Let \mathcal{R} be a ring with a privileged sequence $(\omega_1, \omega_2, \omega_4, \dots)$ of roots of unity. Let $f, g \in \mathcal{R}[z_1, \dots, z_d]$ be polynomials of degrees $< r$ and let $s = \binom{r+d-1}{r}$. Then the truncated product of f and g at degree $< r$ can be computed using $O(s \log^2 s)$ operations in \mathcal{R} .*

REMARK 8. In practice, if the coefficients f_{i_1, \dots, i_d} have different growths in i_1, \dots, i_d , then it may be useful to consider truncations along more general degrees of the form

$$\deg_{\alpha} f = \max \{ \alpha_1 i_1 + \dots + \alpha_d i_d : f_{i_1, \dots, i_d} \neq 0 \}.$$

The “slicing technique” from section 6.3.5 in [vdH02] may then be used in order to obtain complexity bounds of the same type.

REMARK 9. Using remark 4, the polynomial and truncated multiplication algorithms can be used in combination with the strategy of relaxed evaluation [vdH97, vdH02, vdH03b] for solving partial differential equations in multivariate power series with an additional overhead of $O(\log r)$. A recent technique [vdH03a] allows to reduce this overhead even further and it would be interesting to study more precisely what happens in the multivariate case.

BIBLIOGRAPHY

- [CK91] D.G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991.
- [CKL89] J. Canny, E. Kaltofen, and Y. Lakshman. Solving systems of non-linear polynomial equations faster. In *Proc. ISSAC '89*, pages 121–128, Portland, Oregon, A.C.M., New York, 1989. ACM Press.
- [CT65] J.W. Cooley and J.W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Computat.*, 19:297–301, 1965.
- [HQZ00] Guillaume Hanrot, Michel Quercia, and Paul Zimmermann. Speeding up the division and square root of power series. Research Report 3973, INRIA, July 2000. Available from <http://www.inria.fr/RRRT/RR-3973.html>.
- [HQZ02] Guillaume Hanrot, Michel Quercia, and Paul Zimmermann. The middle product algorithm I. speeding up the division and square root of power series. Submitted, 2002.
- [HZ02] Guillaume Hanrot and Paul Zimmermann. A long note on Mulders' short product. Research Report 4654, INRIA, December 2002. Available from <http://www.loria.fr/~hanrot/Papers/mulders.ps>.
- [LS01] G. Lecerf and É. Schost. Fast multivariate power series multiplication in characteristic zero. Technical Report 2001-1, GAGE, École polytechnique, 91228 Palaiseau, France, 2001.
- [SS71] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. *Computing* 7, 7:281–292, 1971.
- [vdH97] J. van der Hoeven. Lazy multiplication of formal power series. In W. W. Kuchlin, editor, *Proc. ISSAC '97*, pages 17–20, Maui, Hawaii, July 1997.
- [vdH02] Joris van der Hoeven. Relax, but don't be too lazy. *JSC*, 34:479–542, 2002.
- [vdH03a] Joris van der Hoeven. New algorithms for relaxed multiplication. Technical Report 2003-44, Univ. d'Orsay, 2003.
- [vdH03b] Joris van der Hoeven. Relaxed multiplication using the middle product. In Manual Bronstein, editor, *Proc. ISSAC '03*, pages 143–147, Philadelphia, USA, August 2003.