1

Through the Kaleidoscope; Symmetries, Groups and Chebyshev Approximations from a Computational Point of View

H. Z. Munthe-Kaas, M. Nome and B. N. Ryland

Dept. of Mathematics, Univ. of Bergen, P.O. Box 7800, 5020 Norway

Abstract

In this paper we survey parts of group theory, with emphasis on structures that are important in design and analysis of numerical algorithms and in software design. In particular, we provide an extensive introduction to Fourier analysis on locally compact abelian groups, and point towards applications of this theory in computational mathematics. Fourier analysis on non-commutative groups, with applications, is discussed more briefly. In the final part of the paper we provide an introduction to multivariate Chebyshev polynomials. These are constructed by a kaleidoscope of mirrors acting upon an abelian group, and have recently been applied in numerical Clenshaw–Curtis type numerical quadrature and in spectral element solution of partial differential equations, based on triangular and simplicial subdivisions of the domain.

1.1 Introduction

Group theory is the mathematical language of symmetry. As a mature branch of mathematics, with roots going almost two centuries back, it has evolved into a highly technical discipline. Many texts on group theory and representation theory are not readily accessible to applied mathematicians and computational scientists, and the relevance of group theoretical techniques in computational mathematics is not widely recognized.

Nevertheless, it is our conviction that knowledge of central parts of group theory and harmonic analysis on groups is invaluable also for computational scientists, both as a language to unify, analyze and generalize computational algorithms and also as an organizing principle of mathematical software construction.

The first part of the paper provides a quite detailed self contained introduction to Fourier analysis in the language of compact abelian groups. Abelian groups are fundamental structures of major importance in computations as they describe spaces with commutative shift operations. Both continuous and discrete abelian groups are omnipresent as computational domains in applied mathematics. Classical Fourier analysis can be understood as the theory of linear operators which commute with translations in space; examples being differential operators with constant coefficients, and other linear operators which are invariant under change of time or space. A central theme in computation is the interplay between continuous and discrete structures, between continuous mathematical models and their discretizations. Many aspects of discretizations can be understood via subgroups and sub-lattices of continuous abelian groups. The duality between time/space and frequency/wave numbers provided by the Fourier transform is another central topic which mathematically is expressed via Pontryagin duality of abelian groups.

Symmetries is a core topic of all group theory. Translation invariance of linear operators is one type of symmetry. Another kind of symmetry is invariance of operators under isometries acting upon the domain, e.g. reflection symmetries. In the second part of this paper we will discuss the importance of certain kaleidoscopic groups generated by a set of mirrors acting upon a domain. We will review basic properties of multivariate versions of Chebyshev polynomials, which are constructed by the folding of exponential functions under the action of kaleidoscopes. Theoretically these have remarkable properties, both from an approximation theoretical point of view and also with respect to computational complexity. However, until recently they have not been applied in computations, probably due to the quite complicated underlying theory. We will briefly review recent work where multivariate Chebyshev polynomials are applied in numerical quadrature and in spectral element solution of PDEs, and discuss some advantages and difficulties in this line of work.

1.2 Fourier analysis on groups

Classical Fourier analysis is intimately connected with the theory of locally compact abelian groups. In mathematical analysis this view allows a unified presentation of the Fourier transform, Fourier series and the discrete Fourier transform [31]. In computational science, some understanding of the group structures underlying discrete and continuous Fourier analysis is invaluable both as a language to discuss computational aspects of Fourier transforms and sampling theory, as well as an organizing principle of mathematical software [14]. In this section we provide a self-contained review of Fourier analysis on abelian groups. We conclude this section with a brief discussion on generalized Fourier transforms on non-abelian groups and point to some applications of representation theory in computational science. We refer to [5, 7, 11, 15, 16, 27, 31, 33] for more details on the topics of this chapter.

1.2.1 Locally compact abelian groups

A Locally Compact Abelian group (LCA) is a locally compact topological space G which is also an abelian group. Thus G has an identity element 0, a group operation $+$ and a unary $-$ operation, such that $a+b=b+a$, $a+(b+c)=(a+b)+c$ and $a+(-a)=0$ for all $a,b,c \in G$. For our discussion it suffices to consider the *elementary* LCAs and groups isomorphic to one of these:

Definition 1.1 The *elementary LCAs* are:

- The *reals* $\mathbb{R}$ under addition, with the standard definition of open sets.
- The *integers* $\mathbb{Z}$ under addition (with the discrete topology). This is also known as the infinite cyclic group.
- The 1-dimensional *torus*, or *circle* $T=\mathbb{R}/2\pi\mathbb{Z}$ defined as $[0,2\pi)\subset\mathbb{R}$ under addition modulo 2π, with the circle topology.
- The *cyclic group* of order k, $\mathbb{Z}_k=\mathbb{Z}/k\mathbb{Z}$, which consists of the integers $0,1,\ldots,k-1$ under addition modulo k (with the discrete topology).
- *Direct products* of the above spaces, $G\times H$, in particular $\mathbb{R}^n$ (real n-space), T^n (the n-torus) and all finitely generated abelian groups.

It is often convenient to consider these groups in their isomorphic *multiplicative form.* E.g. the multiplicative group of complex numbers of modulus 1, denoted $\mathbb{T}=\{\, z\in\mathbb{C} \mid |z|=1 \,\}$ is isomorphic to T via the exponential map $T \ni \theta \mapsto \exp(i\theta) \in \mathbb{T}$. More generally, if G is an additive abelian group with elements λ, we denote a corresponding multiplicative group as e^G with elements e^λ, multiplicative group operation $e^\lambda e^{\lambda'} = e^{\lambda+\lambda'}$ and unit $1 = e^0$. The exponential notation is here just a formal change of notation, a 'syntactic sugar' to simplify certain calculations.

An LCA with a finite subset of generators is called a *Finitely Generated Abelian group* (FGA). These, and in particular the finite ones, are the basic domains of computational Fourier analysis and the Fast Fourier Transform. A complete classification of these is simple and very useful:

Theorem 1.2 (Classification of FGA) *If G is an FGA, then G is isomorphic to a group of the form*

$$\mathbb{Z}^n \times \mathbb{Z}_{p_0^{n_0}} \times \mathbb{Z}_{p_1^{n_1}} \times \cdots \times \mathbb{Z}_{p_{k-1}^{n_{k-1}}} \tag{1.1}$$

where p_i are primes, $p_0 \leq p_1 \leq \cdots \leq p_{n-1}$ and $n_i \leq n_{i+1}$ whenever $p_i = p_{i+1}$. Furthermore G is also isomorphic to a group of the form

$$\mathbb{Z}^n \times \mathbb{Z}_{m_0} \times \mathbb{Z}_{m_1} \times \cdots \times \mathbb{Z}_{m_{\ell-1}} \tag{1.2}$$

where m_i divides m_{i+1}. In both forms the representation is unique, i.e. two FGA are isomorphic iff they can be transformed into the same canonical form[1].

1.2.2 Dual groups and the Fourier transform

Group homomorphisms are fundamental in describing sampling theory, fast Fourier transforms and convolution theorems.

Definition 1.3 For two LCA H and G, we let $\hom(H, G)$ denote the set of continuous homomorphisms from H to G, i.e. all continuous maps $\phi\colon H \to G$ such that

$$\phi(h + h') = \phi(h) + \phi(h') \quad \text{for all } h, h' \in H.$$

In particular $\chi \in \hom(G, \mathbb{T})$ are homomorphisms into the multiplicative group of complex numbers with modulus 1, defined as $\chi(x) = e^{i\theta(x)}$, where $\theta \in \hom(G, \mathbb{R})$. These homomorphisms χ are called the *characters* of G. The characters are always given by exponential maps and form the basis for the Fourier transform. For two characters χ, χ' it is clear that the product $(\chi\chi')(g) = \chi(g)\chi'(g)$ is also a character, and also the complex conjugate $\overline{\chi}$ is a character, which is the multiplicative inverse of χ. Thus $\hom(G, \mathbb{T})$ is an abelian group called the *dual group* of G,

[1] The isomorphisms are provided by the *Chinese remainder theorem.*

denoted $\widehat{G}$. With the compact-open topology $\widehat{G}$ is also an LCA, thus, we can define the double dual $\widehat{\widehat{G}}$. Pontryagin's duality theorem states that G and $\widehat{\widehat{G}}$ are isomorphic LCAs. Due to this symmetry, we adopt a slightly different definition of dual LCAs. In the following, we consider both G and $\widehat{G}$ as additive abelian groups and define the duality between these via a dual pairing.

Definition 1.4 (Dual pair) Two LCAs G and $\widehat{G}$ are called a *dual pair of LCAs* if there exists a continuous function $\langle\cdot,\cdot\rangle : \widehat{G}\times G \to \mathbb{T}$ such that

$$\begin{aligned}\hom(G,\mathbb{T}) &= \left\{ \langle k,\cdot\rangle \mid k\in\widehat{G} \right\}\\ \hom(\widehat{G},\mathbb{T}) &= \{ \langle\cdot,x\rangle \mid x\in G \}.\end{aligned}$$

The definition implies that

$$\langle k+k',x\rangle = \langle k,x\rangle\cdot\langle k',x\rangle \tag{1.3}$$

$$\langle k,x+x'\rangle = \langle k,x\rangle\cdot\langle k,x'\rangle \tag{1.4}$$

$$\langle 0,x\rangle = \langle k,0\rangle = 1 \tag{1.5}$$

$$\overline{\langle k,x\rangle} = \langle -k,x\rangle = \langle k,-x\rangle. \tag{1.6}$$

Furthermore, if $\langle k,x\rangle = 1$ for all k, then $x = 0$, and if $\langle k,x\rangle = 1$ for all x then $k = 0$. If we work with pairings on different groups G, H, we write $\langle\cdot,\cdot\rangle_G$ and $\langle\cdot,\cdot\rangle_H$ to distinguish them.

Every LCA has a translation invariant Haar measure, yielding an integral $\int_G f(x)d\mu$ which is uniquely defined up to a scaling. For $\mathbb{R}^n$ and T^n, this is the standard integral, and for the discrete $\mathbb{Z}_k$ and $\mathbb{Z}$, it is the sum over the elements. The characters are orthogonal under the inner product defined by the Haar measure. $L^2(G)$ denotes square integrable functions on G.

Definition 1.5 (Fourier transform) The Fourier transform is a unitary[2] map $\widehat{\ }: L^2(G) \to L^2(\widehat{G})$ given as

$$\widehat{f}(k) = \int_G \langle -k,x\rangle f(x)dx. \tag{1.7}$$

There exists a constant C such that the inverse transform is given as

$$f(x) = \frac{1}{C}\int_{\widehat{G}} \langle k,x\rangle \widehat{f}(k)dk. \tag{1.8}$$

Sometimes we will write $\mathcal{F}(f)$ or $\mathcal{F}_G(f)$ instead of $\widehat{f}$.

[2] Unitary up to a scaling.

Note that $\widehat{G}$ is discrete if and only if G is compact. In this case the inversion formula becomes a sum, and $C = \int_G 1dx = \mathrm{vol}(G)$. The following table presents the dual pairs for the elementary groups $\mathbb{R}$, T, $\mathbb{Z}$ and $\mathbb{Z}_n$.

G	$\widehat{G}$	$\langle \cdot, \cdot \rangle$	$\widehat{f}(\cdot)$	$f(\cdot)$
$x \in \mathbb{R}$	$\omega \in \mathbb{R}$	$e^{2\pi i\omega x}$	$\int_{-\infty}^{\infty} e^{-2\pi i\omega x} f(x)dx$	$\int_{-\infty}^{\infty} e^{2\pi i\omega x} \widehat{f}(\omega)d\omega$
$x \in T$	$k \in \mathbb{Z}$	$e^{2\pi ikx}$	$\int_0^1 e^{-2\pi ikx} f(x)dx$	$\sum_{k=-\infty}^{\infty} e^{2\pi ikx} \widehat{f}(k)$
$j \in \mathbb{Z}_n$	$k \in \mathbb{Z}_n$	$e^{\frac{2\pi ikj}{n}}$	$\sum_{j=0}^{n-1} e^{\frac{-2\pi ikj}{n}} f(j)$	$\frac{1}{n} \sum_{k=0}^{n-1} e^{\frac{2\pi ikj}{n}} \widehat{f}(k)$

Multidimensional versions are given by the componentwise formulae:

$$
\begin{aligned}
x &= (x_1, x_2) \in G = G_1 \times G_2 \\
k &= (k_1, k_2) \in \widehat{G} = \widehat{G_1} \times \widehat{G_2} \\
\langle k, x \rangle &= \langle k_1, x_1 \rangle \cdot \langle k_2, x_2 \rangle \\
\widehat{f}(k_1, k_2) &= \int_{G_1} \int_{G_2} \langle -k_1, x_1 \rangle \langle -k_2, x_2 \rangle f(x_1, x_2) dx_1 dx_2 \\
f(x_1, x_2) &= \frac{1}{C_1 C_2} \int_{\widehat{G_1}} \int_{\widehat{G_2}} \langle k_1, x_1 \rangle \langle k_2, x_2 \rangle f(k_1, k_2) dk_1 dk_2.
\end{aligned}
$$

We end this section with a brief discussion of shifts and the convolution theorem. For a finite group G, we let $\mathbb{C}G$ denote the *group algebra* (or group ring). This is the complex vector space of dimension $|G|$, where each element in G is a basis vector, and with a product given by convolution. The convolution is most easily computed if we write the basis vector for $\lambda \in G$ in the multiplicative form e^λ. Then $f \in \mathbb{C}G$ is represented by the vector $f = \sum_{\lambda \in G} f(\lambda) e^\lambda$ and we obtain the convolution $f * g$ of $f, g \in \mathbb{C}G$ by computing their product and collecting equal terms

$$fg = \sum_{\lambda \in G} f(\lambda) e^\lambda \sum_{\lambda' \in G} g(\lambda') e^{\lambda'} = \sum_{\lambda \in G} (f * g)(\lambda) e^\lambda,$$

where

$$(f * g)(\lambda) = \sum_{\lambda' \in G} f(\lambda') g(\lambda - \lambda'). \tag{1.9}$$

For a continuous group G, we can similarly understand $\mathbb{C}G$ as a function space of complex valued functions on G. Care must be taken in order to define a suitable function space for which convolutions make sense, but these issues will not be discussed in this paper, see [31]. *All results tacitly assume that we define an appropriate function space, e.g. $L^2(G)$, where operations such as convolutions and Fourier transforms are well-defined.*

The continuous convolution is given for $f, g \in \mathbb{C}G$ as

$$(f * g)(x) = \int_{x' \in G} f(x')g(x - x')dx'. \tag{1.10}$$

Arguably the most important property of the Fourier transform is that it diagonalizes the convolution, i.e.

$$\widehat{(f * g)}(\xi) = \widehat{f}(\xi)\widehat{g}(\xi) \quad \text{for all } \xi \in \widehat{G}. \tag{1.11}$$

The proof is a straightforward computation. It is worthwhile to notice that it relies upon the property that the group characters are homomorphisms, i.e. $\langle \xi, x+x' \rangle = \langle \xi, x \rangle \langle \xi, x' \rangle$, and upon the translation invariance of the integral.

