

APPLICATIONS OF THE GENERALIZED FOURIER TRANSFORM IN NUMERICAL LINEAR ALGEBRA*

KRISTER ÅHLANDER¹ and HANS MUNTHE-KAAS^{2,**}

¹*Department of Information Technology, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden. email: krister@it.uu.se*

²*Department of Mathematics, University of Bergen, Joh. Brunsgate 12, N-5008 Bergen, Norway. email: hans@munthe-kaas.no*

Abstract.

Equivariant matrices, commuting with a group of permutation matrices, are considered. Such matrices typically arise from PDEs and other computational problems where the computational domain exhibits discrete geometrical symmetries. In these cases, group representation theory provides a powerful tool for block diagonalizing the matrix via the Generalized Fourier Transform (GFT). This technique yields substantial computational savings in problems such as solving linear systems, computing eigenvalues and computing analytic matrix functions such as the matrix exponential.

The paper is presenting a comprehensive self contained introduction to this field. Building upon the familiar special (finite commutative) case of circulant matrices being diagonalized with the Discrete Fourier Transform, we generalize the classical convolution theorem and diagonalization results to the noncommutative case of block diagonalizing equivariant matrices.

Applications of the GFT in problems with domain symmetries have been developed by several authors in a series of papers. In this paper we elaborate upon the results in these papers by emphasizing the connection between equivariant matrices, block group algebras and noncommutative convolutions. Furthermore, we describe the algebraic structure of projections related to non-free group actions. This approach highlights the role of the underlying mathematical structures, and provides insight useful both for software construction and numerical analysis. The theory is illustrated with a selection of numerical examples.

AMS subject classification (2000): 43A30, 65T99, 20B25.

Key words: non commutative Fourier analysis, equivariant operators, block diagonalization.

1 Introduction.

As a motivation for the general theory, we briefly recapture the theory of circulant matrices. Let $\mathcal{I} = \{0, 1, \dots, n-1\}$ be a set of indices. A matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$

* Submitted July 2004. Revised April 2005. Communicated by Lars Eldén.

** This work has been partially supported by a grant from the Göran Gustafsson foundation and by the Center for Advanced Study, Oslo, through the Special Year in Geometric Integration 2002–2003.

is *circulant* if there exists a function $A : \mathcal{I} \rightarrow \mathbb{C}$ such that

$$\mathbf{A}_{i,j} = \mathbf{A}_{i-j \bmod n,0} = A(i - j \bmod n) \quad \text{for all } i, j \in \mathcal{I}.$$

Circulants form one of the most important classes of matrices in computational science. Their omnipresence is due to the fact that circulants represent discretizations of shift invariant linear operators, and their computational usefulness arise because they can be diagonalized via the Discrete Fourier Transform (DFT). Employing the Fast Fourier Transform, important operations on circulants such as computing products, inverses and matrix exponentials can be achieved in just $\mathcal{O}(n \log n)$ operations. This is fundamental for many algorithms in digital filter theory, spectral methods for differential equations, variable precision arithmetic and computational number theory.

A point of departure for understanding these remarkable properties of circulants is noting that \mathbf{A} is circulant if and only if it is *equivariant with respect to cyclic shifts of the domain*. That is, $\mathbf{A}\mathbf{S} = \mathbf{S}\mathbf{A}$, where \mathbf{S} is the cyclic shift matrix,

$$\mathbf{S} = \begin{pmatrix} & & 1 & & \\ & & & \ddots & \\ & & & & 1 \\ 1 & & & & \end{pmatrix}.$$

The cyclic shift matrix \mathbf{S} can be defined by a permutation of vector indices

$$(\mathbf{S}\mathbf{x})_i = \mathbf{x}_{i+1 \bmod n} \quad \text{for all } i \in \mathcal{I}.$$

The eigenvectors of \mathbf{S} are given as $\{\boldsymbol{\chi}^k\}_{k \in \mathcal{I}}$, where $\boldsymbol{\chi}_j^k = \exp(2\pi i k j / n)$ for $j \in \mathcal{I}$. The DFT computes the expansion of a vector in this basis. Note that \mathbf{A} can be written as a polynomial in \mathbf{S} ,

$$(1.1) \quad \mathbf{A} = \sum_{i \in \mathcal{I}} A(i) \mathbf{S}^{-i},$$

thus $\boldsymbol{\chi}^k$ are also eigenvectors of \mathbf{A} , and this explains why circulants are diagonalized by the DFT.

A slightly different explanation is in terms of *convolutions*. The product of circulants \mathbf{A} and \mathbf{B} is a circulant matrix $\mathbf{C} = \mathbf{A}\mathbf{B}$ represented by a function $C : \mathcal{I} \rightarrow \mathbb{C}$ given by the (circular) convolution

$$C(j) = \sum_{\ell \in \mathcal{I}} A(j - \ell \bmod n) B(\ell).$$

Defining Fourier coefficients as

$$(1.2) \quad \hat{A}(k) = \sum_{j \in \mathcal{I}} A(j) \boldsymbol{\chi}_j^k$$

we find, using the following *homomorphism property* of the eigenfunctions

$$(1.3) \quad \chi_{j+\ell \bmod n}^k = \chi_j^k \chi_\ell^k,$$

that the Fourier transform *diagonalizes convolutions*, $\hat{C}(k) = \hat{A}(k)\hat{B}(k)$.

All the above theory is readily generalized to the multidimensional case. For example, on a 2D grid, periodic in both directions, we say that \mathbf{A} is 2D circulant if it satisfies the equivariance $\mathbf{A}\mathbf{S}_x = \mathbf{S}_x\mathbf{A}$ and $\mathbf{A}\mathbf{S}_y = \mathbf{S}_y\mathbf{A}$, where \mathbf{S}_x and \mathbf{S}_y represent cyclic shifts in x and y -directions. Since shifts commute, $\mathbf{S}_x\mathbf{S}_y = \mathbf{S}_y\mathbf{S}_x$, they have a common set of eigenvectors, which are also eigenvectors for \mathbf{A} . The 2D DFT is an expansion in this basis, which diagonalizes any 2D circulant matrix, or equivalently any 2D circulant convolution.

The 2D DFT is a well-known example of classical Fourier analysis and circulant matrix theory, which is concerned with equivariance under *commutative* or *abelian* groups. In this case the existence of a complete set of common eigenvectors for all elements in the group provides an adequate tool for diagonalizing circulants.

This paper is concerned with generalizations of circulant matrix theory to *noncommutative* groups \mathcal{G} , in which case a complete set of common eigenvectors for all $\mathbf{G} \in \mathcal{G}$ cannot be found. Instead there always exists a complete set of *irreducible representations* for \mathcal{G} . This brings us into the beautiful topic of representation theory and noncommutative Fourier analysis, a very useful mathematical tool which has not received the attention it deserves in the computational science communities. Our goal is to provide a self contained introduction to this field, with matrix computations as our main application. The *Generalized Fourier Transform* (GFT), expands vectors in terms of the basis given by the irreducible representations. Under this change of basis, any \mathcal{G} -equivariant matrix \mathbf{A} is *block-diagonalized* rather than diagonalized. For a wide range of computational problems, such as solving linear equations, computing matrix functions and eigenvalue computations, the computational savings obtained by GFT based block diagonalizations are substantial.

Group representation theory was invented by Frobenius in the late 19th century, and has since grown to a major branch of pure mathematics. For our applications in numerical linear algebra, it is, however, sufficient to understand the very basics of this theory. The basic algorithm for solving linear systems with symmetries, presented in Section 5.1, is due to Allgower et al. [6], and the extension to the case of symmetries with fixed points in Section 5.2 was first done in [5]. The numerical algorithms, including eigenvalue computations and symmetric preconditioners, have been developed in a series of papers [3, 4, 19, 20, 21, 35]. Related to these techniques is the theory of *domain reduction* of PDEs, described in [9, 10, 16]. A group theoretic development of symmetric FFTs on domains with discrete symmetries is developed in [29], and algebraic properties of symmetric transforms are developed in [17].

In our current exposition we have aimed at presenting the mathematical theory in such a way that the basic underlying structures become visible. Our development differs from the above references by its emphasis on the equivalence of

\mathcal{G} -equivariant matrices and convolutional operators in the group algebra. One advantage with our reformulation is that we can more easily relate to the theory of fast GFT algorithms. See [27, 32] for an overview of this field and [34] for parallel computing aspects. We have also simplified the formulas for handling fixed points.

Other important applications of the GFT include statistical analysis of ranked data [14, 15] and the construction of symmetrized fast cosine-type transforms [17]. The mathematical machinery and basic matrix decompositions employed in these works are related to the present theory, but the problems studied there differ in important aspects, and thus these algorithms are not readily applicable to the problems at hand. In the case of ranked data, the algorithms are coined at situations where the symmetry group is much *larger* than the number of points in the domain. The notion of equivariance employed in [17] is a natural generalization of the present definition. The main focus of their study is finding symmetries and symmetric factorizations of matrices related to generalized cosine-type transforms.

