

Structured FFT and TFT: symmetric and lattice polynomials

Joris van der Hoeven*
Laboratoire d'informatique
UMR 7161 CNRS
École polytechnique
91128 Palaiseau Cedex, France
vdhoeven@lix.polytechnique.fr

Romain Lebreton
LIRMM
UMR 5506 CNRS
Université de Montpellier II
Montpellier, France
lebreton@lirmm.fr

Éric Schost
Computer Science Department
Western University
London, Ontario
Canada
eschost@uwo.ca

ABSTRACT

In this paper, we consider the problem of efficient computations with structured polynomials. We provide complexity results for computing Fourier Transform and Truncated Fourier Transform of symmetric polynomials, and for multiplying polynomials supported on a lattice.

1. INTRODUCTION

Fast computations with multivariate polynomials and power series have been of fundamental importance since the early ages of computer algebra. The representation is an important issue which conditions the performance in an intrinsic way; see [23, 31, 8] for some historical references.

It is customary to distinguish three main types of representations: dense, sparse, and functional. A *dense representation* is made of a compact description of the support of the polynomial and the sequence of its coefficients. The main example concerns *block supports* – it suffices to store the coordinates of two opposite vertices. In a dense representation all the coefficients of the considered support are stored, even if they are zero. If a polynomial has only a few non-zero terms in its bounding block, we shall prefer to use a *sparse representation* which stores only the sequence of the non-zero terms as pairs of monomials and coefficients. Finally, a *functional representation* stores a function that can produce values of the polynomials at any given point. This can be a pure blackbox (which means that its internal structure is not supposed to be known) or a specific data structure such as *straight-line programs* (see Chapter 4 of [4], for instance, and [10] for a concrete library for the manipulation of polynomials with a functional representation).

For dense representations with block supports, it is classical that the algorithms used for the univariate case can be naturally extended: the naive algorithm, Karatsuba's algorithm, and even Fast Fourier Transforms [7, 30, 6, 28, 16] can be applied recursively in each variable, with good performance. Another classical approach is the Kronecker substitution which reduces the multivariate product to one variable only; for all these questions, we refer the reader to classical books such as [29, 14]. When the number of variables is fixed and the partial degrees tend to infinity, these techniques lead to softly linear costs.

After the discovery of sparse interpolation [2, 24, 25, 26, 11], probabilistic algorithms with a quasi-linear complexity have been developed for sparse polynomial multiplication [5]. It has recently been shown that such asymptotically

fast algorithms may indeed become more efficient than naive sparse multiplication [19].

In practice however, it frequently happens that multivariate polynomials with a dense flavor do not admit a block support. For instance, it is common to consider polynomials of a given total degree. In a recent series of works [15, 17, 18, 21, 20], we have studied the complexity of polynomial multiplication in this “semi-dense” setting. In the case when the supports of the polynomials are initial segments for the partial ordering on \mathbb{N}^n , the truncated Fourier transform is a useful device for the design of efficient algorithms.

Besides polynomials with supports of a special kind, we may also consider what will call “structured polynomials”. By analogy with linear algebra, such polynomials carry a special structure which might be exploited for the design of more efficient algorithms. In this paper, we turn our attention to a first important example of this kind: polynomials which are invariant under the action of certain matrix groups. We consider only two special cases: finite subgroups of \mathfrak{S}_n and finite groups of diagonal matrices. These cases are already sufficient to address questions raised *e.g.* in celestial mechanics [12]; it is hoped that more general groups can be dealt with using similar ideas.

In the limited scope of this paper, our main objective is to prove complexity results that demonstrate the savings induced by a proper use of the symmetries. Our complexity analyses take into account the number of arithmetic operations in the base field, and as often, we consider that operations with groups and lattices take a constant number of operations. A serious implementation of our algorithms would require an improved study of these aspects.

Of course, there already exists an abundant body of work on some of these questions. *Crystallographic* FFT algorithms date back to [32], with contributions as recent as [27], but are dedicated to crystallographic symmetries. A more general framework due to [1] was recently revisited under the point of view of high-performance code generation [22]; our treatment of permutation groups is in a similar spirit, but to our understanding, these previous papers do not prove results such as those we give below (and they only consider the FFT, not its truncated version).

Similarly, our results on diagonal groups, which fit in the general context of FFTs over lattices, use similar techniques as in a series of papers initiated by [13] and continued as recently as [33, 3], but the actual results we prove are not in those references.

*. This work has been partly supported by the Digiteo 2009-36HD grant of the Rgion Ile-de-France, NSERC and the CRC program.

2. THE CLASSICAL FFT

2.1 Notation

We will work over an effective base field (or ring) \mathbb{K} with sufficiently many roots of unity; the main objects are polynomials $\mathbb{K}[\mathbf{x}] := \mathbb{K}[x_1, \dots, x_n]$ in n variables over \mathbb{K} . For any $\mathbf{k} = (\mathbf{k}_1, \dots, \mathbf{k}_n) \in \mathbb{N}^n$ and $P \in \mathbb{K}[\mathbf{x}]$, we denote by $\mathbf{x}^{\mathbf{k}}$ the monomial $x_1^{\mathbf{k}_1} \dots x_n^{\mathbf{k}_n}$ and by $P_{\mathbf{k}}$ the coefficient of P in $\mathbf{x}^{\mathbf{k}}$. The support $\text{supp}(P)$ of a polynomial $P \in \mathbb{K}[\mathbf{x}]$ is the set of exponents $\mathbf{k} \in \mathbb{N}^n$ such that $P_{\mathbf{k}} \neq 0$.

For any subset S of \mathbb{N}^n , we define $\mathbb{K}[\mathbf{x}]_S$ as the polynomials with support included in S . As an important special case, when $S = \mathbb{N}_d^n := \{0, \dots, d-1\}^n$, we denote by $\mathbb{K}[\mathbf{x}]_d := \mathbb{K}[\mathbf{x}]_S$ the set of polynomials with partial degree less than d in all variables.

Let $\omega \in \mathbb{K}$ be a primitive d th root of unity (in Section 4, it will be convenient to write this root $e^{2\pi i/d}$). For any $\mathbf{k} = (\mathbf{k}_1, \dots, \mathbf{k}_n) \in \mathbb{N}^n$, we define $\omega^{\mathbf{i}} = (\omega^{\mathbf{k}_1}, \dots, \omega^{\mathbf{k}_n})$. One of our aims is to compute efficiently the map

$$\text{FFT}_{\omega}: \begin{cases} \mathbb{K}[\mathbf{x}]_d & \longrightarrow & \mathbb{K}^{\mathbb{N}_d^n} \\ P & \longmapsto & (P(\omega^{\mathbf{i}}))_{\mathbf{i} \in \mathbb{N}_d^n} \end{cases}$$

when P is a ‘‘structured polynomial’’. Here, $\mathbb{K}^{\mathbb{N}_d^n}$ denotes the set of vectors with entries in \mathbb{K} indexed by \mathbb{N}_d^n ; in other words, $\mathbf{v} \in \mathbb{K}^{\mathbb{N}_d^n}$ implies that $\mathbf{v}_{\mathbf{i}} \in \mathbb{K}$ for all $\mathbf{i} \in \mathbb{N}_d^n$. Likewise, $\mathbb{K}^{\mathbb{N}_d^n \times \mathbb{N}_d^n}$ denotes the set of matrices with indices in \mathbb{N}_d^n , that is, $\mathbf{M} \in \mathbb{K}^{\mathbb{N}_d^n \times \mathbb{N}_d^n}$ implies that $\mathbf{M}_{\mathbf{i}, \mathbf{j}} \in \mathbb{K}$ for all $\mathbf{i}, \mathbf{j} \in \mathbb{N}_d^n$.

Most of the time, we will take $d = 2^\ell$ with $\ell \in \mathbb{N}$, although we will also need more general d of mixed radices $d = p_1 \dots p_\ell$.

We denote by $\mathbf{i} \cdot \mathbf{j}$ the inner product of two vectors in \mathbb{N}^n . We also let $\mathbf{e}_k := (0, \dots, \overset{k}{1}, 0, \dots, 0)$, for $1 \leq k \leq n$, be the k th element of the canonical basis of \mathbb{K}^n . At last, for $\ell \in \mathbb{N}$ and $k \in \mathbb{N}_{2^\ell}$ we denote by $[k]_\ell$ the bit reversal of k in length n and we extend this notation to vectors by $[\mathbf{k}]_s = ([\mathbf{k}_1]_s, \dots, [\mathbf{k}_n]_s)$.

2.2 The classical multivariate FFT

Let us first consider the in-place computation of a full n -dimensional FFT of length $d = 2^\ell$ in each variable. We first recall the notations and operations of the decimation in time variant of the FFT, and we refer the reader to [9] for more details. In what follows, ω is a d th primitive root of unity.

We start with the FFT of a univariate polynomial $P \in \mathbb{K}[x]_d$. Decimation in time amounts to decomposing P into its even and odd parts, and proceeding recursively, by means of ℓ decimations applied to the variable x . For $0 \leq k < d$, we will write $\mathbf{c}_k^0 = P_k$ for the input and denote by $\mathbf{c}^s = (\mathbf{c}_k^s)_{0 \leq k < d}$ the result after s decimation steps, for $s \in \{1, \dots, \ell\}$.