Similar results hold for more general non-commutative groups, in which case the characters are replaced by group representations. Irreducible group representations are matrix valued group homomorphisms which form an orthogonal basis for $L^2(G)$, stated by the Frobenius theorem for finite G and Peter–Weyl theorem for compact G. Applications of such generalized Fourier transforms in numerical linear algebra are discussed in [2].

1.2.3 Subgroups, lattices and sampling

A subgroup of G, written $H < G$, is a topologically closed subset $H \subset G$ that is closed under the group operations $+$ and $-$.

Definition 1.6 For a subgroup $H < G$ we define the *annihilator subgroup* $H^{\perp} < \widehat{G}$ as

$$H^{\perp} = \left\{ k \in \widehat{G} \mid \langle k, h \rangle_G = 1 \text{ for all } h \in H \right\}.$$

Note that if $k, k' \in \widehat{G}$ are such that $k - k' \in H^{\perp}$, then $\langle k, h \rangle_G = \langle k', h \rangle_G$ for all $h \in H$. This phenomenon is called *aliasing* in signal processing, the two characters corresponding to k and k' are indistinguishable when restricted to H.

Since G is abelian we can always form the quotient group $K = G/H$, where the elements of K are the cosets $H + g$ with the group operation $(H + g) + (H + g') = H + g + g'$, and the identity element of K is H. Similarly, we can form the quotient $\widehat{H}/H^{\perp}$, where each coset $H^{\perp} + k$ consists of a set of characters aliasing on H.

Definition 1.7 A *lattice* in G is a discrete subgroup $H < G$ such that G/H is compact.

An example is $G = \mathbb{R}$, $H = \mathbb{Z}$ and $K = \mathbb{R}/\mathbb{Z} = T$. Also $\mathbb{Z}\times\mathbb{Z} < \mathbb{R}\times\mathbb{R}$ is a lattice. However, $\mathbb{Z}\times 0 < \mathbb{R}\times\mathbb{R}$ is *not* a lattice since the quotient $T\times\mathbb{R}$ is not compact. If G is a finite group, then any $H < G$ is a lattice.

As another example, consider $G = \mathbb{Z}$ with the lattice $2\mathbb{Z} < G$ consisting of all even integers. Theorem 1.2 states that $2\mathbb{Z}$ is isomorphic to $\mathbb{Z}$ as an abstract group, thus it seems natural to define $H = \mathbb{Z}$ and identify H with a subgroup of G via a group homomorphism $\phi_0 \in \hom(H, G)$ given as $\phi_0(j) = 2j$. Similarly the quotient is $K = \mathbb{Z}_2$, identified with the cosets of H in G via $\phi_1 \in \hom(G, K)$ given as $\phi_1(j) = j \bmod 2$, where two elements $j, j' \in G$ belong to the same coset iff $\phi_1(j - j') = 0$.

It is in general easy to define which properties of $\phi_1 \in \hom(H, G)$ and $\phi_2 \in (G, K)$ are necessary and sufficient for H to be (isomorphic to) a subgroup of G with quotient (isomorphic to) K. Recall that the kernel and image of $\phi \in \hom(G_1, G_2)$ are defined as $\ker(\phi) = \{\, x \in G_1 \mid \phi(x) = 0 \,\}$ and $\mathrm{im}(\phi) = \phi(G_1) \subset G_2$. A *(co)chain complex* is a sequence $\{G_j, \phi_j\}$ of homomorphisms between abelian groups $\phi_j \in \hom(G_j, G_{j+1})$ such that $\mathrm{im}(\phi_j) \subset \ker(\phi_{j+1})$ for all j, in other words $\phi_{j+1}\circ\phi_j = 0$ for all j. The sequence is *exact* if $\ker(\phi_{j+1}) = \mathrm{im}(\phi_j)$. A *short exact sequence* is an exact sequence of five terms of the form

$$\mathbf{0} \longrightarrow H \xrightarrow{\phi_0} G \xrightarrow{\phi_1} K \longrightarrow \mathbf{0}\ . \tag{1.12}$$

The leftmost and rightmost arrows are the trivial maps $0 \mapsto 0$ and $K \mapsto 0$. A short exact sequence defines a subgroup $H < G$ with quotient $K = G/H$, or more precisely ϕ_0 is a monomorphism (injective homomorphism) identifying H with a subgroup $\phi_0(H) < G$ and ϕ_1 an epimorphism (surjective homomorphism) identifying $G/\phi_0(H)$ with K. Henceforth we will always define H as a subgroup of G with quotient K by explicitly defining a short exact sequence and the maps ϕ_0 and ϕ_1.

Although this homological algebra point of view is ubiquitous in many areas of pure mathematics, it is not a commonly used language in applied and computational mathematics. However, this presentation is also very important from a computational point of view. First of all, this language allows for a general and unified discussion of sampling, interpolation and Fast Fourier Transforms. Furthermore, from an object oriented programming point of view, it is an advantage to characterize mathematical objects in terms of categorical diagrams. Classes in an object oriented program consist of an internal representation of a certain abstraction as well as an interface defining the interaction and relationship between different objects. Category theory and homological algebra is thus a language which is important in defining classes in object oriented programming. We will not discuss implementations further here, but refer to [1, 14] for examples of this line of ideas within numerical analysis. To understand sampling theory and the FFT in a group language, we need to define *adjoint homomorphisms*, similarly to adjoints of linear operators.

Definition 1.8 Given two LCAs H and G with dual pairings $\langle\cdot,\cdot\rangle_H$ and $\langle\cdot,\cdot\rangle_G$. The adjoint of a homomorphism $\phi \in \hom(H,G)$ is $\widehat{\phi} \in \hom(\widehat{G},\widehat{H})$ defined such that

$$\langle \xi, \phi(x)\rangle_G = \langle \widehat{\phi}(\xi), x\rangle_H \quad \text{for all } \xi \in \widehat{G} \text{ and } x \in H.$$

The following fundamental theorem is proven by standard techniques in homological algebra.

Theorem 1.9 *A short sequence of LCAs*

$$\mathbf{0} \longrightarrow H \xrightarrow{\phi_0} G \xrightarrow{\phi_1} K \longrightarrow \mathbf{0}\ .$$

is exact if and only if the adjoint sequence

$$\mathbf{0} \longleftarrow \widehat{H} \xleftarrow{\widehat{\phi}_0} \widehat{G} \xleftarrow{\widehat{\phi}_1} \widehat{K} \longleftarrow \mathbf{0}\ .$$

is exact.

Note that $\widehat{K}$ is a subgroup of $\widehat{G}$ with quotient $\widehat{G}/\widehat{K} = \widehat{H}$. Furthermore, for any $x \in H$ and for any $k \in \widehat{K}$, we have that $\langle\widehat{\phi}_1(k), \phi_0(h)\rangle_G = \langle k, \phi_1 \circ \phi_0(x)\rangle_K = \langle k, 0\rangle_K = 1$, since the composition of any two adjacent arrows is 0. The exactness of the adjoint sequence implies that if $\xi \in \widehat{G}$ is such that $\langle \xi, \phi_0(h)\rangle_G = 0$ for all $h \in H$, then $\xi = \widehat{\phi}_1(k)$ for some $k \in \widehat{K}$. Thus Theorem 1.9 implies the fundamental dualities

$$\begin{aligned} H &< G, & K &= G/H, \\ \widehat{H} &= \widehat{G}/\widehat{K}, & \widehat{K} &= H^{\perp} < \widehat{G}. \end{aligned}$$

For many computational problems it is necessary to choose representative elements from each of the cosets in the quotient groups G/H and $\widehat{G}/\widehat{K}$. E.g. in sampling theory, all characters in a coset $\widehat{K} + \xi \subset \widehat{G}$ alias on H, but physical relevance is usually given to the character $\xi' \in \widehat{K} + \xi$ which is closest to 0 (the lowest frequency mode). Similarly, we often represent G/H by picking a representative from each coset, e.g. $\mathbb{R}/2\pi\mathbb{Z}$ can be represented by $[0, 2\pi) \subset \mathbb{R}$. The quotient map $\phi_1 \colon G \to K$ assigns each coset to a unique element in K, and we need to decide on a right inverse of this map.

Definition 1.10 (Transversal of coset map) Given the short exact sequence (1.12), a function $\sigma\colon K \to G$ is called a *transversal* of the quotient map $\phi_1\colon G \to K$ if $\phi_1 \circ \sigma = \mathrm{Id}_K$.

Note that in general we cannot choose σ as a group homomorphism (only if $G = H \times K$), but it can be chosen as a continuous function. In most applications G has a natural norm (e.g. Euclidean distance) and we can choose σ such that the coset representatives are as close to the origin as possible, i.e. such that $||\sigma(k)|| \leq ||\sigma(k) - h||$ for all $h \in \phi_0(H)$. This choice is called a *Voronoi transversal.* It is usually not uniquely defined on the boundary, and the treatment of points on the boundary must be done with some care in many applications. If H is a lattice in a continuous group G, then the closure of the image $\sigma(K) \subset G$ is a polyhedron limited by hyperplanes halfway between 0 and its neighbouring lattice points. In sampling theory one usually picks out the coset representatives for aliasing characters in $\widehat{G}/\widehat{K}$ by letting $\widehat{\sigma}\colon \widehat{H} \to \widehat{G}$ be the Voronoi transversal of $\widehat{\phi}_0$ with respect to the L^2 norm on $\widehat{G}$.

In the rest of this section, we assume that H, G and K form a short exact sequence as in (1.12), where H is a lattice, i.e. H is discrete and K is compact. Then $\widehat{H}$ is compact and $\widehat{K} = H^\perp$ is discrete, so $H^\perp$ is a lattice in $\widehat{G}$, called the *reciprocal lattice.* For a function $f \in \mathbb{C}G$, we let $f_H = f \circ \phi_0 \in \mathbb{C}H$ denote the function f *downsampled* to the lattice H. Similarly, $\widehat{f}_{H^\perp} = \widehat{f} \circ \widehat{\phi}_1 \in \mathbb{C}H^\perp$.

Lemma 1.11 (Poisson summation formula) *Given a lattice $H < G$ with reciprocal lattice $H^\perp < \widehat{G}$, there exists a constant C such that*

$$\sum_H f_H = \frac{1}{C} \sum_{H^\perp} \widehat{f}_{H^\perp}. \tag{1.13}$$

If G is compact, then $C = \mathrm{vol}(G/H)$. With our normalization of the Fourier transform on $\mathbb{R}$, $C = \mathrm{vol}(G/H)$ also when $G = \mathbb{R}^n$.

Proof Consider the group G/H via its set of coset representatives $V := \sigma(K) \subset G$. The characters of this group are $\{\langle \widehat{\phi}_1(k), \cdot \rangle_G\}_{k \in \widehat{K}}$. Thus by Fourier inversion in G/H there exists a constant C such that

$$f(0) = C \sum_{k \in \widehat{K}} \int_V f(x) \overline{\langle \widehat{\phi}_1(k), x \rangle_G} dx.$$

We write $x \in G$ as $x = y + \phi_0(h)$, where $y \in V$, use $\langle \widehat{\phi}_1(k), \phi_0(h) \rangle \equiv 1$

and the result above to obtain:

$$\begin{aligned}\sum_{k\in\widehat{K}} \widehat{f}\circ\widehat{\phi}_1(k) &= \sum_{k\in\widehat{K}} \int_G f(x)\overline{\langle\widehat{\phi}_1(k), x\rangle_G} dx \\ &= \sum_{k\in\widehat{K}} \sum_{h\in H} \int_V f(y+\phi_0(h))\overline{\langle\widehat{\phi}_1(k), \phi_0(h)+y\rangle_G} dy \\ &= \sum_{h\in H} \sum_{k\in\widehat{K}} \int_V f(y+\phi_0(h))\overline{\langle\widehat{\phi}_1(k), y\rangle_G} dy \\ &= C\sum_{h\in H} f\circ\phi_0(h).\end{aligned}$$

If G is compact, we set $f = 1$ and compute $\widehat{f} = \mathrm{vol}(G)\delta_0$. Since $|H| = \mathrm{vol}(G)/\mathrm{vol}(V)$ we find $C = \mathrm{vol}(V)$. The constant is computed for $G = \mathbb{R}^n$ by considering the Fourier transform of $f(x) = e^{-x^T x}$, which under an appropriate scaling, is invariant under the Fourier transform on $\mathbb{R}^n$.

□

1.2.4 Heisenberg groups and the FFT

More material on topics related to this section is found in [4, 5, 35].

