Contents

1 Basic Arithmetic **2**

1.1 The School Method for Integer Multiplication 2

1.2 The Toom-Cook Algorithm 5

1.3 Approximate Computation 9

1.3.1 Fixed Point Arithmetic 9

1.3.2 Interval Arithmetic 12

1.3.3 Floating point arithmetic 15

1.4 Division 16

2 The Fast Fourier Transform and Fast Polynomial Arithmetic **19**

2.1 Schönhage-Strassen Multiplication 19

2.1.1 The Algorithm in a Nutshell 19

2.1.2 Fast Fourier Transform 21

2.1.3 Fast Multiplication in $\mathbb{Z}$ and $\mathbb{Z}[x]$. 26

2.1.4 Fast Multiplication over arbitrary Rings* 30

2.2 Fast Polynomial Division and Applications 32

2.3 Fast Polynomial Arithmetic in $\mathbb{C}[x]$ 38

Chapter 1

Basic Arithmetic

In this section, we present an efficient algorithm due to Toom and Cook for multiplying two integers, which already considerably improves upon the method that most people have learned in school. We further investigate in methods for carrying out approximate computations on fixed-point and floating-point numbers, and we derive bounds on the occurring error when using approximate instead of exact arithmetic. In addition, we introduce the concepts of interval arithmetic and box-functions and show that these concepts yield a powerful and very practical approach for carrying out approximate arithmetic. This is due to the fact that adaptive bounds on the error can directly be computed "on the fly", and that these bounds are often much better than any a priori bounds obtained by a worst-case error analysis. Finally, we give an efficient method to compute an arbitrary good approximation of the quotient of two integers or, more generally, two arbitrary complex values.

1.1 The School Method for Integer Multiplication

We represent integers $a \in \mathbb{Z}$ as digit strings with respect to a fixed base $B \in \mathbb{N}_{\geq 2}$. That is,

$$a = (-1)^s \cdot \sum_{i=0}^{n-1} a_i \cdot B^i, \text{ with } s \in \{0,1\} \text{ and } a_i \in \{0,\dots,B-1\} \text{ for all } i = 0,\dots,n-1.$$

We call the a_i's the *digits* and s the *sign digit* of a with respect to B. For convenience, we also write (if B is fixed)

$$a = (-1)^s \; a_{n-1}a_{n-2}\dots a_0$$

if the base B is fixed.

Example: Important bases are $B = 2, 10, 16$, and 2^k for some $k \in \mathbb{N}$. The integer 29 writes as

$$29 = 1 \cdot 2^0 + 0 \cdot 2^1 + 1 \cdot 2^2 + 1 \cdot 2^3 + 1 \cdot 2^4 = 11101.$$

The *length* (or *bitsize* for $B = 2$) of an integer a with respect to B is defined as the number of digits needed to represent a. For convenience, we use the term *n-digit number* to denote an integer of length n. Notice that any n-digit number can always be considered as an N-digit number for arbitrary $N \geq n$. This is advantageous in the analysis of many algorithms as it allows us to assume that the length of the input is a power of 2 (or some other value k).

Algorithm 1: School Method for Addition

Input : Two non-negative n-digit integers $a = a_{n-1} \dots a_0$ and $b = b_{n-1} \dots b_0$.
Output: An $(n+1)$-digit integer $c = c_n \dots c_0$ with $c = a + b$.

1 $\gamma_0 := 0$
2 **for** $i = 0, \dots, n-1$ **do**
3 Recursively define
4 $\gamma_{i+1} \cdot B + c_i = a_i + b_i + \gamma_i$ with $c_i, \gamma_i \in \{0, \dots, B-1\}$
5 $c_n := \gamma_n$
6 **return** $c_n \dots c_0$

We mainly consider two different ways of measuring the efficiency of an algorithm. The first one is to count the number of additions and multiplications between integers that an algorithm needs to return a result. This is referred to as the *arithmetic complexity* of an algorithm. Notice that the arithmetic complexity might be unrelated to the actual running time of an algorithm as the involved integers can be arbitrarily large. Hence, a more meaningful and precise way of measuring the efficiency of an algorithm is to count instead the number of *primitive operations* (or bit operations if the base B equals 2) that are carried out by the algorithm, often referred to as the *bit complexity* of an algorithm. Notice that the result of a primitive operations is always a one- or two-digit number.

Example: A prominent example is Gaussian elimination for solving a linear system in n unknowns. It is easy to see that the method uses $O(n^3)$ arithmetic operations, hence the arithmetic complexity of Gaussian elimination is polynomial in the input size. However, a straight forward analysis does NOT guarantee that the intermediate results as computed by the algorithm (which are rationals if the input matrix has integer entries) have size that is polynomial in the size of the input, thus it is not obvious that Gaussian elimination actually constitutes a polynomial time algorithm for solving linear systems. A more refined argument however shows that by recursively removing common factors of the intermediate results, it can be guaranteed that all intermediate results have polynomial size. We will go into more detail in one of the exercises. Later, we will also consider a different approach based on modular computation that does not come with any of these drawbacks.

We now review and analyze the school method for adding and multiplying two non-negative n-digit integers $a = a_{n-1} \dots a_0$ and $b = b_{n-1} \dots b_0$. We first start with addition; see Algorithm 1. The γ_i's are called *carries*. Using induction, it is easy to see that $\gamma_i \in \{0, 1\}$ for all i. Further notice that γ_{i+1} is non-zero if and only if the sum of the two digits a_i and b_i and the previous carry γ_i is larger than the base B. We also remark that, for subtraction (i.e. the computation of $a - b$), we can assume that $a \geq b$. The recursion for c_i and γ is then almost identical. More specifically, we have

$$-\gamma_{i+1} \cdot B + c_i = a_i - b_i - \gamma_i \text{ with } c_i, \gamma_i \in \{0, \dots, B-1\}.$$

The proof of the following theorem is straight-forward.

Theorem 1.1.1. *The school method for adding (or subtracting) two $n-$digit numbers requires at most $2n$ primitive operations. The addition of an m-digit number and an n-digit number uses at most $m + n + 2$ primitive operations.*

Algorithm 2: School Method for Multiplication

Input : Two non-negative n-digit integers $a = a_{n-1} \dots a_0$ and $b = b_{n-1} \dots b_0$.
Output: A $2n$-digit integer $c = c_{2n-1} \dots c_0$ with $c = a \cdot b$.

```
1  P_0 := 0
2  for j = 0,...,n-1 do
3      for i = 0,...,n-1 do
4          Define
5          a_i · b_j = c_ij · B + d_ij with c_ij, d_ij ∈ {0,...,B-1}
6      c_j := c_{n-1,j} ... c_{0,j} 0
7      d_j := d_{n-1,j} ..., d_{0,j}
8      p_j = p_{n,j} ... p_{0,j} := c_j + d_j
       // * Notice that p_j = a · b_j, and thus a · b = Σ_{j=0}^{n-1} p_j · B^j *//
9      P_{j+1} := P_j + p_j · B
10 return P_n
```

In the next step, we consider the school method for multiplying integers; see Algorithm 2. Let us count the number of primitive operations that Algorithm 2 needs:

- The computation of each product $a_i \cdot b_j$ requires one primitive operations, thus n^2 many primitive operations in total.
- Computing each of the integers p_j amounts for adding two $(n+1)$-digit numbers. Hence, in total, we need $2n(n+1)$ primitive operations.
- For computing P_n we need n additions each involving $2n$-digit numbers. Thus, we need $2n^2$ many primitive operations for this step.

We now obtain the following result. For the second claim on the complexity of computing the product of an m-digit number and an n-digit number, a completely analogous argument applies.

Theorem 1.1.2. *Using the school method, we need at most $5n^2 + 2n = O(n^2)$ primitive operations to multiply two n-digit numbers. Multiplication of an n-digit number and an m-digit number needs $O(mn)$ primitive operations.*

Exercise 1.1.3. *Let $f = a_0 + \cdots + a_d \cdot x^d \in \mathbb{Z}[x]$ be a polynomial of degree d with integer coefficients of length at most L, and let $m \in \mathbb{Z}$ be an ℓ-digit number. Show that*

(a) $f(m)$ is a $O(d\ell + L)$-digit number.

(b) Computing $f(m)$ using Horner's method

$$f(m) = a_0 + m \cdot (a_1 + m \cdot (a_2 + \cdots m \cdot (a_{d-1} + m \cdot a_d))$$

and the school method for multiplication uses $O(d \cdot (d\ell^2 + \ell \cdot L)))$ primitive operations.

We will later see that it is even possible to compute $f(m)$ in only $\tilde{O}(d \cdot (\ell + L))$ primitive operations, where $\tilde{O}(.)$ means that poly-logarithmic factors are suppressed, that is, $\tilde{O}(T) = O(T \cdot (\log T)^c)$ for some constant c.

Exercise 1.1.4. *Let $A = (a_{i,j})_{i,j=1,\ldots,n} \in \mathbb{Z}^{n\times n}$ be an $n \times n$-matrix with integer entries $a_{i,j}$ of length at most L.*

(a) Derive an upper bound on the number of primitive operations that are needed to compute the inverse A^{-1} of A.

(b) Show that the entries of A^{-1} are rational numbers with numerators and denominators of length $O(n(L + \log n))$.

(c) Suppose that Gaussian elimination with pivoting is used to compute the determinant of A. Further suppose that, after each iteration, we reduce all intermediate entries $a'_{i,j} = \frac{p}{q} \in \mathbb{Q}$, that is, we ensure that $\gcd(p, q) = 1$. Show that p and q can be represented using $O(n^2(L+\log n))$ digits and conclude that Gaussian elimination constitutes a polynomial time algorithm for computing determinants.*

Hints: For (a), consider Gaussian elimination to compute A^{-1} and derive a bound on the numerators and denominators of the rational entries of the matrices produced after each iteration. For (b), use Cramer's Rule to write the entries of A^{-1} as fractions of determinants of suitable $n \times n$-matrices and use the definition of the determinant to bound the size of the numerator and denominator. For (c), show that, in each iteration, the pivot element can be written as the quotient of the determinants of two sub-matrices of A.

1.2 The Toom-Cook Algorithm

We now investigate in algorithms for multiplying integers that are considerably faster than the school method. We start with a simple algorithm due to Karatsuba [AY62] (from 1960). Its running time $O(n^{\log_2 3})$ already constitutes a considerable improvement upon the running time $O(n^2)$ of the school method. Then, we show how to generalize the approach to achieve a running time $O(n^{1+\epsilon})$ for arbitrary $\epsilon > 0$.

Let $a = a_{n-1} \ldots a_0$ and $b = b_{n-1} \ldots b_0$ be integers of length n. We first write

$$a = a_{n-1} \ldots a_0 = a' \cdot B^{\lceil n/2\rceil} + a'', \text{ and}$$
$$b = b_{n-1} \ldots b_0 = b' \cdot B^{\lceil n/2\rceil} + b'',$$

with integers a', a'', b', b'' of length $\lceil n/2\rceil$. Then, it holds that

$$\begin{aligned} a \cdot b &= (a' \cdot B^{\lceil n/2\rceil} + a'') \cdot (b' \cdot B^{\lceil n/2\rceil} + b'') \\ &= a'b' \cdot B^{2\lceil n/2\rceil} + (a'b'' + a'' \cdot b') \cdot B^{\lceil n/2\rceil} + a''b'' \\ &= \underbrace{a'b'}_{=:P_1} \cdot B^{2\lceil n/2\rceil} + [\underbrace{(a' + a'')(b' + b'')}_{=:P_2} - (a'b' + a''b'')] \cdot B^{\lceil n/2\rceil} + \underbrace{a''b''}_{=:P_3} \end{aligned} \tag{1.1}$$

What have we gained in the last step? They crucial point is that, when passing from the second line to the last line, we reduced the problem to three (instead of four!) multiplications and six (instead of three) additions. Notice that there are actually five multiplications, however, each of the products P_1 and P_2 appears twice, and thus only 3 different products need to be computed. So the total number of additions and multiplication has increased, however, additions are much cheaper than multiplications. We can now recursively use the

Algorithm 3: Karatsuba Multiplication (1960)

Input : Two non-negative n-digit integers $a = a_{n-1} \dots a_0$ and $b = b_{n-1} \dots b_0$.
Output: A $2n$-digit integer $c = c_{2n-1} \dots c_0$ with $c = a \cdot b$.

1 **if** $n \leq 4$ **then**
2 Compute $c = a \cdot b$ using Algorithm 2
3 **else**
4 Define

$$\begin{aligned} a &= a_{n-1} \dots a_0 = a' \cdot B^{\lceil n/2 \rceil} + a'', \text{ and} \\ b &= b_{n-1} \dots b_0 = b' \cdot B^{\lceil n/2 \rceil} + b'', \end{aligned}$$

with integers a', a'', b', b'' of length $n/2$.
5 $A := a' + a''$
6 $B := b' + b''$
7 Compute $P_1 := a' \cdot b'$, $P_2 := A \cdot B$, and $P_3 := a'' \cdot b''$ by recursively calling Algorithm 3.
8 $P := P_1 \cdot B^{2\lceil n/2 \rceil} + (P_2 - P_1 - P_3) \cdot B^{\lceil n/2 \rceil} + P_3$
9 **return** P

above approach for multiplication until all remaining multiplications are numbers with four or less digits; see Algorithm 3

Theorem 1.2.1. *Using Karatsuba multiplication, we need $O(n^{\log 3}) = O(n^{1.58\dots})$ primitive operations to multiply two n-digit numbers.*

Proof. Let $T(n)$ denote the maximal number of operations needed to multiply two n-digit numbers using the Karatsuba algorithm. If $n \leq 4$, Theorem 1.1.2 yields that $T(n) \leq 5n^2 + 2n \leq 88$. For $n \geq 5$, it holds that

$$T(n) \leq 3 \cdot T(\lceil n/2 \rceil + 1) + 6 \cdot (4n).$$

as we need to compute 3 products involving $\lceil n/2 \rceil$- or $(\lceil n/2 \rceil + 1)$-digit numbers and 6 additions involving $2n$-digit numbers. Now, a general version of the Master Theorem (e.g. see [MS08, Sec. 2.6]) yields a total running time of size $O(n^{\log_2 3})$. □

Remark. For readers who are not familiar with the general Master Theorem, we give the following direct argument from [MS08], which also yields an explicit bound for $T(n)$. For $\ell \in \mathbb{N}_{\geq 1}$, we first prove that

$$T(2^\ell + 2) \leq 33 \cdot 3^\ell + 12 \cdot (2^{\ell+1} + 2\ell - 2)$$

using induction on ℓ. For $\ell = 1$, the claim is obviously true as $T(4) \leq 88$. For $\ell \geq 2$, we thus conclude from the induction hypothesis and the above recursive formula for $T(n)$ that

$$\begin{aligned} T(2^\ell + 2) &\leq 3 \cdot T(2^\ell + 2) + 12 \cdot (2^\ell + 2) \\ &\leq 3 \cdot [33 \cdot 3^{\ell-1} + 12 \cdot (2^\ell + 2(\ell - 1) - 2)] + 12 \cdot (2^\ell + 2) \\ &= 33 \cdot 3^\ell + 12 \cdot (2^{\ell+1} + 2\ell - 2). \end{aligned}$$

Notice that our special choice for n (i.e. $n = 2^\ell + 2$) guarantees that $\lceil n/2 \rceil + 1 = 2^{\ell-1} + 2$ is again of the same form, and thus we can recursively apply the induction hypothesis on $T(\lceil n/2 \rceil + 1)$. It remains to derive a bound on $T(n)$ for arbitrary n. Setting $\ell := \lceil \log n \rceil \le 1 + \log n$, we have

$$\begin{aligned} T(n) \le T(2^\ell) &\le 33 \cdot 3^\ell + 12 \cdot (2^{\ell+1} + 2\ell - 2) \\ &\le 33 \cdot 3 \cdot 3^{\log n} + 12 \cdot (4 \cdot 3^{\log n} + 2(1 + \log n) - 2) \\ &\le 99 \cdot n^{\log 3} + 48 \cdot \cdot n + 24 \cdot \log n. \end{aligned}$$

We now consider the following approach due to Toom and Cook (1966),[1] which extends Karatsuba's idea; see Algorithm 4. The first step is similar as in Karatsuba's method, however, instead of splitting each of the input numbers into two almost equally sized parts, we now consider a split into k parts, where $k \in \mathbb{N}_{\ge 2}$ is an arbitrary but fixed constant. That is, with $m := \lceil n/k \rceil$, we write

$$\begin{aligned} a &= a^{(0)} + a^{(1)} \cdot B^m + \cdots + a^{(k-1)} \cdot B^{(k-1)\cdot m}, \text{ and} \\ b &= b^{(0)} + b^{(1)} \cdot B^m + \cdots + b^{(k-1)} \cdot B^{(k-1)\cdot m}, \end{aligned}$$

such that each integer $a^{(i)}$ and $b^{(i)}$ has length at most m. Now, let $f(x) := \sum_{i=0}^{k-1} a^{(i)} \cdot x^i$ and $g(x) := \sum_{i=0}^{k-1} b^{(i)} \cdot x^i$ be corresponding polynomials of degree $k-1$ with coefficients $a^{(i)}$ and $b^{(i)}$. Then, it holds that $a \cdot b = f(B^m) \cdot g(B^m) = h(B^m)$, where

$$h(x) = \sum_{i=0}^{2k-2} c^{(i)} \cdot x^i := f(x) \cdot g(x).$$

Notice that the coefficients $c^{(i)}$ of h are integers of length at most $O(m)$. Now, suppose that we know these coefficients, then we can easily compute $a \cdot b$ by shifting each of the coefficients $c^{(i)}$ by $i \cdot m$ digits and adding up the resulting integers. The cost for these additions (there are only constantly many!) is then bounded by $O(n)$. Hence, we have reduced the problem of computing the product $a \cdot b$ of two integers of length n to the problem of computing a product $g(x) \cdot h(x)$ of polynomials of degree less than k and with coefficients of length at most $\lceil n/k \rceil$. For the latter problem, we consider an *evaluation/interpolation approach*, that is, we first evaluate f and g at $2k-1$ many different points $x_0, \ldots, x_{2k-2} \in \mathbb{Z}$ of constant length. Typically, we consider $x_j := j$ for $j = 0, \ldots, 2k-2$ but also other choices are possible. Then, the resulting integer values $f_j = f(x_j)$ and $g_j := g(x_j)$ are of length $O(m)$ according to Exercise 1.1.3. For computing the k products $h_j := f_j \cdot g_j = f(x_j) \cdot g(x_j) = h(x_j)$, we call the multiplication algorithm recursively. In the third step, we interpolate $h(x)$ from its values h_j at the points x_j. Notice that

$$\underbrace{\begin{pmatrix} 1 & x_0 & \cdots & x_0^{2k-2} \\ 1 & x_1 & \cdots & x_1^{2k-2} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & x_{2k-2} & \cdots & x_{2k-2}^{2k-2} \end{pmatrix}}_{=:V} \cdot \begin{pmatrix} c^{(0)} \\ c^{(1)} \\ \vdots \\ c^{(2k-2)} \end{pmatrix} = \begin{pmatrix} h(x_0) \\ h(x_1) \\ \vdots \\ h(x_{2k-2}) \end{pmatrix} = \begin{pmatrix} h_0 \\ h_1 \\ \vdots \\ h_{2k-2} \end{pmatrix},$$

[1] In his Phd Thesis (http://cr.yp.to/bib/1966/cook.html), Cook improves upon Toom's original approach [Too63] from 1963

Algorithm 4: Toom-Cook-k Algorithm

Input : Two non-negative integers a and b of length at most n.
Output: The product $c = a \cdot b$.