At stage s , the decimation is computed using butterflies of span $\delta := 2^{\ell-s}$. If $\mathbf{i} \in \mathbb{N}_{2^s}$ is even and \mathbf{j} belongs to \mathbb{N}_δ then these butterflies are given by

$$\begin{pmatrix} \mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^s \\ \mathbf{c}_{(\mathbf{i}+1)\delta+\mathbf{j}}^s \end{pmatrix} = \begin{pmatrix} 1 & \omega^{[\mathbf{i}]_s \delta} \\ 1 & -\omega^{[\mathbf{i}]_s \delta} \end{pmatrix} \begin{pmatrix} \mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^{s-1} \\ \mathbf{c}_{(\mathbf{i}+1)\delta+\mathbf{j}}^{s-1} \end{pmatrix}.$$

Putting the coefficients of all these linear relations in a matrix $\mathbf{B}^s \in \mathbb{K}^{\mathbb{N}_d \times \mathbb{N}_d}$, we get $\mathbf{c}^s = \mathbf{B}^s \mathbf{c}^{s-1}$. The matrix \mathbf{B}^s is sparse, with at most two non-zero coefficients on each row and each column; up to permutation, it is block-diagonal, with blocks of size 2. After ℓ stages, we get the evaluations of P in the bit reversed order:

$$\mathbf{c}_k^\ell = P(\omega^{[k]_\ell}).$$

We now adapt this notation to the multivariate case. The computation is still divided in ℓ stages, each stage doing one decimation in every variable x_1, \dots, x_n . Therefore, we will denote by $\mathbf{c}_k^0 := P_{\mathbf{k}}$ the coefficients of the input for $\mathbf{k} \in \mathbb{N}_d^n$ and by $\mathbf{c}_k^{s,t}$ the coefficients obtained during the s th stage, after the decimations in x_1, \dots, x_t are done, with $s \in \{1, \dots, \ell\}$ and $t \in \{0, \dots, n\}$, so that $\mathbf{c}_k^{s-1, n} = \mathbf{c}_k^{s, 0}$. We abbreviate $\mathbf{c}_k^{s-1} := \mathbf{c}_k^{s, 0}$ for the coefficients after $s-1$ stages.

The intermediate coefficients \mathbf{c}_k^s can be seen as evaluations of intermediate polynomials: for every $s \in \{0, \dots, \ell\}$, $\mathbf{i} \in \mathbb{N}_{2^s}^n$ and $\mathbf{j} \in \mathbb{N}_\delta^n$, one has

$$\mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^s = P_j^s((\omega^\delta)^{[\mathbf{i}]_s}) \quad (1)$$

where $P_j^s = \sum_{\mathbf{i}' \in \mathbb{N}_{2^s}^n} P_{\mathbf{i}'\delta+\mathbf{j}} \mathbf{x}^{\mathbf{i}'}$ are obtained through an s -fold decimation of P . Equivalently, the coefficients \mathbf{c}_k^s satisfy

$$\mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^s = \sum_{\mathbf{i}' \in \mathbb{N}_{2^s}^n} P_{\mathbf{i}'\delta+\mathbf{j}} \omega^{[\mathbf{i}]_s \cdot \mathbf{i}' \delta}. \quad (2)$$

Thus, $\mathbf{c}_j^\ell = P(\omega^{[j]_\ell})$ yields $\text{FFT}_{\omega}(P)$ in bit reversed order.

For the concrete computation of the coefficients \mathbf{c}_k^s at stage s from the coefficients \mathbf{c}_k^{s-1} at stage $s-1$, we use n so called ‘‘elementary transforms’’ with respect to each of the variables x_1, \dots, x_n . For any $t \in \{1, \dots, n\}$, the coefficients $\mathbf{c}_k^{s,t}$ are obtained from the coefficients $\mathbf{c}_k^{s,t-1}$ through butterflies of span $\delta = 2^{\ell-s}$ with respect to the variable x_t ; this can be rewritten by means of the formula

$$\begin{pmatrix} \mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^{s,t} \\ \mathbf{c}_{(\mathbf{i}+\mathbf{e}_t)\delta+\mathbf{j}}^{s,t} \end{pmatrix} = \begin{pmatrix} 1 & \omega^{[\mathbf{i}]_s \delta} \\ 1 & -\omega^{[\mathbf{i}]_s \delta} \end{pmatrix} \begin{pmatrix} \mathbf{c}_{\mathbf{i}\delta+\mathbf{j}}^{s,t-1} \\ \mathbf{c}_{(\mathbf{i}+\mathbf{e}_t)\delta+\mathbf{j}}^{s,t-1} \end{pmatrix} \quad (3)$$

for any $\mathbf{i} \in \mathbb{N}_{2^s}^n$ with even t -th coordinate i_t and $\mathbf{j} \in \mathbb{N}_\delta^n$. Equation (3) can be rewritten more compactly in matrix form $\mathbf{c}^{s,t} = \mathbf{B}^{s,t} \mathbf{c}^{s,t-1}$ where $\mathbf{B}^{s,t}$ is a sparse matrix in $\mathbb{K}^{\mathbb{N}_d^n \times \mathbb{N}_d^n}$ which is naturally indexed by pairs of multi-indices in \mathbb{N}_d^n . Setting $\mathbf{B}^s = \mathbf{B}^{s,n} \dots \mathbf{B}^{s,1} \in \mathbb{K}^{\mathbb{N}_d^n \times \mathbb{N}_d^n}$, we also obtain a short formula for \mathbf{c}^s as a function of \mathbf{c}^{s-1} :

$$\mathbf{c}^s = \mathbf{B}^s \mathbf{c}^{s-1} \quad (4)$$

Remark that the matrices $\mathbf{B}^{s,1}, \dots, \mathbf{B}^{s,n}$ commute pairwise. Notice also that each of the rows and columns of $\mathbf{B}^{s,t}$ has at most 2 non zero entries and consequently those of \mathbf{B}^s contains at most 2^n non zero entries. For this reason, we can apply the matrices $\mathbf{B}^{s,t}$ (resp. \mathbf{B}^s) within $\mathcal{O}(|\mathbb{N}_d^n|) = \mathcal{O}(d^n)$ (resp. $\mathcal{O}(n d^n)$) operations in \mathbb{K} . Finally, the full n -dimensional FFT of length $d = 2^\ell$ costs $F(d, n) := 3/2 \ell n d^n = 3/2 d^n \log(d^n)$ operations in \mathbb{K} (see [14, Section 8.2]).

Example 1. Let us make these matrices explicit for polynomials in $n = 2$ variables and degree less than $d = 4$, so that $\ell = 2$. We start with the first decimation in x_1 whose butterflies of span δ are captured by $\mathbf{B}^{s,t}$ with $\delta = 2^{\ell-s} = 2^1$, $s = 1$ and $t = 1$. It takes as input $\mathbf{c}_k^0 = \mathbf{c}_k^{1,0} := P_{\mathbf{k}}$ and outputs $\mathbf{c}_k^{1,1}$ for $\mathbf{k} \in \mathbb{N}_4^2 = \{0, \dots, 3\}^2$. For any $\mathbf{j}_1 \in \mathbb{N}_2 = \{0, 1\}$ and $\mathbf{k}_2 \in \{0, \dots, 3\}$, we let

$$\begin{pmatrix} \mathbf{c}_{(\mathbf{j}_1, \mathbf{k}_2)}^{1,1} \\ \mathbf{c}_{(2+\mathbf{j}_1, \mathbf{k}_2)}^{1,1} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \mathbf{c}_{(\mathbf{j}_1, \mathbf{k}_2)}^{1,0} \\ \mathbf{c}_{(2+\mathbf{j}_1, \mathbf{k}_2)}^{1,0} \end{pmatrix}.$$

So, $\mathbf{B}^{1,1}$ is a 16 by 16 matrix, which is made of diagonal blocks of $\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ in a suitable basis. The decimation in x_2 during stage 1 is similar :

$$\begin{pmatrix} \mathbf{c}_{(\mathbf{k}_1, \mathbf{j}_2)}^{1,2} \\ \mathbf{c}_{(\mathbf{k}_1, 2+\mathbf{j}_2)}^{1,2} \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \mathbf{c}_{(\mathbf{k}_1, \mathbf{j}_2)}^{1,1} \\ \mathbf{c}_{(\mathbf{k}_1, 2+\mathbf{j}_2)}^{1,1} \end{pmatrix}$$

for $\mathbf{j}_2 \in \mathbb{N}_2$ and $\mathbf{k}_1 \in \mathbb{N}_4$. Consequently, $\mathbf{B}^{1,2}$ is made of diagonal blocks $\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ (in another basis than $\mathbf{B}^{1,1}$). Their product $\mathbf{B}^1 := \mathbf{B}^{1,2} \mathbf{B}^{1,1}$ corresponds to the operations

$$\begin{pmatrix} \mathbf{c}_{(\mathbf{j}_1, \mathbf{j}_2)}^1 \\ \mathbf{c}_{(2+\mathbf{j}_1, \mathbf{j}_2)}^1 \\ \mathbf{c}_{(\mathbf{j}_1, 2+\mathbf{j}_2)}^1 \\ \mathbf{c}_{(2+\mathbf{j}_1, 2+\mathbf{j}_2)}^1 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{c}_{(\mathbf{j}_1, \mathbf{j}_2)}^0 \\ \mathbf{c}_{(2+\mathbf{j}_1, \mathbf{j}_2)}^0 \\ \mathbf{c}_{(\mathbf{j}_1, 2+\mathbf{j}_2)}^0 \\ \mathbf{c}_{(2+\mathbf{j}_1, 2+\mathbf{j}_2)}^0 \end{pmatrix}$$

for $\mathbf{j}_1, \mathbf{j}_2 \in \mathbb{N}_2$. Thus \mathbf{B}^1 is made of such 4 by 4 matrices on the diagonal (in yet another basis). Note that in general, the matrices $\mathbf{B}^{s,t}$ and \mathbf{B}^s are still made of diagonal blocks, but these blocks vary along the diagonal.

We can sum it all up in the two following algorithms.