We can act upon $f \in \mathbb{C}G$ with a time-shift $S_x f(t) := f(t+x)$ and with a frequency shift $\chi_\xi f(t) := \langle\xi, t\rangle f(t)$. These two operations are dual under the Fourier transform, but do not commute:

$$\widehat{S_x f}(\xi) = \chi_x \widehat{f}(\eta) \tag{1.14}$$

$$\widehat{\chi_\xi f}(\eta) = S_{-\xi}\widehat{f}(\eta) \tag{1.15}$$

$$(S_x\chi_\xi f)(t) = \langle\xi, x\rangle \cdot (\chi_\xi S_x f)(t). \tag{1.16}$$

The full (non-commutative) group generated by time and frequency shifts on $\mathbb{C}G$ is called the *Heisenberg group* of G.

The Heisenberg group of $\mathbb{R}^n$ is commonly defined as the multiplicative group of matrices of the form

$$\begin{pmatrix} 1 & x^T & s \\ 0 & I_n & \xi \\ 0 & 0 & 1 \end{pmatrix},$$

where $\xi, x \in \mathbb{R}^n$, $s \in \mathbb{R}$. This group is isomorphic to the semidirect product $\mathbb{R}^n \times \mathbb{R}^n \rtimes \mathbb{R}$ where

$$(\xi', x', s') \cdot (\xi, x, s) = (\xi' + \xi, x' + x, s' + s + x'^T\xi).$$

We prefer to instead consider $\mathbb{R}^n \times \mathbb{R}^n \rtimes \mathbb{T}$ (where $\mathbb{T}$ is the multiplicative group consisting of $z \in \mathbb{C}$ such that $|z| = 1$) with product

$$(\xi', x', z') \cdot (\xi, x, z) = (\xi' + \xi, x' + x, z' z e^{2\pi i x'^T \xi}).$$

More generally:

Definition 1.12 For an LCA G we define the Heisenberg group $\mathcal{H}_G = \widehat{G} \times G \rtimes \mathbb{T}$ with the semidirect product

$$(\xi', x', z') \cdot (\xi, x, z) = (\xi' + \xi, x' + x, z' \cdot z \cdot \langle \xi, x' \rangle).$$

We define a *left action* $\mathcal{H}_G \times \mathbb{C}G \to \mathbb{C}G$ as follows

$$(\xi, x, z) \cdot f = z \cdot \chi_\xi S_x f. \tag{1.17}$$

To see that this defines a left action, we check that $(0, 0, 1) \cdot f = f$ and

$$(\xi', x', z') \cdot ((\xi, x, z) \cdot f) = ((\xi', x', z') \cdot (\xi, x, z)) \cdot f.$$

Lemma 1.13 *Let $\mathcal{H}_G = \widehat{G} \times G \rtimes \mathbb{T}$ and $\mathcal{H}_{\widehat{G}} = G \times \widehat{G} \rtimes \mathbb{T}$ act upon $f \in \mathbb{C}G$ and $\widehat{f} \in \mathbb{C}\widehat{G}$ as in (1.17). Then*

$$\mathcal{F}((\xi, x, z) \cdot f) = z \cdot \langle -\xi, x \rangle \cdot \chi_x S_{-\xi} \widehat{f} = (x, -\xi, z \cdot \langle -\xi, x \rangle) \cdot \widehat{f}$$

Proof This follows from (1.14)–(1.16). □

We will henceforth assume that H, G and K form a short exact sequence as in (1.12), with H discrete and K compact.

Definition 1.14 (Weil–Brezin map) The *Weil–Brezin map* $\mathcal{W}_G^H$ is defined for $f \in \mathbb{C}G$ and $(\xi, x, s) \in \mathcal{H}_G$ as

$$\mathcal{W}_G^H f(\xi, x, z) = \sum_H ((\xi, x, z) \cdot f)_H .$$

A direct computation shows that the Weil–Brezin map satisfies the following symmetries for all $(h', h, 1) \in H^\perp \times H \times 1 \subset \mathcal{H}_G$ and all $z \in \mathbb{T}$:

$$\mathcal{W}_G^H f\left((h', h, 1) \cdot (\xi, x, s)\right) = \mathcal{W}_G^H f(\xi, x, s) \tag{1.18}$$

$$\mathcal{W}_G^H f(\xi, x, z) = z \cdot \mathcal{W}_G^H f(\xi, x, 1). \tag{1.19}$$

Lemma 1.15 $\Gamma = H^\perp \times H \times 1$ *is a subgroup of $\mathcal{H}_G$. It is not a normal subgroup, so we cannot form the quotient group. However, as a manifold the set of right cosets is*

$$\Gamma \backslash \mathcal{H}_G = \widehat{H} \times K \times \mathbb{T}.$$

The Heisenberg group has a right and left invariant volume measure given by the direct product of the invariant measures of $\widehat{G}$, G and $\mathbb{T}$. Thus we can define the Hilbert spaces $L^2(\mathcal{H}_G^H)$ and $L^2(\widehat{H} \times K \times \mathbb{T})$. By Fourier decomposition in the last variable (z-transform), $L^2(\widehat{H} \times K \times \mathbb{T})$ splits into an orthogonal sum of subspaces $\mathcal{V}_k$ for $k \in \mathbb{Z}$, consisting of those $g \in L^2(\widehat{H} \times K \times \mathbb{T})$ such that

$$g(\xi, x, z) = z^k g(\xi, x, 1) \quad \text{for all } z = e^{2\pi i\theta}.$$

It can be verified that $\mathcal{W}_G^H$ is unitary with respect to the L^2 inner product. Together with (1.18)–(1.19) this implies:

Lemma 1.16 *The Weil–Brezin map is a unitary transform*

$$\mathcal{W}_G^H : L^2(G) \to \mathcal{V}_1 \subset L^2(\widehat{H} \times K \times \mathbb{T}).$$

Note that the Weil–Bezin map on $\widehat{G}$, with respect to the reciprocal lattice $H^\perp$, is

$$\mathcal{W}_{\widehat{G}}^{H^\perp} : L^2(G) \to \mathcal{V}_1 \subset L^2(K \times \widehat{H} \times \mathbb{T}).$$

The Poisson summation formula (Lemma 1.11) together with Lemma 1.13 implies that these two maps are related via

$$\mathcal{W}_G^H f(\xi, x, z) = \mathcal{W}_{\widehat{G}}^{H^\perp} \widehat{f}(x, -\xi, z \cdot \langle \xi, x \rangle).$$

Defining the unitary map $J \colon L^2 \subset L^2(\widehat{H} \times K \rtimes \mathbb{T}) \to L^2(K \times \widehat{H} \times \mathbb{T})$ as

$$Jf(x, -\xi, z \cdot \langle \xi, x \rangle) = f(\xi, x, z), \tag{1.20}$$

we obtain the following fundamental theorem.

Theorem 1.17 (Weil–Brezin factorization) *Given an LCA G and a lattice $H < G$. The Fourier transform on G factorizes in a product of three unitary maps*

$$\mathcal{F}_G = \left(\mathcal{W}_{\widehat{G}}^{H^\perp}\right)^{-1} \circ J \circ \mathcal{W}_G^H. \tag{1.21}$$

The Zak transform. Given a lattice $H < G$ and transversals $\sigma \colon K \to G$ and $\widehat{\sigma} \colon \widehat{H} \to \widehat{G}$. The *Zak transform* is defined as

$$\mathcal{Z}_G^H f(\xi, x) := \mathcal{W}_G^H f(\xi, x, 1) \quad \text{for } \xi \in \widehat{\sigma}(\widehat{H}),\ x \in \sigma(K). \tag{1.22}$$

The Zak transform can be computed as a collection of Fourier transforms on H of f shifted by x, for all $x \in \sigma(K)$. The definition of the Fourier transform yields:

$$\mathcal{Z}_G^H f(-\xi, x) = \mathcal{F}_H\left((S_x f)_H\right)(\widehat{\phi}_0(\xi)). \tag{1.23}$$

We see that the Zak transform is invertible when $\mathcal{Z}_G^H f(-\xi, x)$ is computed for all $\xi \in \widehat{\sigma}(\widehat{H})$ and all $x \in \sigma(K)$. Written in terms of the Zak transform, the Weil–Brezin factorization (1.21) becomes

$$\mathcal{Z}_{\widehat{G}}^{H^\perp} \widehat{f}(x, \xi) = \langle \xi, x \rangle \mathcal{Z}_G^H f(-\xi, x). \tag{1.24}$$

The factor $\langle \xi, x \rangle$ is called a *twiddle factor* in the computational FFT literature.

In the special case where $G = H \times K$, then also $K < G$ and $\widehat{H} = K^\perp$. Thus in this case it is possible to choose σ and $\widehat{\sigma}$ as group homomorphisms, resulting in $\langle \xi, x \rangle \equiv 1$ for all $\xi \in \widehat{\sigma}(\widehat{H})$ and $x \in \sigma(K)$. This choice is called a *twiddle free* factorization. However, by other choices of the transversals, the twiddle factors also enter into the formula in this case.

Due to the symmetries (1.18)–(1.19), the Weil–Brezin map is trivially recovered from the Zak transform. The Zak transform is the practical way of computing the Weil–Brezin map and its inverse. However, since the invertible Zak transform cannot be defined canonically, independently of the transversals σ and $\widehat{\sigma}$, the Weil–Brezin formulation is more fundamental.

The Fast Fourier Transform. Cooley–Tukey style FFT algorithms are based on recursive use of (1.24), where a Fourier transform on G is computed by a collection of Fourier transforms on H composed with inverse Fourier transforms on $H^\perp$. We choose transversals σ and $\widehat{\sigma}$. If G is finite, then $\widehat{\sigma}(\widehat{H})$ and $\sigma(K)$ are finite. The Cooley–Tukey factorization follows from (1.24):

- For each $x \in \sigma(K)$ compute:

$$\mathcal{Z}_G^H f(-\xi, x) = \mathcal{F}_H(S_x f)(\widehat{\phi}_0(\xi)) \quad \text{for all } \xi \in \widehat{\sigma}(\widehat{H}) \ .$$

- For each $\xi \in \widehat{\sigma}(\widehat{H})$ compute:

$$\widehat{f}(\xi + \widehat{\phi}_1(\kappa)) = \mathcal{F}_{H^\perp}^{-1} \left(\langle \xi, \sigma(\cdot) \rangle \mathcal{Z}_G^H f(-\xi, \sigma(\cdot)) \right)(\kappa) \quad \text{for all } \kappa \in H^\perp,$$

  where the inverse Fourier transform $\mathcal{F}_{H^\perp}^{-1}$ is with respect to the variable $\cdot \in K$.

The Fast Fourier Transform is obtained by recursive application of this splitting. This general formulation allows for Cooley–Tukey kind FFTs based on any decomposition of G with respect to a lattice H. In particular this is useful for functions with symmetries, in which case it is important to choose lattices H that preserve the symmetries in order to take advantage of all the symmetries in the FFT. We return to this issue in Section 1.3.4.

Shannon's sampling theorem. By setting $x = 0$ in (1.24), we obtain the important dual relationship between downsampling and periodization:

$$\mathcal{F}_H(f_H)(\widehat{\phi}_0(\xi)) = \sum_{k \in \widehat{\phi}_1(\widehat{K})} \mathcal{F}_G(f)(k+\xi). \tag{1.25}$$

A function $f \in \mathbb{C}G$ is *band limited* with respect to the reciprocal lattice $H^\perp$ if its Fourier transform is zero outside the Voronoi polyhedron, i.e. if $\text{supp}(\widehat{f}) \subset \widehat{\sigma}(\widehat{H})$, where $\widehat{\sigma}$ is a Voronoi transversal of $\widehat{\phi}_0$. If f is band limited, the terms on the right hand side of (1.25) are zero for $k \neq 0$ and $\xi \in \widehat{\sigma}(\widehat{H})$. This yields

$$\mathcal{F}_G(f)(\xi) = \mathcal{F}_H(f_H)(\widehat{\phi}_0(\xi)),$$

thus we obtain Shannon's celebrated result that a band limited f can be exactly recovered from its downsampling f_H.

Lattice rules. For general functions $f \in \mathbb{C}G$, the error between the Fourier transform of the true and the sampled function is given as

$$\mathcal{F}_H(f_H)(\widehat{\phi}_0(\xi)) - \mathcal{F}_G(f)(\xi) = \sum_{k \in \widehat{\phi}_1(\widehat{K}) \setminus \{0\}} \mathcal{F}_G(f)(k+\xi).$$

The game of Lattice rules is, given f with specific properties, to find a lattice $H < G$ such that the error is minimised. We now assume that the original domain is periodic $G = T^n$. Lattice rules are designed such that the nonzero points in $H^\perp$ neighbouring 0 are pushed as far out as possible with respect to a given norm, depending on f. If f is spherically symmetric, H should be chosen as a *densest lattice packing* (with respect to the 2-norm) [10], e.g. hexagonal lattice in $\mathbb{R}^2$ and face centred cubic packing in $\mathbb{R}^3$ (as the orange farmers know well). In dimensions up to 8, these are given by certain root lattices [30]. The savings, compared to standard tensor product lattices, are given by the factors 1.15, 1.4, 2.0, 2.8 4.6, 8.0 and 16.0 in dimensions $n = 2, 3, \ldots, 8$. This is important, but not dramatic, e.g. a camera with 8.7 megapixels arranged in a hexagonal lattice has approximately the same sampling error as a 10 megapixel camera with a standard square pixel distribution. However, these alternative lattices have other attractive features, such as larger spatial symmetry groups, yielding more isotropic discretizations.

A hexagonal lattice picture can be rotated more uniformly than a square lattice picture.

Dramatic savings can be obtained for functions belonging to the *Korobov spaces*. This is a common assumption in much work on high dimensional approximation theory. Korobov functions are functions whose Fourier transforms have energy concentrated along the axis directions in $\widehat{G}$, the so-called hyperbolic cross mass distribution. Whereas the tensor product lattice with $2d$ points in each direction contains $(2d)^n$ lattice points in T^n, the optimal lattice with respect to the Korobov norm contains only $\mathcal{O}(2^n d(\log(d))^{n-1})$ points, removing exponential dependence on d.