1 Write

$$a = a^{(0)} + a^{(1)} \cdot B^m + \cdots + a^{(k-1)} \cdot B^{(k-1)\cdot m}, \text{ and}$$
$$b = b^{(0)} + b^{(1)} \cdot B^m + \cdots + b^{(k-1)} \cdot B^{(k-1)\cdot m},$$

with $m := \lceil n/k \rceil$ and integers $a^{(i)}$, $b^{(i)}$ of length at most m.

2 $f(x) := a^{(0)} + a^{(1)} \cdot x + \cdots + a^{(k-1)} \cdot x^{k-1}$

3 $g(x) := b^{(0)} + b^{(1)} \cdot x + \cdots + b^{(k-1)} \cdot x^{k-1}$

4 **for** $j = 0, \ldots, 2k-2$ **do**

5 Define $x_j = j$

// * *We can also choose other values for x_j unless the x_j's are pairwise distinct and of constant length* *//

6 Compute $f_j := f(x_j)$ and $g_j := g(x_j)$

7 Compute $h_j := f_j \cdot g_j$ by calling the Algorithm 4 recursively.

8 Compute the inverse V^{-1} of the *Vandermonde Matrix*

$$V := \text{Vand}(x_0, \ldots, x_{2k-2}) := \begin{pmatrix} 1 & x_0 & \cdots & x_0^{2k-2} \\ 1 & x_1 & \cdots & x_1^{2k-2} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & x_{2k-2} & \cdots & x_{2k-2}^{2k-2} \end{pmatrix}$$

Compute

$$\begin{pmatrix} c^{(0)} \\ c_{(1)} \\ \vdots \\ c^{(2k-2)} \end{pmatrix} = V^{-1} \cdot \begin{pmatrix} h_0 \\ h_1 \\ \vdots \\ h_{2k-2} \end{pmatrix}$$

$C_0 = c^{(0)}$

9 **for** $j = 1, \ldots, 2k-2$ **do**

10 $C_j = C_j + B^{mj} \cdot c^{(j)}$

11 **return** $C_{2k-2} = c = a \cdot b$

where $V = \text{Vand}(x_0, \ldots, x_{2k-2})$ is the so-called *Vandermonde-Matrix* of $x_0, \ldots, x_{2k-2}$. Hence, we can compute the coefficients $c^{(i)}$ of $h(x)$ from its values h_j at the $2k-1$ points x_j as

$$\begin{pmatrix} c^{(0)} \\ c^{(1)} \\ \vdots \\ c^{(2k-2)} \end{pmatrix} = V^{-1} \cdot \begin{pmatrix} h_0 \\ h_1 \\ \vdots \\ h_{2k-2} \end{pmatrix}$$

Since k is a constant and since each entry of V is of constant size, only a constant number

of primitive operations is needed to compute V^{-1}. Computing the product of V^{-1} and the vector $(h_0, \dots, h_{2k-2})^t$ needs $O(n)$ primitive operations as each h_j has length $O(n)$. Finally, we compute $c = a \cdot b$ as the sum of the $2k-1$ integers $c_j \cdot B^j$, for $j = 0, \dots, 2k-2$, which also uses $O(n)$ primitive operations.

In summary, we thus obtain the following recursion for the computation time $T(n)$ of the Toom-Cook-k Algorithm:

$$T(n) \leq (2k-1) \cdot T(\lceil n/k \rceil) + O(n).$$

Again, the Master Theorem yields the following result:

Theorem 1.2.2. *For a fixed integer $k \in \mathbb{N}_{\geq 2}$, the Toom-Cook-$k$ Algorithm uses $O(n^{\frac{\log(2k-1)}{\log k}})$ primitive operations to multiply two n-digit numbers.*

From the above theorem and the fact that $\lim_{k \mapsto \infty} \frac{\log(2k-1)}{\log k} = 1$, we conclude that, for any fixed $\epsilon > 0$, there exists an algorithm with running time $O(n^{1+\epsilon})$ to multiply to n-digit numbers. In the next chapter, will discuss a method due to Schönhage and Strassen (1971) that even yields a running time of size $O(n \cdot \log^c(n))$, with some constant $c > 1$. The method is similar to the Toom-Cook approach in the sense that it considers the input integers as polynomials and then computes the product of the polynomials using an evaluation/interpolation-approach. The main difference however is that n-digits numbers are considered as polynomials of degree $n-1$ (and not k for some fixed constant k) and that the interpolation points are chosen to be the $2n$-th roots of unity. Here, the crucial point is that evaluating and interpolating a polynomial at the roots of unity can be done in a very efficient way.

Exercise 1.2.3. *Show that Karatsuba's method can be considered as a special case of Toom-Cook-2. For this, you need to choose suitable interpolation points x_0, x_1, x_2 in the Toom-Cook-2 algorithm.*

Hint: You may choose $x_0 = \infty$ as one of the interpolation points, where we define $P(\infty) := P_d$ for a polynomial $P(x) = P_0 + \dots + P_d \cdot x^d$. For the interpolation step, you cannot use the Vandermonde matrix any more but need a more direct approach instead.

Exercise 1.2.4. *For two integers $a = a^{(0)} + a^{(1)} \cdot B^{\lceil n \rceil} + a^{(3)} \cdot B^{2\lceil n \rceil}$ and $b = b^{(0)} + b^{(1)} \cdot B^{\lceil n \rceil} + b^{(3)} \cdot B^{2\lceil n \rceil}$ of length n, use the Toom-Cook-3 approach to derive a relation between the values $a^{(i)}$ and $b^{(i)}$ that is similar to the relation in (1.1) as considered in Karatsuba's method.*

1.3 Approximate Computation

1.3.1 Fixed Point Arithmetic

A common approach when dealing with non-integer values a (e.g. $1/3$, $\sqrt{2}$, or π) is to approximate them by rational numbers $\tilde{a} = m \cdot B^{-\rho}$, with B the working base, $m \in \mathbb{Z}$ and $\rho \in \mathbb{N}$, such that $|a - \tilde{a}| \leq B^{-\rho+1}$. That is, $\tilde{a}$ constitutes the best approximation of a among all *fixed-point number* with base B and precision ρ:

$$\mathbb{F}_{B,\rho} := \{a = (-1)^s \cdot B^{-\rho} \cdot \sum_{i=0}^{n-1} a_i B^i \text{ with } n \in \mathbb{N},\ s \in \{0,1\}, \text{ and } a_i \in \{0, \dots, B-1\}\}$$

If B and ρ are clear from the context, we also write $\mathbb{F} = \mathbb{F}_{B,\rho}$. For convenience, we also write

$$a = (-1)^s \ a_{n-1} \dots a_{\rho+1} a_\rho, a_{\rho-1} \dots a_0$$

for an arbitrary element $a = (-1)^s \cdot B^{-\rho} \cdot \sum_{i=0}^{n-1} a_i B^i \in \mathbb{F}_{B,\rho}$. The *length of* a (with respect to B) is defined as the number n of digits that is needed to represent a. It is common to consider the base $B = 2$ and to work with so called *dyadic* numbers (also called dyadic rationals). These are exactly the fixed point numbers with respect to base 2 and arbitrary but finite precision:

$$\mathbb{D} := \bigcup_{\rho=0}^{\infty} \mathbb{F}_{2,\rho} = \{p \cdot 2^{-\rho} : p \in \mathbb{Z} \text{ and } \rho \in \mathbb{N}\}.$$

In what follows, we always assume that the base B and the precision ρ is fixed. For an arbitrary real value x, we define

$$\operatorname{flu}(x) := \min\{a \in \mathbb{F} : x \leq a\}$$

and

$$\operatorname{fld}(x) := \max\{a \in \mathbb{F} : x \geq a\}.$$

the two *rounding functions* to the nearest fixed-point number that is larger/smaller than or equal to a. $\operatorname{fl}(.)$ defines the *rounding to nearest*, that is, $\operatorname{fl}(x) = \operatorname{flu}(x)$ if $|\operatorname{flu}(x) - x| < |\operatorname{fld}(x) - x|$ and $\operatorname{fl}(x) = \operatorname{fld}(x)$ if $|\operatorname{fld}(x) - x| < |\operatorname{flu}(x) - x|$. In case of ties (i.e. $|\operatorname{fld}(x) - x| = |\operatorname{flu}(x) - x|$), we round to even, that is, $\operatorname{fl}(x) = \operatorname{flu}(x)$ if the last digit of $\operatorname{flu}(x)$ is even, otherwise $\operatorname{fl}(x) = \operatorname{fld}(x)$. For each arithmetic operations $\circ \in \{+, -, \cdot\}$, we now consider a corresponding approximate variant $\tilde{\circ}$, where we use $\operatorname{fl}(.)$ to round the exact result to a nearby number in $\mathbb{F}$:

Definition 1.3.1. *For $x, y \in \mathbb{R}$ and $\circ \in \{+, -, \cdot\}$, we define*

$$x \tilde{\circ} y := \operatorname{fl}(\operatorname{fl}(x) \circ \operatorname{fl}(y)).$$

In particular, we have $x \tilde{\circ} y := \operatorname{fl}(x \circ y)$ for $x, y \in \mathbb{F}$.

Notice that the above definition yields a canonical way of approximately evaluating a polynomial $f(x) = a_0 + \cdots a_d \cdot x^d \in \mathbb{R}[x]$ at an arbitrary real value x. More precisely, we consider some evaluation method (e.g. Horner Evaluation) and replace each of the occurring arithmetic operations $\circ$ by the corresponding fixed point variant $\tilde{\circ}$. We denote the so-obtained result by $f_{\mathbb{F}}(x)$. We remark at this point that the result may crucially depend on the chosen evaluation method. That is, we might get completely different values when using Horner Evaluation instead of the "classical" way of evaluating the polynomial, that is, by first computing all powers x^i of x, then multiplying each power with the corresponding coefficient a_i, and finally summing up the obtained values. In other terms, it does not necessarily hold that

$$a_0 \tilde{+} x \tilde{\cdot} (a_1 \tilde{+} \cdots (a_{d-1} \tilde{+} x \tilde{\cdot} a_d) \dots) = a_0 \tilde{+} a_1 \tilde{\cdot} \tilde{x} \tilde{+} \cdots \tilde{+} a_d \tilde{\cdot} x \tilde{\cdot} x \cdots x \tilde{\cdot} x$$

Exercise 1.3.2. *Give an example where Horner Evaluation and classical evaluation give different results for $f_{\mathbb{F}}(x)$.*

The above approach for approximately evaluating a univariate polynomial at a point then further extends to polynomials $F(\mathbf{x}) \in \mathbb{R}[\mathbf{x}] = \mathbb{R}[x_1, \dots, x_n]$ in several variables. Since each complex number z can be written as $z = x + \mathbf{i} \cdot y$ with $x, y \in \mathbb{R}$, and since each addition and

multiplication in $\mathbb{C}$ amounts for a constant number of additions and multiplications in $\mathbb{R}$, we may further extend the approach to polynomials with complex coefficients. In this case, the set of *complex fixed point numbers* is given as

$$\mathbb{F}_{\mathbb{C}} := \mathbb{F} + \mathbf{i} \cdot \mathbb{F},$$

and the set of *complex dyadic numbers* is given as

$$\mathbb{D}_{\mathbb{C}} := \mathbb{D} + \mathbf{i} \cdot \mathbb{D}.$$

In the next step, we investigate the error when performing a series of additions and multiplications using fixed point arithmetic. Assume that we are given approximations $\tilde{x}, \tilde{y} \in \mathbb{F}_{\mathbb{C}}$ of two complex numbers $x, y \in \mathbb{C}$ with $|x - \tilde{x}| < \epsilon_x$ and $|x - \tilde{x}| < \epsilon_y$. Then, it holds that

$$|(\tilde{x} \mathbin{\tilde{+}} \tilde{y}) - (x + y)| \leq \sqrt{2} \cdot B^{-(\rho+1)} + |(\tilde{x} + \tilde{y}) - (x + y)| < B^{-\rho} + \epsilon_x + \epsilon_y, \tag{1.2}$$

and the same error bound holds true for subtraction. For multiplication, we have

$$|\tilde{x} \mathbin{\tilde{\cdot}} \tilde{y} - x \cdot y| \leq \sqrt{2} \cdot B^{-(\rho+1)} + |(\tilde{x} \cdot \tilde{y}) - x \cdot y| < B^{-\rho} + \epsilon_x \cdot |y| + \epsilon_y \cdot |x| + \epsilon_x \cdot \epsilon_y. \tag{1.3}$$

From the above error bounds, we can now derive a bound on the error $|f(x_0) - f_{\mathbb{F}}(x_0)|$ that we obtain when using Horner evaluation and fixed point arithmetic to compute the value of a polynomial f at a complex point x_0.

Theorem 1.3.3. *For any $x_0 \in \mathbb{C}$ and any polynomial $f \in \mathbb{C}[x]$ of degree d with coefficients of absolute value less than 2^L, with $L \in \mathbb{Z}_{\geq 0}$, it holds that*

$$|f(x_0) - f_{\mathbb{F}}(x_0)| < 4(d+1)^2 \cdot 2^L \cdot B^{-\rho} \cdot \max(1, |x_0|)^d.$$

if Horner Evaluation and fixed point arithmetic with a precision $\rho \geq \log d$ is used for the evaluation of f at x_0.

Proof. We argue by induction on the degree d of $f = a_0 + \cdots + a_d \cdot x^d$. Obviously, the error bound is true for $d = 0$ as

$$|a_0 - \mathrm{fl}(a_0)| \leq \sqrt{2} \cdot B^{-(\rho+1)},$$

When using Horner evaluation to evaluate a polynomial f of degree $d \geq 1$ at x_0, we first evaluate $\hat{f} := a_1 + a_2 \cdot x + \cdots + a_d \cdot x^{d-1}$ at x_0, then multiply the result by x_0 and eventually add a_0. Using fixed point arithmetic with precision ρ, our induction hypotheses yields that

$$|\hat{f}_{\mathbb{F}}(x_0) - \hat{f}(x_0)| < \epsilon := 4d^2 \cdot 2^L \cdot B^{-\rho} \cdot \max(1, |x_0|)^{d-1}$$

Since $|\hat{f}(x_0)| \leq d \cdot 2^L \cdot \max(1, |x_0|)^{d-1}$ and $|x_0 - \mathrm{fl}(x_0)| \leq \sqrt{2} \cdot B^{-(\rho+1)} < B^{-\rho}$, we conclude from (1.3) that

$$\begin{aligned}
|x_0 \cdot \hat{f}(x_0) - \mathrm{fl}(x_0) \mathbin{\tilde{\cdot}} \hat{f}_{\mathbb{F}}(x_0))| &< \sqrt{2} \cdot B^{-(\rho+1)} + \epsilon \cdot |x_0| + B^{-\rho} \cdot |\hat{f}(x_0)| + B^{-\rho} \cdot \epsilon \\
&< B^{-\rho} + \epsilon \cdot \max(1, |x_0|) + B^{-\rho} \cdot |\hat{f}(x_0)| + \frac{\epsilon \cdot \max(1, |x_0|)}{d} \\
&< B^{-\rho} \cdot [1 + 5d \cdot 2^L \cdot \max(1, |x_0|)^{d-1} + 4d^2 \cdot 2^L \cdot \max(1, |x_0|)^d]. \\
&\leq B^{-\rho} \cdot \max(1, |x_0|)^d \cdot 2^L \cdot (1 + 5d + 4d^2)
\end{aligned}$$

Adding the constant a_0 increases the error by less than $2 \cdot B^{-\rho}$ due to (1.2). Hence, the total error is bounded by

$$B^{-\rho} \cdot 2^L \cdot (3 + 5d + 4d^2) \cdot \max(1, |x_0|)^d \leq 4(d+1)^2 \cdot 2^L \cdot B^{-\rho} \cdot \max(1, |x_0|)^d.$$

Hence, the claim follows. □

1.3.2 Interval Arithmetic

Instead of computing an approximation of the value $f(x_0)$ that a function $f : \mathbb{R} \mapsto \mathbb{R}$ (or more general, $f : \mathbb{C} \mapsto \mathbb{C}$) takes at a specific point $x_0 \in \mathbb{R}$ (or $x_0 \in \mathbb{C}$), it is often useful to compute an approximation of the image $f([a, b])$ (or $f([a, b] + \mathbf{i} \cdot [c, d])$) of an interval $[a, b]$ (rectangle $[a, b] + \mathbf{i} \cdot [c, d]$) under the mapping f.

Definition 1.3.4 (Interval Extensions and Box Functions)**.** *Let $f : \mathbb{R} \mapsto \mathbb{R}$ be an arbitrary function. An* interval extension *$\Box f : H \mapsto H$ of f is a function from the halfplane $H := \{[a, b] : a, b \in \mathbb{R}$ with $a \leq b\}$ of intervals $X = [a, b]$ to itself such that $f(x) \in \Box f(X)$ for all $x \in X$. For continuous f, $\Box f$ is a* continuous interval extension *(or* box-function*) if*

$$\bigcap_{i=1}^{\infty} \Box f(X_i) = f(x_0)$$

for any sequence $X_1 \supset X_2 \supset \cdots$ such that $\bigcap_{i=1}^{\infty} X_i$ contains only a single point x_0.

In simpler terms, an interval extension $\Box f$ of f is a function that maps an interval $[a, b]$ to an interval $[A, B]$ such that $f(x) \in [A, B]$ for any $x \in [a, b]$. Notice that this is not a very restricting condition as we can simply choose $\Box f$ as the function that maps any interval to $(-\infty, +\infty)$. However, for a box function, it must also hold that $[A, B]$ shrinks to one point $(f(x_0))$ if $[a, b]$ shrinks to one point (x_0).

We further remark that Definition 1.3.4 further generalizes to complex valued functions $f : \mathbb{C} \mapsto \mathbb{C}$. Then, an interval extension $\Box f : H_{\mathbb{C}} \mapsto H_{\mathbb{C}}$ computes for each rectangle $R = [a, b] + \mathbf{i} \cdot [c, d] \in H_{\mathbb{C}} := H + \mathbf{i} \cdot H$ a rectangle $\Box f(B) \in H_{\mathbb{C}}$ with $f(B) \subset \Box f(B)$. The definition of a box function is also completely analogous to the real case. We now show how to compute a box-function for a polynomial. For this, we introduce the concept of interval-arithmetic.