The paper is organized as follows. In Section 2 we define groups acting on domains and arrive at the definition of \mathcal{G} -equivariant matrices, which is a generalization of circulants. Section 3 establishes the equivalence between equivariant matrices and convolutions. In Section 4 we introduce fundamentals of group representation theory and Fourier analysis on (finite noncommutative) groups. Finally, in Section 5 we discuss applications of Fourier based block diagonalizations in numerical linear algebra, and in Section 6 a few numerical examples are presented.

2 Equivariance and domain symmetries.

2.1 The Triangle Example.

An important class of computational problems where noncommutative equivariance groups appear, is equations with *domain symmetries*. As a motivating example, we consider the simplest noncommutative case, the symmetries of an equilateral triangle. There are six linear transformations that map the triangle onto itself, three pure rotations and three rotations combined with reflections. In Figure 2.1a we indicate the two generators α (rotation 120° clockwise) and β (right-left reflection). These satisfy the algebraic relations $\alpha^3 = \beta^2 = e$, $\beta\alpha\beta = \alpha^{-1}$, where e denotes the identity transform. The whole group of symmetries is known as the *dihedral group* of order 6 with the elements $\mathcal{D}_3 = \{e, \alpha, \alpha^2, \beta, \alpha\beta, \alpha^2\beta\}$.

Suppose we want to solve a linear elliptic PDE $\mathcal{L}u = f$ on the triangle, where both α and β are symmetries of the elliptic operator, i.e. $(\mathcal{L}\alpha)u = (\alpha\mathcal{L})u$ and $(\mathcal{L}\beta)u = (\beta\mathcal{L})u$ for any u satisfying the appropriate boundary conditions on the triangle.

Let the domain be discretized with a *symmetry respecting discretization*, see Figure 2.1b. In this example we consider a finite difference discretization represented by the nodes $\mathcal{I} = \{1, 2, \dots, 10\}$, such that both α and β map nodes to nodes. In finite element discretizations one would use basis functions mapped

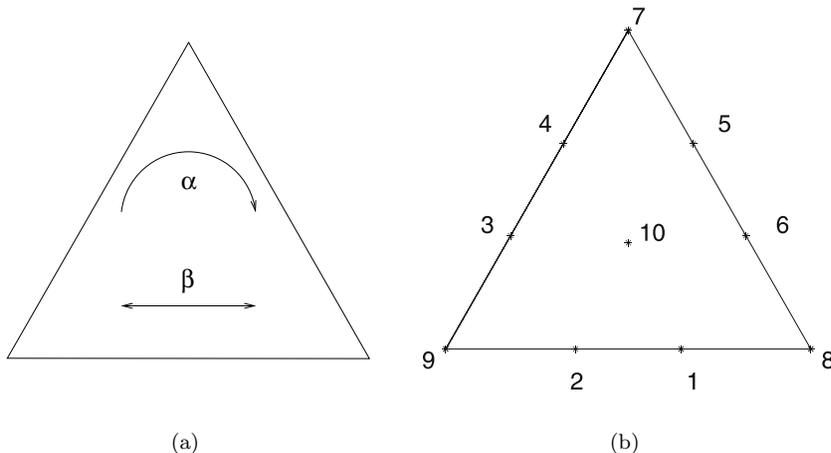


Figure 2.1: (a) An equilateral triangle is mapped onto itself by a rotation α , a reflection β and all combinations of α and β . (b) The discretization with nodes $1, \dots, 10$ preserves the symmetry.

to other basis functions by the symmetries. On the grid, the linear maps α and β are represented by two permutation matrices \mathbf{G}_α and \mathbf{G}_β acting on a vector $\mathbf{x} \in \mathbb{C}^{10}$ as

$$\mathbf{G}_\alpha(x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10})^T = (x_5, x_6, x_1, x_2, x_3, x_4, x_9, x_7, x_8, x_{10})^T,$$

$$\mathbf{G}_\beta(x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10})^T = (x_2, x_1, x_6, x_5, x_4, x_3, x_7, x_9, x_8, x_{10})^T.$$

The operator \mathcal{L} is discretized as a matrix $\mathbf{A} \in \mathbb{C}^{10 \times 10}$ satisfying the equivariances $\mathbf{A}\mathbf{G}_g = \mathbf{G}_g\mathbf{A}$ for $g \in \{\alpha, \beta\}$. It is important to note that since \mathbf{G}_g are permutation matrices, the equivariance of \mathbf{A} can be expressed in terms of symmetries of the matrix indices. Thus we have e.g.

$$\mathbf{A}_{1,6} = \mathbf{A}_{3,2} = \mathbf{A}_{5,4} = \mathbf{A}_{4,5} = \mathbf{A}_{2,3} = \mathbf{A}_{6,1}.$$

The above relations can also be understood by considering Figure 2.1b. For example, by rotating the dependence between nodes 1 and 6 (i.e. $\mathbf{A}_{1,6}$) we see that $\mathbf{A}_{3,2}$, the dependence between nodes 3 and 2, must be identical.

In the following, we will present a systematic way of representing and utilizing symmetries of this kind in numerical linear algebra. This Triangle Example will recur throughout the text.

2.2 Groups and actions.

A *group* is a set \mathcal{G} with a binary operation $g, h \mapsto gh$, inverse $g \mapsto g^{-1}$ and identity element e , such that $g(ht) = (gh)t$, $eg = ge = g$ and $gg^{-1} = g^{-1}g = e$ for all $g, h, t \in \mathcal{G}$. We let $|\mathcal{G}|$ denote the number of elements in the group. In this paper most groups are finite, $|\mathcal{G}| < \infty$. Important examples of groups are *groups of matrices*, where the group operation is the matrix product and e is the

identity matrix, and *groups of permutations*, where the group operation is the composition of two permutations and e is the identity permutation.

Let \mathcal{I} denote the set of indices used to enumerate the nodes in the discretization of a computational domain. We consider permutations acting on the indices \mathcal{I} and the corresponding linear transformations acting on the vector space $\mathbb{C}^{|\mathcal{I}|}$.

We say that a group \mathcal{G} *acts on* a set \mathcal{I} (from the right) if there exists a product $(i, g) \mapsto ig : \mathcal{I} \times \mathcal{G} \rightarrow \mathcal{I}$ such that

$$(2.1) \quad ie = i \quad \text{for all } i \in \mathcal{I},$$

$$(2.2) \quad i(gh) = (ig)h \quad \text{for all } g, h \in \mathcal{G} \text{ and } i \in \mathcal{I}.$$

The map $i \mapsto ig$ is a permutation of the set \mathcal{I} , with the inverse permutation being $i \mapsto ig^{-1}$.

An action partitions \mathcal{I} into disjoint *orbits*

$$\mathcal{O}_i = \{j \in \mathcal{I} \mid j = ig \text{ for some } g \in \mathcal{G}\}, \quad i \in \mathcal{I}.$$

We let $\mathcal{S} \subset \mathcal{I}$ denote a selection of *orbit representatives*, i.e. one element from each orbit. The action is called *transitive* if \mathcal{I} consists of just a single orbit, $|\mathcal{S}| = 1$.

For any $i \in \mathcal{I}$ we let the *isotropy subgroup at i* , \mathcal{G}_i be defined as

$$\mathcal{G}_i = \{g \in \mathcal{G} \mid ig = i\}.$$

The action is *free* if $\mathcal{G}_i = \{e\}$ for every $i \in \mathcal{I}$, i.e., there are no fixed points under the action of \mathcal{G} . We remark that the relation $|\mathcal{G}_i||\mathcal{O}_i| = |\mathcal{G}|$ holds.

EXAMPLE 2.1. In the Triangle Example we may pick orbit representatives as $\mathcal{S} = \{1, 7, 10\}$. The action of the symmetry group is free on the orbit $\mathcal{O}_1 = \{1, 2, 3, 4, 5, 6\}$, while the points in the orbit $\mathcal{O}_7 = \{7, 8, 9\}$ have isotropy subgroups of size 2, and finally $\mathcal{O}_{10} = \{10\}$ has isotropy of size 6.

2.3 Equivariant matrices.

Consider a group of permutations \mathcal{G} acting on \mathcal{I} . Let $n = |\mathcal{I}|$. For any permutation $g \in \mathcal{G}$ there corresponds a *permutation matrix* $\mathbf{P}(g) \in \mathbb{C}^{n \times n}$ defined by permutations of vector indices

$$(2.3) \quad (\mathbf{P}(g)\mathbf{x})_i = \mathbf{x}_{ig} \quad \text{for all } \mathbf{x} \in \mathbb{C}^n \text{ and } i \in \mathcal{I}.$$

Letting $\{\mathbf{k}^i\}_{i=1}^n$ denote the standard basis for \mathbb{C}^n , this implies that

$$(2.4) \quad \mathbf{P}(g)\mathbf{k}^i = \mathbf{k}^{ig^{-1}} \quad \text{for all } i \in \mathcal{I}.$$

Note that $\mathbf{P}(g)\mathbf{P}(h) = \mathbf{P}(gh)$ for all $g, h \in \mathcal{G}$ and \mathbf{P} is thus called a *permutation representation* of \mathcal{G} .