The group theoretical understanding of lattice rules makes software implementation very clean and straightforward. In [29], numerical experiments are reported on lattice rules for FFT-based spectral methods for PDEs. Note that whereas the choice of transversal $\widehat{\sigma}\colon \widehat{H} \to \widehat{G}$ is irrelevant for lattice integration rules, it is essential for pseudospectral derivation. The Laplacian $\nabla^2 f$ is computed on $\widehat{G}$ as $\widehat{f}(\xi) \mapsto c|\xi|^2\widehat{f}(\xi)$, whereas the corresponding computation on $\widehat{H}$ must be done as $\mathcal{F}_H(f_H)(\eta) \mapsto c|\widehat{\sigma}(\eta)|^2\mathcal{F}_H(f_H)(\eta)$ for $\eta \in \widehat{H}$, and we must choose the Voronoi transversal to minimise aliasing errors.

Polyhedral Dirichlet kernels. The theoretical understanding of lattice sampling rules depends on the analytical properties of *polyhedral Dirichlet kernels*. Let $\widehat{\sigma}\colon \widehat{H} \to \widehat{G}$ be the Voronoi transversal. The perfect low-pass filter $\widehat{\mathcal{D}}_H \in \mathbb{C}\widehat{G}$ is defined as

$$\widehat{\mathcal{D}}_H(\xi) = \begin{cases} 1 & \text{if } \xi \in \widehat{\sigma}(\widehat{H}) \\ 0 & \text{otherwise} \end{cases} .$$

The polyhedral Dirichlet kernel $\mathcal{D}_H \in \mathbb{C}G$ is defined as

$$\mathcal{D}_H = \mathcal{F}_G^{-1}(\widehat{\mathcal{D}}_H).$$

This function plays the same role as the classical Dirichlet kernel in the 1–dimensional sampling theory, e.g. low-pass reconstruction of a down sampled function is done by convolution with $\mathcal{D}_H$. Detailed analysis of these functions is done in [34, 36]. In particular it is important that they in general have Lebesque constants scaling like $\mathcal{O}(\log^n(N))$, where n is the dimension of G, and N measures the number of sampling points in H.

1.2.5 Fourier analysis on non-commutative groups

In this section we will briefly discuss computational aspects of Fourier techniques on non-commutative groups. We will be much less detailed than in the previous section, since this material is covered in detail elsewhere, e.g. [3, 2, 30].

A starting point of the LCA discussion was the definition of the group ring $\mathbb{C}G$ and the existence of a translation invariant measure, which led to convolutions in the group ring. For non-commutative groups the situation is a bit more complicated, since invariance with respect to left and right translations might not yield the same measure. Groups for which there exist a (unique up to scaling) measure which is both left and right invariant is called *unimodular*. For such groups a lot of the previous theory carries over, with some modifications. Groups which are not unimodular are considerably more complicated and will not be discussed here.

Important examples of unimodular groups are:

- Abelian groups.
- Finite groups.
- Compact groups.
- Semidirect product of compact and abelian groups, e.g. the Euclidean motion group consisting of translations and rotations.
- Semisimple and nilpotent Lie groups.

Finite groups. Let $\mathbb{C}G$ denote the group ring, the complex vector space of dimension $|G|$, where each element in G is a basis vector, so, as before, $f \in \mathbb{C}G$ is given as $f = \sum_{g\in G} f(g)g$, where $f(g) \in \mathbb{C}$. The right and left invariant Haar measure is given as the sum over the elements

$$\int_G f d\mu = \sum_{x\in G} f(x).$$

The product in G yields a convolution product in $\mathbb{C}G$

$$(f * g)(y) = \sum_{x\in G} f(x)g(x^{-1}y) = \sum_{x\in G} f(yx)g(x^{-1}).$$

This is, however, not a commutative product on $\mathbb{C}G$, $f * g \neq g * f$. In the abelian case, the Fourier transform diagonalizes the convolution because the exponential basis consists of group homomorphisms (into $\mathbb{T}$). In the non-commutative case, it cannot be possible to diagonalize the convolution using just $\hom(G, \mathbb{T})$ because the convolution is not commutative. The idea of Schur and Frobenius in the late 19th century was to look for a basis for $\mathbb{C}G$ in terms of *group representations*, defined as elements of $\hom(G, U(n))$, where $U(n)$ is the set of unitary $n \times n$ matrices. (Note that $U(1) = \mathbb{T}$.) An n-dimensional representation is thus a function $R\colon G \to U(n)$ satisfying $R(xy) = R(x)R(y)$ and $R(x^{-1}) = R(x)^{-1} = R(x)^{\dagger}$, where $R(x)^{\dagger}$ is the complex conjugate transpose. Let $d_R = n$ denote the dimension of the representation. For each representation R, we may define Fourier coefficients of a function $f \in \mathbb{C}G$ as a complex $d_R \times d_R$ matrix defined as

$$\widehat{f}(R) = \sum_{x \in G} f(x) R(x)^{\dagger}.$$

A computation using the homomorphism property and shift invariance of the sum, shows that the representations may be used to (block-)diagonalize the convolution:

$$\widehat{f * g}(R) = \widehat{g}(R)\widehat{f}(R). \tag{1.26}$$

However, we need a basis for $\mathbb{C}G$, and we need an inversion formula for this generalised Fourier transform.

The concepts of equivalent and reducible representations are crucial for constructing a suitable basis. Two representations R and $\tilde{R}$ are equivalent if there exists an invertible matrix V such that $R(x) = V\tilde{R}(x)V^{-1}$ for all x. A representation is reducible if it is equivalent to a representation which is block-diagonal. In that case the representation can be seen as a direct sum of smaller representations (one for each diagonal block). Frobenius found that there always exists a complete list of non-equivalent irreducible representations which forms an orthogonal basis for $\mathbb{C}G$.

Theorem 1.18 (Frobenius) *For a finite group G there exists a complete list of non-equivalent irreducible representations $\mathcal{R} = \{R_1, \ldots, R_k\}$ such that $\sum_{R \in \mathcal{R}} d_R^2 = |G|$. Define the generalized Fourier transform of $f \in \mathbb{C}G$ as*

$$\widehat{f}(R) = \sum_{x \in G} f(x) R(x)^{\dagger} \quad \text{for all } R \in \mathcal{R}.$$

Then f is reconstructed by the formula

$$f(x) = \frac{1}{|G|} \sum_{R \in \mathcal{R}} d_R \operatorname{trace}(\widehat{f}(R)R(x)).$$

As an example, we consider the computation of the convolution of $f, g \in \mathbb{C}G$ when $G < O(3)$ is the icosahedral group, the collection of the 120 orthogonal matrices which leave the icosahedron in $\mathbb{R}^3$ invariant. This group has a complete list of irreducible representations of dimensions $\{1, 1, 3, 3, 3, 3, 4, 4, 5, 5\}$. A direct computation of the convolution involves 120^2 multiplications. Instead, computing the convolution in Fourier space involves the multiplication of matrices of size $1, 1, 3, 3, \ldots$, which requires only 120 multiplications, saving a factor 120. For the computation of matrix exponentials and eigenvalues, the savings are more dramatic; a direct computation costs 120^3 operations, while the equivalent computation in Fourier space costs $2 + 4 \cdot 3^3 + 2 \cdot 4^3 + 2 \cdot 5^3$ operations, which is cheaper by a factor of about 3500.

A source of computational problems leading to group convolutions is linear problems with spatial symmetries. Given a linear operator $\mathcal{L}$ which commutes with a finite group of isometries acting upon the domain (e.g. the Laplacian on the sphere commutes with any group of isometries, e.g. the discrete icosahedral group). Let $\mathcal{L}$ be discretized in a symmetry preserving manner, such that the discrete L commutes with the isometries in G. Then L can be described as a block-convolution, i.e. L belongs to a group ring of the form $\mathbb{C}^{m \times m}G$. The blocks represent the interaction between different orbits of the action of G on the domain. As an example, if the space of spherical functions is discretized with 12,000 degrees of freedom, the full space splits into about $12,000/120 = 100$ different orbits under the action of the icosahedral group. The Laplacian can then be represented as a block-convolution in $\mathbb{C}^{100 \times 100}G$. Under the generalized Fourier transform, the matrix becomes block-diagonal, with blocks of sizes $100 d_R$ for $d_R \in \{1, 1, 3, 3, \ldots, 5\}$.

The use of the generalized Fourier transform is important in various computational tasks, such as eigenvalue problems, solutions of linear equations and computations of matrix exponentials. Experience shows that symmetry preserving discretizations and algorithms are not only much faster than direct algorithms, but they are also often more accurate, since preservation of symmetry tends to diminish the effect of numerical round-off errors.

Compact groups. Compact groups, such as, for example, the group of rotations SO(3), are in many respects quite similar to finite groups. There exists a bi-invariant Haar measure and a space of functions $L^2(G)$ with a convolution product. The Peter–Weyl theorem guarantees that there exist a discrete, infinite family of irreducible unitary representations forming an orthogonal basis for $L^2(G)$. The Fourier transform becomes an integral over G and the inversion formula is similar to the finite case, although the sum here is over an infinite list $\mathcal{R}$.

The representation theory of orthogonal groups has numerous applications in computational science and technology, an example being recent work on Cryo-Electron microscopy [18, 19]. The basic problem here is the reconstruction of a 3D molecular structure from a large collection of 2D projections of the molecule seen from unknown angles. Representation theory provides an important tool to analyze numerical algorithms for this problem.

Euclidean motions. The Euclidean motion group of translations and rotations on $\mathbb{R}^3$ is important in many technological applications. The group $E(3) = \mathrm{SO}(3) \rtimes \mathbb{R}^3$ is the semidirect product of a compact group and an abelian group. Such groups are always unimodular, and the representation theory is relatively simple. The irreducible representations on $E(3)$ are induced from the representations of the compact part, $SO(3)$, by a standard method called the method of small groups [33]. An example of an application of Fourier analysis on $E(3)$ is the problem of medical image registration; *Find the Euclidean motion which best matches two different 3-D images of an object.* This can be phrased as the question of computing the maximum of the cross correlation (or phase correlation) of the two images. The cross correlation is very similar to a convolution and can be cheaply computed in Fourier space.

1.3 Multivariate Chebyshev polynomials in computations

Univariate (classical) Chebyshev polynomials are ubiquitous in numerical analysis and computational science, due to their in many ways optimal approximation properties and the tight relationship between Chebyshev approximations and fast cosine transforms. First and second kind Chebyshev polynomials $\{T_k\}_{k=0}^{\infty}$ and $\{U_k\}_{k=0}^{\infty}$ are defined as

$$\begin{aligned} x &= \cos(\theta) \\ T_k(x) &= \cos(k\theta) \\ U_k(x) &= \sin((k+1)\theta)/\sin(\theta). \end{aligned}$$

In this section we will discuss the connection between Chebyshev approximations and group theory. Once the group theoretical view is established, it will become clear that there exist certain interesting multivariate generalizations of Chebyshev polynomials. These share most of the favorable properties of the univariate polynomials, and they are orthogonal on domains related to triangles and simplices. Bivariate Chebyshev polynomials were constructed independently by Koornwinder [25] and Lidl [26] by folding exponential functions. Multidimensional generalizations (the A_2 family) appeared first in [13]. In [21] a general folding construction was presented. Characterization of such polynomials as eigenfunctions of differential operators is found in [6, 25].

Our interest in these polynomials originates from their potential applications in computational approximation methods, in particular spectral element methods and multidimensional quadrature. We have developed the theory of their discrete orthogonality, triangular based Clenshaw–Curtis type quadrature formulae, recursion formulae for computing spectral derivations as well as software for the application of multivariate Chebyshev polynomials in approximation, quadrature and PDE solution. Our exposition here will aim at giving an overview of the main ideas. More detailed presentations are found in [30, 32, 9].

Spectral element methods are computational techniques for solving PDEs where the domain of the equation is divided into a fixed collection of regularly shaped subdomains. On each subdomain a high order polynomial space is constructed, and a global solution is obtained by patching together local solutions, either in a strong sense by imposing continuity conditions across subdomain boundaries, or in a weak sense by variational formulations (discontinuous Galerkin methods). The advantage of spectral element methods compared to its competitors (finite elements, finite differences and finite volume methods) is the phenomenon called *spectral convergence*. When an analytic function is approximated by an Nth order polynomial, one may achieve errors decaying as e^{-N}. Thus spectral methods are particularly attractive when high accuracy is important.

The drawbacks of spectral element methods are that high order polynomial approximations must be constructed with care. Jim Wilkinson famously demonstrated that high order polynomial interpolation in equispaced points is a highly unstable process, due to the fact that the Lebesgue constant of equi-spaced interpolation points grows exponentially in N. Interpolation in Chebyshev zeros, or Chebyshev extremal points, is on the other hand near optimally stable, as the Lebesque constant in such points grows as $\mathcal{O}(\log(N))$. Another problem with spectral element methods is inflexibility with respect to sub-domain divisions. High order polynomial approximations are easy to construct on rectangular subdomains (by tensor products of univariate polynomials), but high order approximation theory based on triangular and tetrahedral subdivisions is far less developed. The most common practice is therefore rectangular subdivision methods. Triangular and tetrahedral subdivision schemes are far more flexible, if they can be implemented in an efficient and stable manner. A singular mapping technique from squares to triangles [12] is a possible solution, but has drawbacks in breaking of triangular symmetries as well as other problems. Nodal spectral Galerkin methods are another approach where good collocation nodes are computed by numerical optimization (e.g. Fekete points). But in this approach one has no direct connection to Fourier analysis, and fast transforms are not available [17, 20]. This makes spectral element methods based on multivariate Chebyshev polynomials an attractive alternative. We start with a discussion of particular eigenfunctions of the Laplacian on simplices.