<p>Algorithm Butterfly INPUT: n, degree d, stage s, index $(2i', j)$ of the butterfly with $i' \in \mathbb{N}_{2^{s-1}}^n$, $j \in \mathbb{N}_2^n$, root of unity ω and coefficients $\mathbf{c} \in \mathbb{K}^{\mathbb{N}_2^n}$ for the butterfly ($\mathbf{c}_b := \mathbf{c}_{(2i'+b)\delta+j}^{s-1}$ for $\mathbf{b} \in \mathbb{N}_2^n$). OUTPUT: the output of the butterfly $\mathbf{d} \in \mathbb{K}^{\mathbb{N}_2^n}$</p> <hr/> <p>For $t=1, \dots, n$ do //decimation in x_t For $\mathbf{b} \in \mathbb{N}_2^n = \{0, 1\}^n$ such that $\mathbf{b}_t = 0$ do $r = \omega^{[2i']_s \delta \mathbf{c}_{b+e_t}}$ //butterfly in one variable $\mathbf{d}_b = \mathbf{c}_b + r$ $\mathbf{d}_{b+e_t} = \mathbf{c}_b - r$ Return \mathbf{d}_b</p>
--

<p>Algorithm FFT INPUT: n, degree d, ω and coefficients $\mathbf{c}^0 \in \mathbb{K}^{\mathbb{N}_2^n}$ of P OUTPUT: coefficients $\mathbf{c}^\ell \in \mathbb{K}^{\mathbb{N}_2^n}$ in bit reversed order</p> <hr/> <p>For $s=1, \dots, \ell$ do //stage s $\delta := d/2^s$ For $i' \in \mathbb{N}_{2^{s-1}}^n$, $j \in \mathbb{N}_{2^{\ell-s}}^n$ do //pick a butterfly For $\mathbf{b} \in \mathbb{N}_2^n$ do $\mathbf{c}_b := \mathbf{c}_{(2i'+b)\delta+j}^{s-1}$ $\mathbf{d}_b = \mathbf{Butterfly}(n, d, s, i', j, \mathbf{c}_b)$ For $\mathbf{b} \in \mathbb{N}_2^n$ do $\mathbf{c}_{(2i'+b)\delta+j}^s := \mathbf{d}_b$ Return \mathbf{c}^ℓ</p>
--

In the last section, we will also use the ring isomorphism

$$\mathbb{K}[\mathbf{x}]_d / (x_1^d - 1, \dots, x_n^d - 1) \rightarrow [\mathbb{K}[\mathbf{x}]_\delta / (x_1^\delta - 1, \dots, x_n^\delta - 1)]^{2^{n-s}}$$

$$P \mapsto Q^s = (Q_i^s(\mathbf{x}))_{i \in \mathbb{N}_{2^s}^n}$$

with $Q_i^s(\mathbf{x}) = \sum_{j \in \mathbb{N}_2^n} \left(\sum_{i' \in \mathbb{N}_{2^s}^n} P_{i'\delta+j} \omega^{i \cdot i' \delta} \right) (\omega^i \mathbf{x})^j$ for $i \in \mathbb{N}_{2^s}^n$, and $(\omega^i \mathbf{x})^j = (\omega^{i_1} x_1)^{j_1} \dots (\omega^{i_n} x_n)^{j_n}$. These polynomials generalize the decomposition of a univariate polynomials P into $P \bmod (x^{d/2} - 1)$ and $P \bmod (x^{d/2} + 1)$. We could obtain them through a decimation in frequency, but it is enough to remark that we can reconstruct them from the coefficients \mathbf{c}^s thanks to the formula

$$Q_i^s(\mathbf{x}) = \sum_{j \in \mathbb{N}_2^n} \mathbf{c}_{[i]_s \delta + j}^s (\omega^i \mathbf{x})^j. \quad (5)$$

3. THE SYMMETRIC FFT

In this section, we let $G \subseteq \mathfrak{S}_n$ be a permutation group. The group G acts on the polynomials $\mathbb{K}[\mathbf{x}]$ via the map

$$\varphi_g: \begin{cases} \mathbb{K}[\mathbf{x}] & \longrightarrow \mathbb{K}[\mathbf{x}] \\ P(x_1, \dots, x_n) & \longmapsto P^g(\mathbf{x}) := P(x_{g(1)}, \dots, x_{g(n)}) \end{cases}$$

We denote by $\mathbb{K}[\mathbf{x}]^G := \text{stab}_G \mathbb{K}[\mathbf{x}]$ the set of polynomials invariant under the action of G . Our main result here is that one can save a factor (roughly) $|G|$ when computing the FFT of an invariant polynomial. Our approach is in the same spirit as the one in [1], where similar statements on the savings induced by (very general) symmetries can be found. However, we are not aware of published results similar to Theorem 7 below.

3.1 The bivariate case

Let $P \in \mathbb{K}[x_1, x_2]$ be a symmetric bivariate polynomial of partial degrees less than $d = 8$, so that $\ell = 3$; let also $\omega \in \mathbb{K}$ be a primitive eighth root of unity. We detail on this easy example the methods used to exploit the symmetries to decrease the cost of the FFT.

The coefficients of P are placed on a 8×8 grid. The bivariate classical FFT consists in the application of butterflies of size $\delta := 2^{\ell-s}$ for s from 1 to 3, as in Figure 1. When some butterflies overlap, we draw only the shades of all but one of them. The result of stage $s=3$ is the set of evaluations $(P(\omega^{i_1}, \omega^{i_2}))_{i \in \mathbb{N}_8^2}$ in bit reversed order.




Figure 1. Classical bivariate FFT

We will remark below that each stage s preserves the symmetry of the input. In particular, since our polynomial P is symmetric, the coefficients \mathbf{c}_k at stage 1 are symmetric too, so we only need to compute the coefficients $\mathbf{c}_{(\mathbf{k}_1, \mathbf{k}_2)}^1$ for $(\mathbf{k}_1, \mathbf{k}_2)$ in the *fundamental domain* (sometimes called asymmetric unit) $0 \leq k_2 \leq k_1 < 8$. We choose to compute at stage s only the butterflies which involves at least an element of F ; the set of indices that are included in a butterfly of span $\delta = 2^{\ell-s}$ that meets F will be called the extension F_δ of F .




Figure 2. Symmetric bivariate FFT

Every butterfly of stage 1 meets the fundamental domain, so we do not save anything there. However, we save 4 out of 16 butterflies at stage 2 and 6 out of 16 butterflies at the third stage. Asymptotically in d , we gain a factor two in the number of arithmetic operations in \mathbb{K} compared to the classical FFT; this corresponds to the order of the group.

3.2 Notation

The group G also acts on \mathbb{N}^n with $g(\mathbf{i}) = (\mathbf{i}_{g(1)}, \dots, \mathbf{i}_{g(n)})$ for any $\mathbf{i} \in \mathbb{N}^n$. This action is consistent with the action on polynomials since $\varphi_g(\mathbf{x}^{\mathbf{k}}) = \mathbf{x}^{g^{-1}(\mathbf{k})}$.

Because G acts on polynomials, it acts on the vector of its coefficients. More generally, if $I \subseteq \mathbb{N}^n$ and $\mathbf{v} \in \mathbb{K}^{G(I)}$, we denote by $\mathbf{v}^g \in \mathbb{K}^I$ the vector defined by $\mathbf{v}_i^g = \mathbf{v}_{g(\mathbf{i})}$ for all $\mathbf{i} \in I$. If $I \subseteq \mathbb{N}^n$ is stable by G , i.e. $G(I) \subseteq I$, we define the set of invariant vectors $\mathbb{K}^{I,G}$ by $\mathbb{K}^{I,G} := \{\mathbf{v} \in \mathbb{K}^I : \forall g \in G, \mathbf{v}^g = \mathbf{v}\}$.

For any two sets I, J satisfying $J \subseteq I \subseteq \mathbb{N}^n$, we define the restriction $\mathbf{v}_J \in \mathbb{K}^J$ of $\mathbf{v} \in \mathbb{K}^I$ by $(\mathbf{v}_J)_j = \mathbf{v}_j$ for all $j \in J$. Recall the definition of the lexicographical order \leq_{lex} : for all $\mathbf{i}, \mathbf{j} \in \mathbb{N}^n$, $\mathbf{i} <_{\text{lex}} \mathbf{j}$ if there exists $k \in \{0, \dots, n-1\}$ such that $\mathbf{i}_t = \mathbf{j}_t$ for any $1 \leq t \leq k$ and $\mathbf{i}_{k+1} < \mathbf{j}_{k+1}$. Finally for any subset $S \subseteq \mathbb{N}^n$ stable by G , we define a fundamental domain S^G of the action of G on S by $S^G := \{\mathbf{i} \in S : \forall g \in G, \mathbf{i} \geq_{\text{lex}} g(\mathbf{i})\}$ together with the projection $\pi^G: \mathbb{N}^n \rightarrow \mathbb{N}^n$ such as $\{\pi^G(\mathbf{i})\} := G(\{\mathbf{i}\}) \cap (\mathbb{N}^n)^G$.

3.3 Fundamental domains

Let $F = F^{[d]} := (\mathbb{N}_d^n)^G$ be the fundamental domain associated to the action of G on \mathbb{N}_d^n . Any G -symmetric vector $\mathbf{c} \in \mathbb{K}^{\mathbb{N}_d^n}$ can be reconstructed from its restriction $\mathbf{c}_F \in \mathbb{K}^F$ using the formula $\mathbf{c}_i = (\mathbf{c}_F)_{\pi^G(\mathbf{i})}$. As it turns out, the in-place FFT algorithm from Section 2.2 admits the important property that the coefficients at all stage are still G -symmetric.

LEMMA 2. Let $P \in \mathbb{K}[\mathbf{x}]_d$ and the vectors $\mathbf{c}^0, \dots, \mathbf{c}^\ell \in \mathbb{K}^{\mathbb{N}_d^n}$ be as in Section 2.2. Then if $P \in \mathbb{K}[\mathbf{x}]_d^G$, $\mathbf{c}^0, \dots, \mathbf{c}^\ell \in \mathbb{K}^{\mathbb{N}_d^n, G}$.