Definition 1.3.5 (Interval Arithmetic)**.** *Let $[a, b]$ and $[c, d]$ be arbitrary intervals and λ a non-negative real number. Then, we define*

$$\begin{aligned}
\lambda \cdot [a, b] &:= [\lambda \cdot a, \lambda \cdot b] \\
-[a, b] &:= [-b, -a] \\
[a, b] \boxplus [c, d] &:= [a + c, b + d] \\
[a, b] \boxminus [c, d] &:= [a, b] \boxplus [-c, -d], \text{ and} \\
[a, b] \boxdot [c, d] &:= [\min(ab, bd, ad, bc), \max(ab, bd, ad, bc)]
\end{aligned}$$

The above rules then extend to arithmetic operations on rectangles in $\mathbb{C}$ in a straight forward way. In particular, for $R = [a, b] + \mathbf{i} \cdot [c, d]$ and $R' := [a', b'] + \mathbf{i} \cdot [c', d']$, we have

$$\begin{aligned}
R \boxplus R' &:= [a, b] \boxplus [a', b'] + \mathbf{i} \cdot ([c, d] \boxplus [c', d']), \\
R \boxdot R' &:= [a, b] \boxdot [a', b'] \boxminus [c, d] \boxdot [c', d'] + \mathbf{i} \cdot ([a, b] \boxdot [c', d'] \boxplus [a', b'] \boxdot [c, d]).
\end{aligned}$$

Often, we have to restrict to fixed point arithmetic instead of exact arithmetic. Similar to the definition of fl(.), which rounds a real (or complex) value to its best approximation in $\mathbb{F}$ (or $\mathbb{F}_{\mathbb{C}}$), we introduce the following rounding function for intervals (rectangles in $\mathbb{C}$):

$$\mathrm{Fl} : H_{\mathbb{C}} \mapsto H_{\mathbb{C}} : \mathrm{Fl}([a, b] + \mathbf{i} \cdot [c, d]) := [\mathrm{fld}(a), \mathrm{flu}(b)] + \mathbf{i} \cdot [\mathrm{fld}(c), \mathrm{flu}(d)]$$

Hence, Fl(.) rounds each of the vertices of a rectangle B to the nearest corresponding approximations in $\mathbb{F}_{\mathbb{C}}$ such that $\text{Fl}(B)$ contains B. We can now define arithmetic operations on intervals (rectangles) using fixed point arithmetic.

Definition 1.3.6 (Fixed Point Interval Arithmetic)**.** *Let $[a,b]$ and $[c,d]$ be arbitrary intervals with $a,b,c,d \in \mathbb{F}$ and $\lambda \in \mathbb{R}$ a non-negative real number. Then, we define*

$$\begin{aligned}
[a,b] \,\tilde{\boxplus}\, [c,d] &:= \text{Fl}([a,b] \boxplus [c,d]) \\
[a,b] \,\tilde{\boxminus}\, [c,d] &:= \text{Fl}([a,b] \boxplus [-d,-c]), \\
[a,b] \,\tilde{\boxdot}\, [c,d] &:= \text{Fl}([a,b] \boxdot [c,d]), \text{ and} \\
\lambda \,\tilde{\boxdot}\, [a,b] &:= [\text{fld}(\lambda), \text{flu}(\lambda)] \,\tilde{\boxdot}\, [a,b]
\end{aligned}$$

Again, the above rules for arithmetic operations on intervals extend in a straight forward manner to rectangles in $\mathbb{C}$. In addition, they induce interval extensions $\square f$ and $\tilde{\square} f$ for a polynomial $f \in \mathbb{R}[x]$. For this, we replace each arithmetic operation $\circ$ in the evaluation of f (e.g. when using Horner Evaluation) by the corresponding interval variant $\boxcircle$ and $\tilde{\boxcircle}$, respectively. Notice that $\square f$ is a box-function, whereas this is not true for $\tilde{\square} f$.

Exercise 1.3.7. *For any $x \in \mathbb{R}$ with $0 \le x \le 1$ and $k \in \mathbb{N}$, there exists a $\xi \in [0,x]$ such that*

$$\cos(x) = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \cdots + \frac{x^{4k}}{(4k)!} \cdot \cos(\xi) \quad \text{(Taylor Series Expansion with Remainder Term)}$$

Use the above formula to derive a box function $\square \cos$ for $\cos$ for intervals $[a,b] \subset [0,1]$! Can you extend your approach to derive a box function for $\sin x$ and e^x.

We now investigate a bound on the size of the intervals (rectangles) that are obtained when performing a series of additions and multiplication according to the above rules. Notice that there are similarities to our considerations in the previous section, where we derived bounds on the error that occurs when adding or multiplying numbers using fixed point arithmetic. Namely, you might think of two rectangles $R := [a,b] + \mathbf{i} \cdot [c,d]$ and $R' := [a',b'] + \mathbf{i} \cdot [c',d']$ as approximations of its centers $m_R := \frac{a+b}{2} + \mathbf{i} \cdot \frac{c+d}{2}$ and $m_{R'} := \frac{a'+b'}{2} + \mathbf{i} \cdot \frac{c'+d'}{2}$ up to an error of size at most $\epsilon := \sqrt{2} \cdot w(R)$ and $\epsilon' := \sqrt{2} \cdot w(R')$, respectively, where $w(R) = \max(b-a, d-c)$ and $w(R') = \max(b'-a', d'-c')$ are defined as the *width* of R and R'. Then, the output of an arithmetic operation between R and R' can again be considered as an approximation of the corresponding arithmetic operation between m_R and $m_{R'}$. Hence, similarly to the bounds in (1.2) and (1.3), we obtain for any two rectangles R and R' with vertices in $\mathbb{F} + \mathbf{i} \cdot \mathbb{F}$ that

$$w(R \boxplus R') \le w(R \,\tilde{\boxplus}\, R') \le w(R) + w(R') + 2 \cdot B^{-\rho} \tag{1.4}$$

and

$$w(R \boxdot R') \le w(R) \cdot w(R') + |m_R| \cdot w(R') + |m_{R'}| \cdot w(R) + 2 \cdot B^{-\rho}. \tag{1.5}$$

Exercise 1.3.8. *Prove correctness of the inequalities in (1.4) and (1.5).*

Exercise 1.3.9. *Let $f \in \mathbb{C}[x]$ be a polynomial of degree d with coefficients of absolute value less than 2^L, with $L \in \mathbb{Z}_{\ge 0}$, let $\rho \in \mathbb{N}$ be a precision with $\rho > \log d$, and let $\mathbb{F}$ the corresponding set of fixed point numbers with precision ρ. Let $R = [a,b] + \mathbf{i} \cdot [c,d]$ be a rectangle of width*

$w(R) < \frac{1}{d}$ with vertices in $\mathbb{F} + \mathbf{i} \cdot \mathbb{F}$, and suppose that we compute $\square f$ (and $\tilde{\square} f$) using Horner Evaluation and fixed point interval arithmetic with a precision ρ. Then, it holds

$$w(\square f(R)) \le w(\tilde{\square} f(R)) < 8 \cdot (d+1)^2 \cdot 2^L \cdot \max(1, |m_R|)^d \cdot w(R). \tag{1.6}$$

Hint: Consider a similar argument as in the proof of Theorem 1.3.3.

Notice that the bound (1.6) on $w(\square f(R))$ and $w(\tilde{\square} f(R))$ tends to zero if we consider a rectangle (square) R of width $c \cdot B^{-\rho}$, for some constant c, and the precision ρ tends to ∞. Hence, in order to compute an approximation of $f(x_0)$ for some complex value x_0, we may first approximate x_0 by some fixed point number $\tilde{x}_0 = \tilde{x}_{0,\Re} + \mathbf{i} \cdot \tilde{x}_{0,\Im} \in \mathbb{F}_{B,\rho} + \mathbf{i} \cdot \mathbb{F}_{B,\rho}$ such that $|x_0 - \tilde{x}_0| \le B^{-\rho}$ and consider a rectangle

$$R := [\tilde{x}_{0,\Re} - B^{-\rho}, \tilde{x}_{0,\Re} + B^{-\rho}] + \mathbf{i} \cdot [\tilde{x}_{0,\Im} - B^{-\rho}, \tilde{x}_{0,\Im} + B^{-\rho}]$$

of width $2B^{-\rho}$ whose vertices are obtained by adding and subtracting $B^{-\rho}$ from the real and complex part of $\tilde{x}_0$. Then, R contains x_0 and we can use use interval arithmetic to compute the rectangle $\tilde{\square} f(R)$, which contains $f(x_0)$. Its center m constitutes an approximation of $f(x_0)$ with $|m - f(x_0)| < w(\tilde{\square} f(R))$. Hence, for computing an approximation m with $|m - f(x_0)| < \epsilon$, we can iteratively compute $\tilde{\square} f(R)$ with increasing precision $\rho = 1, 2, 4, 8, \ldots$ until $w(\tilde{\square} f(R)) < \epsilon$, and then return the center of $\tilde{\square} f(R)$. Exercise 1.3.9 guarantees that we must succeed as soon as the precision ρ fulfills the inequality

$$\rho > \rho_\epsilon := \log_B[16(d+1)^2 \cdot 2^L \cdot \max(1, |x_0|)^d \cdot \epsilon^{-1}] = O(\log d + d \log \max(1, |x_0|) + L + |\log \epsilon|),$$

where we used that

$$\max(1, |m_R|)^d \le \max(1, |x_0| + B^{-\rho}) \le \max(1, |x_0|)^d \cdot (1 + 1/d^2)^d \le 2 \max(1, |x_0|)^d$$

for any $\rho > 2 \log d$. Since we double ρ in each step, this shows that we succeed for a precision $\rho < 2\rho_\epsilon$. We fix this result, which will turn out to be useful at several places in the following considerations.

Theorem 1.3.10. *Let $f \in \mathbb{C}[x]$ be a polynomial of degree d with coefficients of absolute value less than 2^L, with $L \in \mathbb{Z}_{\ge 0}$, and let x_0 be an arbitrary complex value. For any non-negative integer ℓ, we can compute an approximation $\tilde{y}_0$ of $y_0 = f(x_0)$ with $|y_0 - \tilde{y}_0| < 2^{-\ell}$ using fixed point interval arithmetic with a precision ρ bounded by*

$$O(\log d + d \log \max(1, |x_0|) + L + \ell).$$

Notice that the above bound on ρ that is needed in the worst-case is also a (worst-case) bound on the input precision as, in each iteration, we need approximations of the coefficients of f as well as of x_0 to an error less than $B^{-\rho}$. We further remark that, as an alternative to the above approach, one could also use fixed point arithmetic directly to compute an approximation of $f(x_0)$, and to estimate the occurring error using Theorem 1.3.3. This yields a comparable bound on the needed precision in the worst case. However, the main drawback of this approach is that one has to work with an a priori computed worst-case error bound, which means that the needed precision is always of size $\Omega(\log d + d \log \max(1, |x_0|) + L + \ell)$. In contrast, when using interval arithmetic with increasing precision, we might already succeed with a much smaller precision.

Exercise 1.3.11. *Suppose that a polynomial $f \in \mathbb{R}[x]$ as well as a real value x_0 is given by means of an oracle that returns arbitrary good dyadic approximations of the coefficients of f and x_0. Under the assumption that $f(x_0) \neq 0$, formulate an algorithm that computes an $\ell \in \mathbb{Z}$ such that $2^{-\ell} < |f(x_0)| < 2^{\ell+2}$. How does its running time depend on $|f(x_0)|$?*

1.3.3 Floating point arithmetic

When actually implementing algorithms, the standard approach for the approximate computation with real (complex) numbers is NOT fixed point arithmetic but *floating point arithmetic*. However, a corresponding error analysis is more delicate, and thus, for the seek of simplicity, we decided to use fixed point arithmetic as our main tool for approximate computation. For a short but self-contained introduction, we refer to the appendix of [MOS11].

As already suggested by the name, we use floating point numbers to approximate an arbitrary real numbers a to a fixed *relative* error, whereas fixed point numbers are used for approximations to a fixed *absolute* error. By abuse of notation, we also use $\mathbb{F}$ to denote the set of floating point numbers with precision ρ and base B:

$$\mathbb{F} := \{-1^s \cdot m \cdot B^e : s \in \{0,1\},\ m \in \{1,\dots,B^\rho - 1\} \text{ and } e \in \{-e_{\min},\dots,e_{\max}\}\}$$

We call m the mantissa and e the exponent. Many hardware floating point units use the so-called IEEE 754 standard (1985) with *single* or *double precision*. For double precision, we have $B = 2$, $\rho = 52$, and $e \in \{-1023,\dots,0,\dots,1023\}$. There also exist several libraries for arbitrary precision arithmetic (e.g. MPFR[2]) that allow us to compute with floating point numbers of arbitrary length and (almost) no restrictions on the exponent e. In a completely analogous way as for fixed point numbers, we define the rounding functions fl(.), fld(.), and flu(.). While we have $|x - \mathrm{fl}(x)| \le B^{-(\rho+1)}$ for fixed point numbers with precision ρ (and base B), it now holds for floating point numbers of corresponding length that

$$|x - \mathrm{fl}(x)| \le 2^{-(\rho+1)} \cdot \min(|x|, |\,\mathrm{fl}(x)|).$$

Arithmetic operations of floating point numbers are defined in exactly the same way as for fixed point numbers. Furthermore, all concepts as introduced in the previous section for fixed point arithmetic also carry over to floating point arithmetic. Here, we only give the following result from [MOS11], which provides a bound on the error when using floating point arithmetic to evaluate a multivariate polynomial at a specific point. You may consider this as the counterpart of Theorem 1.3.3, which gives a comparable bound for evaluating a univariate polynomial using fixed point arithmetic.

Theorem 1.3.12. *Let $f = \sum_\alpha c_\alpha \cdot \mathbf{x}^\alpha \in \mathbb{R}[\mathbf{x}] = \mathbb{R}[x_1,\dots,x_n]$ be a polynomial of total degree d, and let $c_f := \sum_\alpha \max(1,|c_\alpha|)$ and $m_f := |\{\alpha : c_\alpha \neq 0\}|$. Then, for arbitrary real values $x_1,\dots,x_n$ of absolute value at most 2^Γ, with $\Gamma \in \mathbb{N}$, it holds that*

$$|f(x_1,\dots,x_n) - f_{\mathbb{F}}(x_1,\dots,x_n)| \le c_f \cdot (m_f + 2d) \cdot 2^{d\Gamma} \cdot 2^{-\rho},$$

where $f_{\mathbb{F}}(x_1,\dots,x_n)$ denotes the result of evaluating f at $\mathbf{x} = (x_1,\dots,x_n)$ with floating point arithmetic with precision ρ (and base $B = 2$).

The above theorem also generalizes to complex values x_i and polynomials defined over the complex numbers. The obtained error bound is comparable, that is, it only differs by a multiplicative constant from the above bound.

[2] http://www.mpfr.org/

1.4 Division

In the previous sections, we have shown how to efficiently carry out additions and multiplications on integers. We also considered corresponding operations on fixed-point numbers and intervals and estimated the error that occurs when using approximate instead of exact arithmetic. So far, any such treatment for the division of integers or fixed-/floating-point numbers a and b is missing. We will first show how to compute an arbitrary good dyadic approximation $\tilde{q} \in \mathbb{D}$ of a rational number $q := \frac{a}{b} \in \mathbb{Q}$ using only additions and multiplications of integers. We start with the special case, where $a = 1$ and b is a positive integer of length less than n. The crucial idea underlying the approach is to consider q as the unique solution of the equation $f(x) := \frac{1}{x} - b = 0$ and to use the Newton-Raphson method to derive an approximation of q. That is, with $x_0 := 2^{-\lceil \log b \rceil} \in \mathbb{D}$, we define

$$x_{i+1} := x_i - \frac{f(x_i)}{f'(x_i)} = x_i - \frac{\frac{1}{x_i} - b}{-\frac{1}{x_i^2}} = 2 \cdot x_i - b \cdot x_i^2 = x_i \cdot (2 - b \cdot x_i) \in \mathbb{D} \quad \text{for } i \in \mathbb{N}_{\geq 1}. \quad (1.7)$$

The first part of the following exercise shows that the sequence x_i converges *quadratically* to q. Roughly speaking, this means that the number of correct digits doubles in each iteration. We then conclude that, after $\lceil \log L \rceil$ iterations, we have computed a dyadic approximation $\tilde{q}$ of $q = 1/b$ with $|q - \tilde{q}| < 2^{-L}$. However, there is a small problem with this approach, namely, the lengths of the dyadic numbers x_i double in each iteration, and since x_0 has length $\lceil \log B \rceil \leq n$, we end up with dyadic numbers of length $O(nL)$ after $\lceil \log L \rceil$ iterations. In Part (c) of the exercise, we show that we can improve upon this approach by rounding the result obtained in the i-th iteration to the ρ_i-th digit after the binary point, with $\rho_i := 2^{i+1} + 2n$. As a result, we can reduce the length of the occurring numbers from $O(nL)$ to $O(n + L)$.

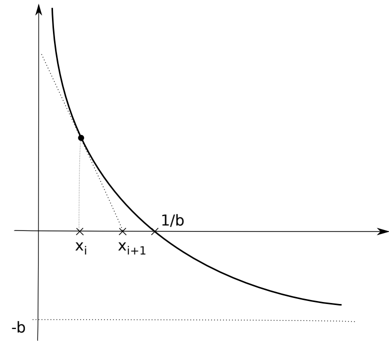

Figure 1.1: The graph of the function $f(x) = \frac{1}{x} - b$. The value x_{i+1} results from applying one step of the Newton-Raphson method to x_i.

Exercise 1.4.1. *Let $(x_i)_i$ be defined as above and L be an arbitrary positive number. Show that, for all i, it holds that*

(a) $\left|x_{i+1} - \frac{1}{b}\right| \leq b \cdot \left|x_i - \frac{1}{b}\right|^2$ *and*

(b) $\left|x_i - \frac{1}{b}\right| < \frac{1}{b} \cdot 2^{-2^i}$. *In particular, it holds that* $\left|x_i - \frac{1}{b}\right| < 2^{-L}$ *for all* $i \geq \log L$.

(c) Suppose now that we start with $y_0 := x_1 = 2^{-\lceil \log b \rceil} \cdot (2 - b \cdot 2^{-\lceil \log b \rceil})$ *and define*

$$y_{i+1} := \mathrm{fl}(y_i \cdot (2 - b \cdot y_i)) \quad \text{for } i \in \mathbb{N}_{\geq 1},$$

where we consider rounding to the nearest fixed-point number of precision $\rho_i := 2^{i+1} + 2n$. *Then, it holds* $\left|y_i - \frac{1}{b}\right| < \frac{1}{b+1} \cdot 2^{-2^i}$ *for all* i.

Hint: For (c), use that the error $2^{-\rho_i - 1}$ that is induced by the rounding in the $(i+1)$-st iteration is smaller than $\frac{2^{-2^{i+1}}}{(b+1)^2}$. Then, use induction on i to prove the claim.

Algorithm 5: Division

Input : Two non-negative n-digit integers a and b and a non-negative integer L.
Output: A dyadic number $\tilde{q} \in \mathbb{D}$ of length $O(n+L)$ such that $|\tilde{q} - a/b| < 2^{-L}$.