DEFINITION 2.1. Consider a group \mathcal{G} acting on \mathcal{I} . A matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ is \mathcal{G} -equivariant if

$$(2.5) \quad \mathbf{P}(g)\mathbf{A} = \mathbf{A}\mathbf{P}(g) \quad \text{for all } g \in \mathcal{G}.$$

PROPOSITION 2.1. A matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ is \mathcal{G} -equivariant iff

$$(2.6) \quad \mathbf{A}_{i,j} = \mathbf{A}_{ig,jg} \quad \text{for all } i, j \in \mathcal{I} \text{ and all } g \in \mathcal{G}.$$

PROOF. For a permutation matrix $\mathbf{P} = \mathbf{P}(g)$ we have $\mathbf{P}^{-1} = \mathbf{P}^T$. Assuming $\mathbf{A}\mathbf{P} = \mathbf{P}\mathbf{A}$, we find

$$\mathbf{A}_{ig,jg} = \mathbf{k}^{ig^T} \mathbf{A} \mathbf{k}^{jg} = (\mathbf{P}^{-1} \mathbf{k}^i)^T \mathbf{A} \mathbf{P}^{-1} \mathbf{k}^j = \mathbf{k}^{iT} \mathbf{P} \mathbf{A} \mathbf{P}^{-1} \mathbf{k}^j = \mathbf{k}^{iT} \mathbf{A} \mathbf{k}^j = \mathbf{A}_{i,j}.$$

The implication in the opposite direction follows similarly. □

EXAMPLE 2.2. In the Triangle Example we define the action of \mathcal{D}_3 on \mathcal{I} as

$$\begin{aligned} (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)\alpha &= (5, 6, 1, 2, 3, 4, 9, 7, 8, 10), \\ (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)\beta &= (2, 1, 6, 5, 4, 3, 7, 9, 8, 10), \end{aligned}$$

and extend to all of \mathcal{D}_3 using (2.2). Thus we obtain $\mathbf{G}_\alpha = \mathbf{P}(\alpha)$ and $\mathbf{G}_\beta = \mathbf{P}(\beta)$. \mathcal{G} -equivariance of \mathbf{A} is here equivalent to $\mathbf{A}_{i,j} = \mathbf{A}_{i\alpha,j\alpha} = \mathbf{A}_{i\beta,j\beta}, \forall i, j \in \mathcal{I}$.

3 Equivariance and convolutions.

To establish a relationship between \mathcal{G} -equivariant matrices and *convolutions*, we find it convenient to coordinate the nodes in \mathcal{I} via two indices. Recall that \mathcal{S} is a selection of orbit representatives. Define a map $\lambda : \mathcal{S} \times \mathcal{G} \rightarrow \mathcal{I}$ as

$$(3.1) \quad \lambda(i, g) = ig \quad \text{for all } i \in \mathcal{S}, g \in \mathcal{G}.$$

The theory will be developed in three stages:

1. If the action is both free and transitive, then \mathcal{S} consists of a single element $0 \in \mathcal{I}$, and λ establishes a natural 1–1 correspondence between \mathcal{G} and \mathcal{I} . The tool for studying equivariant matrices is convolutions in the *group algebra*, a vector space where the elements of \mathcal{G} are the basis vectors.
2. If the action is free, but not transitive, then \mathcal{I} splits into $m = |\mathcal{S}|$ orbits, each of size $|\mathcal{G}|$, and λ is a 1–1 identification of $\mathcal{S} \times \mathcal{G}$ and \mathcal{I} . In this case we will obtain matrix block convolution formulae, where the size of each block is m . In short, we will identify $\mathbb{C}^{|\mathcal{I}|} = \mathbb{C}^{|\mathcal{S}| \times |\mathcal{G}|}$.
3. In the general case where the action is not free, then \mathcal{I} splits into $m = |\mathcal{S}|$ orbits, each of size $|\mathcal{O}_i| = |\mathcal{G}|/|\mathcal{G}_i|$. In this case $\lambda(i, g) = \lambda(i, hg)$ for all $h \in \mathcal{G}_i$. We will identify $\mathbb{C}^{|\mathcal{I}|} \subset \mathbb{C}^{|\mathcal{S}| \times |\mathcal{G}|}$ as the linear subspace consisting of all the vectors being invariant under the actions of $h \in \mathcal{G}_i$ on $\mathcal{S} \times \mathcal{G}$ given by $(i, g) \mapsto (i, hg)$.

3.1 Free transitive actions and the group algebra.

Assume that \mathbf{A} is equivariant with respect to a permutation group \mathcal{G} acting transitively and freely on \mathcal{I} . The adequate tool for studying equivariance in this case is the *group algebra*. Since \mathcal{G} and \mathcal{I} are identified through λ in (3.1), we may also identify the basis $\{\mathbf{k}^i\}_{i \in \mathcal{I}}$ with \mathcal{G} , via $\mathbf{k}^{\lambda(0,g)} \mapsto g$, where $\mathcal{S} = \{0\}$ is the selection of an arbitrary point $0 \in \mathcal{I}$.

DEFINITION 3.1. *The group algebra $\mathbb{C}\mathcal{G}$ is the complex vector space $\mathbb{C}^{|\mathcal{G}|}$ where each $g \in \mathcal{G}$ is a basis vector. $\mathbb{C}\mathcal{G}$ is equipped with a convolution product defined below. A vector $a \in \mathbb{C}\mathcal{G}$ can be written as*

$$a = \sum_{g \in \mathcal{G}} a(g)g \quad \text{where } a(g) \in \mathbb{C}.$$

The convolution product $* : \mathbb{C}\mathcal{G} \times \mathbb{C}\mathcal{G} \rightarrow \mathbb{C}\mathcal{G}$ is induced from the product in \mathcal{G} as follows. For basis vectors $g, h \in \mathcal{G} \subset \mathbb{C}\mathcal{G}$, we set $g * h \equiv gh$, and in general if $a = \sum_{g \in \mathcal{G}} a(g)g$ and $b = \sum_{h \in \mathcal{G}} b(h)h$, then

$$a * b = \left(\sum_{g \in \mathcal{G}} a(g)g \right) * \left(\sum_{h \in \mathcal{G}} b(h)h \right) = \sum_{g, h \in \mathcal{G}} a(g)b(h)(g * h) = \sum_{g \in \mathcal{G}} (a * b)(g)g,$$

where

$$(3.2) \quad (a * b)(g) = \sum_{h \in \mathcal{G}} a(gh^{-1})b(h) = \sum_{h \in \mathcal{G}} a(h)b(h^{-1}g).$$

The convolution is associative, i.e., $a * (b * c) = (a * b) * c$, but generally not commutative, i.e., $a * b \neq b * a$.

Now consider a \mathcal{G} -equivariant $\mathbf{A} \in \mathbb{C}^{n \times n}$. Recall that when the the action is free and transitive, we have a 1-1 correspondance between the indices and the group, and we may identify \mathcal{I} with \mathcal{G} . Corresponding to \mathbf{A} there is a unique $A \in \mathbb{C}\mathcal{G}$, given as $A = \sum_{g \in \mathcal{G}} A(g)g$, where A is the first column of \mathbf{A} , i.e.,

$$(3.3) \quad A(gh^{-1}) = \mathbf{A}_{gh^{-1}, e} = \mathbf{A}_{g, h}.$$

Similarly, any vector $\mathbf{x} \in \mathbb{C}^n$ corresponds uniquely to $x = \sum_{g \in \mathcal{G}} x(g)g \in \mathbb{C}\mathcal{G}$, where $x(g) = \mathbf{x}_g$ for all $g \in \mathcal{G}$. Consider the matrix vector product:

$$(\mathbf{A}\mathbf{x})_g = \sum_{h \in \mathcal{G}} \mathbf{A}_{g, h} \mathbf{x}_h = \sum_{h \in \mathcal{G}} A(gh^{-1})x(h) = (A * x)(g).$$

If \mathbf{A} and \mathbf{B} are two equivariant matrices, then $\mathbf{A}\mathbf{B}$ is the equivariant matrix where the first column is given as

$$(\mathbf{A}\mathbf{B})_{g, e} = \sum_{h \in \mathcal{G}} \mathbf{A}_{g, h} \mathbf{B}_{h, e} = \sum_{h \in \mathcal{G}} A(gh^{-1})B(h) = (A * B)(g).$$

We have shown that if \mathcal{G} acts freely and transitively, then the algebra of \mathcal{G} -equivariant matrices acting on \mathbb{C}^n is isomorphic to the group algebra $\mathbb{C}\mathcal{G}$ acting on itself by convolutions from the left.