PROOF. Given $P \in \mathbb{K}[\mathbf{x}]_d^G$, $\mathbf{k} \in \mathbb{N}_d^n$ and $g \in G$, we clearly have $\mathbf{c}_{g(\mathbf{k})}^0 = \mathbf{c}_{\mathbf{k}}^0$. For any $s \in \{0, \dots, \ell\}$, $\mathbf{i}, \mathbf{i}' \in \mathbb{N}_d^n$, $\mathbf{j} \in \mathbb{N}_d^{2^\ell - s}$ and $g \in G \subseteq \mathfrak{S}_n$, we also notice that

$$\begin{aligned} g(\mathbf{i} 2^{\ell-s} + \mathbf{j}) &= g(\mathbf{i}) 2^{\ell-s} + g(\mathbf{j}) \\ [g(\mathbf{i})]_s &= g([\mathbf{i}]_s) \\ g(\mathbf{i}) \cdot g(\mathbf{i}') &= \mathbf{i} \cdot \mathbf{i}'. \end{aligned}$$

Hence, using Equation (2), we get

$$\begin{aligned} \mathbf{c}_{g(\mathbf{i} 2^{\ell-s} + \mathbf{j})}^s &= \mathbf{c}_{g(\mathbf{i}) 2^{\ell-s} + g(\mathbf{j})}^s \\ &= \sum_{\mathbf{i}' \in \mathbb{N}_d^{2^s}} P_{\mathbf{i}' 2^{\ell-s} + g(\mathbf{j})} \omega^{[g(\mathbf{i})]_s \cdot \mathbf{i}' 2^{\ell-s}} \\ &= \sum_{\mathbf{i}' \in \mathbb{N}_d^{2^s}} P_{g(\mathbf{i}') 2^{\ell-s} + g(\mathbf{j})} \omega^{[g(\mathbf{i})]_s \cdot g(\mathbf{i}') 2^{\ell-s}} \\ &= \sum_{\mathbf{i}' \in \mathbb{N}_d^{2^s}} P_{g(\mathbf{i}' 2^{\ell-s} + \mathbf{j})} \omega^{g([\mathbf{i}]_s) \cdot g(\mathbf{i}') 2^{\ell-s}} \\ &= \sum_{\mathbf{i}' \in \mathbb{N}_d^{2^s}} P_{\mathbf{i}' 2^{\ell-s} + \mathbf{j}} \omega^{[\mathbf{i}]_s \cdot \mathbf{i}' 2^{\ell-s}} \\ &= \mathbf{c}_{\mathbf{i} 2^{\ell-s} + \mathbf{j}}^s. \end{aligned}$$

Thus, $\mathbf{c}_{g(\mathbf{k})}^s = \mathbf{c}_{\mathbf{k}}^s$ for all $\mathbf{k} \in \mathbb{N}_d^n$, whence $\mathbf{c}^0, \dots, \mathbf{c}^\ell \in \mathbb{K}^{\mathbb{N}_d^n, G}$. \square

This lemma implies that for the computation of the FFT of a G -symmetric polynomial P , it suffices to compute the projections $\mathbf{c}_F^0, \dots, \mathbf{c}_F^\ell$.

In order to apply formula (4) for the computation of \mathbf{c}_F^s as a function of \mathbf{c}_F^{s-1} , it is not necessary to completely reconstruct \mathbf{c}^{s-1} , due to the sparsity of the matrix \mathbf{B}^s . Instead, we define the δ -expansion of the set F by

$$F_\delta = \{\mathbf{k} \boxplus \boldsymbol{\epsilon} \delta : \mathbf{k} \in F, \boldsymbol{\epsilon} \in \mathbb{N}_2^n\},$$

where $\mathbf{i} \boxplus \mathbf{j}$ stands for the ‘‘bitwise exclusive or’’ of \mathbf{i} and \mathbf{j} . For any $\mathbf{k} \in \mathbb{N}_d^n$, the set $\{\mathbf{k} \boxplus \boldsymbol{\epsilon} \delta : \boldsymbol{\epsilon} \in \mathbb{N}_2^n\}$ describes the vertices of the butterfly of span δ that includes the point \mathbf{k} . Thus, F_δ is indeed the set of indices of \mathbf{c}^{s-1} that are involved in the computation of \mathbf{c}^s via the formula $\mathbf{c}^s = \mathbf{B}^s \mathbf{c}^{s-1}$.

LEMMA 3. For any $\delta \in \{1, 2, 4, \dots, 2^{\ell-1}\}$, let

$$F_{[\delta]} = \{\mathbf{k} \boxplus \boldsymbol{\epsilon} : \mathbf{k} \in F, \boldsymbol{\epsilon} \in \mathbb{N}_2^n\},$$

so that $F_\delta \subseteq F_{[2\delta]}$. Then we have $F_{[\delta]} = \delta F^{[d/\delta]} + \mathbb{N}_2^n$.

PROOF. Assume that $\mathbf{i} \in F^{[d/\delta]}$. Then clearly, $\delta \mathbf{i} \in F$, whence $\delta \mathbf{i} + \mathbb{N}_2^n \subseteq F_{[\delta]}$. Conversely, if $\mathbf{i} \in F_{[\delta]}$, then there exists a $\mathbf{j} \in F$ with $\mathbf{j} = \mathbf{i} \boxplus \boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon} \in \mathbb{N}_2^n$. Let $\mathbf{k} = \lfloor \mathbf{j} / \delta \rfloor$. For any $g \in G$, we have $g(\mathbf{k}) = \lfloor g(\mathbf{j}) / \delta \rfloor \leq_{\text{lex}} \lfloor \mathbf{j} / \delta \rfloor = \mathbf{k}$. Consequently, $\lfloor \mathbf{j} / \delta \rfloor \in F^{[d/\delta]}$, whence $\mathbf{i} \in \delta \lfloor \mathbf{j} / \delta \rfloor + \mathbb{N}_2^n$. \square

For the proof of the next lemma, for $(j, k) \in \{1, \dots, n\}^2$, define $\Delta_{j,k} = \{\mathbf{i} \in \mathbb{N}^n : \mathbf{i}_j = \mathbf{i}_k\}$ and $\Upsilon = \mathbb{N}_d^n \setminus \bigcup_{k \neq j} \Delta_{j,k}$.

LEMMA 4. There exists a constant C (depending on n) such that

$$\left| |F^{[d]}| - \frac{d^n}{|G|} \right| \leq C d^{n-1}.$$

PROOF. With the notation above, we have $|\Delta_{j,k} \cap \mathbb{N}_d^n| = d^{n-1}$. Taking $C = \binom{n}{2}$, it follows that $\left| \left(\bigcup_{k \neq j} \Delta_{j,k} \right) \cap \mathbb{N}_d^n \right| \leq C d^{n-1}$. On the other hand, the orbit of any element in Υ under G contains exactly $|G|$ elements and one element in F . In other words, $|F \cap \Upsilon| = |\Upsilon| / |G|$ and so $0 \leq |F| - |\Upsilon| / |G| \leq C d^{n-1}$. Finally, $0 \leq (d^n - |\Upsilon|) / |G| \leq C / |G| d^{n-1}$ which implies $-C / |G| d^{n-1} \leq |F| - d^n / |G| \leq C d^{n-1}$. \square

3.4 The symmetric FFT

Let $\mathbf{c}_{F_\delta}^{s-1}$ be the restriction of \mathbf{c}^{s-1} to F_δ and notice that \mathbb{K}^{F_δ} is stable under \mathbf{B}^s , i.e. $\mathbf{B}^s(\mathbb{K}^{F_\delta}) \subseteq \mathbb{K}^{F_\delta}$. Therefore the restriction $\mathbf{B}_{\mathbb{K}^{F_\delta}}^s$ of the map $\mathbf{B}^s: \mathbb{K}^{\mathbb{N}_d^n} \rightarrow \mathbb{K}^{\mathbb{N}_d^n}$ to vectors in \mathbb{K}^{F_δ} is a \mathbb{K} -algebra morphism from \mathbb{K}^{F_δ} to \mathbb{K}^{F_δ} . By construction, we now have $\mathbf{c}_{F_\delta}^s = \mathbf{B}_{\mathbb{K}^{F_\delta}}^s \mathbf{c}_{F_\delta}^{s-1}$. This allows us to compute \mathbf{c}_F^s as a function of \mathbf{c}_F^{s-1} using

$$\mathbf{c}_F^s = (\mathbf{B}_{\mathbb{K}^{F_\delta}}^s \xi_\delta(\mathbf{c}_F^{s-1}))_F \quad (6)$$

where $\xi_\delta(\mathbf{c}_F^{s-1})$ denotes the δ -expansion $\mathbf{c}_{F_\delta}^{s-1}$ of \mathbf{c}_F^{s-1} .

Remark 5. We could have saved a few more operations by only computing the coefficients \mathbf{c}_F^s . To do so, we would have performed just the operations of a butterfly corresponding to vertices inside F . However, the potential gain of complexity is only linear in the numbers of monomials d^n .

The formula (6) yields a straightforward way to compute the direct FFT of a G -symmetric polynomial:

Algorithm Symmetric-FFT
 INPUT: n, d and coefficients $\mathbf{c}_F^0 \in \mathbb{K}^{\mathbb{N}_d^G}$ of $P \in \mathbb{K}[\mathbf{x}]_d^G$ in F
 OUTPUT: $\mathbf{c}_F^\ell \in \mathbb{K}^{\mathbb{N}_d^G}$ in bit reversed order

For $s = 1, \dots, \ell$ do
 $\delta := d/2^s$
 For $\mathbf{i}' \in \mathbb{N}_{2^{\ell-s}}^n, \mathbf{j} \in \mathbb{N}_{2^{\ell-s}}^n$ s.t. $2\mathbf{i}'\delta + \mathbf{j} \in F_\delta$ do
 For $\mathbf{b} \in \mathbb{N}_2^G$ do $\mathbf{c}_b := \mathbf{c}_{\pi^s \circ ((2\mathbf{i}'+\mathbf{b})\delta + \mathbf{j})}^{s-1}$
 $\mathbf{d}_b = \text{Butterfly}(n, d, s, \mathbf{i}', \mathbf{j}, \mathbf{c}_b)$
 For $\mathbf{b} \in \mathbb{N}_2^G$ do $\mathbf{c}_{\pi^s \circ ((2\mathbf{i}'+\mathbf{b})\delta + \mathbf{j})}^s := \mathbf{d}_b$
 Return \mathbf{c}^ℓ

Remark 6. The main challenge for an actual implementation of our algorithms consists in finding a way to iterate on sets like F_δ without too much overhead. We give a hint on how to iterate over F in Lemma 8 (see also [22] for similar considerations).

The inverse FFT can be computed classically by unrolling the loops in the inverse order and inverting the operations of the butterflies. The main result of this section is then the following theorem, where $\mathbb{F}(d, n)$ denotes the cost of a full n -dimensional FFT of multi-degree d .