1.3.1 What is the sound of a triangular drum?

Bases for high order approximation spaces are usually obtained as eigenfunctions of Sturm–Liouville problems, truncated to a given order. Can we find Sturm–Liouville problems on triangles and simplices, that yield good approximation spaces? It is known that the eigenfunctions of the Laplacian (with Dirichlet or Neumann boundary conditions) can be explicitly constructed on certain triangular domains in 2D and some particular simplices in all higher dimensions. An illustrative example is the construction of Laplacian eigenfunctions on an equilateral triangle, with Dirichlet or Neumann boundary conditions.

The equilateral triangle has the particular property that if we set up a kaleidoscope with three mirrors at the three edges, then the reflections of the triangle tile the plane in a periodic pattern, shown as the shaded domain in the right part of Figure 1.2 (labelled A_2). Without loss of generality, we assume that the triangle has corners in the origin $[0;0]$, $\lambda_1 = [1/\sqrt{2}; 1/\sqrt{6}]$ and $\lambda_2 = [0; \sqrt{2}/\sqrt{3}]$. Let $\{s_j\}_{j=1}^3$ denote reflections of $\mathbb{R}^2$ about the edges of the triangle. Let $\widetilde{W}$ denote the full group of isometries of $\mathbb{R}^2$ generated by $\{s_j\}_{j=1}^3$. This is an example of a *crystallographic group*[3], a group of isometries of $\mathbb{R}^d$ where the subgroup

[3] More specifically it is an *affine Weyl group*, to be defined below.

of translations form a lattice in $\mathbb{R}^d$. From this fact we will derive the Laplacian eigenfunctions on the triangle, which we denote by $\triangle$.

The translation lattice of $\widetilde{W}$ is $L = \mathrm{span}_{\mathbb{Z}}\{\alpha_1, \alpha_2\} < \mathbb{R}^2$ generated by the vectors $\alpha_1 = (\sqrt{2}, 0)$ and $\alpha_2 = (-1/\sqrt{2}, \sqrt{3}/\sqrt{2})$. The unit cell of L can be taken as either the rhombus spanned by α_1 and α_2 or the hexagon ⬡ indicated in Figure 1.2. The hexagon is the Voronoi cell of the origin in the lattice L, its interior consists of the points in $\mathbb{R}^2$ that are closer to the origin than to any other lattice points. As a first step in our construction of triangular eigenfunctions, we consider the L-periodic eigenfunctions of the Laplacian ∇^2. Since $(\lambda_j, \alpha_k) = \delta_{j,k}$, the reciprocal lattice is $L^\perp = \mathrm{span}_{\mathbb{Z}}\{\lambda_1, \lambda_2\}$ and the periodic eigenfunctions are

$$\nabla_t^2 \langle \lambda, t\rangle = -(2\pi)^2 ||\lambda||^2 \langle \lambda, t\rangle \quad \text{for all } \lambda \in L^\perp,\ t \in \mathbb{R}^2/L,$$

where $\langle \lambda, t\rangle = e^{2\pi i(\lambda,t)}$ is the dual pairing on $\mathbb{R}^2$. We continue to find the Laplacian eigenfunctions on $\triangle$ by folding the exponentials. Let $W < \widetilde{W}$ be the subgroup which leaves the origin fixed (the symmetries of ⬡):

$$W = \{e, s_1, s_2, s_1s_2, s_2s_1, s_1s_2s_1\},$$

where e is the identity and s_i, $i \in \{1,2\}$ act on $v \in \mathbb{R}^2$ as $s_i v = v - 2(\alpha_i, v)/(\alpha_i, \alpha_i)$. We define even and odd (cosine and sine-type) foldings of the exponentials

$$c_\lambda(t) = \frac{1}{|W|}\sum_{w\in W} \langle \lambda, wt\rangle = \frac{1}{|W|}\sum_{w\in W} \left\langle w^T\lambda, t\right\rangle \tag{1.27}$$

$$s_\lambda(t) = \frac{1}{|W|}\sum_{w\in W} \det(w) \langle \lambda, wt\rangle = \frac{1}{|W|}\sum_{w\in W} \det(w) \left\langle w^T\lambda, t\right\rangle, \tag{1.28}$$

where, in our example, $|W| = 6$. The Laplacian commutes with any isometry, in particular $\nabla^2(f\circ w) = (\nabla^2 f)\circ w$ for $w \in W$. Hence, the reflected exponentials and the functions $c_\lambda(t)$ and $s_\lambda(t)$ are eigenfunctions with the same eigenvalue:

Lemma 1.19 *The eigenfunctions of ∇^2 on the equilateral triangle $\triangle$ with Dirichlet ($f = 0$) and Neumann boundary conditions ($\nabla f \cdot \vec{n} = 0$) are given respectively as*

$$\nabla_t^2 c_\lambda(t) = -(2\pi)^2 ||\lambda||^2 c_\lambda(t)$$
$$\nabla_t^2 s_\lambda(t) = -(2\pi)^2 ||\lambda||^2 s_\lambda(t),$$

for $\lambda \in \mathrm{span}_{\mathbb{N}}\{\lambda_1, \lambda_2\}$, the set of all non-negative integer combinations of $\lambda_1 = [1/\sqrt{2}; 1/\sqrt{6}]$ and $\lambda_2 = [0; \sqrt{2}/\sqrt{3}]$.

The set $\mathrm{span}_{\mathbb{N}}\{\lambda_1, \lambda_2\} \subset L^\perp$ contains exactly one point from each W-orbit, and is called the positive *Weyl chamber*. An important question is how good (or bad!) these eigenfunctions are as bases for approximating analytic functions on $\triangle$. It is well known from the univariate case that the similar construction, yielding the eigenfunctions $\cos(k\theta)$ and $\sin(k\theta)$ on $[0, \pi]$ are *not* giving a basis converging at a super-algebraic speed. For an analytic function $f(\theta)$ where $f(0) = f(\pi) = 0$, the even 2π-periodic extension is piecewise smooth, with only C^0 continuity at $\theta \in \{0, \pi\}$. Hence the Fourier-cosine series converges only as $\mathcal{O}(k^{-2})$. There are several different ways to achieve the desired $\mathcal{O}(\exp(-ck))$ spectral convergence rate for analytic functions. One possible solution is to approximate $f(\theta)$ in a *frame* (not linearly independent) consisting of both $\{\cos(k\theta)\}_{k\in\mathbb{Z}}$ and $\{\sin(k\theta)\}_{k\in\mathbb{Z}^+}$ [22]. Another possibility is to employ a change of variable $x = \cos(\theta)$, yielding (univariate) Chebyshev polynomials with spectral convergence. Here we will discuss a generalization of the latter approach leading to the multivariate Chebyshev polynomials.

1.3.2 Through the kaleidoscope

Recall that the main trick for finding eigenfunctions of the Laplacian on the triangle was in using the fact that the domain is a polyhedron with the property that the group generated by the boundary reflections is a crystallographic group, defined as a group of isometries on $\mathbb{R}^n$ such that the subgroup of translations form a lattice $L < \mathbb{R}^n$. Crystallographic groups is a classical topic of mathematics, physics and chemistry. A classification of such groups starts with the fundamental *crystallographic restriction*: For any crystallographic group, the only allowed rotations are 2-fold, 3-fold, 4-fold and 6- fold[4]. A general group generated just by reflections, as in a kaleidoscope, is called a *Coxeter* group, and the Coxeter groups which comply with the crystallographic restriction are called affine Weyl groups. These are classified in terms of their root systems, where the roots are orthogonal to the mirrors passing through the origin.

Root systems. Let $\mathcal{V}$ be a finite dimensional real Euclidean vector space with standard inner product $(\cdot, \cdot)$. The construction above can

[4] The 5-fold symmetry is not allowed in proper crystals, but is seen in quasicrystals, related to Penrose tilings. This discovery was the topic of the 2011 Nobel prize in chemistry.

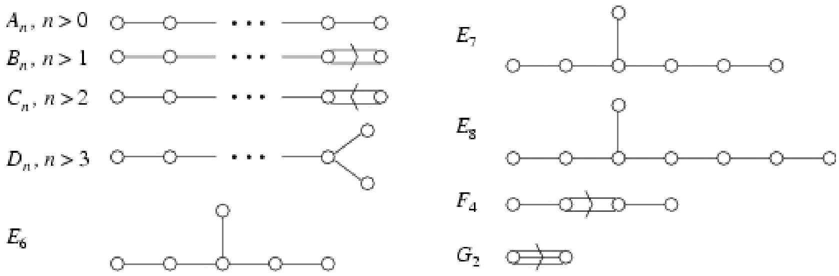

Figure 1.1 Dynkin diagrams for irreducible root systems

be generalized to all those simplices $\triangle \subset \mathcal{V}$ with the property that the group of isometries $\widetilde{W}$ generated by reflecting $\triangle$ about its faces is a crystallographic group. All such simplices are determined by a *root system*, a set of vectors in $\mathcal{V}$ which are perpendicular to the reflection planes of W passing through the origin. We review some basic definitions and results about root systems. For more details we refer to [8].

Definition 1.20 A **root system** in $\mathcal{V}$ is a finite set Φ of non-zero vectors, called roots, that satisfy:

1. The roots span $\mathcal{V}$.
2. The only scalar multiples of a root $\alpha \in \Phi$ that belong to Φ are α and $-\alpha$.
3. For every root $\alpha \in \Phi$, the set Φ is invariant under reflection through the hyperplane perpendicular to α. I.e. for any two roots α and β, the set Φ contains the reflection of β,
$$s_\alpha(\beta) := \beta - 2\frac{(\alpha, \beta)}{(\alpha, \alpha)}\alpha \in \Phi.$$
4. (Crystallographic restriction): For any $\alpha, \beta \in \Phi$ we have
$$2\frac{(\alpha, \beta)}{(\alpha, \alpha)} \in \mathbb{Z}.$$

Condition 4 implies that the obtuse angle between two different reflection planes must be either 90°, 120°, 135° or 150°. The *rank* of the root system is the dimension d of the space $\mathcal{V}$. Any root system contains a subset (not uniquely defined) $\Sigma \subset \Phi$ of so-called *simple positive roots*. This is a set of d linearly independent roots $\Sigma = \{\alpha_1, \ldots, \alpha_d\}$ such that any root $\beta \in \Phi$ can be written either as a linear combination of α_j with non-negative integer coefficients, or as a linear combination with non-positive integer coefficients. We call Σ a *basis* of the root system Φ. A root system is conveniently represented by its *Dynkin diagram*. This is a graph with d nodes corresponding to the simple positive roots. Between two nodes j and k no line is drawn if the angle between α_j and α_k is $90°$, a single line if it is $120°$, a double line for $135°$ and a triple line for $150°$. It is only necessary to understand the geometry of *irreducible root systems*, where the Dynkin diagram is connected. Disconnected Dynkin diagrams (reducible root systems) are trivially understood in terms of products of irreducible root systems. For irreducible root systems, the roots are either all of the same length or have just two different lengths. In the latter case a marker $<$ or $>$ on an edge indicates the separation of long and short roots (short $<$ long). Since the work of W. Killing and E. Cartan in the late 19th century it has been known that Dynkin diagrams of irreducible root systems must belong to one of four possible infinite cases A_n $(n > 0)$, B_n $(n > 1)$, C_n $(n > 2)$, D_n $(n > 3)$ or five special cases E_6, E_7, E_8, F_4, G_2 shown in Figure 1.1. We say that two root systems are equivalent if they differ only by a scaling or an isometry. Up to equivalence there corresponds a unique root system to each Dynkin diagram.

A root system Φ is associated with a dual root system $\Phi^\vee$ defined such that a root $\alpha \in \Phi$ corresponds to a co-root $\alpha^\vee := 2\alpha/(\alpha,\alpha) \in \Phi^\vee$. If all roots have equal lengths then $\Phi^\vee = \Phi$ (upto equivalence), i.e. the root system is *self dual*. For the cases with two root lengths we have $B_2^\vee = B_2$, $B_n^\vee = C_n$ $(n > 2)$, $F_4^\vee = F_4$ and $G_2^\vee = G_2$.

Weyl groups and affine Weyl groups. Given a d-dimensional root system Φ with dual root system $\Phi^\vee$, for the roots $\alpha \in \Phi$, consider the reflection $s_\alpha \colon \mathcal{V} \to \mathcal{V}$ given by

$$s_\alpha(t) = t - \frac{2(t,\alpha)}{(\alpha,\alpha)}\alpha = t - (t,\alpha^\vee)\alpha.$$

For the dual roots $\alpha^\vee \in \Phi^\vee$, we define translations $\tau_{\alpha^\vee} : \mathcal{V} \to \mathcal{V}$ as

$$\tau_{\alpha^\vee}(t) = t + \alpha^\vee.$$

The *Weyl group* of Φ is the finite group of isometries on $\mathcal{V}$ generated by the reflections s_α for $\alpha \in \Sigma$:

$$W = \langle \{s_\alpha\}_{\alpha\in\Sigma} \rangle .$$

An important example of Weyl groups is the infinite family A_n. The

Weyl group W of A_n is isomorphic to the symmetric group S_{n+1}. This is a group of order $|S_{n+1}| = (n+1)!$, consisting of all permutations of $n+1$ objects, and is also isomorphic to the symmetry group of the regular n-simplex. In particular the Weyl group of A_2 has order 6 and can be identified with the symmetries of the regular triangle.