1 $L' := \lceil \log a \rceil + L + 1$
2 $N := \lceil \log L' \rceil$
3 $x_0 := 2^{-\lceil \log b \rceil} \cdot (2 - b \cdot 2^{-\lceil \log b \rceil})$
4 **for** $i = 1, \ldots, N-1$ **do**
5 Recursively define
6
$$x_{i+1} := \mathrm{fl}(x_i \cdot (2 - b \cdot x_i)),$$
where fl(.) is defined as "rounding to the nearest element" in $\mathbb{F}_{2,\rho_i}$ and $\rho_i := 2^{i+1} + 2n$.
7 Compute $\tilde{q} := \mathrm{fl}(a \cdot x_{N-1})$, where fl(.) is defined as rounding to the nearest in $\mathbb{F}_{2,L}$.
8 **return** $\tilde{q}$

From the above consideration, we conclude that we can compute a dyadic number $\tilde{q}$ with $|\tilde{q} - 1/b| < 2^{-L}$ using $O(\log L)$ additions and multiplications of integers of length $O(L+n)$. Now, computing a corresponding approximation $\tilde{q}$ of $q := \frac{a}{b}$, with integers a and b of length less than n, is straightforward; see Algorithm 5. Namely, we first compute a dyadic q' of length $O(L+n)$ such that $|q' - 1/b| < 2^{-L-\lceil a \rceil - 1}$ and then determine the product $a \cdot q'$. The result is eventually rounded to the L-th digit after the binary point. The so-obtained $\tilde{q} = \mathrm{fl}(a \cdot q')$ has length $O(n+L)$ and it holds that $|\tilde{q} - q| < 2^{-L}$. We fix this result:

Theorem 1.4.2. *Let a and b be integers of length n. For any non-negative L, Algorithm 5 computes a dyadic approximation $\tilde{q} \in \mathbb{D}$ of length $O(n+L)$ such that $|\tilde{q} - q| < 2^{-L}$. For this, it uses $O(\log(n+L))$ additions and multiplications of $O(n+L)$-digit integers.*

We can now go one step further and derive a bound on the cost for computing an approximation of the quotient of two arbitrary complex numbers $a = a_0 + \mathbf{i} \cdot a_1$ and $b = b_0 + \mathbf{i} \cdot b_1$. Here, we assume that, for any $L' \in \mathbb{N}$, we can ask for dyadic approximations $\tilde{a}, \tilde{b} \in \mathbb{D}$ such that $|a - \tilde{a}|, |b - \tilde{b}| < 2^{-L'}$. Notice that

$$\frac{a}{b} = \frac{a_0 + \mathbf{i} \cdot a_1}{b_0 + \mathbf{i} \cdot b_1} = \frac{(a_0 + \mathbf{i} \cdot a_1) \cdot (b_0 + \mathbf{i} \cdot b_1)}{(b_0 + \mathbf{i} \cdot b_1) \cdot (b_0 - \mathbf{i} \cdot b_1)} = \frac{(a_0 b_0 - a_1 b_1) + \mathbf{i} \cdot (a_1 b_0 + a_0 b_1)}{|b|^2},$$

thus we can restrict to quotients of real numbers $a, b \in \mathbb{R}_{\neq 0}$. Suppose that dyadic approximations $\tilde{a}, \tilde{b} \in \mathbb{R}_{\neq 0}$ with $|a - \tilde{a}|, |b - \tilde{b}| < 2^{-L'} < |b|/2$ are given. Then, we have

$$\left| \frac{\tilde{a}}{\tilde{b}} - \frac{a}{b} \right| = \left| \frac{b\tilde{a} - a\tilde{b}}{b\tilde{b}} \right| = \frac{\left| b(\tilde{a} - a) - a(b - \tilde{b}) \right|}{\left| b^2 + b(b - \tilde{b}) \right|} < 2^{-L'+1} \cdot \frac{|a| + |b|}{|b|^2} \leq 2^{-L'+2} \cdot \frac{\max(|a|, |b|)}{\min(1, |b|)^2}.$$

For $L' > L + \lceil \log \max(1, |a|) \rceil + 3\lceil |\log b| \rceil + 3$, this implies that $\left| \frac{\tilde{a}}{\tilde{b}} - \frac{a}{b} \right| < 2^{-L-1}$. Hence, we may first consider L'-digit approximations $\tilde{a}, \tilde{b} \in \mathbb{D}$ of a and b, and then compute an $(L+1)$-digit approximation $\tilde{q} \in \mathbb{D}$ of their quotient $q = \frac{\tilde{a}}{\tilde{b}}$ using the method from above. Then, it holds that $|\tilde{q} - a/b| < 2^{-L}$. We fix this result:

Theorem 1.4.3. *Let $a, b \in \mathbb{C}$ be arbitrary complex numbers and $L \in \mathbb{N}$. Then, there exists a positive integer L' of size*

$$L' := O(L + \lceil \log \max(1, |a|) \rceil + \lceil |\log |b|| \rceil)$$

such that we can compute a fixed point number $\tilde{q} \in \mathbb{F} + \mathbf{i} \cdot \mathbb{F}$ of length L' with $|\tilde{q} - q| < 2^{-L}$ using $O(\log L')$ additions and multiplications of $O(L')$-digit integers. The values a and b need to be approximated to an error of size $2^{-L'}$.

Exercise 1.4.4. *For arbitrary $x \in \mathbb{R}$ with $0 \leq x \leq 1$, it holds that*

$$\arctan(x) = x - \frac{x^3}{3} + \frac{x^5}{5} - \frac{x^7}{7} + \cdots \tag{1.8}$$

Now, for given $L \in \mathbb{N}$, use the above formula and the fact (due to Euler) that

$$\pi = 20 \cdot \arctan(1/7) + 8 \cdot \arctan(3/79)$$

to derive an efficient algorithm (i.e. with a running time polynomial in L) for computing a fixed point approximation $\tilde{\pi}$ (wrt. base 2) of π to an error less than 2^{-L}.

Hint: Estimate the error when considering only the first k summands in (1.8). Then, proceed with a suitably truncated series.

Exercise 1.4.5. *For arbitrary $x \in \mathbb{R}$ with $0 \leq x \leq 1$, we have*

$$\cos(x) = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \cdots$$

For fixed $n \in \mathbb{N}_{\geq 8}$ and arbitrary $L \in \mathbb{N}$, formulate an efficient method to compute an L-digit approximation $\tilde{\omega}$ of $\omega := \cos(2\pi/n)$.

Hint: Proceed similar as in Exercise 1.4.4 and use a sufficiently good approximation $\tilde{\pi}$ of π. For the evaluation of the truncated series at $x = \tilde{\pi}$, use Theorem 1.3.3.

Chapter 2

The Fast Fourier Transform and Fast Polynomial Arithmetic

2.1 Schönhage-Strassen Multiplication

In the previous chapter, we have seen that the cost $M(n)$ for computing the product of two integers of length n is bounded by $O(n^{1+\epsilon})$, where ϵ is a an arbitrary but fixed positive real value. For sufficiently large k, this bound is achieved by the Toom-Cook-k algorithm. In this section, we present a method [SS71] due to Schönhage and Strassen whose running time is bounded by[1] $O(n \log n \cdot M(\log n)) = O(n \log^{2+\epsilon} n)$. Before we go into detail, we give an overview of the main steps.

2.1.1 The Algorithm in a Nutshell

In the first step, we split a and b into n blocks $a^{(i)}$ and $b^{(i)}$, that is, we write

$$a = a^{(0)} + a^{(1)} \cdot B + \cdots + a^{(n-1)} \cdot B^{n-1}, \quad \text{and}$$
$$b = b^{(0)} + b^{(1)} \cdot B + \cdots + b^{(n-1)} \cdot B^{n-1}$$

with one-digit numbers a^i, $b^{(i)} \in \{0, \ldots, B-1\}$. Notice the difference to the Toom-Cook algorithm, where we split a and b into only constantly many (i.e. k) blocks of size $\lceil n/k \rceil$. Similar to the Toom-Cook method, we now consider corresponding polynomials

$$f(x) := a^{(0)} + a^{(1)} \cdot x + \cdots + a^{(n-1)} \cdot x^{n-1}, \quad \text{and}$$
$$g(x) := b^{(0)} + b^{(1)} \cdot x + \cdots + b^{(n-1)} \cdot x^{n-1} \tag{2.1}$$

of degree $n-1$ (instead of k as in the Toom-Cook method) with coefficients $a^{(i)}$ and $b^{(i)}$, and reduce the computation of $a \cdot b$ to the problem of computing the product $h = \sum_{i=0}^{2n-2} c^{(i)} \cdot x^i := f \cdot g$ of the polynomials f and g, followed by the evaluation of h at $x = B$. For the computation of h, we again use an evaluation/interpolation approach, that is, we first evaluate f and g at $2n$ points $x_0, \ldots, x_{2n-1}$, compute each of the products $f(x_i) \cdot g(x_i) = h(x_i)$, and then

[1] We remark that there exists a slightly more involved variant of the Schönhage-Strassen method that needs only $O(n \log n \log\log n)$ primitive operations. For the seek of simplicity, we decided to only present the variant with slightly worse running time but hint to the faster approach when discussing the corresponding steps in more detail.

reconstruct h from its values at the points x_i. The crucial part of the algorithm is the special choice of the points x_i, that is, instead of considering arbitrary distinct values for the points x_i, we now choose $x_i = \omega^i$ for $i = 0, \dots, 2n-1$, where $\omega \in \mathbb{C}$ is a *primitive* $2n$-th root of unity. That is, ω is a solution of the equation $x^{2n} - 1 = 0$, and it holds that $\omega^i \neq 1$ for any integer i with $1 \leq i < 2n$. For convenience, we choose $\omega := e^{\frac{\pi \mathbf{i}}{n}} = \cos(\pi/n) + \mathbf{i} \cdot \sin(\pi/n)$, even though other choices are possible. We will see that, for n a power of two, there exists a very efficient method, called *Fast Fourier Transform* (FFT for short) due to Cooley and Tukey (1965), that needs only $O(n \log n)$ additions and multiplications of complex numbers in order compute the so-called *Discrete Fourier Transform* (DFT for short)

$$\mathrm{DFT}_\omega(f) := (f(1), f(\omega), \dots, f(\omega^{2n-1})).$$

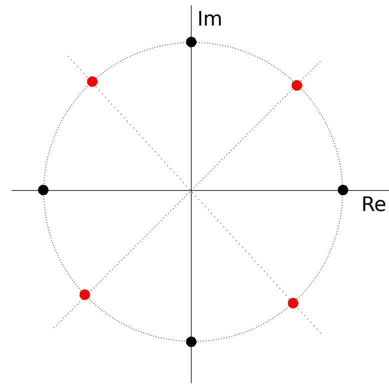

Figure 2.1: The dots on the unit circle are the 8-th roots of unity. The red dots are primitive.

The efficiency of the method is based on the fact that there are only $2n$ different values for x_i^j for any i, j if $x_i = \omega$, whereas, for a general choice of x_i, there are $2n^2$ different values for x_i^j. We will further show that the fast convolution method can also be used to interpolate h from the values $h(x_i)$ in a comparably efficient manner.

One problem of the approach is that, since ω is not a rational number in general, the computations involving ω can only be carried out with approximate arithmetic. However, we will show that the total (absolute) error that occurs during the computation is less than $1/2$ if we use fixed point arithmetic with a precision $\rho > \rho_0$ in each step, where ρ_0 is some computable number of size $O(\log n)$. In addition, we will show that all occurring numbers in the intermediate results have length bounded by $O(\log n)$, and thus we may conclude that, using $O(n \log n)$ arithmetic operations on fixed-point numbers of length $O(\log n)$, we can compute approximations $\tilde{c}^{(i)}$ of the coefficients $c^{(i)}$ of h with $|c^{(i)} - \tilde{c}^{(i)}| < 1/2$. Since each coefficient $c^{(i)}$ is an integer, we can thus derive the exact value $c^{(i)}$ from its approximation $\tilde{c}^{(i)}$. We give the following example to illustrate the last step: Suppose that our approach yields the approximation

$$\tilde{h} = 2.34 \cdot x^{10} - 0.14 \cdot x^9 + 0.98 \cdot x^8 + \cdots + 0.67 \cdot x + 1.11 \tag{2.2}$$

for the product $h = f \cdot g$ of two integer polynomials f and g. In addition, according to the choice of our precision ρ, we can guarantee that the absolute error is less than $1/2$. Now, since the coefficients of h are integers and since they differ from the corresponding approximations by less than $1/2$, we conclude that $h = 2 \cdot x^{10} + x^8 + \cdots + x + 1$.

It remains to show how to recover the product $c = a \cdot b$ from the polynomial h. For this, we evaluate h at $x = B$, which amounts for shifting each coefficient $c^{(i)}$ by i digits and summing up the so obtained numbers. Here, it is crucial that each $c^{(i)}$ has length $O(\log n)$, and thus each summation uses only $O(\log n)$ primitive operations. We conclude that the total cost is bounded by $O(n \log n \cdot M(\log n))) = O(n(\log n)^{2+\epsilon})$ primitive operations, where ϵ is an arbitrary fixed positive number. Instead of using the Toom-Cook algorithm for the occurring multiplications in the Schönhage-Strassen method, we could instead call the Schönhage-Strassen method recursively. This yields the running time

$$O(n \log n M(n)) = O(n(\log n)(\log \log n)M(\log \log n)) = O(n(\log n)^2(\log \log n)^2 M(\log \log \log n))$$

and so on. As already mentioned above, it is possible to slightly improve upon this approach. This is achieved by splitting the initial numbers not into $\approx n/\log n$ blocks of size $\approx \log n$. Then, recursively calling the algorithm even yields the complexity bound $O(nM(\log n)))$. We now give details in the following two sections.

2.1.2 Fast Fourier Transform

Even though we are mainly interested in solving problems defined over the real or complex numbers, it will turn out to be useful to work over an arbitrary ring R (or a field $\mathbb{K}$). In what follows, we always assume that R is a commutative ring with $1 = 1_R$.

We start with the following definition:

Definition 2.1.1 (Convolution). *Let $f = a_0 + \cdots + a_{N-1} \cdot x^{N-1}$ and $g = b_0 + \cdots + b_{N-1} \cdot x^{N-1}$ be two polynomials of degree less than N in $R[x]$. We define*

$$f \star_N g := \sum_{k=0}^{N-1} c_k \cdot x^k := \sum_{k=0}^{N-1} \left(\sum_{i,j: i+j=k \mod N} a_i \cdot b_j \right) \cdot x^k$$

as the convolution *of f and g.*

Example. Let $f = 1 + x + x^2 \in \mathbb{Z}[x]$ and $g := 2 - x$, then $f \cdot g = 2 + x + x^2 - x^3$, and

$$f \star_3 g = (2-1) + 1 \cdot x + 1 \cdot x^2 = 1 - x + x^2.$$

Notice that, in general, $f \star_N g = f \cdot g \mod (x^N - 1)$. In particular, if we consider two polynomials f and g of degree less than n as polynomials of degree less than $2n-1$ (by setting $a_n = \cdots = a_{2n-1} = b_n = \cdots = b_{2n-1} = 0$), then it holds that $f \star_{2n} g = f \cdot g$.

In our overview of the Schönhage-Strassen multiplication for n-digit numbers, we mentioned that the method considers an evaluation/interpolation approach using the $2n$-th complex roots of unity. Again, we generalize this approach to arbitrary rings.

Definition 2.1.2 (Root of Unity and Discrete Fourier Transform (DFT)). *Let $\omega \in R$, and $N \in \mathbb{N}$. We call ω an N-th root of unity if $\omega^N = 1$. We further call ω primitive if $\omega^{N/i} - 1$ is not a zero-divisor[2] in R for any divisor i of N. For fixed ω, the* Discrete Fourier Transform *of a polynomial $f \in R[x]$ is defined as*

$$\mathrm{DFT}_\omega(f) := (f(1), f(\omega), \ldots, f(\omega^{N-1})).$$

For a vector $a = (a_0, \ldots, a_{N-1})^t \in R^N$, we define $\mathrm{DFT}_\omega(a) := \mathrm{DFT}_\omega(\sum_{i=0}^{N-1} a_i x^i)$.

We remark that there does not always exist a primitive N-th root of unity in a ring R. For instance, this is the case for $R = \mathbb{Z}$ or $R = \mathbb{R}$. The following exercise (taken from [GG03, Sec. 8]) gives a necessary and sufficient condition on the existence of a primitive root of unity in the finite field $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$.

[2] An element $a \in R$ is a zero divisor if there exists an $r \in R$ with $a \cdot r = 0 = 0_R$ or $r \cdot a = 0$. A zero-divisor does not have to be zero. For instance, $a = \bar{3} \in R = \mathbb{Z}/6\mathbb{Z}$ is a zero divisor in R as $\bar{2} \cdot \bar{3} = \bar{0}$.

Exercise 2.1.3. *Denote by* $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ *the finite field with* p *elements for some prime* p*, and let* $N \in \{1, \dots, p-1\}$*. Show that* $\mathbb{F}_p$ *contains a primitive* N*-th root of unity if and only if* N *divides* $p-1$*, and conclude that the multiplicative group* $\mathbb{F}_p^\times$ *of* $\mathbb{F}_p$ *is cyclic.*

Hints:

1. *Use (without proof)* ***Fermat's little theorem:*** For arbitrary $a \in \mathbb{Z}$ arbitrary, it holds

$$a^p \equiv a \mod p.$$

   In particular, if $a \in \{1, \dots, p-1\}$, then

$$a^{p-1} \equiv 1 \mod p.$$

2. *Let* $q \in \mathbb{N}$ *be a divisor of* $p-1$ *and* $q = q_1^{e_1} \cdots q_r^{e_r}$ *its prime factorization. For* $a \in \mathbb{F}_p^\times$*, we denote by* $\operatorname{ord}(a) := \min\{i \in \mathbb{N}_{>0} : a^i = 1\}$ *the order of* a *in* $\mathbb{F}_p^\times$*.*

   Prove the following facts:

   - $\operatorname{ord}(a) = q$ *if and only if* $a^q = 1$ *and* $a^{q/q_i} \neq 1$ *for* $i = 1, \dots, r$*.*
   - *For each* i*,* $\mathbb{F}_p^\times$ *contains an element* a_i *with* $q_i^{e_i} \mid \operatorname{ord}(a_i)$*. Conclude that there is an element* b_i *with* $\operatorname{ord}(b_i) = q_i^{e_i}$*.*
   - *If* $a, b \in \mathbb{F}_p^\times$ *are elements of coprime orders, then* $\operatorname{ord}(ab) = \operatorname{ord}(a)\operatorname{ord}(b)$*.*
   - $\mathbb{F}_p^\times$ *contains an element of order* q*.*

Lemma 2.1.4. *For* $N \in \mathbb{N}$*, suppose that there exists a primitive* N*-root of unity* ω *in* R*. For any two polynomials* $f, g \in R[x]$ *of degree less than* N*, it holds that*

$$\mathrm{DFT}_\omega(f \star_N g) = \mathrm{DFT}_\omega(f) \cdot \mathrm{DFT}_\omega(g) = (f(1) \cdot g(1), f(\omega) \cdot g(\omega), \dots, f(\omega^{N-1}) \cdot g(\omega^{N-1})).$$

Proof. There exists a polynomial $q \in R[x]$ with $f \star_N g = f \cdot g + q \cdot (x^N - 1)$. Thus, we have

$$\begin{aligned}(f \star_N g)(\omega^i) &= f(\omega^i) \cdot g(\omega^i) + q(\omega^i) \cdot ((\omega^i)^N - 1) = f(\omega^i) \cdot g(\omega^i) + q(\omega^i) \cdot ((\omega^N)^i - 1) = \\ &= f(\omega^i) \cdot g(\omega^i) + q(\omega^i) \cdot (1^i - 1) = f(\omega^i) \cdot g(\omega^i).\end{aligned}$$

□

In our overview of the Schönhage-Strassen method, one step is to compute the Discrete Fourier Transforms $\mathrm{DFT}_\omega(f)$ and $\mathrm{DFT}_\omega(g)$ of two polynomials of degree at most $n-1$, where ω is an N-th root of unity in $\mathbb{C}$, with $N := 2n$. Now from the above lemma and the fact that $f \star_N g = f \cdot g = h$, we conclude that

$$\mathrm{DFT}_\omega(h) = \mathrm{DFT}_\omega(f \cdot g) = \mathrm{DFT}_\omega(f \star_N g) = \mathrm{DFT}_\omega(f) \cdot \mathrm{DFT}_\omega(g). \tag{2.3}$$

Notice that the mapping $\mathrm{DFT}_\omega : R^N \mapsto R^N$ is given by the Vandermonde matrix

$$V_\omega := \mathrm{Vand}(1, \omega, \dots, \omega^{N-1}) = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & \omega & \cdots & \omega^{N-1} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & \omega^{N-1} & \cdots & \omega^{N(N-1)} \end{pmatrix}.$$

That is, the coefficient vector $a := (a_0, \ldots, a_{N-1})^t$ of a polynomial $f = \sum_{i=0}^{N-1} a_i \cdot x^i \in R[x]$ is mapped to the vector $v := (f(1), f(\omega), \ldots, f(\omega^{N-1}))^t = V_\omega \cdot a$. Vice versa, if v is known, then the coefficients a_i of f can be reconstructed as $a = V_\omega^{-1} \cdot v$. It turns out that a multiple of V_ω^{-1} can be easily computed.