THEOREM 7. *For fixed n and G , and for $d \rightarrow \infty$, the direct (resp. inverse) symmetric FFT can be computed in time*

$$\mathbb{T}(d, n) = \frac{1}{|G|} \mathbb{F}(d, n) + \mathcal{O}(d^n).$$

PROOF. At stage s of the computation, the ratio between the computation of \mathbf{c}_F^s as a function of \mathbf{c}_F^{s-1} and \mathbf{c}^s as a function of \mathbf{c}^{s-1} is $|F_{2^{\ell-s}}|/|\mathbb{N}_2^G|$, so that

$$\frac{\mathbb{T}(d, n)}{\mathbb{F}(d, n)} \leq \frac{|F_{2^{\ell-1}}| + |F_{2^{\ell-2}}| + \dots + |F_2|}{\ell |\mathbb{N}_2^G|}.$$

Lemmas 3 and 4 thus imply

$$|F_\delta| \leq |F_{2\delta}| \leq |F^{[d/(2\delta)]}| (2\delta)^n \leq d^n/|G| + 2C\delta d^{n-1},$$

whence

$$\frac{\mathbb{T}(d, n)}{\mathbb{F}(d, n)} \leq \frac{\ell d^n/|G| + 2C d^n}{\ell d^n} = \frac{1}{|G|} + \frac{2C}{\ell}.$$

The result follows for $d \rightarrow \infty$. \square

Finally, the following lemma gives a description of an “open” subset of F , characterized only by simple inequalities. Although we do not describe implementation questions in detail here, we point out that this simple characterization would naturally be used in order to iterate over the sets F_δ , possibly using tree-based techniques as in [17].

LEMMA 8. *Let Π be the set of $(j, k) \in \{1, \dots, n\}^2$ such that $j < k$ and $\exists g \in G$ such that $g(i) = i$ for $i \leq j$ and $g(j) = k$. For $(j, k) \in \{1, \dots, n\}^2$, define $H_{j,k} = \{\mathbf{i} \in \mathbb{N}^n : \mathbf{i}_j > \mathbf{i}_k\}$. Then*

$$F \cap \Upsilon = \bigcap_{(j,k) \in \Pi} H_{j,k} \cap \mathbb{N}_d^G.$$

PROOF. Assume that $\mathbf{i} \in \Upsilon$ does not lie in F . Then $g(\mathbf{i}) >_{\text{lex}} \mathbf{i}$ for some $g \in G$. Let $j \in \{1, \dots, n\}$ be minimal with $k = g(j) \neq j$, whence $k > j$ and $(j, k) \in \Pi$. Since $\mathbf{i}_j \neq \mathbf{i}_k$, it follows that $\mathbf{i}_k > \mathbf{i}_j$, so $\mathbf{i} \notin H_{j,k}$. Inversely, assume that $\mathbf{i} \in \mathbb{N}_d^G$ does not lie in $\bigcap_{(j,k) \in \Pi} H_{j,k}$, so that there exists $(j, k) \in \Pi$ with $\mathbf{i} \notin H_{j,k}$. Let $g \in G$ be such that $g(i) = i$ for $i \leq j$ and $g(j) = k$. Then, $g(\mathbf{i}) >_{\text{lex}} \mathbf{i}$. \square

4. THE LATTICE FFT

In this section, we deal with polynomials supported on lattices; our main result is an algorithm for the multiplication of such polynomials using Fourier transforms.

The first subsection introduces the main objects we will need (a lattice Λ and its dual Γ); then, we will give an algorithm, based on Smith normal form computation, for the FFT on what we will call a *basic domain*, and we will finally deduce a multiplication algorithm for the general case.

The techniques we use, based on Smith normal form computation, can be found in several papers, originating from [13]; in particular, the results in Section 4.2 are essentially in [33] (in a more general form).

4.1 Lattice polynomials

Assume that \mathbb{K} admits a primitive k th root of unity for any order $k > 1$, which we will denote by $e^{2\pi i/k}$. Let Λ be a free \mathbb{Z} -submodule of \mathbb{Z}^n of rank n , generated by the vectors $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_n \in \mathbb{N}^n$. Then

$$\mathbb{K}[\mathbf{x}]_\Lambda = \{P \in \mathbb{K}[\mathbf{x}] : \text{supp } P \subseteq \Lambda\}$$

is a subring of $\mathbb{K}[\mathbf{x}]$, and we will call elements of $\mathbb{K}[\mathbf{x}]_\Lambda$ *lattice polynomials*. First, we show that these polynomials are the invariant polynomial ring for a diagonal matrix group.

The set \mathbb{Q}^n acts on $\mathbb{K}[\mathbf{x}]$ via, for any $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_n) \in \mathbb{Q}^n$, the isomorphism $\varphi_\boldsymbol{\gamma}$ of $\mathbb{K}[\mathbf{x}]$ given by

$$\varphi_\boldsymbol{\gamma} : P(x_1, \dots, x_n) \mapsto P(e^{2\pi i \gamma_1} x_1, \dots, e^{2\pi i \gamma_n} x_n).$$

Note that $\varphi_{\boldsymbol{\gamma} + \mathbf{v}} = \varphi_\boldsymbol{\gamma}$ for any $\mathbf{v} \in \mathbb{Z}^n$. The action of $\varphi_\boldsymbol{\gamma}$ on the monomial $\mathbf{x}^{\boldsymbol{\lambda}_i}$ is given by $\varphi_\boldsymbol{\gamma}(\mathbf{x}^{\boldsymbol{\lambda}_i}) = e^{2\pi i(\boldsymbol{\lambda}_i \cdot \boldsymbol{\gamma})} \mathbf{x}^{\boldsymbol{\lambda}_i}$.

In particular, all elements of $\mathbb{K}[\mathbf{x}]_\Lambda$ are invariants under the action of $\varphi_\boldsymbol{\gamma}$ if and only if

$${}^t \boldsymbol{\Lambda} \boldsymbol{\gamma} \in \mathbb{Z}^n, \tag{7}$$

where $\boldsymbol{\Lambda} \in \mathbb{N}^{n \times n}$ is the matrix with columns $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_n$, that is $\boldsymbol{\Lambda}_{i,j} = (\boldsymbol{\lambda}_j)_i$. Let Γ be the dual (or reciprocal) lattice of Λ , that is the set of all $\boldsymbol{\gamma} \in \mathbb{Q}^n$ satisfying Equation (7). A basis of Γ is given by the columns of $\boldsymbol{\Gamma} = {}^t \boldsymbol{\Lambda}^{-1} \in \mathbb{Q}^{n \times n}$.

Let G be the group of actions $\{\varphi_\boldsymbol{\gamma} : \boldsymbol{\gamma} \in \boldsymbol{\Gamma}\}$. It is generated by $\{\varphi_{\boldsymbol{\gamma}_i}\}_{1 \leq i \leq n}$, where $\boldsymbol{\gamma}_i$ are the columns of $\boldsymbol{\Gamma}$. From Equation (7), we deduce that G is the diagonal matrix group of all actions $\varphi_\boldsymbol{\gamma}$ which leave $\mathbb{K}[\mathbf{x}]_\Lambda$ invariant. Conversely, because monomials are mapped to monomials by elements of G , the ring of invariants is spanned by the monomials $\mathbf{x}^\boldsymbol{\lambda}$ with ${}^t \boldsymbol{\Lambda} \boldsymbol{\lambda} \in \mathbb{Z}^n$, i.e. $\boldsymbol{\lambda}$ belongs to the dual of Γ . Since the dual of a lattice is the lattice itself, we deduce that $\boldsymbol{\lambda} \in \Lambda$ and $\mathbb{K}[\mathbf{x}]_\Lambda = \mathbb{K}[\mathbf{x}]^G$. Note that only $\boldsymbol{\Gamma}$ modulo $\mathbb{Z}^{n \times n}$ matters to determine the group G .

Example 9. Consider the lattice Λ generated by $\lambda_1 = (2, 0)$ and $\lambda_2 = (1, 1)$, and let $\mathbf{\Lambda} = \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$. The lattice polynomials $\mathbb{K}[\mathbf{x}]_\Lambda$ are the polynomials $\mathbb{K}[x_1^2, x_1 x_2, x_2^2]$. We have $\mathbf{\Gamma} = \begin{pmatrix} 1/2 & 0 \\ -1/2 & 1 \end{pmatrix}$, so G is the group generated by $\varphi_{\gamma_1}: P(x_1, x_2) \mapsto P(-x_1, -x_2)$ and $\varphi_{\gamma_2} = \text{Id}$. The lattice polynomials $\mathbb{K}[\mathbf{x}]_\Lambda = \mathbb{K}[x_1^2, x_1 x_2, x_2^2]$ are those polynomials invariant under the symmetry $P(x_1, x_2) \mapsto P(-x_1, -x_2)$.

4.2 The lattice FFT on a basic domain

Given $\mathbf{d} = (d_1, \dots, d_n)$, we define

$$\begin{aligned} \mathbb{K}[\mathbf{x}]_{\mathbf{d}} &= \{P \in \mathbb{K}[\mathbf{x}] : \deg_{x_1}(P) < d_1, \dots, \deg_{x_n}(P) < d_n\} \\ \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{d}} &= \mathbb{K}[\mathbf{x}]_\Lambda \cap \mathbb{K}[\mathbf{x}]_{\mathbf{d}}. \end{aligned}$$

For $i \in \{1, \dots, n\}$, let $p_i > 0$ be minimal such that $x_i^{p_i} \in \mathbb{K}[\mathbf{x}]_\Lambda$. We call the block $\mathbb{N}_{\mathbf{p}} := \mathbb{N}_{p_1} \times \dots \times \mathbb{N}_{p_n}$ a *basic domain* for Λ . In this subsection, we consider the computation of the FFT at order \mathbf{p} of a polynomial $P \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$.

In what follows, we set $\omega_i = e^{2\pi i/p_i}$ for each i , so that we have to show how to compute the set of evaluations $P(\omega^{\mathbf{k}}) = P(\omega_1^{k_1}, \dots, \omega_n^{k_n})$ for $\mathbf{k} \in \mathbb{N}_{\mathbf{p}}$. If we proceeded directly, we would compute more evaluations than the number of monomials of P , which is $|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$. We show how to reduce the problem to computing exactly $|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$ evaluations.