The *dual root lattice* $L^\vee$ is the lattice spanned by the translations $\tau_{\alpha^\vee}$ for $\alpha^\vee \in \Sigma^\vee$. We identify this with the abelian group of translations on $\mathcal{V}$ generated by the dual roots

$$L^\vee = \langle \{\tau_{\alpha^\vee}\}_{\alpha^\vee \in \Sigma^\vee} \rangle .$$

The *affine Weyl group* $\widetilde{W}$ is the infinite crystallographic symmetry group of $\mathcal{V}$ generated by the reflections s_α for $\alpha \in \Sigma$ and the translations $\tau_{\alpha^\vee}$ for $\alpha^\vee \in \Sigma^\vee$, thus it is the semidirect product of the Weyl group W with the dual[5] root lattice $L^\vee$

$$\widetilde{W} = \langle \{s_\alpha\}_{\alpha \in \Sigma}, \{\tau_{\alpha^\vee}\}_{\alpha^\vee \in \Sigma^\vee} \rangle = W \rtimes L^\vee .$$

Let $\Lambda = (L^\vee)^\perp$ denote the reciprocal lattice of $L^\vee$. The lattice Λ is spanned by vectors $\{\lambda_j\}_{j=1}^d$ such that $(\lambda_j, \alpha_k^\vee) = \delta_{j,k}$ for all $\alpha_k^\vee \in \Sigma^\vee$. The vectors λ_j are called the *fundamental dominant weights* of the root system Φ, and Λ is called the *weights lattice*.

The *positive Weyl chamber* $\mathcal{C}_+$ is defined as the closed conic subset of $\mathcal{V}$ containing the points with nonnegative coordinates with respect to the dual basis $\{\lambda_1, \ldots, \lambda_d\}$, in other words

$$\mathcal{C}_+ = \{t \in \mathcal{V} \colon (t, \alpha_j) \geq 0\}.$$

This is a fundamental domain for the Weyl group acting on $\mathcal{V}$. The boundary of $\mathcal{C}_+$ consists of the hyperplanes perpendicular to $\{\alpha_1, \ldots, \alpha_d\}$. The affine Weyl group contains reflection symmetries about affine planes perpendicular to the roots, shifted a half integer multiple of the length of a co-root away from the origin, i.e. for each $\alpha^\vee \in \Phi^\vee$ and each $k \in Z$ there is an affine plane consisting of the points $P_{k,\alpha^\vee} = \{t \in \mathcal{V} \colon 2(t, \alpha^\vee) = k(\alpha^\vee, \alpha^\vee)\} = \{t \in \mathcal{V} \colon (t, \alpha) = k\}$, and this affine plane is invariant under the affine reflection $\tau_{k\alpha^\vee} \cdot s_\alpha$. A connected closed subset of $\mathcal{V}$ limited by such affine planes is called an *alcove* and is a fundamental domain for the affine Weyl group $\widetilde{W}$.

The situation is particularly simple for irreducible root systems, where

[5] Since the Weyl groups of Φ and of $\Phi^\vee$ are identical, it is no problem to instead define $\widetilde{W} = W \rtimes L$ as the semidirect product of the Weyl group with the *primal* root lattice. We have, however, chosen to follow the most common definition here, which leads to a slightly simpler notation for the Fourier analysis.

the alcoves are always d-simplices. Recall that all roots $\alpha \in \Phi$ can be written as $\alpha = \sum_{k=1}^{d} n_k \alpha_k$ where all $n_k = 2(\alpha, \lambda_k)/(\alpha_k, \alpha_k)$ are either non-negative or all non-positive integers. A root $\tilde{\alpha}$ strictly dominates another root α, written $\tilde{\alpha} \succ \alpha$, if $\tilde{n}_k \geq n_k$ for all k, with strict inequality for at least one k. Irreducible root systems have a unique *dominant root* $\tilde{\alpha} \in \Phi$ such that $\tilde{\alpha} \succ \alpha$ for all $\alpha \neq \tilde{\alpha}$. The dominant root $\tilde{\alpha}$ is the unique long root in the Weyl chamber (possibly on the boundary). The basic geometric properties of affine Weyl groups are summarized by the following lemma:

Lemma 1.21

1 If Φ is irreducible with dominant root $\tilde{\alpha}$ then a fundamental domain for $\widetilde{W}$ is the simplex $\triangle \subset \mathcal{V}$ given as

$$\triangle = \{t \in \mathcal{V} \colon (t, \tilde{\alpha}) \leq 1 \text{ and } (t, \alpha_j) \geq 0 \text{ for all } \alpha_j \in \Sigma\},$$

where $\triangle$ has corners in the origin and in the points $\lambda_j/(\lambda_j, \tilde{\alpha})$ for $j = 1, \ldots, d$.

2 The affine Weyl group is generated by the affine reflections about the boundary faces of the fundamental domain $\triangle$. For irreducible Φ these are

$$\widetilde{W} = \langle \{s_{\alpha_j}\}_{j=1}^{d}, \tau_{\tilde{\alpha}} \cdot s_{\tilde{\alpha}} \rangle.$$

3 If Φ is reducible then a fundamental domain for the affine Weyl group is given as the Cartesian product of the fundamental domains for each of its irreducible components.

The simplest rank d root system is the reducible system $A_1 \times \cdots \times A_1$, where the Dynkin diagram consists of d non-connected dots. Figure 1.2 shows $A_1 \times A_1$. The solid black square is the fundamental domain of the root lattice, and the small shaded square the fundamental domain of the affine Weyl group $\widetilde{W}$.

The right part of Figure 1.2 and Figure 1.3 shows the irreducible 2-d cases with roots α (large dots), dominant root $\tilde{\alpha}$ (circle), simple positive roots (α_1, α_2), and fundamental dominant weights (λ_1, λ_2). The roots are normalized such that the longest roots have length $\sqrt{2}$, thus for long roots $\alpha^\vee = \alpha$. For short roots we have for B_2 that $\alpha^\vee = 2\alpha$ and for G_2 that $\alpha^\vee = 3\alpha$. The fundamental domain of the dual root lattice (Voronoi region of $L^\vee$) is indicated by ⬡, □ and ⬡, and the fundamental domain for the affine Weyl group (an alcove) is indicated by shaded triangles.

Figure 1.4 shows the A_3 case (self dual), where the fundamental domain (Voronoi region) of the root lattice is a rhombic dodecahedron,

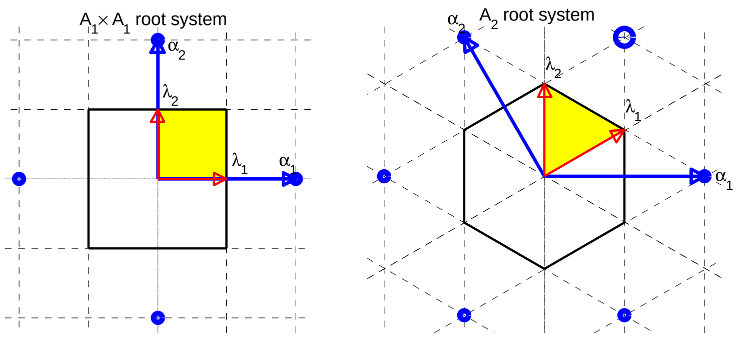

Figure 1.2 Reducible root system $A_1 \times A_1$ and irreducible system A_2

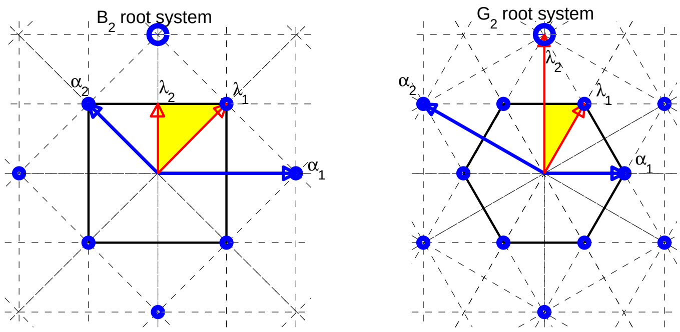

Figure 1.3 The irreducible root systems B_2 and G_2

a convex polyhedron with 12 rhombic faces. Each of these faces (composed of two triangles) is part of a plane perpendicular to one of the 12 roots, halfway out to the root (roots are not drawn). The fundamental domain of the affine Weyl group is the tetrahedron with an inscribed octahedron. The corners of this tetrahedron constitute the origin and the fundamental dominant weights λ_1, λ_2, λ_3. The regular octahedron drawn inside the Weyl chamber is important for applications of multivariate Chebyshev polynomials in domain decompositions [9].

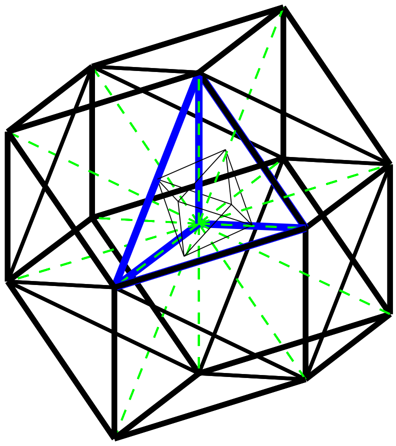

Figure 1.4 Fundamental domains of A_3 root system

1.3.3 Laplacian eigenfunctions on triangles, tetrahedra and simplices.

In this section we will consider real or complex valued functions on $\mathcal{V}$ respecting symmetries of an affine Weyl group. Consider first $L^2(\mathrm{T})$, the space of complex valued L^2-integrable periodic functions on the torus $\mathrm{T} = \mathcal{V}/L^\vee$, i.e. functions f such that $f(y+\alpha^\vee) = f(y)$ for all $y \in \mathcal{V}$ and $\alpha^\vee \in L^\vee$. Since the weights lattice Λ is reciprocal to $L^\vee$, the Fourier transform and its inverse are given for $f \in L^2(\mathrm{T})$ as

$$\widehat{f}(\lambda) = \mathcal{F}(f)(\lambda) = \frac{1}{\mathrm{vol(T)}} \int_{\mathrm{T}} f(t)\langle -\lambda, t\rangle dt, \tag{1.29}$$

$$f(t) = \mathcal{F}^{-1}(\widehat{f})(t) = \sum_{\lambda\in\Lambda} \widehat{f}(\lambda)\langle \lambda, t\rangle. \tag{1.30}$$

We are interested in functions which are both periodic under translations in $L^\vee$ and also respect the other symmetries in $\widetilde{W}$, e.g. functions with odd or even symmetry[6] with respect to the reflections in $\widetilde{W}$. Due to the semi-direct product structure $\widetilde{W} = W \rtimes L^\vee$ it follows that any $\widetilde{W} \in \widetilde{W}$ can be written $\widetilde{W} = w \cdot \tau_{\alpha^\vee}$, where $w \in W$ and $\alpha^\vee \in L^\vee$. Thus on the space of periodic functions $L^2(\mathrm{T})$, the action of $\widetilde{W}$ and the finite group W are identical. We define subspaces of symmetric and skew-symmetric

[6] There are other possible symmetries as well, related to other representations of the Weyl group [9].

periodic functions as follows

$$L^2_\vee(\mathrm{T}) = \{f \in L^2(\mathrm{T})\colon f(wt) = f(t) \text{ for all } w \in W,\, t \in T\}$$
$$L^2_\wedge(\mathrm{T}) = \{f \in L^2(\mathrm{T})\colon f(wt) = (-1)^{|w|} f(t) \text{ for all } w \in W,\, t \in T\},$$

where $|w|$ denotes the *length*, defined as $|w| = \ell$, where $w = s_{\alpha_{j_1}} \cdots s_{\alpha_{j_\ell}}$ is written in the shortest possible way as a product of reflections about the simple positive roots $\alpha_j \in \Sigma$. Thus $(-1)^{|w|} = \det(w) = \pm 1$ depending on whether w is a product of an even or an odd number of reflections. We define L^2 orthogonal projections $\pi^W_\vee$ and $\pi^W_\wedge$ on these subspaces as

$$\pi^W_\vee f(t) = \frac{1}{|W|} \sum_{w \in W} f(wt), \tag{1.31}$$

$$\pi^W_\wedge f(t) = \frac{1}{|W|} \sum_{w \in W} (-1)^{|w|} f(wt). \tag{1.32}$$

Orthogonal bases for these subspaces are obtained by projecting the exponentials, yielding the cosine and sine type basis functions

$$c_\lambda(t) := \pi^W_\vee \exp 2\pi i(\lambda, t) \tag{1.33}$$

$$s_\lambda(t) := \pi^W_\wedge \exp 2\pi i(\lambda, t). \tag{1.34}$$

Note that these are not all distinct functions. Since they possess symmetries

$$c_\lambda(t) = c_{w\lambda}(t),$$
$$s_\lambda(t) = (-1)^{|w|} s_{w\lambda}(t),$$

for every $w \in W$, we need only one λ from each orbit of W. The weights in the Weyl chamber, $\Lambda_+ = \mathcal{C}_+ \cap \Lambda$, form a natural index set of orbit representatives, and we find L^2 orthogonal bases by taking the corresponding $c_\lambda(t)$ and $s_\lambda(t)$. Lemma 1.19 holds also in this more general case.

Lemma 1.22 *Given a rank d root system ϕ with weights lattice Λ, let W be the Weyl group and $\triangle$ denote the fundamental domain of the affine Weyl group $\widetilde{W} = W \rtimes L^\vee$. The functions $\{c_\lambda(t)\}_{\lambda \in \Lambda_+}$ and $\{s_\lambda(t)\}_{\lambda \in \Lambda_+}$ form two distinct L^2 orthogonal bases for $L^2(\triangle)$. These basis functions are eigenfunctions of the Laplacian ∇^2 on $\triangle$ satisfying homogeneous Neumann and Dirichlet boundary conditions, as in Lemma 1.19.*

Truncations of these bases do not, unfortunately, form spectrally convergent approximation spaces for analytic functions on $\triangle$. Approximation of a function f, defined on $\triangle$, in terms of $\{c_\lambda(t)\}$ is equivalent to

Fourier approximation of the even extension of f in $L^2_\vee(\mathrm{T})$, and we do in general only observe quadratic convergence due to discontinuity of the gradient across the boundary of $\triangle$. A route to spectral convergence is by a change of variables which turns the trigonometric polynomials $\{c_\lambda(t)\}$ and $\{s_\lambda(t)\}$ into multivariate Chebyshev polynomials of first and second kind.