Also the action of the permutation matrices $P(g)$ on \mathbb{C}^n can be understood as convolutions in $\mathbb{C}\mathcal{G}$,

$$(P(g)\mathbf{x})_i = x(ig) = (x * g^{-1})(i),$$

i.e. convolutions from the right by g^{-1} . Note that $\mathbf{A}P(g) = P(g)\mathbf{A}$ follows from associativity of the convolution,

$$\mathbf{A}(P(g)\mathbf{x}) \leftrightarrow A * (x * g^{-1}) = (A * x) * g^{-1} \leftrightarrow P(g)(\mathbf{A}\mathbf{x}).$$

In fact, the reason for choosing a right group action $i \mapsto ig$ in (2.3) is that this corresponds to the equivariant matrices acting as convolutions from the left. Alternatively, one can define the group action from left, in which case left matrix multiplication will correspond to right convolutions in $\mathbb{C}\mathcal{G}$.

EXAMPLE 3.1. In the Triangle Example, a free transitive action would correspond to a discretization which only uses nodes 1 to 6 of Figure 2.1b. A \mathcal{D}_3 -equivariant $\mathbf{A} \in \mathbb{C}^{6 \times 6}$ maps to $A \in \mathbb{C}\mathcal{G}$ via (3.3). The vectors \mathbf{x} and $\mathbf{b} = \mathbf{A}\mathbf{x} \in \mathbb{C}^6$ correspond directly to x and $b = A * x \in \mathbb{C}\mathcal{G}$ via the identification of \mathcal{I} and \mathcal{G} , e.g., $b(g) = \mathbf{b}_g$.

3.2 Free non-transitive actions and block convolutions.

In the case where \mathbf{A} is \mathcal{G} -equivariant w.r.t. a free, but not transitive, action of \mathcal{G} on \mathcal{I} , we need a block version of the above theory.

Let $\mathbb{C}^{m \times \ell}\mathcal{G} \equiv \mathbb{C}^{m \times \ell} \otimes \mathbb{C}\mathcal{G}$ denote the space of vectors consisting of $|\mathcal{G}|$ matrix blocks, each block of size $m \times \ell$, thus $A \in \mathbb{C}^{m \times \ell}\mathcal{G}$ can be written as

$$(3.4) \quad A = \sum_{g \in \mathcal{G}} A(g) \otimes g \quad \text{where } A(g) \in \mathbb{C}^{m \times \ell}.$$

The convolution product (3.2) generalizes to a block convolution $*$: $\mathbb{C}^{m \times \ell}\mathcal{G} \times \mathbb{C}^{\ell \times k}\mathcal{G} \rightarrow \mathbb{C}^{m \times k}\mathcal{G}$ given as

$$A * B = \left(\sum_{g \in \mathcal{G}} A(g) \otimes g \right) * \left(\sum_{h \in \mathcal{G}} B(h) \otimes h \right) = \sum_{g \in \mathcal{G}} (A * B)(g) \otimes g,$$

where

$$(3.5) \quad (A * B)(g) = \sum_{h \in \mathcal{G}} A(gh^{-1})B(h) = \sum_{h \in \mathcal{G}} A(h)B(h^{-1}g),$$

and $A(h)B(h^{-1}g)$ denotes a matrix product.

We will establish an isomorphism between the algebra of \mathcal{G} -equivariant matrices acting on \mathbb{C}^n and the block-convolution algebra $\mathbb{C}^{m \times m}\mathcal{G}$ acting on $\mathbb{C}^m\mathcal{G}$.

We use λ in (3.1) to define the mappings (the subscript “f” stands for “free”; general actions will be treated below) $\mu_f : \mathbb{C}^n \rightarrow \mathbb{C}^m \mathcal{G}$, $\nu_f : \mathbb{C}^{n \times n} \rightarrow \mathbb{C}^{m \times m} \mathcal{G}$ as:

$$(3.6) \quad \mu_f(\mathbf{y})_i(g) = y_i(g) = \mathbf{y}_{ig} \quad \forall i \in \mathcal{S}, g \in \mathcal{G},$$

$$(3.7) \quad \nu_f(\mathbf{A})_{i,j}(g) = A_{i,j}(g) = \mathbf{A}_{ig,j} \quad \forall i, j \in \mathcal{S} \quad g \in \mathcal{G}.$$

Note the indexing convention: $\mathbf{y} \in \mathbb{C}^n$ corresponds to $y \in \mathbb{C}^m \mathcal{G}$, and we have $y(g) \in \mathbb{C}^m$, $y_i \in \mathbb{C} \mathcal{G}$, $y_i(g) \in \mathbb{C}$. A is indexed accordingly. The introduced mappings reflect the equivariance of \mathbf{A} since only a fraction of \mathbf{A} is stored in A ; the number of elements in \mathbf{A} is $n^2 = m^2 |\mathcal{G}|^2$ and the number of elements in A is $m^2 |\mathcal{G}|$.

PROPOSITION 3.1. *Let \mathcal{G} act freely on \mathcal{I} . Then μ_f is invertible and ν_f is invertible on the subspace of \mathcal{G} -equivariant matrices. Furthermore, if $\mathbf{A}, \mathbf{B} \in \mathbb{C}^{n \times n}$ are \mathcal{G} -equivariant, and $\mathbf{y} \in \mathbb{C}^n$, then*

$$(3.8) \quad \mu_f(\mathbf{A}\mathbf{y}) = \nu_f(\mathbf{A}) * \mu_f(\mathbf{y}),$$

$$(3.9) \quad \nu_f(\mathbf{A}\mathbf{B}) = \nu_f(\mathbf{A}) * \nu_f(\mathbf{B}).$$

PROOF. Invertability of μ_f follows from invertability of λ in (3.1). Invertability of ν_f for equivariant \mathbf{A} follows from

$$\mathbf{A}_{ig,jh} = \mathbf{A}_{igh^{-1},j} = \mathbf{A}_{i,j}(gh^{-1}).$$

Equation (3.8) follows from the following computation, where $i \in \mathcal{S}$ and $g \in \mathcal{G}$:

$$\begin{aligned} \mu_f(\mathbf{A}\mathbf{y})_i(g) &= (\mathbf{A}\mathbf{y})_{ig} = \sum_{j \in \mathcal{S}, h \in \mathcal{G}} \mathbf{A}_{ig,jh} \mathbf{y}_{jh} \\ &= \sum_{j \in \mathcal{S}, h \in \mathcal{G}} A_{i,j}(gh^{-1}) y_j(h) = (A * y)_i(g) = (\nu_f(\mathbf{A}) * \mu_f(\mathbf{y}))_i(g), \end{aligned}$$

and (3.9) follows from a similar computation. □

EXAMPLE 3.2. In the Triangle Example, a free non-transitive action would correspond to a discretization with no fixed points and with $m > 1$ orbits. A \mathcal{D}_3 -equivariant $\mathbf{A} \in \mathbb{C}^{6m \times 6m}$ maps to $A \in \mathbb{C}^{m \times m} \mathcal{G}$ via (3.7). The vectors \mathbf{x} and $\mathbf{b} = \mathbf{A}\mathbf{x} \in \mathbb{C}^{6m}$ correspond x and $b = A * x \in \mathbb{C}^m \mathcal{G}$ via (3.6).

3.3 Convolutions in the general case.

The proof of $\mu_f(\mathbf{A}\mathbf{x}) = \nu_f(\mathbf{A}) * \mu_f(\mathbf{x})$ above does not hold in the case of general actions, since the double sum $\sum_{j \in \mathcal{S}, h \in \mathcal{G}} \mathbf{A}_{ig,jh} \mathbf{y}_{jh}$ covers each orbit \mathcal{O}_j exactly $|\mathcal{G}_j|$ times. This is easily addressed by scaling the mappings ν_f and μ_f . One possible choice is

$$(3.10) \quad \mu(\mathbf{x})_i(g) = x_i(g) = \mathbf{x}_{ig} \quad \forall i \in \mathcal{S}, g \in \mathcal{G},$$

$$(3.11) \quad \nu(\mathbf{A})_{i,j}(g) = A_{i,j}(g) = \frac{1}{|\mathcal{G}_j|} \mathbf{A}_{ig,j} \quad \forall i, j \in \mathcal{S} \quad g \in \mathcal{G}.$$

Note that the general mappings ν and μ equals ν_f and μ_f , if the action is free. Now since

$$\sum_{h \in \mathcal{G}} \frac{1}{|\mathcal{G}_j|} \mathbf{A}_{ig,jh} \mathbf{x}_{hj} = \sum_{k \in \mathcal{O}_j} \frac{|\mathcal{G}_j|}{|\mathcal{G}_j|} \mathbf{A}_{ig,k} \mathbf{x}_k = \sum_{k \in \mathcal{O}_j} \mathbf{A}_{ig,k} \mathbf{x}_k$$

we obtain, with a computation similar to the proof of Proposition 3.1, the following result.