For any $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_n) \in \Gamma$, notice that

$$\begin{aligned} P(\omega^{\mathbf{k}}) &= P(e^{2\pi i k_1/p_1}, \dots, e^{2\pi i k_n/p_n}) \\ &= P(e^{2\pi i(k_1/p_1 + \gamma_1)}, \dots, e^{2\pi i(k_n/p_n + \gamma_n)}). \end{aligned}$$

Therefore we would only need to consider the evaluations with multi-indices in $(1/p_1 \mathbb{Z}, \dots, 1/p_n \mathbb{Z})$ modulo Γ . There are only $|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$ such evaluations but, as in Section 3.4, we would have to expand the fundamental domain of evaluations at each stage of the FFT to compute the butterflies.

Instead, we propose a more direct method with no domain expansion. We show that, regarding the evaluations, polynomials P in $\mathbb{K}[\mathbf{x}]_\Lambda$ can be written $Q(\mathbf{x}^{\lambda'_1}, \dots, \mathbf{x}^{\lambda'_n})$ where $\lambda'_1, \dots, \lambda'_n$ is a basis of Λ and the evaluation can be done directly in this rewritten form.

We introduce the notation $\bar{k}^p \in \{0, \dots, p-1\}$ for the remainder of $k \in \mathbb{Z}$ by $p \in \mathbb{N}^*$. If $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{Z}^n$ and $\mathbf{p} = (p_1, \dots, p_n) \in (\mathbb{N}^*)^n$, we let $\bar{\mathbf{k}}^{\mathbf{p}} := (\bar{k}_1^{p_1}, \dots, \bar{k}_n^{p_n})$. This notation is motivated by the remark that $\mathbf{x}^{\lambda}(\omega^{\mathbf{k}})$ depends only on the class of λ and \mathbf{k} modulo $(p_1 \mathbb{Z}, \dots, p_n \mathbb{Z})$.

LEMMA 10. *There exists a basis $(\lambda'_1, \dots, \lambda'_n)$ of Λ and a basis $(\mathbf{k}_1, \dots, \mathbf{k}_n)$ of \mathbb{Z}^n such that*

$$\mathbf{y}_i(\omega^{\mathbf{k}_j}) = e^{2\pi i \delta_{i,j}/q_i} \quad (8)$$

where $\mathbf{y}_i = \mathbf{x}^{\bar{\lambda}'_i \mathbf{p}}$ and the q_i 's are positive integers satisfying $q_1 | q_2 | \dots | q_n$.

PROOF. Let us consider the lattice of the exponents of $(\mathbf{x}^{\lambda'_i}(\omega^{\mathbf{k}}))_{1 \leq i \leq n}$ for $\mathbf{k} \in \mathbb{N}^n$. If $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{N}^n$, then

$$\mathbf{x}^{\lambda'_i}(\omega^{\mathbf{k}}) = e^{2\pi i[(k_1/p_1, \dots, k_n/p_n) \cdot \lambda'_i]} = e^{2\pi i[\mathbf{k} \cdot ((\lambda'_i)_1/p_1, \dots, (\lambda'_i)_n/p_n)]}.$$

We define the lattice Δ spanned by the columns of $\mathbf{\Delta} = {}^t \mathbf{\Lambda} \begin{pmatrix} 1/p_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & 1/p_n \end{pmatrix}$, that is $\Delta_{i,j} = (\lambda_i)_j/p_j$. We want to take the Smith normal form of $\mathbf{\Delta}$ but the coefficients of $\mathbf{\Delta}$ are not necessarily integers. So we multiply the matrix by $\ell := \text{LCM}(p_1, \dots, p_n)$, take the Smith normal form and divide by ℓ . Therefore there exists ${}^t \mathbf{H}, \mathbf{K} \in \text{GL}_n(\mathbb{Z})$ and integers $d_n | \dots | d_2 | d_1$ such that

$${}^t \mathbf{H} \mathbf{\Delta} \mathbf{K} = \begin{pmatrix} d_1/\ell & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & d_n/\ell \end{pmatrix}. \quad (9)$$

Let us prove that $\mathbb{Z}^n \subseteq \Delta$. By definition of p_1, \dots, p_n , there exists $\mathbf{S} \in \mathbb{Z}^{n \times n}$ such that $\mathbf{\Lambda} \mathbf{S} = \begin{pmatrix} p_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & p_n \end{pmatrix}$. Thus we have ${}^t \mathbf{S} {}^t \mathbf{\Lambda} \begin{pmatrix} 1/p_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & 1/p_n \end{pmatrix} = \text{Id}$, implying that $\mathbf{\Delta} {}^t \mathbf{S} = \text{Id}$ and our result. Because $\mathbb{Z}^n \subseteq \Delta$, we have $d_i | \ell$ and by setting $q_i := \ell/d_i \in \mathbb{N}^*$, we get $q_1 | q_2 | \dots | q_n$.

With a geometrical point of view, the equality of Equation (9) gives the existence of the two required bases. The columns of \mathbf{K} give the basis $(\mathbf{k}_1, \dots, \mathbf{k}_n)$ of \mathbb{Z}^n and the columns of $\mathbf{\Lambda}' := \mathbf{\Lambda} \mathbf{H}$ give the basis $(\lambda'_1, \dots, \lambda'_n)$ of Λ . To sum up, we have

$${}^t \mathbf{\Lambda}' \begin{pmatrix} 1/p_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & 1/p_n \end{pmatrix} \mathbf{K} = \begin{pmatrix} 1/q_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & 1/q_n \end{pmatrix} \quad (10)$$

which is the matricial form of Equation (8). \square

PROPOSITION 11. *The following ring morphism Φ*

$$\begin{aligned} \mathbb{K}[\mathbf{y}]/(y_1^{q_1} - 1, \dots, y_n^{q_n} - 1) &\longrightarrow \mathbb{K}[\mathbf{x}]/(x_1^{p_1} - 1, \dots, x_n^{p_n} - 1) \\ \mathbf{y}_i &\longmapsto \mathbf{x}^{\bar{\lambda}'_i \mathbf{p}} \end{aligned}$$

is well-defined, injective and its image is $\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$. Moreover, if $P = \Phi(Q)$ then

$$P(\omega^{\ell_1 \mathbf{k}_1 + \dots + \ell_n \mathbf{k}_n}) = Q(e^{2\pi i \ell_1/q_1}, \dots, e^{2\pi i \ell_n/q_n}). \quad (11)$$

PROOF. First, we prove that Φ is well defined. For this matter we have to check that $\Phi(y_i^{q_i} - 1) = \mathbf{x}^{q_i \bar{\lambda}'_i \mathbf{p}} - 1 = 0$ modulo $(x_1^{p_1} - 1, \dots, x_n^{p_n} - 1)$. It is sufficient to prove that $q_i \bar{\lambda}'_i \mathbf{p} \in (p_1 \mathbb{Z}, \dots, p_n \mathbb{Z})$, which follows from Equation (10).

Then we prove Equation (11). Let $Q \in \mathbb{K}[\mathbf{y}]$ and $P(\mathbf{x}) := \Phi(Q)(\mathbf{x})$. As a consequence of Lemma 10, one has

$$P(\omega^{\ell_1 \mathbf{k}_1 + \dots + \ell_n \mathbf{k}_n}) = Q(e^{2\pi i \ell_1/q_1}, \dots, e^{2\pi i \ell_n/q_n}).$$

Now, we prove that Φ is injective. Indeed if $Q \in \mathbb{K}[\mathbf{y}]$ satisfies $\Phi(Q) = 0$ then for all $\mathbf{k} \in \mathbb{N}^n$, one has $\Phi(Q)(\omega^{\mathbf{k}}) = 0$. Using Equation (11), we get that for all $\ell_1, \dots, \ell_n \in \mathbb{Z}$, $Q(e^{2\pi i \ell_1/q_1}, \dots, e^{2\pi i \ell_n/q_n}) = 0$. As a result, Q belongs to the ideal generated by $(y_1^{q_1} - 1, \dots, y_n^{q_n} - 1)$.

Finally it is trivial to see that the image of Φ is included in $\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$. Reciprocally, let $\mathbf{x}^\lambda \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$. We write $\lambda = \sum_{i=1}^n \ell_i \bar{\lambda}'_i \mathbf{p}$ with $\ell_i \in \mathbb{Z}$ and define $\bar{\lambda} = \sum_{i=1}^n \bar{\ell}_i^{q_i} \bar{\lambda}'_i \mathbf{p}$. Now because the lattice $q_1 \lambda'_1 \mathbb{Z} \oplus \dots \oplus q_n \lambda'_n \mathbb{Z}$ is included in $(p_1 \mathbb{Z}, \dots, p_n \mathbb{Z})$, we have that $\Phi(\mathbf{y}^{(\bar{\ell}_1^{q_1}, \dots, \bar{\ell}_n^{q_n})}) = \mathbf{x}^{\bar{\lambda}} = \mathbf{x}^\lambda$ modulo $(x_1^{p_1} - 1, \dots, x_n^{p_n} - 1)$. \square

In particular, the FFT of P is uniquely determined by its restriction to evaluations of the form $P(\omega^{\mathbf{k}})$ with $\mathbf{k} = \ell_1 \mathbf{k}_1 + \dots + \ell_n \mathbf{k}_n$ and $\ell = (\ell_1, \dots, \ell_n) \in \mathbb{N}_{\mathbf{q}}$. Notice that Proposition 11 implies that the number $q_1 \dots q_n$ of evaluations equals the numbers $|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$ of monomials in $\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$.

We define the *lattice FFT* of P to be the vector of values $(P(\omega^{\ell_1 \mathbf{k}_1 + \dots + \ell_n \mathbf{k}_n}))_{\ell \in \mathbb{N}_{\mathbf{q}}}$. We have thus shown that the computation of the lattice FFT of P reduces to the computation of an ordinary multivariate FFT of order \mathbf{q} .

4.3 The general lattice FFT multiplication

Assume that $\mathbf{d} \in (\mathbb{N}^*)^n$ is such that $p_i \mid d_i$ for each i and consider the computation of the FFT at order \mathbf{d} of a lattice polynomial $P \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{d}}$. In practice, one often has $\mathbf{d}_i = p_i 2^{\ell_i}$, and this is what we will assume from now on. For simplicity, we also suppose that $\ell_1 = \dots = \ell_n$ and denote by ℓ this common value. The results are valid in general but we would have to do more decimations in some variables than in others, which would not match the notations of Section 2.