1.3.4 Multivariate Chebyshev polynomials

Recall that classical Chebyshev polynomials of first and second kind, $T_k(x)$ and $U_k(x)$ are obtained from $\cos(k\theta)$ and $\sin(k\theta)$ by a change of variable $x = \cos(\theta)$ as

$$T_k(x) = \cos(k\theta),$$
$$U_k(x) = \frac{\sin((k+1)\theta)}{\sin(\theta)}.$$

We want to understand this construction in the context of affine Weyl groups. We recognize $\cos(k\theta)$ and $\sin(k\theta)$ as the symmetrized and skew-symmetrized exponentials. The $\cos(\theta)$ used in the change of variables is the 2π-periodic function that is symmetric, non-constant and has the longest wavelength (as such, uniquely defined up to a constant). In other words $\cos(\theta) = \pi_\vee \exp(\lambda_1\theta)$, where $\lambda_1 = 1$ is the generator of the weights lattice. Any periodic band limited even function f has a symmetric Fourier series of finite support on the weights lattice, and must hence be a polynomial in the variable $x = \cos(\theta)$. The denominator $\sin(\theta)$ is similarly the odd function of longest possible wavelength. Any periodic band limited odd function f has a skew-symmetric Fourier series on the weights lattice. Dividing out by $\sin(\theta)$ results in a band limited even function which again must map to a polynomial under our change of variables. The denominator, which in this special case is $\sin(\theta)$, is called the *Weyl denominator*. It plays an important role in representation theory of compact Lie groups as the denominator in *Weyl's* character formula, a cornerstone of representation theory. We will detail these constructions in the sequel.

As before, we let Φ be a rank d root system on $\mathcal{V} = \mathbb{R}^d$, with Weyl group W, co-root lattice $L^\vee$ and affine Weyl group $\widetilde{W} = W \rtimes L^\vee$. Let $\mathrm{T} = \mathcal{V}/L^\vee$ be the torus of periodicity and $\Lambda = \mathrm{span}_{\mathbb{Z}}\{\lambda_1, \ldots, \lambda_d\}$ the reciprocal lattice of $L^\vee$. It is convenient to write the group Λ in multiplicative form, where we let $\{e^\lambda\}_{\lambda\in\Lambda}$ denote the elements of the multiplicative group, understood as formal symbols such that for $\lambda, \mu \in \Lambda$ we have $e^\lambda \cdot e^\mu = e^{\lambda+\mu}$.

Let $\mathcal{E} = \mathcal{E}(\mathbb{C}) \subset L^2(\Lambda)$ denote the free complex vector space over the symbols e^λ. This consists of all formal sums $a = \sum_{\lambda\in\Lambda} a(\lambda)e^\lambda$ where the coefficients $a(\lambda) \in \mathbb{C}$ and all but a finite number of these are non-zero. An element $a \in \mathcal{E}$ is identified with a trigonometric polynomial $f(t) = \mathcal{F}^{-1}(a)(t)$ on the torus T (i.e. a band limited periodic function) through the Fourier transforms given in (1.29)–(1.30).

Let $\mathcal{E}_\vee^W \subset \mathcal{E}$ denote the symmetric subalgebra of those elements that are invariant under the action of the Weyl group W on $\mathcal{E}$. This consists of those $a \in \mathcal{E}$ where $a(\lambda) = a(w\lambda)$ for all λ and all $w \in W$. Similarly, $\mathcal{E}_\wedge^W \subset \mathcal{E}$ denotes those $a \in \mathcal{E}$ that are alternating sign under reflections s_α, i.e. where the coefficients satisfy $a(\lambda) = (-1)^{|w|}a(w\lambda)$. Projections $\pi_\vee^W : \mathcal{E} \to \mathcal{E}_\vee$ and $\pi_\wedge^W : \mathcal{E} \to \mathcal{E}_\wedge$ are defined as in (1.31)–(1.32).

The algebra $\mathcal{E}$ is generated by $\{e^{\lambda_j}\}_{j=1}^d \cup \{e^{-\lambda_j}\}_{j=1}^d$ where λ_j are the fundamental dominant weights. $\mathcal{E}_\vee^W$ is the subalgebra generated by the symmetric generators $\{z_j\}_{j=1}^d$ defined as

$$z_j = \pi_\vee^W e^{\lambda_j} = \frac{1}{|W|}\sum_{w\in W} e^{w\lambda_j} = \frac{2}{|W|}\sum_{w\in W^+} e^{w\lambda_j}, \tag{1.35}$$

where W^+ denotes the even subgroup of W containing those w such that $|w|$ is even. The latter identity follows from $s_{\alpha_j}\lambda_j = \lambda_j$, thus it is enough to consider only w of even length. The action of W^+ on λ_j is free and effective.

It can be shown that $\mathcal{E}_\vee^W$ is a unique factorization domain over the generators $\{z_j\}$, i.e. any $a \in \mathcal{E}_\vee^W$ can be expressed uniquely as a polynomial in $\{z_j\}_{j=1}^d$.

The skew subspace $\mathcal{E}_\wedge^W$ does not form an algebra, but this can be corrected by dividing out the *Weyl denominator*. Define the *Weyl vector* $\rho \in \Lambda$ as

$$\rho = \sum_{j=1}^d \lambda_j = \frac{1}{2}\sum_{\alpha\in\Phi^+} \alpha.$$

We define the Weyl denominator $D \in \mathcal{E}_\wedge^W$ as

$$D = \sum_{w\in W} (-1)^{|w|} e^{w\rho}.$$

Proposition 1 Any $a \in \mathcal{E}_\wedge^W$ is divisible by D, i.e. there exist a unique $b \in \mathcal{E}_\vee^W$ such that $a = bD$.

Proof See [8] Prop. 25.2. □

Any $a \in \mathcal{E}_\vee^W$ can be written as a polynomial in $z_1, \ldots, z_d$, hence the following polynomials are well defined.

Definition 1.23 For $\lambda \in \Lambda$ we define multivariate Chebyshev polynomials of first and second kind T_λ and U_λ as the unique polynomials that satisfy

$$T_\lambda(z_1, \ldots, z_d) = \pi_\vee^W e^\lambda = \frac{1}{|W|} \sum_{w \in W} e^{w\lambda}, \tag{1.36}$$

$$U_\lambda(z_1, \ldots, z_d) = \frac{|W| \pi_\wedge^W e^{\lambda+\rho}}{D} = \frac{\sum_{w \in W} (-1)^{|w|} e^{w(\rho+\lambda)}}{\sum_{w \in W} (-1)^{|w|} e^{w\rho}}. \tag{1.37}$$

By a slight abuse of notation, we will also consider z_j as W-invariant functions in $\mathbb{C}T$ as

$$z_j(t) = \mathcal{F}^{-1}(z_j)(t) = \frac{2}{|W|} \sum_{w \in W^+} \langle w\lambda_j, t \rangle_T. \tag{1.38}$$

The functions $z_j(t)$ may be real or complex. If there exist an $w \in W^+$ such that $w\lambda_j = \lambda_j$ then $z_j = \overline{z_j}$ is real. Otherwise there must exist an index $\overline{j} \neq j$ and a $w \in W^+$ such that $w\lambda_j = \lambda_{\overline{j}}$ and we have $\overline{z_j} = z_{\overline{j}}$. In the latter case we can replace these with d real coordinates $x_j = \frac{1}{2}(z_j + z_{\overline{j}})$, $x_{\overline{j}} = \frac{1}{2i}(z_j - z_{\overline{j}})$.

We remark that (1.37) is exactly the same formula as Weyl's character formula, giving the trace of all the irreducible characters on a semisimple Lie group [8]. These characters form an L^2 orthogonal basis for the space of class functions on the Lie group. Thus, expansion in terms of second kind multivariate Chebyshev polynomials is equivalent to expansion in terms of irreducible characters on a Lie group. In a similar way, the basis given by the irreducible representations block-diagonalize equivariant linear operators on a Lie group, it is also such that one may use the irreducible characters to obtain block diagonalizations, see [24]. Thus, our software, which is primarily constructed to deal with spectral element discretizations of PDEs, may also have important applications in computations on Lie groups. This opens up a whole area of possible applications of these approximations.

We will briefly summarize important properties of the multivariate Chebyshev polynomials which are presented in detail in [28, 30, 32, 9], and we will be a bit more detailed on some properties which are not detailed in these references.

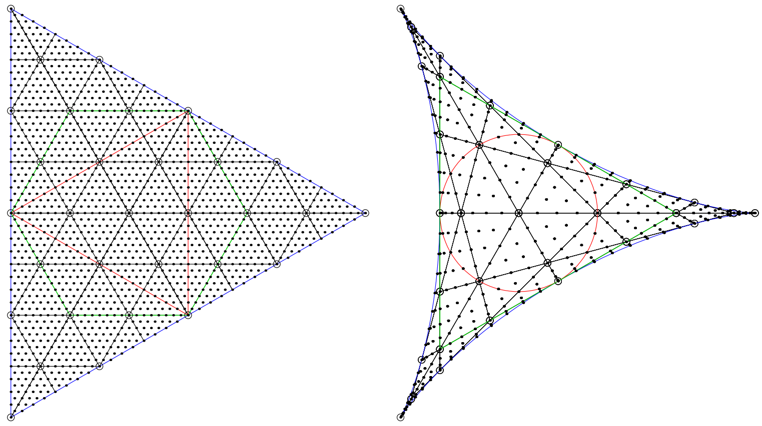

(a) (b)

Figure 1.5 The equilateral domain Δ in (a) maps to the Deltoid δ in (b) under $t \mapsto z(t)$.

Continuous orthogonality. Let Φ be an irreducible root system on $\mathcal{V} = \mathbb{R}^d$ with an alcove $\triangle$ being the simplex defined in Lemma 1.21. The corresponding family of multivariate Chebyshev polynomials are orthogonal on the domain

$$\delta = z(\triangle), \tag{1.39}$$

with respect to the inner product

$$(f, g) = \int_\delta \overline{f(z)} g(z) \frac{1}{\sqrt{D\overline{D}}} dz, \tag{1.40}$$

where D is the Weyl denominator. Figure 1.5 shows $\triangle$ and δ in the A_2 case. Here δ is a *deltoid*, a domain with cusps in each corner. For the A_3 case, $\triangle$ is shown in Figure 1.4 and δ in Figure 1.6. It poses a problem for applications that the domains δ are *not* simplices. However, in the A_2 case, there is a hexagon inside $\triangle$ which maps to an equilateral triangle inscribed in δ, whereas in the A_3 case there is an octagon in $\triangle$ which maps to a tetrahedron inscribed in δ. Note that the tetrahedron is not regular, but has two sides of length 1 and four sides of length $\sqrt{17/18}$. In our spectral element methods we have applied overlapping δ-subdomains such that the inscribed triangles form a simplicial (non-overlapping) subdivision.

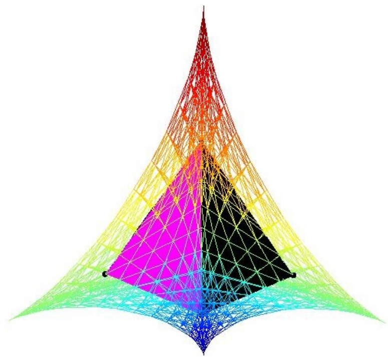

Figure 1.6 The deltoid domain δ in case A_3. The inscribed tetrahedron is the image of the regular octahedron inscribed in the Weyl chamber of $\triangle$, see Figure 1.4.

Discrete orthogonality and sampling. Pulled back to t-coordinates, polynomials in z become band limited symmetric functions $f \in L^2_\vee(\mathrm{T})$. Due to the Shannon sampling theorem, any band limited function $f \in L^2(\mathrm{T})$ can be exactly reconstructed from sampling on a sufficiently fine lattice $S < \mathrm{T}$. To preserve periodicity, we require that

$$S^\perp < \Lambda, \tag{1.41}$$

and in addition, we also require that the sampling lattice is W invariant,

$$WS = S. \tag{1.42}$$

Such a lattice S leads to exact discrete orthogonality formulae and discrete quadrature rules for polynomials such that $f(t)$ satisfies the Shannon criterion

$$\mathrm{supp}(\widehat{f}) \subset \widehat{\sigma}(\widehat{S}), \tag{1.43}$$

where $\widehat{\sigma}$ is a Voronoi transversal. There are several ways to construct such a lattice. In [30, 32] we work with down scalings of the co-root lattice,

$$S = \frac{1}{m} L^\vee,$$

where the integer m is chosen large enough for S to satisfy the Shannon criterion in the polynomial space where we want to perform approximations. This lattice is shown in Figure 1.5 in the A_2 case with circles for $m = 12$ and dots for $m = 48$. The resulting approximation space is the span of the multivariate polynomials up to degree $m/2$, except one particular polynomial of degree $m/2$ which aliases to 0 on S. In addition the space contains some particular polynomials of degree up to $\frac{2}{3}m$.

Another lattice, which in many respects is more elegant, is the downscaling of the weights lattice

$$S = \frac{1}{m}\Lambda, \quad \text{for } m \in \mathbb{N}. \tag{1.44}$$

With this discretization we obtain a perfect symmetry between primal space and Fourier space. The primal space is the periodic domain $T = \mathbb{R}^n/L^\vee$, sampled in $S = \frac{1}{m}\Lambda$. The sampling turns the Fourier space into the periodic domain $\mathbb{R}^n/S^\perp = \mathbb{R}^n/mL^\vee$. On the other hand, periodicity under translation with $L^\vee$ in primal space is equivalent to sampling in Fourier space at the lattice $(L^\vee)^\perp = \Lambda$. Thus, the discretized periodic domain is self dual up to the scaling with m.