Theorem 2.1.5. *Let ω be a primitive N-th root in R. Then, $\omega^{N-1} = \omega^{-1}$ is also a primitive N-th root of unity and $V_\omega \cdot V_{\omega^{-1}} = N \cdot \mathrm{Id}_N$, with Id_N the $N \times N$-identity matrix.*

Proof. We split the proof into four parts:

(1) $\omega^{n-1} = \omega^{-1}$ is a primitive N-th root of unity: Since

$$(\omega^{N-1})^N = (\omega^N)^{N-1} = 1^{N-1} = 1,$$

it follows that ω^{N-1} is an root of unity. Now suppose that there exists a divisor t of N and a $b \in R$ with $((\omega^{N-1})^{N/t} - 1)$. Then, multiplication with $\omega^{N/t}$ implies that

$$0 = \omega^{N/t} \cdot ((\omega^{N-1})^{N/t} - 1) = [(\omega \cdot \omega^{N-1})^{N/t} - \omega^{N/t}) \cdot b = (1 - \omega^{N/t}) \cdot b,$$

and thus $\omega^{N/t} - 1$ is a zero-divisor in R, which contradicts our assumption.

(2) $\omega^\ell - 1$ is not a zero divisor for all $\ell \in \mathbb{N}$ with $1 \leq \ell < N$: Let $g := \gcd(\ell, N)$ be the greatest common divisor of ℓ and $\mathbb{N}$. Then, there exist integers[3] s and t with $s \cdot \ell + t \cdot N = g$. Since $g < n$, there exists a prime divisor p of N that divides N/g, and thus g divides N/p. Hence, we obtain

$$\omega^{N/p} - 1 = (\omega^g)^{\frac{N}{pg}} - 1 = (\omega^g - 1) \cdot \underbrace{\sum_{i=0}^{\frac{N}{pg}-1} \omega^{i \cdot g}}_{=:r}.$$

Now, suppose that there exists a $b \in R$ with $b \cdot (\omega^g - 1) = 0$, then we also have $b \cdot (\omega^{N/p} - 1) = 0$, and thus $b = 0$ as ω is not a zero divisor. This shows that $\omega^g - 1$ is not a zero divisor as well. Notice that $\omega^\ell - 1$ divides $\omega^{s\ell} - 1 = (\omega^\ell - 1) \cdot \sum_{i=0}^{s-1} \omega^{i\ell}$, and since

$$\omega^{s\ell} - 1 = \omega^{s\ell} \cdot (\omega^N)^t - 1 = \omega^{s\ell + tN} - 1 = \omega^g - 1$$

we conclude that $\omega^\ell - 1$ also divides $\omega^g - 1$. It follows that $\omega^\ell - 1$ is not a zero divisor as $b \cdot (\omega^\ell - 1) = 0$ implies that $b \cdot (\omega^g - 1) = 0$, and thus $b = 0$.

(3) It holds that $\sum_{0 \leq j < N} \omega^{\ell j} = 0$ for any $\ell \in \mathbb{N}$ with $1 \leq \ell < N$: It holds that

$$(\omega^\ell - 1) \cdot \sum_{j=0}^{N-1} \omega^{\ell j} = \omega^{\ell N} - 1 = 0,$$

and thus $\sum_{j=0}^{N-1} \omega^{\ell j} = 0$ as $\omega^\ell - 1$ is not a zero divisor.

(4) $V_\omega \cdot V_{\omega^{-1}} = N \cdot \mathrm{Id}_N$: The (i, k)-th entry c_{ij} of $V_\omega \cdot V_{\omega^{-1}}$ is given as

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

where we used (3) for the case $i \neq k$. $\square$

Algorithm 6: Fast Fourier Transform

Input : A polynomial $f = a_0 + \cdots + a_{N-1} \cdot x^{N-1} \in R[x]$, with $N = 2^k$ and $k \in \mathbb{N}_0$, and a primitive N-th root of unity $\omega \in R$.

Output: $\mathrm{DFT}_\omega(f)$.

1 **if** N=*1* **then**

2   **return** a_0

3 Compute $\omega_i := \omega^i$ for $i = 0, \dots, N-1$

4 $f^{\mathrm{ev}} := \sum_{i=0}^{N/2-1} a_{2i} \cdot x^i$ and $f^{\mathrm{odd}} := \sum_{i=0}^{N/2-1} a_{2i+1} \cdot x^i$

5 Call Algorithm 6 recursively to compute

$$(d_0^{\mathrm{ev}}, \dots, d_{N/2-1}^{\mathrm{ev}}) := \mathrm{DFT}_{\omega^2}(f^{\mathrm{ev}})$$

and

$$(d_0^{\mathrm{odd}}, \dots, d_{N/2-1}^{\mathrm{odd}}) := \mathrm{DFT}_{\omega^2}(f^{\mathrm{odd}}).$$

for $i = 1, \dots, N-1$ **do**

6   Let $j = i \mod N/2$. Compute

$$d_i := d_j^{\mathrm{ev}} + \omega_i \cdot d_j^{\mathrm{odd}}.$$

7 **return** $(d_0, \dots, d_{N-1})$

Exercise 2.1.6. *Let* $\mathbb{F} = \mathbb{Z}/29\mathbb{Z}$.

1. *Find a primitive* 4*-th root of unity* $\omega \in \mathbb{F}$ *and compute its inverse* $\omega^{-1} \in \mathbb{F}$.
2. *Check that the product of the two matrices* DFT_ω *and* $\mathrm{DFT}_{\omega^{-1}}$ *equals* $4 \cdot \mathrm{Id}_4$.

Theorem 2.1.5 shows that polynomial interpolation is essentially the same as polynomial evaluation when considering the N-th roots of unity as interpolation points. In particular, applying $\mathrm{DFT}_{\omega^{-1}}$ to both sides of (2.3), we obtain for the coefficient vector $c := (c_0, \dots, c_{N-1})^t$ of $h = \sum_{i=0}^{N-1} c_i x^i$ that

$$N \cdot c = \mathrm{DFT}_{\omega^{-1}}(\mathrm{DFT}_\omega(h)) = \mathrm{DFT}_{\omega^{-1}}(\mathrm{DFT}_\omega(f) \cdot \mathrm{DFT}_\omega(g)). \tag{2.4}$$

Hence, for the evaluation/interpolation step in the Schönhage-Strassen algorithm, we need to carry out three computations of a DFT plus one pointwise multiplication of two DFTs. We next describe an efficient method [CT65] due to Cooley und Tukey (from 1965) for computing the discrete Fourier Transform $\mathrm{DFT}_\omega(f)$ for some polynomial f of degree less than $N-1$ and ω a primitive N-th root of unity.[4] In what follows, we assume that R *supports the* FFT, that is, it contains an N-th root of unity for any $N = 2^k$, with $k \in \mathbb{N}$. In the following considerations, we further assume that N is such a power of two. We can now write a

[3]This follows from the extended Euclidean Algorithm, which we will treat in detail in the next chapter.

[4]In fact, it was Gauss who invented the algorithm already 160 years earlier. Cooley and Tukey rediscovered and popularized the method. The algorithm has a series of applications in engineering, applied mathematics, and the natural sciences. The original paper from 1965 has more than 13400 citations!

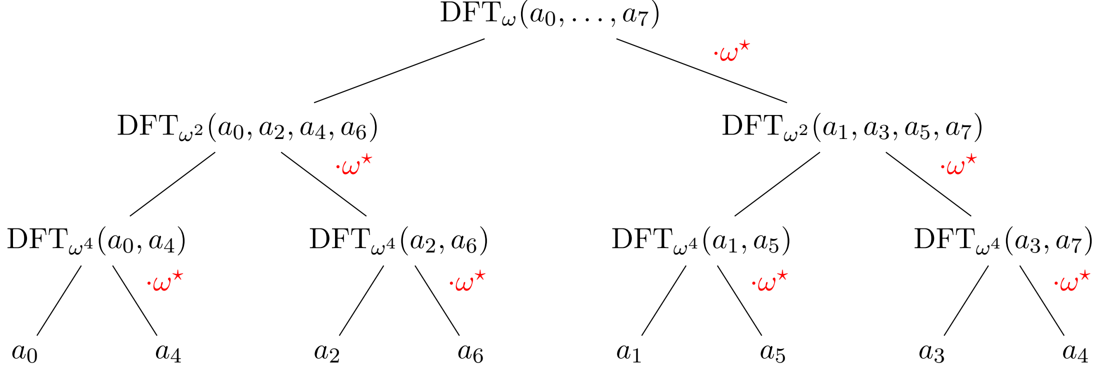

Figure 2.2: Starting with the coefficients $a_i = \mathrm{DFT}_{\omega^8}(a_i)$ of f, we iteratively compute four DFT's of length 2, two DFT's of length 4, and eventually $\mathrm{DFT}_\omega(f)$, which has length 8. In Step ℓ, the i-th entry of a Discrete Fourier Transform of size $N/2^\ell$ is computed as the the sum of the j-th entry of the left child and the j-th entry of the right child multiplied by ω^i (illustrated by the edge labelling "$\cdot\omega^\star$" in the above picture), where $j = i \mod N/2^{\ell+1}$.

polynomial $f(x) = a_0 + \cdots + a_{N-1} \cdot x^{N-1} \in R[x]$ as

$$f(x) = \sum_{i=0}^{N/2-1} a_{2i} \cdot x^{2i} + \sum_{i=0}^{N/2-1} a_{2i+1} \cdot x^{2i+1} = f^{\mathrm{ev}}(x^2) + x \cdot f^{\mathrm{odd}}(x^2),$$

with $f^{\mathrm{ev}} := \sum_{i=0}^{N/2-1} a_{2i} \cdot x^i$ and $f^{\mathrm{odd}} := \sum_{i=0}^{N/2-1} a_{2i+1} \cdot x^i$. Plugging $x = \omega^i$ into the above equation then yields that

$$f(\omega^i) = f^{\mathrm{ev}}(\omega^{2i}) + \omega^i \cdot f^{\mathrm{odd}}(\omega^{2i}). \tag{2.5}$$

Notice that ω^2 is a primitive $N/2$-root, hence the computation of $\mathrm{DFT}_\omega(f) = (d_0, \ldots, d_{N-1})$ can be reduced to the computation of the two Discrete Fourier Transforms $\mathrm{DFT}_{\omega^2}(f^{\mathrm{ev}}) = (d_0^{\mathrm{ev}}, \ldots, d_{N/2-1}^{\mathrm{ev}})$ and $\mathrm{DFT}_{\omega^2}(f^{\mathrm{odd}}) = (d_0^{\mathrm{odd}}, \ldots, d_{N/2-1}^{\mathrm{odd}})$ followed by the computation of $d_i := d_j^{\mathrm{ev}} + \omega^i \cdot d_j^{\mathrm{odd}}$ for all $i = 0, \ldots, N$ and $j = i \mod N/2$; see Algorithm 6.

In terms of complexity, this means that we can compute a Discrete Fourier Transform of size N by computing two Discrete Fourier Transforms of size $N/2$ plus $3N$ additional additions and multiplications (by powers of ω). If we use $T(N)$ to denote the number of arithmetic operations in R that are needed in the worst case to compute the Discrete Fourier Transform $\mathrm{DFT}_\omega(f)$ for a polynomial f of degree less than N and a primitive N-th root of unity ω, the above consideration implies that

$$T(N) \leq 2 \cdot T(N/2) + 3 \cdot N.$$

Hence, we obtain the following result:

Theorem 2.1.7. *Let $f \in R[x]$ be a polynomial of degree less than N and ω be a primitive N-th root of unity ω in R, then Algorithm 6 computes $\mathrm{DFT}_\omega(f)$ using $O(N \log N)$ arithmetic operations in R.*

For an illustration of the FFT Algorithm when applied to a polynomial $f = a_0 + \cdots + a_7 \cdot x^7 \in R[x]$ of degree 7 and ω a primitive 8-th root of unity, see Figure 2.2.

Algorithm 7: Fast Convolution

Input : A commutative ring R, two polynomials $f, g \in R[x]$ of degree less than $N = 2^k$, with $k \in \mathbb{N}_0$, and a primitive N-th root of unity $\omega \in R$.

Output: $f \star_N g$.

1 Compute:
2 $\omega^{-1} = \omega^{N-1}$.
3 $D_f := \mathrm{DFT}_\omega(f)$ and $D_g := \mathrm{DFT}_\omega(g)$
4 $D_h := D_f \cdot D_g$
5 $E := \frac{\mathrm{DFT}_{\omega^{-1}}(D_h)}{N}$
6 **return** E

From (2.4) and the FFT algorithm, we can can now directly derive an efficient algorithm for computing the convolution $f \star_N g$ of two polynomials $f, g \in R[x]$ of degree less than N. Namely, we first compute $\mathrm{DFT}_\omega(f)$ and $\mathrm{DFT}_\omega(g)$ and their pointwise product P. Then, we compute $\mathrm{DFT}_{\omega^{-1}}(P)$ and divide each of its entries by N; see Algorithm 7. Notice that all but the last operation use $O(n \log n)$ arithmetic operations in R. According to Section 1.4, the division by N is relatively cheap in the special case where $R = \mathbb{C}$, however, it might be an entirely non-trivial task for a different ring.

Theorem 2.1.8. *Let $f, g \in R[x]$ be polynomials of degree less than $N = 2^k$ with $k \in \mathbb{N}$. Suppose that a primitive N-th root of unity ω in R is given. Then, Algorithm 7 computes $f \star_N g$ using $O(N \log N)$ arithmetic operations in R plus N divisions by N.*

For two polynomial $f, g \in R[x]$ of degree n or less, it holds that $f \cdot g = f \star_N g$, with $N := 2^{\lceil \log n \rceil + 1}$. Hence, if a primitive N-th root of unity is given, then Algorithm 7 computes the product of f and g using $O(n \log n)$ arithmetic operations in R plus N divisions by N.

Corollary 2.1.9. *Let $f, g \in R[x]$ be polynomials of degree less than n, and $N := 2^{\lceil \log n \rceil + 1}$. If a primitive N-th root of unity ω in R is given, then Algorithm 6 computes $f \cdot g$ using $O(N \log N) = O(n \log n)$ arithmetic operations in R plus N divisions by N.*

2.1.3 Fast Multiplication in $\mathbb{Z}$ and $\mathbb{Z}[x]$.

We are now coming back to our original problem of computing the product of two integer polynomials $f, g \in \mathbb{Z}[x]$ of degree less than n. We further assume that the coefficients of f and g have absolute value less than 2^L. Since $\mathbb{Z}$ does not contain a primitive N-th root of unity for any integer $N > 2$, we cannot directly apply the above approach (with $R = \mathbb{Z}$) to compute the product $f \cdot g$. However, since f, g can also be considered as polynomials with complex coefficients and since $\mathbb{C}$ supports the FFT, Corollary 2.1.9 implies that we can compute the product using $O(n \log n)$ arithmetic operations in $\mathbb{C}$ plus N divisions by N, where $N := 2^{\lceil \log n \rceil + 1}$. As already mentioned in our overview of the Schönhage-Strassen method, we need to address the problem that these operations can only be carried out with approximate arithmetic. Now, suppose that we use fixed point arithmetic with base 2 and a fixed precision ρ in each step of Algorithm 7. Then, we aim to answer the question how large ρ needs to be chosen such that the final error is smaller than $1/2$, which would allow us to derive the exact coefficients of $f \cdot g$ from the computed approximations; see (2.2) for the example we gave at

the beginning of the chapter. Before running Algorithm 7, we first compute an approximation $\tilde{\omega} \in \mathbb{F} = \mathbb{F}_{2,\rho}$ of the N-th root of unity $\omega = \cos(2\pi/N) + \mathbf{i}\cdot\sin(2\pi/N)$ such that $|\tilde{\omega}-\omega| < 2^{-\rho}$. According to Exercise 1.4.4 and Exercise 1.4.5, the cost for this computation is bounded by $O(\rho^c)$ for some constant c. From Theorem 1.3.3, we further conclude that

$$|P(\omega) - P_{\mathbb{F}}(\omega)| < 4N^2 \cdot 2^{-\rho} \cdot \max(1, |\omega|)^{N-1} = 4N^2 \cdot 2^{-\rho}$$

for $P(x) := x^i$ and an arbitrary $i \in \{0, \dots, N-1\}$. Hence, recursively taking powers of the approximation $\tilde{\omega}_1 := \tilde{\omega}$ and using fixed point arithmetic in each step yields approximations $\tilde{\omega}_i$ of $\omega_i := \omega^i$ with $|\tilde{\omega}_i - \omega_i| < 4N^2 \cdot 2^{-\rho}$.