We give here an algorithm for the multiplication of two polynomials $P_1, P_2 \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{d}/2}$.

We start by doing ℓ FFT stages. These stages preserve the lattice Λ , because the butterflies have a span which belongs to Λ . Therefore we compute only the butterflies whose vertices are in Λ .

After these ℓ stages, we twist the coefficients thanks to Formula 5 and obtain the polynomials $Q^\ell = (Q_i^\ell)$, for $\mathbf{i} \in \mathbb{N}_{2^\ell}^n$. As explained in Section 2.2, up to a minor modification (here, the partial degrees are not all the same), these polynomials reduce our multiplication to $2^{n\ell}$ multiplications in $\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}/(x_1^{p_1} - 1, \dots, x_n^{p_n} - 1)$. Thanks to Proposition 11, we are able to perform these latter multiplications as multiplications in $\mathbb{K}[\mathbf{y}]_{\mathbf{q}}/(y_1^{q_1} - 1, \dots, y_n^{q_n} - 1)$.

Algorithm Lattice-partial-FFT
INPUT: n, \mathbf{d}, ω and coefficients $\mathbf{c}^0 \in \mathbb{K}^{\mathbb{N}_{2^\ell}^n}$ of $P \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{d}}$
OUTPUT: the coefficients $\mathbf{c}^\ell \in \mathbb{K}^{\mathbb{N}_{2^\ell}^n}$

For $s=1, \dots, \ell$ do // ℓ decimation steps in each variable
 $\delta := d/2^s$
For $\mathbf{i}' \in \mathbb{N}_{2^{s-1}}^n, \mathbf{j} \in \Lambda \cap \mathbb{N}_{d/2^s}^n$ do // a butterfly in Λ
For $\mathbf{b} \in \mathbb{N}_2^n$ do $\mathbf{c}_{\mathbf{b}} := \mathbf{c}_{(2\mathbf{i}'+\mathbf{b})\delta+\mathbf{j}}^{s-1}$
 $\mathbf{d}_{\mathbf{b}} = \mathbf{Butterfly}(n, d, s, \mathbf{i}', \mathbf{j}, \mathbf{c}_{\mathbf{b}})$
For $\mathbf{b} \in \mathbb{N}_2^n$ do $\mathbf{c}_{(2\mathbf{i}'+\mathbf{b})\delta+\mathbf{j}}^s := \mathbf{d}_{\mathbf{b}}$

Return \mathbf{c}^s

Algorithm Lattice-FFT-multiplication
INPUT: n, \mathbf{d}, ω and $P_1, P_2 \in \mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{d}}$
OUTPUT: the coefficients product $P_1 P_2$

For $i=1, 2$ do
 $\mathbf{c}^\ell = \mathbf{Lattice-partial-FFT}(n, 2\mathbf{d}, \omega, ((P_i)_{\mathbf{k}})_{\mathbf{k} \in \mathbb{N}_{2^\ell}^n})$
Let $Q_i := Q_i^\ell \in (\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}})^{\mathbb{N}_{2^\ell}^n}$ obtained from \mathbf{c}^ℓ
Transform Q_i as a vector of $\mathbb{K}[\mathbf{y}]_{\mathbf{q}}$ using Proposition 11
Multiply $Q = Q_1 Q_2$ in $(\mathbb{K}[\mathbf{y}]/(y_1^{q_1} - 1, \dots, y_n^{q_n} - 1))^{\mathbb{N}_{2^\ell}^n}$
Transform Q as a vector of $\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}}$ using Proposition 11
Recover \mathbf{c}^ℓ from $Q \in (\mathbb{K}[\mathbf{x}]_{\Lambda, \mathbf{p}})^{\mathbb{N}_{2^\ell}^n}$
 $\mathbf{c}^0 = \mathbf{Lattice-partial-FFT}^{-1}(n, 2\mathbf{d}, \omega, (\mathbf{c}_{[\mathbf{k}]_\ell}^\ell)_{\mathbf{k} \in \mathbb{N}_{2^\ell}^n})$
Return \mathbf{c}^0

As for symmetric polynomials, the inverse algorithm **Lattice-partial-FFT**⁻¹ is obtained by reversing the loops and inverting the butterflies of **Lattice-partial-FFT**.

Let us analyze the cost of this algorithm. It starts with ℓ stages of classical n -dimensional FFT, computing only the butterflies whose vertices are in Λ . This amounts to $\frac{3}{2} n \ell |\Lambda \cap \mathbb{N}_{\mathbf{d}}^n|$ arithmetic operations in \mathbb{K} . The second part is made of the transformations of Equation 5 and Proposition 11. Since we consider that operations with the lattice Λ take time $\mathcal{O}(1)$, these transformations take time $\mathcal{O}(|\Lambda \cap \mathbb{N}_{\mathbf{d}}^n|)$. Finally, the componentwise multiplication of $(\mathbb{K}[\mathbf{y}]/(y_1^{q_1} - 1, \dots, y_n^{q_n} - 1))^{\mathbb{N}_{2^\ell}^n}$ costs $2^{n\ell} \mathbf{M}(\mathbf{q}_1, \dots, \mathbf{q}_n)$ where $\mathbf{M}(\mathbf{d})$ stands for the arithmetic complexity of polynomials multiplication in $\mathbb{K}[\mathbf{x}]_{\mathbf{d}}$. Altogether, our multiplication algorithms costs

$$\mathbf{T}(\mathbf{d}) := 3/2 n \ell |\Lambda \cap \mathbb{N}_{\mathbf{d}}^n| + 2^{n\ell} \mathbf{M}(\mathbf{q}) + \mathcal{O}(|\Lambda \cap \mathbb{N}_{\mathbf{d}}^n|).$$

By contrast, a symmetry-oblivious approach would consist in doing ℓ stages of classical n -dimensional FFT and then $2^{n\ell}$ multiplications in $\mathbb{K}[\mathbf{x}]/(x_1^{p_1} - 1, \dots, x_n^{p_n} - 1)$. The cost analysis is similar to the one before and the classical approach's cost is $\mathbf{F}(\mathbf{d}) := 3/2 n \ell d^n + 2^{n\ell} \mathbf{M}(\mathbf{p}) + \mathcal{O}(d^n)$.

The ratio $d^n/|\Lambda \cap \mathbb{N}_{\mathbf{d}}^n|$ is exactly the volume $\text{vol}(\Lambda)$ of the lattice, defined by $\text{vol}(\Lambda) := \det(\mathbf{A}) \in \mathbb{Z}$. Under the superlinearity assumption for the function \mathbf{M} , we get that $\text{vol}(\Lambda) \mathbf{M}(\mathbf{q}) \leq \mathbf{M}(\mathbf{p})$ and we deduce the following theorem.

THEOREM 12. *For a fixed lattice Λ and $d \rightarrow \infty$, the direct (resp. inverse) lattice FFT can be computed in time*

$$\mathbf{T}(\mathbf{d}) = \frac{1}{\text{vol}(\Lambda)} \mathbf{F}(\mathbf{d}) + \mathcal{O}(d^n).$$

5. THE SYMMETRIC TFT

To conclude this paper, we briefly discuss the extensions of the previous results to the Truncated Fourier Transform (TFT). With the notation of Section 2, let $I \subseteq \mathbb{N}_d^n$ be an initial segment subset of \mathbb{N}_d^n : for $\mathbf{i}, \mathbf{j} \in \mathbb{N}_d^n$, if $\mathbf{i} \in I$ and $\mathbf{j} \leq \mathbf{i}$, then $\mathbf{j} \in I$. Given $P \in \mathbb{K}[\mathbf{x}]_I$, the TFT of P is defined to be the vector $\hat{P} \in \mathbb{K}^I$ defined by $\hat{P}_{\mathbf{i}} = P(\omega^{[\mathbf{i}]_\ell})$, where ω is a root of unity of order d .

In [16, 17], van der Hoeven described fast algorithms for computing the TFT and its inverse. For a fixed dimension n , the cost of these algorithms is bounded by $\mathcal{O}(|I| \log |I| + d^n)$ instead of $\mathcal{O}(d^n \log d^n)$. In this section we will outline how a further acceleration can be achieved for symmetric polynomials of the types studied in Sections 3 and 4.

5.1 The symmetric TFT

For each $\delta \in \{1, \dots, 2^\ell\}$, let $I_{[\delta]} = \{\mathbf{i} \boxplus \boldsymbol{\epsilon} : \mathbf{i} \in I, \boldsymbol{\epsilon} \in \mathbb{N}_\delta^n\}$. The entire computation of the TFT and its inverse can be represented schematically by a graph Γ . The vertices of the graph are pairs (s, \mathbf{i}) with $s \in \{0, \dots, \ell\}$ and $\mathbf{i} \in I_{[2^{\ell-s}]}$. The edges are between vertices (s, \mathbf{i}) and $(s+1, \mathbf{j})$ with $\mathbf{i} = \mathbf{j} \boxplus (2^{\ell-s-1} \boldsymbol{\epsilon})$ and $\boldsymbol{\epsilon} \in \mathbb{N}_2^n$. The edge is labeled by a constant that we will simply write $c_{s, \mathbf{i}, \mathbf{j}}$ such that

$$P_j^{s+1} = \sum_{\mathbf{i}} c_{s, \mathbf{i}, \mathbf{j}} P_i^s. \quad (12)$$

For a direct TFT, we are given P_i^0 on input and “fill out” the remaining values P_i^s for increasing values of s using (12). In the case of an inverse TFT, we are given the P_i^ℓ with $\mathbf{i} \in I$ on input, as well as the coefficients $P_i^0 = 0$ with $\mathbf{i} \in \mathbb{N}_d^n \setminus I$. We next “fill out” the remaining values P_i^s using the special algorithm described in [17].