The sampling at $S = \frac{1}{m}\Lambda$ is perfect also in the sense that we, for the A_n Weyl groups, obtain an approximation space consisting of exactly the n-variate polynomials up to degree m and nothing else. The space of n-variate m-degree polynomials has dimension $\binom{m+n}{m}$. The alcove $\triangle$ of the affine Weyl group A_n is a simplex with corners in the origin and in λ_j, the principal dominant weights. Hence S is a downscaling of a simplex containing $\binom{m+n}{m}$ points, in particular these are the triangle numbers for $n = 2$ and the pyramidal numbers for $n = 3$. In Fourier space the situation is the same, the discrete space contains exactly the n-variate Chebyshev polynomials of degree up to and including m.

Lebesgue numbers. The elegant sampling results above are useless unless we can guarantee small Lebesque numbers and hence stable sampling. Fortunately, the following result holds for the A_n family sampled at the downscaled weights lattice.

Theorem 1.24 *Let λ be the weights lattice in the affine Weyl group A_n and let $S = \frac{1}{m}\Lambda$. Let $z\colon \triangle \to \delta$ be the coordinate change in (1.35). The Lebesgue number of the points $z(S)$ grows as*

$$\mathcal{O}(\log^n(m)). \tag{1.45}$$

Similar results hold for other lattices (e.g. downscaled roots lattice) and other affine Weyl groups. To prove this theorem, we must first show that the Lagrangian interpolating polynomials in the nodal set $z(S)$ are given as

$$\ell_j(z(t)) = c \sum_{w \in W} \mathcal{D}_S(w(t - t_j)) \quad \text{for all } t_j \in S,$$

where $\mathcal{D}_S(t)$ is the polyhedral Dirichlet kernel. The result follows from general bounds on Lebesgue numbers of the polyhedral Dirichlet kernels [34, 36]. We omit the details. Similar results are shown for special cases in [37].

Symmetric FFTs and fast Chebyshev expansions. In this paragraph we let Φ be the root system A_n with Weyl group $W \simeq S_{n+1}$ with $|W| = (n+1)!$. Let $S = \frac{1}{m}\Lambda$ and let G be the finite abelian group $G = S/L^\vee$. Under the action of W on G, the fundamental domain is $\triangle_m = \triangle \cup S$, where $\triangle$ is the n-simplex with corners in the origin and in the fundamental dominant weights λ_j. Thus $\triangle_m$ consists of $\binom{m+n}{m}$ points. The Weyl group also acts on $\widehat{G}$ via the adjoint $\langle \xi, wx\rangle = \langle \widehat{w}\xi, x\rangle$, with fundamental domain $\widehat{\triangle}_m$. Note that this case is self dual, so we can identify $\triangle_m$ and $\widehat{\triangle}_m$.

Let $p(z)$ be a real n-variate polynomial of degree m and let $f = p(z(\triangle_m))$ be the sampling of $p(z)$. The Fourier transform is defined as

$$\widehat{f}(\lambda) = \mathcal{F}_G(\pi^W_\vee f)(\lambda) \quad \text{for } \lambda \in \triangle_m. \tag{1.46}$$

The Chebyshev expansion of $p(z)$ is

$$p(z) = \sum_{\lambda \in \widehat{\triangle}_m} \widehat{f}(\lambda) T_\lambda(z). \tag{1.47}$$

How fast can this Chebyshev expansion be computed? One possibility is to extend f to $\mathbb{C}G$ and apply an FFT here. The number of points in G is approximately $N \approx (n+1)!\binom{m+n}{m}$. Counting both multiplications and additions, a complex FFT on N points costs $5N\log_2(N)$ floating point operations (flops).

It is, however, possible to improve the speed of this computation by a factor $2|W| = 2(n+1)!$ by using a symmetric FFT, as in [28]. The factor 2 comes from the assumption that $p(z)$ is real, not complex, and the factor $(n+1)!$ by carefully exploiting all the symmetries. This can be implemented as a Cooley–Tukey style FFT computation. Letting $m = 2^k$ we have a sequence of sub-lattices

$$S_1 < S_2 < S_4 < \cdots < S_{2^k} = S_m.$$

We know that Cooley–Tukey can be formulated with respect to arbitrary decompositions in sub-lattices. All the sub-lattices preserve all the symmetries of W, thus by organizing the computations carefully, it is sufficient to compute in the fundamental domains of the action of W. However, the Cooley–Tukey splitting is not just a splitting in sub-lattices, it also involves computations in the cosets of the sub-lattices. The cosets may lose some symmetries of W, and for such cosets the fundamental domain becomes larger. However, in that case the lost symmetry will map one coset to another coset, and we may discard the cosets that are equivalent to some other coset by a symmetry. In the dual space $\widehat{G}$, the situation is similar, W acts by the adjoint action and we can similarly exploit all the symmetries. In addition we have the real symmetry expressed as $\widehat{f}(\lambda) = \overline{\widehat{f}(-\lambda)}$.

To make a rather complicated story short, by carefully paying attention to exploiting all the symmetries, it can be shown:

Theorem 1.25 *Let π_m^n be the space of n-variate real polynomials of degree m, a real vector space of dimension $N = \binom{m+n}{m}$. The A_n Chebyshev expansion of $p(z) \in \pi_m^n$ can be computed in a stable manner by sampling $p(z)$ in the N points $z(\triangle_m)$. The cost of this computation is $\frac{5}{2}N\log_2(N)$ floating point numbers, counting both additions and multiplications.*

Recurrence formulae. Practical computations with Chebyshev expansions rely on the availability of fast transforms between 'nodal' and 'modal' representation, i.e. between sampled values and Chebyshev expansion coefficients. Products of polynomials can be computed fast in the nodal domain, while other operations such as integration and derivation can be done fast in the modal domain. Also, in the multivariate case, there exist recursion formulae similar to the univariate case.

From the convolution product in $\mathcal{E}_\vee^W$, one finds that the T_λ satisfy the recurrence relations

$$T_0 = 1, \tag{1.48}$$

$$T_{\lambda_j} = z_j, \tag{1.49}$$

$$T_\lambda = T_{w\lambda} \text{ for } w \in W, \tag{1.50}$$

$$T_{-\lambda} = \overline{T_\lambda}, \tag{1.51}$$

$$T_\lambda T_\mu = \frac{1}{|W|}\sum_{w\in W} T_{\lambda+w\mu}. \tag{1.52}$$

These reduce to classical three-term recurrences for A_1 and four term recurrences for A_2, see [30].

There also exist linear recurrence relations which compute the Chebyshev expansion of $\nabla p(z)$ from the Chebyshev expansion of $p(z)$, see [32].

1.4 Conclusions

We have discussed various aspects of group theory as a tool for analyzing numerical algorithms and for developing new computational algorithms. In particular, we have focused on algorithms related to sampling in regular lattices and algorithms related to applications of reflection groups and kaleidoscopes. Group theory provides an important conceptual framework both for understanding the fundamental structure of the algorithms and also for structuring the implementation of numerical software.

Connections between group theory and Chebyshev approximations have been discussed in detail. This is a connection which goes much deeper than what has been presented in this paper. Weyl's character formula for the irreducible characters on simple Lie groups is essentially identical to the definition of second kind multivariate Chebyshev polynomials. Chebyshev approximation theory is therefore intimately linked to approximation theory on Lie groups.

Another topic, which has not been discussed in this paper, is the application of group theory in the development of algorithms for computing time evolution of dynamical systems. Lie group integrators are general numerical methods built around Lie group actions on the phase space of the dynamical system [23]. Lie group integrators rely upon exact computations of operator exponentials. The interplay between Lie group integrators for time evolution and Chebyshev based spectral element methods for spatial discretizations is an active area of research where all these techniques come together.

References

[1] K. Åhlander, M. Haveraaen, and H. Munthe-Kaas. On the role of mathematical abstractions for scientific computing. In R. F. Boisvert and P. T. P. Tang, editors, *The Architecture of Scientific Software*, volume 60 of *IFIP Advances in Information and Communication Technology*, pages 145–158. Citeseer, 2001.

[2] K. Åhlander and H. Munthe-Kaas. Applications of the generalized Fourier transform in numerical linear algebra. *BIT Numerical Mathematics*, 45(4):819–850, 2005.

[3] E. L. Allgower, K. Georg, R. Miranda, and J. Tausch. Numerical exploitation of equivariance. *Zeitschrift für Angewandte Mathematik und Mechanik*, 78:185–201, 1998.

[4] L. Auslander. *Lecture notes on nil-theta functions*. Number 34. Amer Mathematical Society, 1977.

[5] L. Auslander and R. Tolimieri. Is computing with the finite Fourier transform pure or applied mathematics. *Notices of the AMS*, 1(6), 1979.

[6] R. J. Beerends. Chebyshev polynomials in several variables and the radial part of the Laplace–Beltrami operator. *Transactions of the AMS*, 328(2):779–814, 1991.

[7] A. Bossavit. Symmetry, groups, and boundary value poblems. a progressive introduction to noncommutative harmonic analysis of partial differential equations in domains with geometrical symmetry. *Comput. Methods Appl. Mech. and Engrg.*, 56:167–215, 1986.

[8] D. Bump. *Lie Groups*. Springer Verlag, 2004.

[9] S.H. Christiansen, H.Z. Munthe-Kaas, and B. Owren. Topics in structure-preserving discretization. *Acta Numerica*, 20(1):1–119, 2011.

[10] J.H. Conway, N.J.A. Sloane, and E. Bannai. *Sphere packings, lattices, and groups*, volume 290. Springer Verlag, 1999.

[11] C. C. Douglas and J. Mandel. Abstract theory for the domain reduction method. *Computing*, 48:73–96, 1992.

[12] M. Dubiner. Spectral methods on triangles and other domains. *Journal of Scientific Computing*, 6(4):345–390, 1991.

[13] R. Eier and R. Lidl. A class of orthogonal polynomials in k variables. *Math. Ann.*, 260:93–99, 1982.

[14] K. Engo, A. Marthinsen, and H.Z. Munthe-Kaas. Diffman: An object-oriented matlab toolbox for solving differential equations on manifolds. *Applied numerical mathematics*, 39(3-4):323–347, 2001.

[15] A. F. Fässler and E. Stiefel. *Group theoretical methods and their applications*. Birkhäuser, Boston, 1992.

[16] K. Georg and R. Miranda. Exploiting symmetry in solving linear equations. In E. L. Allgower, K. Böhmer, and M. Golubisky, editors, *Bifurcation and Symmetry*, volume 104 of *ISNM*, pages 157–168. Birkhäuser, Basel, 1992.

[17] F. X. Giraldo and T. Warburton. A nodal triangle-based specral element method for the shallow water equations on the sphere. *J. of Comp. Physics*, 2005.

[18] R. Hadani and A. Singer. Representation theoretic patterns in three dimensional cryo-electron microscopy i-the intrinsic reconstitution algorithm. *Arxiv preprint arXiv:0910.3009. To appear in Annals of Mathematics.*, 2009.

[19] R. Hadani and A. Singer. Representation theoretic patterns in three-dimensional cryo-electron microscopy ii—the class averaging problem. *Foundations of Computational Mathematics*, pages 1–28, 2011.

[20] J.S. Hesthaven and T. Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*, volume 54. Springer-Verlag New York Inc, 2008.

[21] M. E. Hoffman and W. D. Withers. Generalized Chebyshev polynomials associated with affine Weyl groups. *Transactions of the AMS*, 308(1):91–104, 1988.

[22] D. Huybrechs. On the Fourier extension of non-periodic functions. *SIAM J. Numer. Anal.*, 47(6):4326–4355, 2010.

[23] A. Iserles, H. Munthe-Kaas, S. P. Nørsett, and A. Zanna. Lie-group methods. In *Acta numerica, 2000*, volume 9 of *Acta Numer.*, pages 215–365. Cambridge Univ. Press, Cambridge, 2000.

[24] G. James and M. Liebeck. *Representations and Characters of Groups.* Cambridge University Press, 2 edition, 2001. ISBN 052100392X.

[25] T. Koornwinder. Orthogonal polynomials in two variables which are eigenfunctions of two algebraically independent partial differential operators I–IV. *Indiag. Math.*, 36:48–66 and 357–381, 1974.

[26] R. Lidl. Tchebyscheffpolynome in mehreren Variabelen. *J. Reine Angew. Math*, 273:178–198, 1975.

[27] J. S. Lomont. *Applications of Finite Groups.* Academic Press, New York, 1959.

[28] H. Munthe-Kaas. Symmetric FFTs; a general approach. In *Topics in linear algebra for vector- and parallel computers, PhD thesis.* NTNU, Trondheim, Norway, 1989. Available at: `http://hans.munthe-kaas.no`.

[29] H. Munthe-Kaas and T. Sørevik. Multidimensional pseudo-spectral methods on lattice grids. *Applied Numerical Mathematics*, 2011.

[30] H.Z. Munthe-Kaas. On group Fourier analysis and symmetry preserving discretizations of PDEs. *Journal of Physics A: Mathematical and General*, 39:5563, 2006.

[31] W. Rudin. *Fourier analysis on groups*, volume 12. Wiley-Interscience, 1990.

[32] B.N. Ryland and H.Z. Munthe-Kaas. On Multivariate Chebyshev Polynomials and Spectral Approximations on Triangles. In J. S. Hesthaven and E. M. Rønquist, editors, *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes Comp. Sci. and Eng.*, pages 19–41. Springer, 2011.

[33] J. P. Serre. *Linear Representations of Finite Groups.* Springer, 1977. ISBN 0387901906.

[34] R.J. Stanton and P.A. Tomas. Polyhedral summability of Fourier series on compact Lie groups. *American Journal of Mathematics*, 100(3):477–493, 1978.

[35] S. Thangavelu. *Harmonic analysis on the Heisenberg group*, volume 159. Birkhauser, 1998.

[36] G. Travaglini. Polyhedral summability of multiple Fourier series. In *Colloq. Math*, volume 65, pages 103–116, 1993.

[37] Y. Xu. Fourier series and approximation on hexagonal and triangular domains. *Constructive Approximation*, 31(1):115–138, 2010.