In the Fast Fourier Transform, the entries of $\mathrm{DFT}_\omega(f) = (c_0, \dots, c_{N-1})$ are recursively computed from the coefficients of $f = a_0 + \cdots + a_{N-1} \cdot x^{N-1}$. That is, at the highest level of the recursion, we start with a suitable permutation of the coefficients a_i and recursively compute corresponding DFT's of size $2, 4, 8, \dots$ until we obtain $\mathrm{DFT}_\omega(f)$. More specifically, at level ℓ of the recursion, the i-th entry d_i of each DFT of size $N/2^{\ell-1}$ is computed as

$$d_i = d_j^{\mathrm{ev}} + \omega^i \cdot d_j^{\mathrm{odd}}$$

where d_j^{ev} and d_j^{odd} are the j-th entries of previously computed DFT's of size $N/2^\ell$ and $j = i \mod N/2^\ell$. Now suppose that we use a precision $\rho > 2(\log N + 1)$ and that we have already computed approximations $\tilde{d_j^{\mathrm{ev}}}$ and $\tilde{d_j^{\mathrm{odd}}}$ of the entries d_j^{ev} and d_j^{odd}, respectively, with $|\tilde{d_j^{\mathrm{ev}}} - d_j^{\mathrm{ev}}|, |\tilde{d_j^{\mathrm{odd}}} - d_j^{\mathrm{odd}}| < \epsilon$. Then $\tilde{d}_i := \tilde{d_j^{\mathrm{ev}}} \,\tilde{+}\, \tilde{w}_i \,\tilde{\cdot}\, \tilde{d_j^{\mathrm{odd}}}$ constitutes an approximation of d_i with

$$\begin{aligned}
|d_i - \tilde{d}_i| &< 2^{-\rho+1} + \epsilon + \epsilon \cdot |\omega_i| + 4N^2 \cdot 2^{-\rho} \cdot |d_j^{\mathrm{odd}}| + 4N^2 \cdot 2^{-\rho} \cdot \epsilon \\
&= \epsilon \cdot (2 + 4N^2 \cdot 2^{-\rho}) + 2^{-\rho} \cdot (2 + 4N^2 \cdot d_j^{\mathrm{odd}}) \\
&< 3\epsilon + 4N^2 \cdot 2^{-\rho} \cdot (1 + |d_j^{\mathrm{odd}}|), \qquad (2.6)
\end{aligned}$$

where we used our bounds (1.2) and (1.3) for the error that occurs when using fixed point arithmetic. Further notice that d_j^{odd} is an entry of $\mathrm{DFT}_{\omega^{N/2^\ell}}(\hat{f})$, where $\hat{f}$ is an integer polynomial of degree less than $N/2^\ell$, whose coefficients form a subset of the set of coefficients of f. Hence, we have $d_j^{\mathrm{odd}} < \frac{N}{2^\ell} \cdot 2^L < \frac{N}{2} \cdot 2^L$, and thus (2.6) yields

$$|d_i - \tilde{d}_i| < 8 \cdot \max(\epsilon, 4N^3 \cdot 2^L \cdot 2^{-\rho})$$

Since there are $\log N$ steps in the recursion, we conclude that the computed approximations of the entries of $\mathrm{DFT}_\omega(f)$ differ from the exact values by at most $8^{\log N}$ times the maximum of the input error[5] for the coefficients a_i and the value $4N^3 \cdot 2^L \cdot 2^{-\rho}$. Hence, the total error is bounded by $4N^6 \cdot 2^L \cdot 2^{-\rho}$. The same bound then also applies to the error that we obtain when computing $\mathrm{DFT}_\omega(g)$ with fixed point arithmetic.

We may now assume that we have computed approximations $\tilde{D}_f = (\tilde{f}_0, \dots, \tilde{f}_{N-1})$ and $\tilde{D}_g = (\tilde{g}_0, \dots, \tilde{g}_{N-1})$ of

$$D_f = (f_0, \dots, f_{N-1}) := \mathrm{DFT}_\omega(f)$$

and

$$D_g = (g_0, \dots, g_{N-1}) := \mathrm{DFT}_\omega(g)$$

[5] Here, the coefficients are given exactly, and thus the input error is zero. However, our analysis also applies to the case where only approximations $\tilde{a}_i$ of the coefficients a_i are given. Then the total error is bounded by $8^{\log N} \cdot \max(4N^3 \cdot 2^L \cdot 2^{-\rho}, \max_i |a_i - \tilde{a}_i|)$.

to an absolute error bounded by $4N^6 \cdot 2^L \cdot 2^{-\rho}$. Pointwise multiplication of $\tilde{D}_f$ and $\tilde{D}_g$ (again using fixed point arithmetic with precision ρ) then yields an approximation $\tilde{D}_h = (\tilde{h}_0, \ldots, \tilde{h}_{N-1}) := \tilde{D}_f \tilde{\cdot} \tilde{D}_g$ of $D_h = \mathrm{DFT}_\omega(h) = (h_0, \ldots, h_{N-1})$, and according to (1.3), the absolute error $|h_i - \tilde{h}_i|$ is bounded by

$$2^{-\rho} + 4N^6 \cdot 2^L \cdot 2^{-\rho} \cdot N \cdot 2^L \cdot (|f_i| + |g_i|) + (4N^6 \cdot 2^L \cdot 2^{-\rho})^2 < 32N^{12} \cdot 2^{2L} \cdot 2^{-\rho}$$

as $|f_i|, |g_i| \leq N \cdot 2^L$ for all $i = 0, \ldots, N-1$.

It remains to estimate the error when computing $\frac{1}{N} \cdot \mathrm{DFT}_{\omega^{-1}}(\tilde{D}_h)$ with fixed point arithmetic. In completely analogous manner as above, one shows that the output error of the computation of $\mathrm{DFT}_{\omega^{-1}}(\tilde{D}_h)$ is bounded by $8^{\log N} \cdot \max_i |h_i - \tilde{h}_i| < 32N^{15} \cdot 2^{2L} \cdot 2^{-\rho}$. The final division by N amounts for a shift by $\log N$ bits as N is a power of two, which shows that the total error is at most $32N^{14} \cdot 2^{2L} \cdot 2^{-\rho}$. Hence, in order to guarantee an output error of less than $1/2$, it suffices to consider a precision

$$\rho > \rho_0 := \log(64N^{14} \cdot 2^{2L}) = 6 + 14 \log N + 2L = O(\log n + L).$$

Each of the intermediate results is an approximation of an entry of some $\mathrm{DFT}_{\omega^{N/2^\ell}}(\hat{f})$, where $\ell \in \{0, \ldots, \log N\}$ and $\hat{f}$ is a polynomial of degree at most N with integer coefficients that form a subset of the set of coefficients of f, g, or $f \cdot g$. Hence, each of these coefficients has absolute value less than $N \cdot 2^L$. It follows that each intermediate result is a fixed point number of length $2^{O(\log N + L + \rho)}$. Since we succeed for $\rho = 2\rho_0$, it follows that the computation of $f \cdot g$ uses $O(n \log n)$ arithmetic operations of fixed numbers of length $O(\log n + L)$. The following result then follows directly.

Theorem 2.1.10. *Let $f, g \in \mathbb{Z}[x]$ be polynomials of degree less than n and with one-digit integer coefficients. Then, the product $f \cdot g$ can be computed using $O(n \log n \cdot M(\log n))$ primitive operations.*

From the above theorem, we can now derive the following result on the cost for multiplying two integers of length less than n:

Theorem 2.1.11. *Given two integers a and b of length less than n, the product $a \cdot b$ can be computed using $O(n \log n \cdot M(\log n)) = O(n(\log n)^{2+\epsilon})$ primitive operations, where ϵ is an arbitrary but fixed constant. Furthermore, we can compute a dyadic approximation $\tilde{q}$ with $|\tilde{q} - a/b| < 2^{-L}$ using $O((n + L) \cdot (\log(n + L))^{3+\epsilon})$ primitive operations.*

Proof. The polynomials $f = a^{(0)} + \cdots + a^{(n-1)} \cdot x^{n-1}$ and $f = b^{(0)} + \cdots + b^{(n-1)} \cdot x^{n-1}$ in (2.1) have one digit coefficients, hence we can compute the product $h = f \cdot g$ using $O(n \log n \cdot M(\log n))$ primitive operations according to Theorem 2.1.10. The computation of $a \cdot b = h(B)$ is bounded by $O(n \log n)$ primitive operations as this step requires $O(n)$ additions, each involving an integer of length $O(n)$ and an integer of length $O(\log n)$. The bound on the cost for the approximate division then follows directly from Theorem 1.4.2. □

You might wonder why we have not given a more general bound in Theorem 2.1.10 that applies to polynomials with integer coefficients of arbitrary length. Namely, if the length of the coefficients is bounded by L, then our above considerations show that the cost for multiplying f and g is bounded by $O(n \log n \cdot M(\log n + L))$ primitive operations *if* a sufficiently good approximation $\tilde{\omega}$ of ω with $|\omega - \tilde{\omega}| = 2^{-\Omega(L + \log n)}$ is already computed. But this is actually

critical as we have only shown that the cost for this step is bounded by $O((\log n + L)^c)$. Hence, in order to derive a bound on the total running time that is near-linear in L, we need a different approach.[6] Here, we consider an approach known as *Kronecker substitution*. The crucial idea is that if an upper bound on the length of the coefficients of a polynomial $F(x) = c_0 + c_1 \cdot x + \cdots c_n \cdot x^n$ is known, then one can recover the coefficients from the value of F at a single point. Namely, suppose that each c_i has length less than L (with respect to some base B), then evaluating F at $x = B^L$ yields

$$F(B^L) = c_0 + B^L \cdot c_1 + B^{2L} \cdot c_2 + \cdot + B^{nL} \cdot c_n.$$

Since each c_i has length less than L and since multiplication by B^{iL} yields a shift of c_i by iL digits, the coefficients can directly be read off the value $F(B^L)$ as there is no overlap. As an example, consider the polynomial $F(x) = 12 + 34 \cdot x + 45 \cdot x^2 + 67 \cdot x^3 + 8x^4$, where we have $f(1000) = \mathbf{8067045034012}$. *Kronecker substitution* now allows us to reduce the problem of computing the product $h = f \cdot g$ of two polynomials $f, g \in \mathbb{Z}[x]$ with coefficients of length less than L to the problem of multiplying two integers of length $O(n(L + \log n))$. This works as follows: Each coefficient of h has length less than $L' := \lceil 2L + \log n \rceil$. Hence, we can directly derive the coefficients of h from the value $h(B^{L'}) = f(B^{L'}) \cdot g(B^{L'})$. Evaluating f (or g) at $x = B^{L'}$ amounts for shifting the corresponding coefficients a_i (or b_i) by iL' digits and summing up the so obtained numbers. This step uses $O(n(L + \log n))$ primitive operations. The values $f(B^{L'})$ and $g(B^{L'})$ are integers of length $O(n(L+\log n))$, and thus we can compute their product using $\tilde{O}(nL)$ primitive operations, where the $\tilde{O}$ - notation indicates that we are omitting poly-logarithmic factors in the input. That is, $\tilde{O}(N) = O(N \cdot \log^c N)$ for some constant c. We fix this result:

Theorem 2.1.12. *Let $f, g \in \mathbb{Z}[x]$ be polynomials of degree less than n and with integer coefficients of length less than L. Then, the product $f \cdot g$ can be computed using $\tilde{O}(nL)$ primitive operations.*

We also state the following complexity bound for the evaluation of a polynomial $f \in \mathbb{Z}[x]$ at a rational point.

Theorem 2.1.13. *Let $f \in \mathbb{Z}[x]$ be a polynomial of degree n with coefficients of length less than 2^L, and let $x_0 = p/q$ be a rational point with integers p, q of length less than ℓ. Then, using Horner Evaluation, we can compute the value $f(x_0)$ using $\tilde{O}(n^2(\ell + L))$ primitive operations.*

Proof. We define $f_0 := a_n$, and

$$f_{i+1} = x \cdot f_i + a_{n-i-1} \in \mathbb{Z}[x] \text{ for } i = 0, \ldots, n-1.$$

Notice that, when using Horner Evaluation, we recursively compute the values $v_i := f_i(x_0)$. Since f_i is a polynomial of degree i, we conclude that $v_i = \frac{p_i}{q_i}$ is a rational number with denominator $q_i = q^i$ and numerator p_i of length less than $\log n + L + i \cdot \ell$. Hence, computing $f_{i+1}(x_0)$ from f_i amount for a constant number of arithmetic operations of integers of length $O(\log n + L + i \cdot \ell) = O(L + n \cdot \ell)$. Each such operations uses $\tilde{O}(L + n \cdot \ell)$ primitive operations, thus the claimed bound follows. □

[6] In fact, one can show that a such an approximation of ω can be computed in a number of primitive operations that is near linear in n and L. However, this requires to introduce some additional tools that we will treat only in one of the following chapters.

In the following exercise, we present a different evaluation method that yields a complexity bound that is near-optimal.

Exercise 2.1.14. *You already know Horner's method for polynomial evaluation. An alternative method is due to* Estrin: *In order to evaluate a polynomial $f(x) = a_0 + \cdots + a_n \cdot x^n$, let $m := 2^{\lceil \log n \rceil - 1}$ and write f as*

$$f(x) = \underbrace{(a_n x^m + a_{n-1} x^{m-1} + \cdots a_m)}_{=:f_H(x)} \cdot x^m + \underbrace{a_{m-1} x^{m-1} + a_{m-2} x^{m-2} + \cdots + a_0}_{=:f_L(x)},$$

where f_H and f_L are polynomials of degree at most m. Recursively evaluate f_H and f_L and reconstruct $f(x) = f_H(x) \cdot x^m + f_L(x)$.

Show that Estrin's method uses only $\tilde{O}(n(L + \ell))$ primitive operations to compute $f(x_0)$ if f has integer coefficients of length L and $x_0 = p/q \in \mathbb{Q}$ is a rational point with integers p, q of length less than ℓ.

Remark. Instead of using fixed-point arithmetic in each step of the Fast Convolution algorithm, we could have used fixed-point interval arithmetic. A corresponding analysis then yields comparable bounds on the needed precision. However, we might again profit from the fact that each interval approximation of some value carries a canonical adaptive bound on the approximation error (i.e. the width of the computed interval), whereas we have to work with a worst-case error bound if fixed point arithmetic is used. For instance, for the Schönhage-Strassen method, this means that we can iteratively increase the precision until the final interval approximations of the coefficients of h have width less than $1/2$ or, alternatively, until they contain only one integer.

2.1.4 Fast Multiplication over arbitrary Rings*

We have already shown that if R supports the FFT and if division by 2 can be carried out in an efficient manner in R, then computing the product of two polynomials $f, g \in R[x]$ of degree less than n uses only $O(n \log n)$ arithmetic operations in R. Can we also give a comparable bound for arbitrary commutative rings that do not support the FFT? The answer is yes, however, we will not give the details here, but only a rough idea of the approach. There are certain cases that need to be distinguished and the actual approach is slightly more involved than what we describe below. The interested reader should have a look into Section 8.3 of the textbook [GG03] "Modern Computer Algebra" from von zur Gathen und Gerhard, which contains a comprehensive description of the algorithm and its analysis.

The crucial idea underlying the approach is to adjoin a so-called *virtual root of unity.* For this, suppose that 2 is a unit in R and that $N = 2^k$ is a power of two. Then, we define $D_N := R[x]/\langle x^N + 1 \rangle$, an extension of the ring R. Since $x^{2N} = (x^N)^2 = 1 \mod (x^N + 1)$, we conclude that $\omega := x \mod (x^N + 1)$ is a $2N$-th root of unity in D_N. Suppose that, for some divisor ℓ of $2N$ and some $b \in D$, we have $b \cdot (\omega^{2N/\ell} - 1) \cdot b = 0$. Since N is a power of two, the same holds for ℓ, and thus we may write $\omega^{2N/\ell} - 1$ as $\omega^{N/\ell'} - 1$ with $\ell' = \ell/2$. Hence, we obtain

$$b \cdot (\omega^N - 1) = b \cdot (\omega^{N/\ell' - 1}) \cdot \sum_{i=0}^{\ell' - 1} \omega^{in/\ell'} = 0.$$

Since $\omega^N - 1 = -2$ is a unit in R, it is also a unit in D_N, and thus we must have $b = 0$. This shows that ω is a primitive $2N$-th root of unity. Now, how does this help to multiply two

polynomials $f, g \in R[x]$ of degree less than n? Remember that, when multiplying two integers of length n using either the Toom-Cook approach or the Schönhage-Strassen method, we first partitioned each integer into k blocks of size n/k and derived corresponding polynomials of degree k whose coefficients are integers of length n/k. We now proceed in a similar way with a suitably chosen k. More specifically, we first partition the coefficients of $f = \sum_{i=0}^{n-1} a_i x^i$ and $g = \sum_{i=0}^{n-1} b_i x^i$ into blocks of size $\sqrt{N}$, where $N := 2^{\lceil \log 2n \rceil}$. That is, we write

$$f(x) = \sum_{j=0}^{\sqrt{N}-1} f_j(x) \cdot x^{\sqrt{N} \cdot j} \text{ and } g(x) = \sum_{j=0}^{\sqrt{N}-1} g_j(x) \cdot x^{\sqrt{N} \cdot j},$$

with polynomials f_j and g_j of degree less than $\sqrt{N}$. Then, we consider polynomials F and G in $R[x][y]$ of degree less than $\sqrt{N}$ (in the variable y) with coefficients in $R[x]$ of degree less than $\sqrt{N}$:

$$F(x) := \sum_{j=0}^{\sqrt{N}-1} f_j(x) \cdot y^j \text{ and } G(x) := \sum_{j=0}^{\sqrt{N}-1} g_j(x) \cdot y^j$$

such that $f(x) = F(x, x^{\sqrt{N}})$ and $g(x) = G(x, x^{\sqrt{N}})$. We now consider the coefficients of F and G as elements in the ring $D_{2\sqrt{N}}$. Computationally, nothing happens at this step, however, in order to distinguish the polynomials F and G, which are contained in $R[x][y]$, from their corresponding images in $D_{2\sqrt{N}}[y]$, we use F^* and G^* to denote these images. Notice that since the coefficients of the product $H := F \cdot G \in R[x][y]$ are polynomials of degree less than $2\sqrt{D}$, they coincide with the corresponding coefficients of the product $H^* := F^* \cdot G^* \in D_{2\sqrt{N}}[y]$. This shows that we can reduce the computation of H (and thus also that of $h = f \cdot g$) to that of H^*. What we have gained with this approach is that since $D_{2\sqrt{N}}$ supports the DFT, we can use the fast convolution algorithm to compute H^*. For the latter computation, we need three FFT computations of size $2\sqrt{N}$ over the ring $D_{2\sqrt{N}}$ plus $2\sqrt{N}$ *essential multiplications* in $D_{2\sqrt{N}}$. Notice that the remaining multiplications in the FFT's are easy as each such multiplication just amounts for a multiplication by x^i modulo $x^{2\sqrt{N}} + 1$. Each essential multiplication amounts for computing the product of two polynomials in $R[x]$ of degree less than $2\sqrt{N}$. For these multiplications, we then call the algorithm recursively. A careful analysis then yields the claimed complexity bound as given in Theorem 2.1.15.

You may notice that $\sqrt{N}$ may not always be an integer, in particular, when calling the algorithm recursively for an N that is different from the initial one. In this case, one has to consider a corresponding rounding to the next power of two. We further remark that there is a variant of the approach that also works for rings, where 3 is a unit. One can further combine the latter two methods to an algorithm to compute the product of $f, g \in R[x]$, where R is an arbitrary commutative ring with 1. Details can be found in [GG03, Sec. 8.3].

Theorem 2.1.15. *Let R be a commutative ring with 1 and $f, g \in R[x]$ polynomials of degree less than n. The product of f and g can be computed using $O(n \log n \log \log n)$ arithmetic operations in R.*

The following Exercise gives an idea of the approach sketched above. Therein, we describe a simplified variant of one of the two algorithms for integer multiplication that Schönhage and Strassen published in their original paper from 1971.