PROPOSITION 3.2. *For a general action it holds that*

$$(3.12) \quad \mu(\mathbf{A} \mathbf{x}) = \nu(\mathbf{A}) * \mu(\mathbf{x}),$$

$$(3.13) \quad \nu(\mathbf{A} \mathbf{B}) = \nu(\mathbf{A}) * \nu(\mathbf{B}).$$

However, if the action is not free, μ is no longer a bijection between \mathbb{C}^n and $\mathbb{C}^m \mathcal{G}$, because their dimensions differ. We introduce a subspace $V \subset \mathbb{C}^m \mathcal{G}$:

$$(3.14) \quad V = \{x \in \mathbb{C}^m \mathcal{G} \mid h * x_i = x_i, \forall i \in \mathcal{S}, h \in \mathcal{G}_i\}.$$

Note that $(h^{-1} * x_i)(g) = x_i(hg)$, thus V consists of those $x \in \mathbb{C}^m \mathcal{G}$ where each component $x_i \in \mathbb{C} \mathcal{G}$ is constant on the *right cosets* $\mathcal{G}_i g = \{hg \in \mathcal{G} \mid h \in \mathcal{G}_i\}$.

PROPOSITION 3.3. *The map*

$$\mu : \mathbb{C}^n \rightarrow V \subset \mathbb{C}^m \mathcal{G}$$

is a bijection.

PROOF. If $x = \mu(\mathbf{x})$ then $x_i(hg) = \mathbf{x}_{ihg} = \mathbf{x}_{ig} = x_i(g)$. Conversely, if $x \in V$ then $x = \mu(\mathbf{x})$ where $\mathbf{x}_{ig} = x_i(g)$. □

We also seek characterisations of the image of $\nu : \mathbb{C}^{n \times n} \rightarrow \mathbb{C}^{m \times m} \mathcal{G}$. Define the projection operator $\pi : \mathbb{C}^m \mathcal{G} \rightarrow V$ as

$$(3.15) \quad (\pi(x))_i(g) = \frac{1}{|\mathcal{G}_i|} \sum_{h \in \mathcal{G}_i} x_i(h^{-1}g) = \left(\left(\frac{1}{|\mathcal{G}_i|} \sum_{h \in \mathcal{G}_i} h \right) * x_i \right)(g),$$

thus π is averaging x_i over each coset $\mathcal{G}_i g$. This projection is studied in detail in Section 5.2, where it is shown to be the orthogonal projection on V with respect to the natural inner product on $\mathbb{C}^m \mathcal{G}$ given in (4.12).

PROPOSITION 3.4. *For $A \in \mathbb{C}^{m \times m} \mathcal{G}$, the following statements are equivalent:*

$$(3.16) \quad \mathbf{i)} \quad A = \nu(\mathbf{A}) \quad \text{for some } G\text{-equivariant } \mathbf{A} \in \mathbb{C}^{n \times n}.$$

$$(3.17) \quad \mathbf{ii)} \quad h * A_{i,j} * t = A_{i,j} \quad \text{for all } h \in \mathcal{G}_i, t \in \mathcal{G}_j.$$

$$(3.18) \quad \mathbf{iii)} \quad A = \pi A \pi.$$

PROOF. Assuming $A = \nu(\mathbf{A})$ we have for every $h \in \mathcal{G}_i$ and $t \in \mathcal{G}_j$

$$A_{i,j}(g) = \frac{1}{|\mathcal{G}_j|} \mathbf{A}_{ig,j} \equiv \frac{1}{|\mathcal{G}_j|} \mathbf{A}_{ihg,jt^{-1}} = A_{i,j}(hgt),$$

thus **i**) \Rightarrow **ii**). Note that

$$(3.19) \quad (\pi A \pi)_{i,j}(g) = \frac{1}{|\mathcal{G}_i||\mathcal{G}_j|} \sum_{h \in \mathcal{G}_i} \sum_{t \in \mathcal{G}_j} h * A_{i,j} * t,$$

thus **ii**) \Rightarrow **iii**). Assuming **iii**), we see from (3.19) that each component $A_{i,j} \in \mathbb{C}\mathcal{G}$ is constant on the *double cosets* $\mathcal{G}_i g \mathcal{G}_j = \{hgt \in \mathcal{G} \mid h \in \mathcal{G}_i, t \in \mathcal{G}_j\}$. Hence $A = \nu(\mathbf{A})$ where $\mathbf{A}_{ig,j} = |\mathcal{G}_j| A_{i,j}(g)$, thus **iii**) \Rightarrow **i**). \square

We have a decomposition $\mathbb{C}^m \mathcal{G} = V \oplus W$ where

$$W = \text{Image}(I - \pi) = \left\{ y \in \mathbb{C}^m \mathcal{G} \mid \sum_{h \in \mathcal{G}_i} h * x_i = 0, \forall i \in \mathcal{S} \right\}.$$

From Proposition 3.4 **iii**), we see that for $A = \nu(\mathbf{A})$ both V and W are invariant subspaces:

$$(3.20) \quad A * v \in V \quad \text{for every } v \in V,$$

$$(3.21) \quad A * w = 0 \quad \text{for every } w \in W.$$

It can be shown that $W = V^\perp$ in the inner product (4.12).

EXAMPLE 3.3. In the Triangle Example, the discretization in Figure 2.1b corresponds to a general action with three orbits, c.f. Example 2.1. A \mathcal{D}_3 -equivariant $\mathbf{A} \in \mathbb{C}^{10 \times 10}$ maps to $A \in \mathbb{C}^{3 \times 3} \mathcal{G}$ via (3.11). The vectors \mathbf{x} and $\mathbf{b} = \mathbf{A}\mathbf{x} \in \mathbb{C}^{10}$ correspond x and $b = A * x \in \mathbb{C}^3 \mathcal{G}$ via (3.10). Note that the image V of μ is a *pure* subspace of $\mathbb{C}^3 \mathcal{G}$, due to the fixed points.

4 The Generalized Fourier Transform.

4.1 Finite noncommutative groups.

In the introductory discussion on circulant matrices, we saw that the shift matrix eigenfunctions χ^k diagonalize convolutions because they satisfy (1.3). A noncommutative group \mathcal{G} has in general too few common eigenvectors for the permutation matrices $\{P(g) \mid g \in \mathcal{G}\}$ to obtain a basis for $\mathbb{C}\mathcal{G}$. To circumvent this problem, Frobenius invented the concept of *group representations*. We recommend the monographs [18, 22, 25, 33] as introductions to representation theory with applications.

DEFINITION 4.1. A *d-dimensional group representation* is a map $R : \mathcal{G} \rightarrow \mathbb{C}^{d \times d}$ such that

$$(4.1) \quad R(gh) = R(g)R(h) \quad \text{for all } g, h \in \mathcal{G}.$$

Generalizing the definition of *Fourier coefficients* (1.2) we define for any $A \in \mathbb{C}^{m \times k} \mathcal{G}$ and any d -dimensional representation R a matrix $\hat{A}(R) \in \mathbb{C}^{m \times k} \otimes \mathbb{C}^{d \times d}$ as:

$$(4.2) \quad \hat{A}(R) = \sum_{g \in \mathcal{G}} A(g) \otimes R(g).$$

PROPOSITION 4.1 (THE CONVOLUTION THEOREM). *For any $A \in \mathbb{C}^{m \times k} \mathcal{G}$, $B \in \mathbb{C}^{k \times \ell} \mathcal{G}$ and any representation R we have*

$$(4.3) \quad \widehat{(A * B)}(R) = \hat{A}(R) \hat{B}(R).$$

PROOF. The statement follows from

$$\begin{aligned} \hat{A}(R) \hat{B}(R) &= \left(\sum_{g \in \mathcal{G}} A(g) \otimes R(g) \right) \left(\sum_{h \in \mathcal{G}} B(h) \otimes R(h) \right) \\ &= \sum_{g, h \in \mathcal{G}} A(g) B(h) \otimes R(g) R(h) = \sum_{g, h \in \mathcal{G}} A(g) B(h) \otimes R(gh) \\ &= \sum_{g, h \in \mathcal{G}} A(gh^{-1}) B(h) \otimes R(g) = \widehat{(A * B)}(R). \end{aligned}$$

□

Let d_R denote the dimension of the representation. For use in practical computations, it is important that $A * B$ can be recovered by knowing $\widehat{(A * B)}(R)$ for a suitable selection of representations, and furthermore that their dimensions d_R are as small as possible. Note that if R is a representation and $X \in \mathbb{C}^{d_R \times d_R}$ is non-singular, then also $\tilde{R}(g) = XR(g)X^{-1}$ is a representation. We say that R and \tilde{R} are equivalent representations. If there exists a similarity transform $\tilde{R}(g) = XR(g)X^{-1}$ such that $\tilde{R}(g)$ has a block diagonal structure, independent of $g \in \mathcal{G}$, then R is called *reducible*, otherwise it is *irreducible*.

Luckily, any group possesses exactly the necessary number of irreducible representations to provide a basis for $\mathbb{C}\mathcal{G}$:

THEOREM 4.2 (FROBENIUS). *For any finite group \mathcal{G} there exists a complete list \mathcal{R} of non-equivalent irreducible representations such that*

$$\sum_{R \in \mathcal{R}} d_R^2 = |\mathcal{G}|.$$

Defining the GFT for $a \in \mathbb{C}\mathcal{G}$ as

$$(4.4) \quad \hat{a}(R) = \sum_{g \in \mathcal{G}} a(g) R(g) \quad \text{for every } R \in \mathcal{R},$$

we may recover a by the inverse GFT (IGFT):

$$(4.5) \quad a(g) = \frac{1}{|\mathcal{G}|} \sum_{R \in \mathcal{R}} d_R \text{trace}(R(g^{-1}) \hat{a}(R)).$$

Complete lists of irreducible representations for a selection of common groups are found in Appendix A.