Now let G be as in Section 3 and assume that I is stable under G . Then each of the $I_{[\delta]}$ is also stable under G . Furthermore, any $\mathbf{i} \in I$ lies in the orbit of some element of $I \cap F$ under the action of G . Let $(I \cap F)_{[\delta]} = \{\mathbf{i} \boxplus \epsilon: \mathbf{i} \in I \cap F, \epsilon \in \mathbb{N}_\delta^s\}$. Given a G -symmetric input polynomial on input, the idea behind the symmetric TFT is to use the restriction Γ' of the above graph Γ by keeping only those vertices (s, \mathbf{i}) such that $\mathbf{i} \in (F \cap I)_{[2^{\ell-s}]}$. The symmetric TFT and its inverse can then be computed by intertwining the above “filling out” process with steps in which we compute $P_{g(\mathbf{i})}^s = P_{\mathbf{i}}^s$ for all $\mathbf{i} \in \mathbb{N}_d^n$ and $g \in G$ such that $P_{\mathbf{i}}^s$ is known but not $P_{g(\mathbf{i})}^s$.

The computational complexities of the symmetric TFT and its inverse are both proportional to the size $|\Gamma'|$ of Γ' . For many initial domains I of interest, it can be shown that $|\Gamma'| \sim |\Gamma|/|G|$, as soon as d gets large. This is for instance the case for $I = \Sigma_k = \{\mathbf{i} \in \mathbb{N}_d^n: i_1 + \dots + i_n \leq k\}$, when $k \rightarrow \infty$, and where $\ell = \lceil \log_2 k \rceil$. Indeed, using similar techniques as in the proof of Theorem 7, we first show that $|(\Sigma_k \cap F)_{[\delta]} - (\Sigma_k)_{[\delta]}|/|G| = \mathcal{O}(\delta d^{n-1})$, and then conclude in a similar way.

5.2 The lattice TFT

Let us now consider the case of a polynomial $P \in \mathbb{K}[\mathbf{x}]_{\Lambda \cap I}$, with Λ as in Section 4. In order to design a fast “lattice TFT”, the idea is simply to replace I by a slightly larger set J which preserves the fundamental domains. More precisely, with p_1, \dots, p_n as in Section 4, we take

$$J = \{(p_1 i_1 + j_1, \dots, p_n i_n + j_n): \mathbf{i} \in K, \mathbf{j} \in \mathbb{N}_{\mathbf{p}}\}$$

$$K = \{(\lfloor i_1/p_1 \rfloor, \dots, \lfloor i_n/p_n \rfloor): \mathbf{i} \in I\}.$$

A lattice TFT of order $(p_1 2^{\ell_1}, \dots, p_n 2^{\ell_n})$ can then be regarded as $|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$ TFTs at order $(2^{\ell_1}, \dots, 2^{\ell_n})$ and initial segment K , followed by $|K|$ TFTs lattice FFTs on a fundamental domain. Asymptotically speaking, we thus gain a factor $|\mathbb{N}_{\mathbf{p}}|/|\Lambda \cap \mathbb{N}_{\mathbf{p}}|$ with respect to a usual TFT with initial segment J . In the case when $I = \Sigma_k = \{\mathbf{i} \in \mathbb{N}_d^n: i_1 + \dots + i_n \leq k\}$, we finally notice that $|J| \sim |I|$ for $k \rightarrow \infty$.

6. CONCLUSION

Let us quickly outline possible generalizations of the results of this paper.

It seems that the most general kinds of finite groups for which the techniques in this paper work are finite subgroups G of $\Pi_n U_n$, where Π_n is the group of $n \times n$ permutation matrices and U_n the group of diagonal matrices whose entries are all roots of unity. Indeed, any such group G both acts on the sets $\mathbb{K}[\mathbf{x}]$ and on the torus \mathbb{U}^n , where $\mathbb{U} = \{e^{2\pi i/n}: i, n \in \mathbb{N}\}$. Even more generally, the results may still hold for closed algebraic subgroups G generated by an infinite number of elements of $\Pi_n U_n$.

Many other interesting groups can be obtained as conjugates $G' = T^{-1} G T$ of groups G of the above kind. For certain applications, such as the integration of dynamical systems, the change of coordinates T can be done “once and for all” on the initial differential equations, after which the results of this paper again apply.

It is classical that the FFT of a polynomial with real coefficients can be computed twice as fast (roughly speaking) as a polynomial with complex coefficient. Real polynomials can be considered as symmetric polynomials for complex conjugation: $P(\bar{z}) = \overline{P(z)}$. Some further extensions of our setting are possible by including this kind of symmetries.

On some simple examples, we have verified that the ideas of this paper generalize to other evaluation-interpolation models for polynomial multiplication, such as Karatsuba multiplication and Toom-Cook multiplication.

We intend to study the above generalizations in more detail in a forthcoming paper.

7. REFERENCES

- [1] L. Auslander, J. R. Johnson and R. W. Johnson. An equivariant Fast Fourier Transform algorithm. Drexel University Technical Report DU-MCS-96-02, 1996.
- [2] M. Ben-Or and P. Tiwari. A deterministic algorithm for sparse multivariate polynomial interpolation. In *STOC '88: Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 301–309. New York, NY, USA, 1988. ACM Press.
- [3] R. Bergmann. The fast Fourier transform and fast wavelet transform for patterns on the torus. *Applied and Computational Harmonic Analysis*, , 2012. In Press.
- [4] P. Bürgisser, M. Clausen and M. A. Shokrollahi. *Algebraic complexity theory*. Springer-Verlag, 1997.
- [5] J. Canny, E. Kaltofen and Y. Lakshman. Solving systems of nonlinear polynomial equations faster. In *Proc. ISSAC '89*, pages 121–128. Portland, Oregon, A.C.M., New York, 1989. ACM Press.
- [6] D.G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28:693–701, 1991.
- [7] J.W. Cooley and J.W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Computat.*, 19:297–301, 1965.
- [8] S. Czapor, K. Geddes and G. Labahn. *Algorithms for Computer Algebra*. Kluwer Academic Publishers, 1992.
- [9] P. Duhamel and M. Vetterli. Fast Fourier transforms: a tutorial review and a state of the art. *Signal Process.*, 19(4):259–299, apr 1990.
- [10] T. S. Freeman, G. M. Imirzian, E. Kaltofen and Y. Lakshman. DAGWOOD a system for manipulating polynomials given by straight-line programs. *ACM Trans. Math. Software*, 14:218–240, 1988.
- [11] S. Garg and É. Schost. Interpolation of polynomials given by straight-line programs. *Theoretical Computer Science*, 410(27-29):2659–2662, 2009.
- [12] M. Gastineau and J. Laskar. TRIP 1.2.26. TRIP Reference manual, IMCCE, 2012. <http://www.imcce.fr/trip/>.
- [13] A. Guessoum and R. Mersereau. Fast algorithms for the multidimensional discrete Fourier transform. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 34(4):937–943, 1986.
- [14] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2-nd edition, 2002.
- [15] J. van der Hoeven. Relax, but don't be too lazy. *JSC*, 34:479–542, 2002.
- [16] J. van der Hoeven. The truncated Fourier transform and applications. In J. Gutierrez, editor, *Proc. ISSAC 2004*, pages 290–296. Univ. of Cantabria, Santander, Spain, July 4–7 2004.
- [17] J. van der Hoeven. Notes on the Truncated Fourier Transform. Technical Report 2005-5, Université Paris-Sud, Orsay, France, 2005.
- [18] J. van der Hoeven. Newton's method and FFT trading. *JSC*, 45(8):857–878, 2010.
- [19] J. van der Hoeven and G. Lecerf. On the bit-complexity of sparse polynomial multiplication. Technical Report, HAL, 2010. <http://hal.archives-ouvertes.fr/hal-00476223>, accepted for publication in JSC.

- [20] J. van der Hoeven and G. Lecerf. On the complexity of blockwise polynomial multiplication. In *Proc. ISSAC '12*, pages 211–218. Grenoble, France, July 2012.
- [21] J. van der Hoeven and É. Schost. Multi-point evaluation in higher dimensions. Technical Report, HAL, 2010. <http://hal.archives-ouvertes.fr/hal-00477658>, accepted for publication in AAECC.
- [22] J. Johnson and X. Xu. Generating symmetric DFTs and equivariant FFT algorithms. In *ISSAC'07*, pages 195–202. ACM, 2007.
- [23] S. C. Johnson. Sparse polynomial arithmetic. *SIGSAM Bull.*, 8(3):63–71, 1974.
- [24] E. Kaltofen and Y. N. Lakshman. Improved sparse multivariate polynomial interpolation algorithms. In *ISSAC '88: Proceedings of the international symposium on Symbolic and algebraic computation*, pages 467–474. Springer Verlag, 1988.
- [25] E. Kaltofen, Y. N. Lakshman and J.-M. Wiley. Modular rational sparse multivariate polynomial interpolation. In *ISSAC '90: Proceedings of the international symposium on Symbolic and algebraic computation*, pages 135–139. New York, NY, USA, 1990. ACM Press.
- [26] E. Kaltofen, W. Lee and A. A. Lobo. Early termination in Ben-Or/Tiwari sparse interpolation and a hybrid of Zippel's algorithm. In *ISSAC '00: Proceedings of the 2000 international symposium on Symbolic and algebraic computation*, pages 192–201. New York, NY, USA, 2000. ACM Press.
- [27] A. Kudlicki, M. Rowicka and Z. Otwinowski. The crystallographic Fast Fourier Transform. recursive symmetry reduction. *Acta Cryst.*, A63:465–480, 2007.
- [28] Victor Y. Pan. Simple multivariate polynomial multiplication. *JSC*, 18(3):183–186, 1994.
- [29] V. Pan and D. Bini. *Polynomial and matrix computations*. Birkhauser, 1994.
- [30] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. *Computing*, 7:281–292, 1971.
- [31] D. R. Stoutemyer. Which polynomial representation is best? In *Proceedings of the 1984 MACSYMA Users' Conference: Schenectady, New York, July 23–25, 1984*, pages 221–243. 1984.
- [32] L. F. Ten Eyck. Crystallographic Fast Fourier Transform. *Acta Cryst.*, A29:183–191, 1973.
- [33] A. Vince and X. Zheng. Computing the Discrete Fourier Transform on a hexagonal lattice. *Journal of Mathematical Imaging and Vision*, 28:125–133, 2007.