Exercise 2.1.16. *Let $n = 2^{2^k}$ with $k \in \mathbb{N}$.*

(a) *Show that* $\omega := 8$ *is a primitive* $\sqrt{n}$*-th root of unity in* $R := \mathbb{Z}/(2^{3\sqrt{n}}+1)$.

(b) *Let* $a = a_{n-1}a_{n-2}\ldots a_0$ *and* $b = b_{n-1}b_{n-2}\ldots b_0$ *be two integers of length* n. *Consider the integer polynomials*

$$f(x) := \sum_{i=0}^{\sqrt{n}-1}(a_{(i+1)\sqrt{n}-1}\ldots a_{i\sqrt{n}+1}a_{i\sqrt{n}}) \cdot x^i$$
$$g(x) := \sum_{i=0}^{\sqrt{n}-1}(b_{(i+1)\sqrt{n}-1}\ldots b_{i\sqrt{n}+1}b_{i\sqrt{n}}) \cdot x^i,$$

and their images $f^* := f \mod (2^{3\sqrt{n}}+1)$ *and* $g^* := f \mod (2^{3\sqrt{n}}+1)$ *in* $R[x]$. *Show that the coefficients of* $h^* = f^* \star_{2\sqrt{n}} g^* \in R[x]$ *equal the coefficients of* $f \cdot g \in \mathbb{Z}[x]$, *and conclude that* h *can be computed with* $O(n \log n)$ *arithmetic operations in* R.

(c) *Notice that, for computing* h^*, *we need only* $2\sqrt{n}$ *essential multiplications in* R, *whereas the remaining multiplications are multiplications by powers of* ω. *Which complexity bound can you derive for the computation of* $a \cdot b$ *when using the approach recursively for the essential multiplications?*

Hint: You should first prove that each of these essential multiplications can be reduced to a constant number of additions and multiplications of integers of length $\sqrt{n}$.

2.2 Fast Polynomial Division and Applications

We start with the following definition of a Euclidean domain.

Definition 2.2.1. *A Euclidean domain is an integral domain*[7] R *together with a function* $d : R \mapsto \mathbb{N} \cup \{-\infty\}$ *if for all* $a, b \in R$, *with* $b \neq 0$, *there exist* $q, r \in R$ *with* $a = q \cdot b + r$ *and* $d(r) < d(b)$.

Exercise 2.2.2. *For* $R = \mathbb{Z}$ *and* $R = F[x]$, *with* F *an arbitrary field, give a function* $d : R \mapsto \mathbb{N} \cup \{-\infty\}$ *such that* R *together with* d *is a Euclidean domain. Does there exist such a function* d *such that* $R = \mathbb{Z}[x]$ *is a Euclidean domain?*

In what follows, we now assume that R is an integral domain and that $R[x]$ together with the degree function $d := \deg$ is a Euclidean domain. Hence, for two polynomials $f = \sum_{i=0}^{n} a_i x^i$ and $g = \sum_{i=0}^{m} b_i x^i$ in $R[x]$, with $n \geq m$, there exist polynomials $q, r \in R[x]$ with

$$f(x) = q(x) \cdot g(x) + r(x) \text{ and } \deg(r) < m. \tag{2.7}$$

Notice that the polynomials q and r in the above representation are uniquely defined if b_m is a unit in R. Namely, $f = q \cdot g + r = q^* \cdot g + r^*(x)$ implies that $r - r^* = g \cdot (q^* - q)$, and thus $r = r^*$ and $q^* = q$ as otherwise $\deg(g \cdot (q^* - q)) > \deg(r - r^*)$. Hence, we can assume that g is monic, that is, $b_m = 1$. We now give an efficient method for computing q and r. If $f = q \cdot g + r$, then

$$f(1/x) = q(1/x) \cdot g(1/x) + r(1/x)$$

[7] An integral domain is a commutative ring with 1 that contains no zero-divisor.

and thus

$$\underbrace{x^n \cdot f(1/x)}_{=:\hat{f}(x)} = \underbrace{x^{n-m} \cdot q(1/x)}_{=:\hat{q}(x)} \cdot \underbrace{x^m \cdot g(1/x)}_{=:\hat{g}(x)} + x^{n-m+1} \cdot (x^{m-1} \cdot r(1/x)).$$

Notice that $\hat{f}, \hat{g}$, and $\hat{q}$ are obtained by just reversing the coefficients of f, g, and q, respectively. In addition, since r has degree less than m, $x^{m-1} \cdot r(1/x)$ is a polynomial. Hence, we obtain

$$\hat{f}(x) = \hat{q}(x) \cdot \hat{g}(x) \mod x^{n-m+1},$$

which shows that, in order to compute $\hat{q}(x)$ (and thus $q(x)$), we can alternatively compute the product of $\hat{f}(x)$ and an inverse of $\hat{g}(x)$ modulo x^{n-m+1}. This does not sound much easier, however, there is a simple way of recursively computing an inverse $\hat{h}_i \in R[x]/\langle x^{2^i} \rangle$ of $\hat{g}(x)$ mod x^{2^i} such that $\hat{h}_i \cdot \hat{g} = 1 \mod x^{2^i}$. Notice that $\hat{g}$ has constant coefficient $\hat{g}_0 = 1$ as g is monic, and thus $\hat{h}_0 := 1$ fulfills the equation $\hat{h}_0 \cdot \hat{g}_0 \mod x$. Now, for $i \geq 0$, we recursively define:

$$\hat{h}_{i+1} := 2\hat{h}_i - \hat{g} \cdot \hat{h}_i^2 \mod x^{2^{i+1}}. \tag{2.8}$$

You might remember that we have already used a similar recursion in (1.7) to compute an approximation of $1/b$ for some integer b based on Newton iteration. The following computation now shoes that $\hat{h}_i$ has indeed the desired property. Using induction, we may assume that $\hat{h}_i \cdot \hat{g} = 1 \mod x^{2^i}$, and thus $\hat{h}_i \cdot \hat{g} = 1 + s_i \cdot x^{2^i}$ for some $s_i \in R[x]$. From (2.8), we further conclude that there exists a polynomial $s \in R[x]$ with $\hat{h}_{i+1} := 2\hat{h}_i - \hat{g} \cdot \hat{h}_i^2 + s \cdot x^{2^{i+1}}$. Hence we obtain

$$\begin{aligned}
\hat{h}_{i+1} \cdot \hat{g} &= [\hat{h}_i \cdot (2 - \hat{g} \cdot \hat{h}_i) + s \cdot x^{2^{i+1}}] \cdot \hat{g} \\
&= \hat{h}_i \cdot \hat{g} \cdot (2 - \hat{g} \cdot \hat{h}_i) && \mod x^{2^{i+1}} \\
&= (1 + s_i \cdot x^{2^i}) \cdot (2 - s_i \cdot x^{2^i}) && \mod x^{2^{i+1}} \\
&= 1 - s^2 \cdot x^{2^{i+1}} && \mod x^{2^{i+1}} \\
&= 1 && \mod x^{2^{i+1}}
\end{aligned}$$

It follows that, for $i_0 := \lceil \log(n - m + 1) \rceil$, we have

$$\hat{h}_{i_0} \cdot \hat{g} = 1 \mod x^{n-m+1}.$$

Since $\hat{q}$ has degree at most $n - m$, we can now immediately compute $\hat{q}$ from $\hat{h}_{i_0}$ as

$$\hat{q} = \hat{f} \cdot \hat{h}_{i_0} \mod x^{n-m+1}.$$

This further yields the polynomial $q(x) = x^{n-m} \cdot \hat{q}(1/x)$, and eventually the remainder

$$r(x) = f(x) - q(x) \cdot g(x).$$

We now estimate the computational cost of the above approach. The computation of $\hat{h}_i$ amounts for two multiplications and one addition of polynomials in $R[x]$ of degree $2^{2^{i+1}}$. Hence, we conclude that the cost for computing all polynomials $\hat{h}_i$ for $i = 0, \ldots, i_0$ is bounded by

$$4 \cdot [M_P(0) + M_P(2) + M_P(4) + \cdots + M_P(n)] < 8 \cdot M_P(n),$$

Algorithm 8: Fast Polynomial Division

Input : A Euclidean ring $R[x]$, a polynomial $f \in R[x]$ of degree n, and a monic polynomial $g \in R[x]$ of degree m, with $m \leq n$.

Output: Polynomials $q, r \in R[x]$ with $f = q \cdot g + r$ and $\deg r < \deg g$.

1 $\hat{f} := x^n \cdot f(1/x)$ and $\hat{g} := x^m \cdot g(1/x)$.

2 $\hat{h}_0 := 1$

3 $i_0 := \lceil \log(n - m + 1) \rceil$

4 **for** $i = 1, \ldots, i_0$ **do**

5 Recursively define

$$\hat{h}_{i+1} := 2\hat{h}_i - \hat{g} \cdot \hat{h}_i^2 \mod x^{2^{i+1}}$$

6 $\hat{q} := \hat{f} \cdot \hat{h}_{i_0} \mod x^{n-m+1}$

7 $q := x^{n-m} \cdot \hat{q}(1/x)$

8 $r := f - q \cdot g$

9 **return** q, r

where $M_P(N)$ denotes the cost for adding or multiplying two polynomials in $R[x]$ of degree at most N. According to Theorem 2.1.15, we have $M_P(n) = O(n \log n \log \log n)$. The cost for the last two steps is comparable as there are two multiplications and one addition of polynomials of degree n or less. We fix this result:

Theorem 2.2.3. *Let $f \in R[x]$ be a polynomial of degree n, and g a monic polynomial of degree m, with $m \leq n$. Then, we can compute polynomials $q, r \in R[x]$ with*

$$f(x) = q(x) \cdot g(x) + r(x) \text{ and } \deg(r) < m$$

in a number of arithmetic operations in R bounded by $8 \cdot M_P(n) = O(n \log n \log \log n)$.

Exercise 2.2.4. *Let*

$$f = 30x^7 + 31x^6 + 32x^5 + 33x^4 + 34x^3 + 35x^2 + 36x + 37$$

and

$$g = 17x^3 + 18x^2 + 19x + 20$$

be two polynomials in $\mathbb{Z}/101[x]$.

(i) Compute $f^{-1} \bmod x^4$.

(ii) Compute q and r in $\mathbb{Z}/101[x]$ with $f = q \cdot g + r$ and $\deg r < 3 = \deg g$.

Exercise 2.2.5. *Let p be an arbitrary prime and a an integer that is not divisible by p.*

- *Derive an algorithm to compute an integer $b \in \{1, \ldots, p^\ell - 1\}$ with $a \cdot b \equiv 1 \bmod p^\ell$, where $\ell \neq 0$ is an arbitrary given integer.*

  Hint: Use Newton iteration.

- *Compute* $97^{-1} \bmod 4096$.

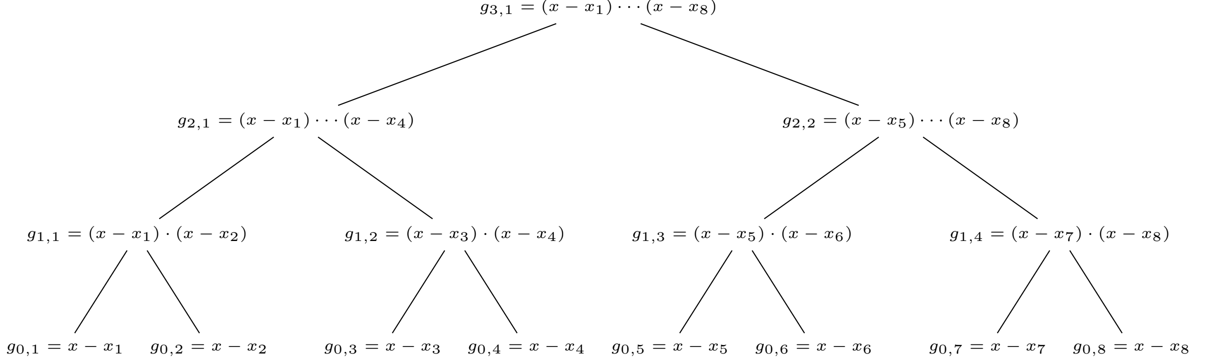

Figure 2.3: Illustration for the computation of all polynomials $g_{i,j}$ for $n = 8$.

There are a series of applications of the fast division algorithm. We start with an algorithm [MB72] due to Moenck and Borodin (from 1972) that allows us to evaluate a polynomial $f \in R[x]$ of degree n at n points $x_1, \ldots, x_n \in R$ in only $O(M_P(n) \cdot \log n) = \tilde{O}(n)$ primitive operations. This can be considered as a generalization of the FFT algorithm. For the seek of a simplified presentation, we again assume that $n = 2^k$ is a power of two. Starting with linear forms $g_{0,j}(x) := x - x_j$, we recursively compute

$$g_{i,j}(x) := g_{i-1,2j-1}(x) \cdot g_{i-1,2j} = (x - x_{(j-1)\cdot 2^i + 1}) \cdots (x - x_{j \cdot 2^i}) \text{ for } i = 1, \ldots, k \text{ and } j = 1, \ldots, \frac{n}{2^i}.$$

Notice that each $g_{i,j}$ is a product of 2^i linear forms, and that $g_{k,1}(x) = \prod_{i=1}^{n}(x - x_i)$; see also Figure 2.3 for an illustration in the case $n = 8$.

In the second step, we start with $r_{k,1} := f$, and recursively compute

$$\begin{aligned} r_{k-i,j} &:= f(x) \mod g_{k-i,j} \text{ for } i = 1, \ldots, k \text{ and } j = 1, \ldots, \frac{n}{2^{k-i}} \\ &= r_{k-i+1,\lceil j/2 \rceil} \mod g_{k-i,j}, \end{aligned}$$

where the latter equality follows from the fact that $g_{k-i,j}$ divides $g_{k-i+1,\lceil j/2 \rceil}$ and

$$\begin{aligned} f(x) &= q_{k-i+1,\lceil j/2 \rceil}(x) \cdot g_{k-i+1,\lceil j/2 \rceil} + r_{k-i+1,\lceil j/2 \rceil} \\ &= \left[q_{k-i+1,\lceil j/2 \rceil}(x) \cdot \frac{g_{k-i+1,\lceil j/2 \rceil}}{g_{k-i,j}} \right] \cdot g_{k-i,j} + r_{k-i+1,\lceil j/2 \rceil} \end{aligned}$$

for some $q_{k-i+1,j} \in R[x]$. Since each $r_{i,j}$ is the remainder of a polynomial division by some $g_{i,j'}$, it follows that $r_{i,j}$ has degree less than 2^i. Further notice that

$$r_{0,j} = f(x) \mod g_{0,j}(x) = f(x) \mod (x - x_j) = f(x_j),$$

thus we have computed all values $f(x_1), \ldots, f(x_n)$; see Algorithm 9 for pseudocode.

It remains to bound the cost for running Algorithm 9. The computation of each $g_{i,j}$ amounts for multiplying two polynomials in $R[x]$ of degree 2^{i-1}. For the computation of each $r_{i,j}$, we need to carry out one division with remainder between a polynomial of degree less than 2^{i+1} and a polynomial of degree 2^i. Hence, from Theorem 2.1.15 and 2.2.3, we conclude that the total cost is bounded by

$$\sum_{i=1}^{k} 2^{k-i} \cdot 8M_P(2^i)) = \sum_{i=1}^{k} 8M_P(n) = 8 \log n \cdot M_P(n).$$

Algorithm 9: Fast Multipoint Evaluation

Input : A Euclidean ring $R[x]$, a polynomial $f \in R[x]$ of degree $n = 2^k$, with $k \in \mathbb{N}$, and $x_1, \ldots, x_n \in R$

Output: $(f(x_1), \ldots, f(x_n))$

1 $g_{0,j}(x) := x - x_j$ for $j = 1, \ldots, n$

2 **for** $i = 1, \ldots, k$ **do**

3 Recursively define

$$g_{i,j}(x) := g_{i-1,2j-1}(x) \cdot g_{i-1,2j} \text{ for } j = 1, \ldots, 2^{k-i}.$$

4 $r_{k,1} := f$

5 **for** $i = 1, \ldots, k$ **do**

6 Recursively define

$$r_{k-i,j} := r_{k-i+1,\lceil j/2 \rceil} \mod g_{k-i,j} \text{ for } j = 1, \ldots, 2^i.$$

7 **return** $(r_{0,1}, \ldots, r_{0,n})$

We fix this result:

Theorem 2.2.6. *Let $f \in R[x]$ be a polynomial of degree n, and $x_1, \ldots, x_n \in R$. Then, Algorithm 9 computes all values $f(x_i)$, for $i = 1, \ldots, n$, using at most $6 \log n \cdot M_P(n) = O(n \log^2 n \log \log n)$ arithmetic operations in R.*

In the next step, we focus on the inverse problem, that is, given n distinct elements $x_1, \ldots, x_n \in R$, with $n = 2^k$, and corresponding values $v_1, \ldots, v_n \in R$, determine a polynomial $f(x) \in R[x]$ of degree less than n such that $f(x_i) = v_i$ for all i. We will now give a very efficient method for interpolation problem under the additional assumption that $x_i - x_j$ is a unit in R for all pairs i, j with $i \neq j$. Using *Lagrange interpolation*, we have

$$f(x) = \sum_{i=1}^{n} v_i \cdot \prod_{j \neq i} \frac{x - x_j}{x_i - x_j} = \sum_{i=1}^{n} v_i \cdot \underbrace{\frac{1}{\prod_{j \neq i}(x_i - x_j)}}_{=:\lambda_i} \cdot \underbrace{\prod_{j \neq i}(x - x_j)}_{=:g_i(x)}.$$

Notice that $\lambda_i' := \lambda_i^{-1} = \prod_{j \neq i}(x_i - x_j) = g_{k,1}'(x_i)$, where $\frac{\partial}{\partial x} g_{k,1}(x)$ is the (formal) derivative of $g_{k,1} = \prod_{j=1}^{n}(x - x_j)$ as defined in the fast multipoint evaluation algorithm above.[8] Hence, we may first compute $g_{k,1}$ and its derivative $g_{k,1}'$, and then use the fast multipoint evaluation algorithm to evaluate $g_{k,1}'$ at the points x_i to compute the values λ_i'. Then, dividing v_i by λ_i' yields the values $\mu_i := v_i / \lambda_i'$. The cost for this step is bounded by $O(\log n \cdot M_P(n))$ arithmetic operations in R plus n divisions in R. Now, in order to compute $f_{k,1}(x) := f(x) =$

[8]For a polynomial $f = \sum_{i=0}^{n} a_i \cdot x^i \in R[x]$, the formal derivative is defined as $\frac{\partial}{\partial x} f := \sum_{i=1}^{n} i \cdot a_i \cdot x^{i-1}$. Then, for any two polynomials $f, g \in R[x]$, it holds that $\frac{\partial}{\partial x}(f \cdot g) = \frac{\partial}{\partial x} f \cdot g + \frac{\partial}{\partial x} g \cdot f$, and thus $g_{k,1}'(x) = \prod_{j \neq i}(x - x_i) + (x - x_i) \cdot \left(\prod_{j \neq i}(x - x_i)\right)'$. It follows that $\frac{\partial}{\partial x} g_{k,1}(x_i) = \prod_{j \neq i}(x_i - x_j)$.