The list \mathcal{R} is not unique, but the number of representations in \mathcal{R} and their sizes d_R are invariants of \mathcal{G} . It is important to note that the list \mathcal{R} can always be chosen such that $R(g)$ are *unitary matrices* for all $R \in \mathcal{R}$ and all $g \in \mathcal{G}$. In this case the GFT can be seen as an orthogonal transform. Define an inner product on $\mathbb{C}\mathcal{G}$ as

$$(4.6) \quad \langle a, b \rangle = \sum_{g \in \mathcal{G}} a(g)^H b(g),$$

where a^H denotes complex conjugate for scalars and complex conjugate transpose for matrices. Let

$$\widehat{\mathbb{C}\mathcal{G}} = \bigoplus_{R \in \mathcal{R}} \mathbb{C}^{d_R \times d_R}$$

be the space of Fourier coefficients, equipped with the inner product

$$(4.7) \quad \langle \hat{a}, \hat{b} \rangle = \sum_{R \in \mathcal{R}} \frac{d_R}{|\mathcal{G}|} \text{trace}(\hat{a}(R)^H \hat{b}(R)).$$

LEMMA 4.3. *If \mathcal{R} is unitary then $\text{GFT} : \mathbb{C}\mathcal{G} \rightarrow \widehat{\mathbb{C}\mathcal{G}}$ is orthogonal,*

$$\langle a, b \rangle = \langle \hat{a}, \hat{b} \rangle \quad \text{for all } a, b \in \mathbb{C}\mathcal{G}.$$

PROOF. For $g, h \in \mathcal{G} \subset \mathbb{C}\mathcal{G}$ we have $\langle g, h \rangle = \delta_{g,h}$. Note that $\hat{g}(R) = R(g)$ and in particular $\hat{e}(R) = I_{d_R}$. From unitarity we have $R(g)^H = R(g^{-1})$, thus

$$\begin{aligned} \langle \hat{g}, \hat{h} \rangle &= \sum_{R \in \mathcal{R}} \frac{d_R}{|\mathcal{G}|} \text{trace}(R(g)^H R(h)) \\ &= \sum_{R \in \mathcal{R}} \frac{d_R}{|\mathcal{G}|} \text{trace}(R(g^{-1}h)) = \text{IGFT}(\hat{e})(hg^{-1}) = \delta_{g,h}. \end{aligned}$$

Thus the GFT sends an orthonormal basis to an orthonormal basis. □

Returning to the the block transform of $A \in \mathbb{C}^{m \times k} \mathcal{G}$ given in (4.2), we see that the GFT and the IGFT are given component-wise as

$$(4.8) \quad \hat{A}_{i,j}(R) = \sum_{g \in \mathcal{G}} A_{i,j}(g) R(g) \in \mathbb{C}^{d_R \times d_R},$$

$$(4.9) \quad A_{i,j}(g) = \frac{1}{|\mathcal{G}|} \sum_{R \in \mathcal{R}} d_R \text{trace}(R(g^{-1}) \hat{A}_{i,j}(R)).$$

Recall that $\mathbb{C}^{m \times k} \mathcal{G} \equiv \mathbb{C}^{m \times k} \otimes \mathbb{C}\mathcal{G}$, similarly we define $\widehat{\mathbb{C}^{m \times k} \mathcal{G}} \equiv \mathbb{C}^{m \times k} \otimes \widehat{\mathbb{C}\mathcal{G}}$. Thus the block transforms (i)GFT_b relate to the scalar transforms (i)GFT_s as

$$(4.10) \quad \text{GFT}_b = (I \otimes \text{GFT}_s) : \mathbb{C}^{m \times k} \otimes \mathbb{C}\mathcal{G} \rightarrow \mathbb{C}^{m \times k} \otimes \widehat{\mathbb{C}\mathcal{G}},$$

$$(4.11) \quad \text{IGFT}_b = (I \otimes \text{IGFT}_s) : \mathbb{C}^{m \times k} \otimes \widehat{\mathbb{C}\mathcal{G}} \rightarrow \mathbb{C}^{m \times k} \otimes \mathbb{C}\mathcal{G}.$$

Similar to Lemma 4.3 it can be shown that if the representations are unitary, then the block transform GFT_b is orthogonal with respect to the following inner products

$$(4.12) \quad \langle A, B \rangle = \sum_{g \in \mathcal{G}} \text{trace}(A(g)^H B(g)) \quad \text{for } A, B \in \mathbb{C}^{m \times k} \mathcal{G},$$

$$(4.13) \quad \langle \hat{A}, \hat{B} \rangle = \sum_{R \in \mathcal{R}} \frac{d_R}{|\mathcal{G}|} \text{trace}(\hat{A}(R)^H \hat{B}(R)) \quad \text{for } \hat{A}, \hat{B} \in \widehat{\mathbb{C}^{m \times k} \mathcal{G}}.$$

The orthogonality of the GFT is important for computational algorithms, since this ensures that the GFT is numerically well-conditioned.

We conclude this section with a noncommutative analog of the important basic result of commutative Fourier analysis; that a 2D transform can be computed by doing 1-D transforms first on the rows and subsequently on the columns of the dataset. The corresponding general statement ([23]) concerns the *direct product* of two groups \mathcal{G}_1 and \mathcal{G}_2 , defined as the set of all pairs

$$\mathcal{G} = \mathcal{G}_1 \times \mathcal{G}_2 = \{(g_1, g_2) \mid g_1 \in \mathcal{G}_1, g_2 \in \mathcal{G}_2\},$$

where the product is componentwise $(g_1, g_2)(h_1, h_2) = (g_1 h_1, g_2 h_2)$.

PROPOSITION 4.4. *Let \mathcal{R}_1 and \mathcal{R}_2 be complete lists of irreducible representations for \mathcal{G}_1 and \mathcal{G}_2 . Then a complete list for $\mathcal{G} = \mathcal{G}_1 \times \mathcal{G}_2$ is*

$$\mathcal{R} = \{R_1 \otimes R_2 \mid R_1 \in \mathcal{R}_1, R_2 \in \mathcal{R}_2\},$$

where $(R_1 \otimes R_2)(g_1, g_2) = R_1(g_1) \otimes R_2(g_2)$. The GFT of $a \in \mathbb{C} \mathcal{G}$ can be computed as

$$\begin{aligned} \hat{a}(R_1 \otimes R_2) &= \sum_{(g_1, g_2) \in \mathcal{G}} a(g_1, g_2) R_1(g_1) \otimes R_2(g_2) \\ &= \sum_{g_2 \in \mathcal{G}_2} \left(\sum_{g_1 \in \mathcal{G}_1} a(g_1, g_2) R_1(g_1) \otimes I \right) (I \otimes R_2(g_2)). \end{aligned}$$

The fact that the GFT on any direct product of groups splits nicely into a ‘first rows then columns’ type transform is important for decreasing the computational complexity of the transform, see Section 5.4.

4.2 Commutative groups.

As a reference we will briefly summarize the Fourier analysis on finite commutative (abelian) groups. Let $\langle \mathbb{Z}_n, + \rangle$ denote the group of integers $\{0, 1, \dots, n - 1\}$ under addition modulo n . A basic result of algebra states that any finite commutative group is isomorphic to a direct product

$$(4.14) \quad \mathcal{G} = \mathbb{Z}_{n_1} \times \mathbb{Z}_{n_2} \times \dots \times \mathbb{Z}_{n_d}.$$

Thus elements $g \in \mathcal{G}$ are represented as $g = (g_1, \dots, g_d)$, where $0 \leq g_i < n_i$, the group operation is

$$g + h = (g_1 + h_1 \bmod n_1, \dots, g_d + h_d \bmod n_d)$$

and the inverse operation is $-g = (-g_1 \bmod n_1, \dots, -g_d \bmod n_d)$. This additive representation has an equivalent multiplicative representation. Let $\mathcal{C}_n = \{\exp(2\pi i j/n) : j \in \mathbb{Z}_n\}$ denote the cyclic group, where the group operation is product of complex numbers. The exponential map is an isomorphism between \mathbb{Z}_n and \mathcal{C}_n . It is common to use \mathcal{C}_n to denote commutative subgroups of noncommutative groups, and \mathbb{Z}_n when non-commutativity is not an issue.

It is well-known that all the irreducible representations of any commutative group \mathcal{G} are 1-dimensional. We write them as $\mathcal{R} \equiv \hat{\mathcal{G}} = \{\chi^k\}$. It might be remarked that also $\hat{\mathcal{G}}$ has the natural structure of a commutative group such that $\chi^k \cdot \chi^\ell = \chi^{k+\ell}$. This is called the dual group of \mathcal{G} . The finite case is particularly simple in which case $\hat{\mathcal{G}} = \mathcal{G}$. Thus $\hat{\mathcal{G}} = \{\chi^k\}_{k \in \mathcal{G}}$, where

$$\chi^k(g) = \exp(2\pi i(k_1 g_1/n_1 + k_2 g_2/n_2 + \dots + k_d g_d/n_d)) \quad \text{for all } k, g \in \mathcal{G}.$$

The (scalar and block) finite commutative version of the Fourier and inverse transforms (4.4)–(4.5) is called the Discrete Fourier Transform:

$$(4.15) \quad \hat{a}(k) = \sum_{g \in \mathcal{G}} a(g) \chi^k(g),$$

$$(4.16) \quad a(g) = \frac{1}{|\mathcal{G}|} \sum_{k \in \mathcal{G}} \hat{a}(k) \chi^k(-g).$$

For $A \in \mathbb{C}^{m \times m} \mathcal{G}$, these can be computed in $\mathcal{O}(m^2 n \log(n))$ operations via the Fast Fourier Transform, where $n = |\mathcal{G}|$. A group theoretical explanation of the FFT algorithm is based on group duality theory, but will not be discussed here.

5 Block diagonalization and numerical algorithms.

We have seen that a \mathcal{G} -equivariant matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ can be represented in terms of a convolutional operator $A \in \mathbb{C}^{m \times m} \mathcal{G}$, where m is the number of orbits in the index set \mathcal{I} under the action of \mathcal{G} . Via the GFT, A can again be transformed into \hat{A} which is a block matrix with blocks $\hat{A}(R) \in \mathbb{C}^{md_R \times md_R}$ for each of the irreducible representations $R \in \mathcal{R}$. There is a number of computational algorithms that may benefit from this block diagonalization of A , e.g.:

- Solving linear systems $\mathbf{A}\mathbf{x} = \mathbf{b}$ by direct or iterative methods.
- Computing matrix functions, e.g. $\mathbf{A} \mapsto \exp(\mathbf{A})$.
- The matrix exponential applied to a vector $\mathbf{x} \mapsto \exp(\mathbf{A})\mathbf{x}$.
- Computing some or all of the eigenvalues and eigenvectors of \mathbf{A} .

We will in this paper mainly focus on the direct solution of linear systems, but also briefly comment on other matrix computations.

5.1 Solving linear systems; free actions.

If \mathcal{G} acts freely on \mathcal{I} then an equivariant \mathbf{A} is isomorphic to a convolutional operator A , which again is isomorphic to $\hat{A} = \text{GFT}(A)$, as illustrated in Figure 5.1.

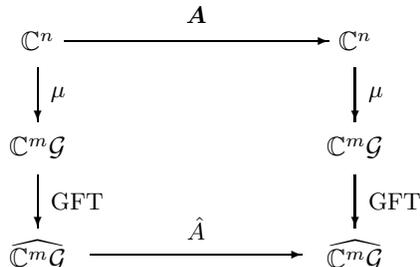


Figure 5.1: The diagram commutes, so $\text{GFT}(\mu(\mathbf{A} \mathbf{x})) = \hat{A} \text{GFT}(\mu(\mathbf{x}))$.

As an example, consider solution of $\mathbf{A} \mathbf{x} = \mathbf{b}$. Using μ and ν , this equation transforms to $A * x = b$, and the GFT further transforms it into a block-diagonal matrix equation:

$$\hat{A}(R) \hat{x}(R) = \hat{b}(R), \quad \text{for every } R \in \mathcal{R},$$

where $\hat{A}(R) \in \mathbb{C}^{md_R \times md_R}$, and both $\hat{x}(R), \hat{b}(R) \in \mathbb{C}^{md_R \times m}$.

In summary, when \mathbf{A} is equivariant w.r.t. a free group action, the GFT can be exploited to solve $\mathbf{A} \mathbf{x} = \mathbf{b}$ as follows:

1. Represent \mathbf{A} as A and \mathbf{b} as b .
2. Fourier transform: $\hat{A} = \text{GFT}(A), \hat{b} = \text{GFT}(b)$.
3. For each $R \in \mathcal{R}$, solve $\hat{A}(R) \hat{x}(R) = \hat{b}(R)$.
4. Inverse transform the solution: $x = \text{IGFT}(\hat{x})$.
5. Finally, construct \mathbf{x} from x .

Note that the mappings $\mathbf{x} \leftrightarrow x$ etc. only permute data, and they can therefore be incorporated in the transforms.

5.2 Solving linear systems; general actions.

In Section 3.3 we saw that if the action is not free, then the map $\mu : \mathbb{C}^n \rightarrow \mathbb{C}^m \mathcal{G}$ is not surjective, and hence the convolutional operator $A = \nu(\mathbf{A})$ becomes singular. In Proposition 3.3 we characterized the subspace $V \subset \mathbb{C}^m \mathcal{G}$ on which μ is an isomorphism.

To solve linear systems we seek characterizations of V and $\hat{V} = \text{GFT}(V)$ in terms of orthogonal projections, and for this purpose we study functions in $\mathbb{C} \mathcal{G}$ invariant under subgroup actions.

Let \mathcal{H} be a subgroup of \mathcal{G} . Consider the subspace $\mathbb{C} \mathcal{G} | \mathcal{H} \subset \mathbb{C} \mathcal{G}$ consisting of all functions invariant under the left action of \mathcal{H} on \mathcal{G} :

DEFINITION 5.1. *Let*

$$\mathbb{C}\mathcal{G}|\mathcal{H} = \{x \in \mathbb{C}\mathcal{G} \mid x(hg) = x(g) \ \forall h \in \mathcal{H}, g \in \mathcal{G}\}.$$

Since

$$(5.1) \quad x(h^{-1}g) = (h * x)(g),$$

an equivalent characterization is

$$(5.2) \quad \mathbb{C}\mathcal{G}|\mathcal{H} = \{x \in \mathbb{C}\mathcal{G} \mid h * x = x, \ \forall h \in \mathcal{H}\}.$$

Associativity $h * (x * y) = (h * x) * y$ yields:

PROPOSITION 5.1. *$\mathbb{C}\mathcal{G}|\mathcal{H}$ is a right ideal of $\mathbb{C}\mathcal{G}$, i.e. it is a linear subspace such that*

$$x * y \in \mathbb{C}\mathcal{G}|\mathcal{H} \quad \text{for all } x \in \mathbb{C}\mathcal{G}|\mathcal{H} \text{ and } y \in \mathbb{C}\mathcal{G}.$$

PROPOSITION 5.2. *The operator $\pi = \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} h$, acting on $\mathbb{C}\mathcal{G}$ as*

$$\pi y = \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} h * y$$

is an orthogonal projection onto $\mathbb{C}\mathcal{G}|\mathcal{H}$.

PROOF. We see from (5.1) that π is averaging y over each orbit of \mathcal{H} , from which it follows that it is a projection. To see orthogonality, we compute the components of π with respect to the orthogonal basis spanned $\mathcal{G} \subset \mathbb{C}\mathcal{G}$:

$$\pi_{g,t} = \langle g, \pi t \rangle = \begin{cases} \frac{1}{|\mathcal{H}|} & \text{if } g \text{ and } t \text{ are in the same } \mathcal{H}\text{-orbit} \\ 0 & \text{otherwise.} \end{cases}$$

Thus $\pi_{g,t} = \pi_{t,g}$. □

In Fourier space we characterize \mathcal{H} -invariant functions as follows:

THEOREM 5.3. *The image $\text{GFT}(\mathbb{C}\mathcal{G}|\mathcal{H}) = \widehat{\mathbb{C}\mathcal{G}|\mathcal{H}}$ is a right ideal given by*

$$\widehat{\mathbb{C}\mathcal{G}|\mathcal{H}} = \{\hat{x} \in \widehat{\mathbb{C}\mathcal{G}} \mid R(h)\hat{x}(R) = \hat{x}(R), \ \forall h \in \mathcal{H}, R \in \mathcal{R}\}.$$

The projection $\hat{\pi} : \widehat{\mathbb{C}\mathcal{G}} \rightarrow \widehat{\mathbb{C}\mathcal{G}|\mathcal{H}}$ is defined for an arbitrary $\hat{y} \in \widehat{\mathbb{C}\mathcal{G}}$ as

$$(5.3) \quad (\hat{\pi}\hat{y})(R) = \hat{\pi}(R)\hat{y}(R) = \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} R(h)\hat{y}(R), \ R \in \mathcal{R}.$$

PROOF. Since the GFT is an algebra isomorphism and since $\mathbb{C}\mathcal{G}|\mathcal{H}$ is a right ideal, the first statement follows immediately. Since $\hat{h}(R) = R(h)$, equation (5.2) yields the description of $\widehat{\mathbb{C}\mathcal{G}|\mathcal{H}}$ as well as of $\hat{\pi}$. □