Algorithm 10: Fast Polynomial Interpolation

Input : A Euclidean ring $R[x]$, points $x_1, \dots, x_n \in R$, with $n = 2^k$ and $k \in \mathbb{N}$, such that $x_i - x_j$ is a unit in R for all $i \neq j$, and values $v_1, \dots, v_n \in R$.

Output: A polynomial $f \in R[x]$ of degree less than n such that $f(x_i) = v_i$ for all $i = 1, \dots, n$.

1 Compute all polynomials $g_{i,j}$, with $i = 0, \dots, k$ and $j = 1, \dots, n/2^i$.
2 $G := \frac{\partial}{\partial x} g_{k,1}$
3 Use Algorithm 9 to compute $\lambda_i := G(x_i)$ for all $i = 1, \dots, n$.
4 Compute $f_{0,j} := \mu_j := v_j/\lambda_j$ for $j = 1, \dots, n$.
5 **for** $i = 1, \dots, k$ **do**
6 Recursively define

$$f_{i,j}(x) := g_{i-1,2j-1}(x) \cdot f_{i-1,2j-1} + g_{i-1,2j}(x) \cdot f_{i-1,2j} \text{ for } j = 1, \dots, 2^{k-i}.$$

7 **return** $f_{k,1}$

$\sum_{i=1}^n \mu_i \cdot \prod_{j \neq i}(x - x_j)$, we write

$$f_{k,1}(x) = g_{k-1,1}(x) \cdot \underbrace{\sum_{i=n/2+1}^{n} \mu_i \cdot \prod_{j=n/2+1, j\neq i}^{n} (x - x_j)}_{=:f_{k-1,1(x)}} + g_{k-1,2}(x) \cdot \underbrace{\sum_{i=1}^{n/2} \mu_i \cdot \prod_{j=1,j\neq i}^{n/2} (x - x_j)}_{=:f_{k-1,2}(x)}.$$

Hence, we can recursively compute the polynomial f from the values μ_i and the polynomials $g_{i,j}$; see Algorithm 10. A completely analogous analysis as for the fast multipoint evaluation then yields the following result:

Theorem 2.2.7. *Let $x_1, \dots, x_n \in R$ be arbitrary points in R such that $x_i - x_j$ is a unit for all $i \neq j$, and let $v_1, \dots, v_n \in R$ be arbitrary points in R. Then, computing the unique polynomial $f \in R[x]$ of degree less than n with $f(x_i) = v_i$ for all i uses $O(\log n \cdot M_P(n)) = O(n \log^2 n \log\log n)$ additions and multiplication plus n divisions in R.*

We give a final application of the fast division algorithm, that is, the computation of a Taylor shift $x \mapsto m + x$ for a polynomial $f = \sum_{i=0}^n a_i \cdot x^i \in R[x]$ of degree n. Given the coefficients of f and a point $m \in R$, we aim to compute the coefficients of $\hat{f}(x) := f(m+x) = \sum_{i=0}^n \hat{a}_i \cdot x^i$. The idea is to reduce the problem to a fast multipoint evaluation followed by an interpolation. Suppose that there exist points $\hat{x}_1, \dots, \hat{x}_n$ such that $\hat{x}_i - \hat{x}_j$ is a unit in R for all $i \neq j$. Then, we evaluate f at the points $x_i := m + \hat{x}_i$, and eventually interpolate $\hat{f}$ from its values $\hat{f}(\hat{x}_i) = f(x_i)$ at the points $\hat{x}_i$. Notice that, if R supports the FFT, we may also choose $\hat{x}_i = \omega^i$, with ω a $2n$-th root of unity. Then, the interpolation step amounts for a single FFT computation. The following result immediately follows from Theorem 2.2.6 and Theorem 2.2.7.

Theorem 2.2.8. *Suppose that R contains elements $x_1, \dots, x_n$ such that $x_i - x_j$ is a unit in R for all $i \neq j$ (or R supports the FFT). Then, for an arbitrary polynomial $f \in R[x]$ and a point $m \in R$, we can compute the coefficients of $f(m + x)$ using $O(\log n \cdot M_P(n)) = O(n \log^2 n \log\log n)$ additions and multiplication plus n divisions in R.*

2.3 Fast Polynomial Arithmetic in $\mathbb{C}[x]$

We finally investigate in fast numerical variants of the algorithms presented in the previous two sections. Here, we assume that the coefficients of the input polynomial $f \in \mathbb{C}[x]$ (or any other input points in $\mathbb{C}$) are only given up to a certain precision, that is, for arbitrary $\rho \in \mathbb{N}$, we may ask for a dyadic approximation in $\mathbb{F} = \mathbb{F}_{2,\rho}$ of each coefficient (or of each point) to an error less than $2^{-\rho}$. For short, we call a corresponding approximation $\tilde{f}$ of f an *(absolute) ρ-bit approximation of f*. We start with a method for the approximate computation of a product of two polynomials.

Theorem 2.3.1. *Let $f, g \in \mathbb{C}[x]$ be polynomials of degree less than n and with coefficients of absolute value less than 2^L. Then, an ℓ - bit approximation $\tilde{h} = \sum_{i=0}^{2n-2} \tilde{h}_i \cdot x^i$ of the product $h = \sum_{i=0}^{2n-2} h_i \cdot x^i := f \cdot g$ (i.e. $|h_i - \tilde{h}_i| < 2^{-\ell}$ for all i) can be computed using $O(n(L + \ell))$ primitive operations. For this, we need ρ-bit approximations of f and g for some ρ of size $O(\log n + \ell + L)$.*

Proof. We reduce the multiplication of f and g to that of integer polynomials. For a non-negative integer ρ, consider ρ-bit approximations $\tilde{f} = \sum_{i=0}^{n-1} \tilde{a}_i \cdot x^i$ and $\tilde{g} = \sum_{i=0}^{n-1} \tilde{b}_i \cdot x^i$ of f and g. Then, $\tilde{h} = \sum_{i=0}^{2n-2} \tilde{c}_i \cdot x^i := \tilde{f} \cdot \tilde{g}$ constitutes an approximation of $h = \sum_{i=0}^{2n-2} c_i x^i := f \cdot g$ with

$$|h_i - \tilde{h}_i| < n \cdot [2^{L+1} \cdot 2^{-\rho} + 2^{-2\rho}] < 2^{L+2+\lceil \log n \rceil - \rho},$$

which follows from the fact that $|\tilde{a}_i \tilde{b}_j - a_i b_j| < (|a_i| + |b_j|) \cdot 2^{-\rho} + 2^{-2\rho}$ for all i, j. Hence, in order to guarantee that $\tilde{h}$ approximate h to an error less than $2^{-\ell}$, it suffices to choose $\rho := L + 2 + \lceil \log n \rceil + \ell$. In order to compute the product $\tilde{f} \cdot \tilde{g}$, we first compute the product $(2^\rho \cdot \tilde{f}) \cdot (2^\rho \cdot \tilde{g})$ of integer polynomials and then shift the coefficients of the result by 2ρ bits. The latter product can be computed in $\tilde{O}(n(\ell + L))$ primitive operations according to Theorem 2.1.12. □

For a corresponding numerical variant of the fast polynomial division, we have to work harder. We start with following lemma:

Lemma 2.3.2. *Let $f \in \mathbb{C}[x]$ be a polynomial of degree n, $g \in \mathbb{C}[x]$ a monic polynomial of degree m, with $m \le n$, and $q, r \in \mathbb{C}[x]$ with $f = q \cdot g + r$ and $\deg r < m$. Then, it holds that*

$$\log \max(\|q\|_\infty, \|r\|_\infty) = O(L + n + (n - m) \cdot \Gamma),$$

where L and Γ are non-negative integers with $\max(\|f\|_\infty, \|g\|_\infty) < 2^L$ and $|z| < 2^\Gamma$ for any complex root z of g.

Proof. Let $f(x) = a_0 + \cdots + a_n \cdot x^n$ and $g(x) = b_0 + \cdots + b_m \cdot x^m$, then we have

$$\begin{aligned} \frac{f(x)}{x^{n-m} \cdot g(x)} &= \frac{q(x) \cdot g(x) + r(x)}{x^{n-m} \cdot g(x)} = \frac{q(x)}{x^{n-m}} + \frac{r(x)}{x^{n-m} \cdot g(x)} \\ &= q_{n-m} + \frac{q_{n-m-1}}{x} + \cdots + \frac{q_0}{x^{n-m}} + \cdots \end{aligned} \tag{2.9}$$

where $q(x) = q_0 + \cdots + q_{n-m} \cdot x^{n-m}$. Here, we use the fact that $r(x)/g(x)$ is a holomorphic function in the domain $D := \{x \in \mathbb{C} : |x| > 2^\Gamma\}$. Using a corresponding result from Complex Analysis, we thus conclude that it can be written as a Laurent series $\sum_{i=-\infty}^{\infty} c_i \cdot x^i$, which converges for all $x \in D$. We further remark that $c_i = 0$ for all $i \ge 0$ as, otherwise,

$\lim_{x\mapsto\infty} \frac{|x\cdot r(x)|}{|g(x)|} = \infty$, which contradicts the fact that $\deg r < m$. Now, from (2.9), we conclude that

$$\frac{f}{x^{n-m}\cdot g(x)}\cdot x^{i-1} = q_{n-m}\cdot x^{i-1} + \cdots + \frac{q_{n-m-i}}{x} + \cdots,$$

and thus the Residue Theorem yields that

$$\frac{1}{2\pi\mathbf{i}}\oint_{|x|=2^{\Gamma+1}} \frac{f(x)}{x^{n-m-i+1}\cdot g(x)}dx = q_{n-m-i} \text{ for all } i = 0,\ldots,n-m$$

or

$$\frac{1}{2\pi\mathbf{i}}\oint_{|x|=2^{\Gamma+1}} \frac{f(x)}{x^{j+1}\cdot g(x)}dx = q_j \text{ for all } j = 0,\ldots,n-m.$$

For $|x| = 2^{\Gamma+1}$, we have $|f(x)| \le (n+1)\cdot 2^L \cdot 2^{n(\Gamma+1)}$ and $|g(x)| \ge 2^{m\Gamma}$. Hence, it follows that the absolute value of the integrand is upper bounded by $B = (n+1)\cdot 2^{L+n+(n-m)\Gamma}$. We conclude that each coefficient q_j of q is bounded by $B\cdot 2^{\Gamma+1} = 2^{O(L+n+(n-m)\Gamma)}$. The claim regarding the size of the coefficients of r then immediately follows from the bound on $\|q\|_\infty$ and the fact that $r = f - q\cdot g$. □

Using the above lemma, we can now derive a bound on the polynomials $\hat{h}_i \in \mathbb{C}[x]$ as computed in Algorithm 8. Namely, let $\hat{h}_i$ be a polynomial of degree less than $2^i - 1$ such that $\hat{h}_i \cdot \hat{g} = 1 \mod x^{2^i}$, with $g(x) = x^m \cdot g(1/x)$ and $g \in \mathbb{C}[x]$ a monic polynomial of degree m. Then, there exists a polynomial $s_i \in \mathbb{C}[x]$ of degree less than m such that

$$\hat{g}(x)\cdot\hat{h}_i = 1 + x^{2^i}\cdot s_i(x) \Rightarrow \underbrace{x^m\cdot\hat{g}(1/x)}_{=g(x)}\cdot\underbrace{x^{2^i}\cdot\hat{h}_i(1/x)}_{=:h_i} = x^{m+2^i} + \underbrace{x^m\cdot s_i(1/x)}_{\hat{s}_i}.$$

Hence, the polynomials $h_i := x^{2^i}\cdot\hat{h}_i(1/x)$ and $\hat{s}_i := x^m\cdot s_i(1/x)$ are the quotient and remainder obtained dividing x^{m+2^i} by $g(x)$. Lemma 2.3.2 then yields that

$$\log\|\hat{h}_i\|_\infty = \log\|h_i\|_\infty = O(\log\|g\|_\infty + n + n\Gamma),$$

with $\Gamma \ge 0$ a bound on $\log|z_i|$ for every complex root of g. In the $(i+1)$ - st iteration step in Algorithm 8, we compute $\hat{h}_{i+1} = 2\hat{h}_i - \hat{g}\cdot\hat{h}_i^2 \mod x^{2^{i+1}}$. Hence, if we use ρ-bit approximations of $\hat{h}_i$ and $\hat{g}$ instead of the exact polynomials h_i and g, then Theorem 2.3.1 shows that we obtain an approximation of $\hat{h}_{i+1}$ to an error less than $2^{-\rho+O(\log\|g\|_\infty+n+n\Gamma)}$. In other words the precision loss in each iteration is bounded by $O(\log\|g\|_\infty + n + n\Gamma)$. Since there are at most $\lceil\log n\rceil$ many iterations, the total precision loss is bounded by $\tilde{O}(\log\|g\|_\infty + n + n\Gamma)$. Hence, we can use fixed point arithmetic with precision $\rho = \ell + \tilde{O}(\log\|g\|_\infty + n + n\Gamma)$ to guarantee an output error of less than $2^{-\ell}$.

Theorem 2.3.3. *Let f and g be polynomials as in Lemma 2.3.2. Then, computing ℓ-bit approximations $\tilde{q}$ and $\tilde{r}$ of q and r uses $\tilde{O}(n(\ell + L + n + n\Gamma))$ primitive operations. For this, we need ρ-bit approximations of the polynomials f and g for some ρ of size $\ell + \tilde{O}(L+n+n\Gamma)$.*

We briefly summarize out findings from Theorems 2.3.1 and 2.3.3: A multiplication of two polynomials $f, g \in \mathbb{C}[x]$ using fixed point arithmetic with precision ρ yields a loss in precision bounded by $O(\log n + \log\max(\|f\|_\infty, \|g\|_\infty))$, whereas the precision loss of a corresponding division with remainder is bounded by $\tilde{O}(n + \log\max(\|f\|_\infty, \|g\|_\infty + n\Gamma)$. Now, what can we conclude about the precision loss in the fast multipoint evaluation algorithm? The polynomials

$g_{i,j}$ are products of linear forms $x - x_s$, hence $\log \|g_{i,j}\|_\infty$ is bounded by $O(n\Gamma^*)$ with $\Gamma^* := \max(1, \log \max_{i=1,\ldots,n} |x_i|)$. Since the depth of the recursion is $\log n$, we conclude that the precision loss is bounded by $O(n\Gamma^* \cdot \log n)$. Now, for the divisions in the algorithm, notice that we start with $r_{k,1} = f$. In each step of the recursion, we divide a previously computed remainder $r_{i,j}$ by some $g_{i',j'}$. Further notice that $r_{i,j} = f \mod g_{i,j}$, and thus $\log \|r_{i,j}\|_\infty = O(L + n\Gamma^*)$ according to Lemma 2.3.2. It follows that the precision loss in each of the considered divisions is bounded by $\tilde{O}(L + n\Gamma^*)$. Now, since the depth of the recursion is bounded by $O(\log n)$, we conclude that the total loss in precision is bounded by $\tilde{O}((L+n\Gamma^*))$. Thus, in order to guarantee an output error of size less than $2^{-\ell}$, it suffices to use fixed point arithmetic with a precision of size $\ell + \tilde{O}(L + n\Gamma^*)$.

Theorem 2.3.4. *Let $f \in \mathbb{C}[x]$ be a polynomial of degree n with coefficients of absolute value bounded by 2^L, with $L \in \mathbb{N}_{\geq 1}$, and let $x_1, \ldots, x_n \in \mathbb{C}$ be arbitrary points of absolute value less than 2^{Γ^*}, with $\Gamma^* \geq 1$. For an arbitrary non-negative number ℓ, we can compute ℓ-bit approximations $\tilde{v}_i$ of all values $v_i := f(x_i)$ using $\tilde{O}(n \cdot (\ell + L + \Gamma^*))$ primitive operations. For this, we need ρ-bit approximations of f and the points x_i for some ρ of size $\ell + \tilde{O}(L + n\Gamma^*)$.*

Corollary 2.3.5. *Let $f \in \mathbb{C}[x]$ be a polynomial of degree n with coefficients of absolute value bounded by 2^L, with $L \in \mathbb{N}_{\geq 1}$, and $m \in \mathbb{C}$ be an arbitrary point of absolute value less than 2^{Γ^*}, with $\Gamma^* \geq 1$. For an arbitrary non-negative number ℓ, we can compute an ℓ-bit approximations of $\hat{F}(x) = F(m + x)$ using $\tilde{O}(n \cdot (\ell + L + \Gamma^*))$ primitive operations. For this, we need ρ-bit approximations of f and m for some ρ of size $\ell + \tilde{O}(L + n\Gamma^*)$.*

Proof. The proof is left as an exercise. □

Exercise 2.3.6. *Let $f \in \mathbb{Z}[x]$ be an integer polynomial of degree less than n with coefficients of absolute value less than 2^L. Furthermore, let $x_1, \ldots, x_n$ be n be distinct rational points in $[0, 1]$ of bitsize ℓ (i.e., $x_i = p_i/q_i \in [0, 1]$ with integers p_i and q_i of absolute value less than 2^ℓ).*

We say that the point x_i is large *for f among $X := \{x_1, \ldots, x_n\}$ if*

$$4 \cdot |f(x_i)| \geq \max_{1 \leq j \leq n} |f(x_j)| =: \lambda.$$

- *Determine the cost of finding a large point in a naive way, that is, by evaluating f at all points x_j exactly.*
- *Show how to find a large point in $\tilde{O}(n(L + \log \max(1, \lambda^{-1})))$ bit operations.*

  Hint: Use approximate multipoint evaluation with increasing precision.

Bibliography

[AY62] Karatsuba A. and Ofman Y. “Multiplication of Many-Digital Numbers by Automatic Computers”. In: *Doklady Akademii Nauk SSSR* 14 (1962), pp. 293–294 (cit. on p. 5).

[CT65] James W. Cooley and John W. Tukey. “An Algorithm for the Machine Calculation of Complex Fourier Series”. In: *Mathematics of Computation* 19.90 (1965), pp. 297–301. ISSN: 00255718, 10886842 (cit. on p. 24).

[GG03] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2003. ISBN: 9780521826464 (cit. on pp. 21, 30, 31).

[MB72] Robert T. Moenck and Allan Borodin. “Fast modular transforms via division”. In: *13th*. 1972, pp. 90–96 (cit. on p. 35).

[MOS11] Kurt Mehlhorn, Ralf Osbild, and Michael Sagraloff. “A general approach to the analysis of controlled perturbation algorithms”. In: *Comput. Geom.* 44.9 (2011), pp. 507–528 (cit. on p. 15).

[MS08] K. Mehlhorn and P. Sanders. *Algorithms and Data Structures: The Basic Toolbox*. SpringerLink: Springer e-Books. Springer, 2008. ISBN: 9783540779773 (cit. on p. 6).

[SS71] A. Schönhage and V. Strassen. “Schnelle Multiplikation großer Zahlen”. In: *Computing* 7.3 (1971), pp. 281–292. ISSN: 1436-5057 (cit. on p. 19).

[Too63] Andrei Toom. “The Complexity of a Scheme of Functional Elements Realizing the Multiplication of Integers”. In: *Soviet Mathematics-Doklady* 7 (1963), pp. 714–716 (cit. on p. 7).