Let $r_R = \text{rank}(\hat{\pi}(R)) = \text{trace}(\hat{\pi}(R))$. To compute a basis for $\widehat{\mathbb{C}\mathcal{G}}|\mathcal{H}$ we factor $\hat{\pi}(R)$ into matrices $Q(R)$ and $P(R)$ of dimension $r_R \times d_R$ and $d_R \times r_R$ such that

$$(5.4) \quad \hat{\pi}(R) = P(R)Q(R),$$

$$(5.5) \quad Q(R)P(R) = I.$$

Thus P defines a basis for $\widehat{\mathbb{C}\mathcal{G}}|\mathcal{H}$. Note that since $\hat{\pi}$ is an orthogonal projection, we may let $Q(R) = P(R)^H$, but other factorizations can also be used. Let \hat{x}' denote $\hat{x} \in \widehat{\mathbb{C}\mathcal{G}}|\mathcal{H}$ expressed in terms of the basis given by P . Thus

$$\begin{aligned} \hat{x}'(R) &= Q(R)\hat{x}(R), \\ \hat{x}(R) &= P(R)\hat{x}'(R). \end{aligned}$$

We seek an efficient computation of $x \mapsto \hat{x}'$ for $x \in \mathbb{C}\mathcal{G}|\mathcal{H}$. Recall that a subgroup \mathcal{H} partitions \mathcal{G} into disjoint subsets $\mathcal{H}g = \{hg \in \mathcal{G} \mid h \in \mathcal{H}\}$ called the (right) *cosets*. Let $\mathcal{C} \subset \mathcal{G}$ denote a selection of coset representatives (one element from each coset). Note that if $x \in \mathbb{C}\mathcal{G}|\mathcal{H}$ then x takes a constant value on each coset, $x(\mathcal{H}g) = x(g)$.

PROPOSITION 5.4. *For $x \in \mathbb{C}\mathcal{G}|\mathcal{H}$ we have*

$$(5.6) \quad \hat{x}'(R) = \sum_{g \in \mathcal{C}} x(\mathcal{H}g)R'(g)$$

where $R'(g) = |\mathcal{H}|Q(R)R(g)$.

PROOF. We have that

$$\begin{aligned} \hat{x}(R) &= \sum_{h \in \mathcal{H}, g \in \mathcal{C}} x(hg)R(hg) = \sum_{h \in \mathcal{H}} R(h) \sum_{g \in \mathcal{C}} x(g)R(g) \\ &= |\mathcal{H}|\hat{\pi}(R) \sum_{g \in \mathcal{C}} x(g)R(g) = |\mathcal{H}|P(R) \sum_{g \in \mathcal{C}} x(g)Q(R)R(g), \end{aligned}$$

from which the result follows. □

We return to the question of solving linear systems $\mathbf{A}\mathbf{x} = \mathbf{b}$ for \mathcal{G} -equivariant \mathbf{A} , and consider the following alternatives:

1. Iterative methods.
2. Regularization via projections.
3. Direct methods using a basis for V .

If iterative methods are employed, the singular part of \hat{A} is not a problem. Since \hat{A} maps V to V , and $\hat{x}_0 = \text{GFT}(\mu(\mathbf{x}_0)) \in V$, we may simply employ standard iterative techniques on $\hat{A}(R)\hat{x}(R) = \hat{b}(R)$.

Regularization is considered by [6]. Briefly, the idea is as follows. If \hat{A} is regular on \hat{V} and $\hat{\pi}$ is the projection onto \hat{V} , then we may replace $\hat{A}(R)$ with the regularized matrix $\hat{A}(R) + \epsilon(I - \hat{\pi}(R))$, where ϵ is a regularization parameter.

Note that if $\hat{x} \in V$ then $(I - \hat{\pi}(R))\hat{x}(R) = 0$, hence we may replace the system with

$$(\hat{A}(R) + \epsilon(I - \hat{\pi}(R)))\hat{x}(R) = \hat{b}(R).$$

For sufficiently small ϵ , the regularized matrix is non-singular on $\widehat{\mathbb{C}^m \mathcal{G}}$. Besides this [6] also considers a projection method that results in a similarity transform of the equivariant matrix.

In our numerical examples, we have concentrated on alternative 3; direct methods using a basis for the space V defined in (3.14). This is motivated by the fact that \mathbf{A} is dense in many applications where equivariance may be exploited, for example boundary element method applications. Note that $x \in V$ is characterized by invariance of each component $x_i \in \mathbb{C}\mathcal{G}$ under the subgroup action of \mathcal{G}_i . Let $\pi_i : \mathbb{C}\mathcal{G} \rightarrow \mathbb{C}\mathcal{G}$ denote the projection onto $\mathbb{C}\mathcal{G}|_{\mathcal{G}_i}$ defined in Proposition 5.2, and let $\hat{\pi}_i(R)$, $P_i(R)$ and $Q_i(R)$ be given as

$$\hat{\pi}_i(R) = \frac{1}{|\mathcal{G}_i|} \sum_{h \in \mathcal{G}_i} R(h) = P_i(R)Q_i(R) \quad \text{where } Q_i(R)P_i(R) = I.$$

By insertion we verify the following result:

THEOREM 5.5. *Solving the equivariant system $\mathbf{Ax} = \mathbf{b}$ is equivalent to solving the block diagonal system $\hat{A}(R)\hat{x}(R) = \hat{b}(R)$ on the subspace V . Expressed in terms of the local basis for V given by $P_i(R)$, this becomes*

$$(5.7) \quad \sum_{j \in \mathcal{S}} \hat{A}''_{i,j}(R)\hat{x}'_j(R) = \hat{b}'_i(R), \quad \text{for } i \in \mathcal{S},$$

where

$$(5.8) \quad \hat{b}'_i(R) = Q_i(R)\hat{b}_i(R),$$

$$(5.9) \quad \hat{A}''_{i,j} = Q_i(R)\hat{A}_{i,j}(R)P_j(R),$$

$$(5.10) \quad \hat{x}'_j(R) = P_j(R)\hat{x}_j(R).$$

For each $R \in \mathcal{R}$, \hat{A}'' in (5.7) has size $r_R \times r_R$, where $r_R = \sum_{i \in \mathcal{S}} \text{rank}(\hat{\pi}_i(R))$.

The solution of a \mathcal{G} -equivariant system $\mathbf{Ax} = \mathbf{b}$ is summarized as:

1. Formulate $\mathbf{Ax} = \mathbf{b}$ as a convolution: compute $A = \nu(\mathbf{A})$ and $b = \mu(\mathbf{b})$.
2. Apply the GFT: Compute \hat{A}'' and \hat{b}' .
3. Solve (5.7) for each $R \in \mathcal{R}$.
4. Inverse transform the solution: $x = \text{IGFT}(\hat{x})$.
5. Finally, construct \mathbf{x} from x .

Note that the computations $b \mapsto \hat{b}'$ and $A \mapsto \hat{A}''$ can be done more efficiently as in Proposition 5.4. For $b_i, A_{i,j} \in \mathbb{C}\mathcal{G}$ we obtain

$$\hat{b}'_i = \sum_{g \in \mathcal{C}'} b_i(\mathcal{G}_i g)R'_i(g),$$

$$\hat{A}''_{i,j} = \sum_{g \in \mathcal{C}'} A_{i,j}(\mathcal{G}_i g)R'_i(g)P_j(R),$$

where $\mathcal{C}' \subset \mathcal{G}$ represents the cosets $\mathcal{G}_i g$, and where $R'_i(g) = |\mathcal{G}_i| Q_i(R) R(g)$. For \hat{A}'' , we can simplify further, since $A_{i,j} \in \mathbb{C}\mathcal{G}$ takes constant values on each of the *double cosets* $\mathcal{G}_i g \mathcal{G}_j \subset \mathcal{G}$ defined as

$$\mathcal{G}_i g \mathcal{G}_j = \{h g k \mid h \in \mathcal{G}_i, k \in \mathcal{G}_j\}.$$

This relation partitions \mathcal{C}' into disjoint sets $\mathcal{D}(g) = \mathcal{C}' \cap \mathcal{G}_i g \mathcal{G}_j$, where $A_{i,j}(\mathcal{D}(g))$ is constant. Let $\mathcal{C}'' \subset \mathcal{C}'$ represent this partition, so $\cup_{g \in \mathcal{C}''} \mathcal{D}(g) = \mathcal{C}'$. By rewriting the summation above we obtain

$$\hat{A}''_{i,j} = \sum_{g \in \mathcal{C}''} A_{i,j}(\mathcal{G}_i g \mathcal{G}_j) R''_{i,j}(g),$$

where $R''_{i,j}(g) = \sum_{h \in \mathcal{D}(g)} R'_i(h) P_j(R)$.

5.3 Other numerical algorithms and issues.

We briefly comment on various issues regarding symmetry based GFT diagonalization of an equivariant \mathbf{A} for various linear algebra computations.

Eigenvalue computations are discussed in [4], some details and numerical results are found in [1]. Commutativity of the diagram in Figure 5.1 yields

$$\mathbf{A} \mathbf{x} = \lambda \mathbf{x} \Rightarrow \mathbf{A} * \mathbf{x} = \lambda \mathbf{x} \Rightarrow \hat{A}(R) \hat{x}(R) = \lambda \hat{x}(R) \quad \text{for all } R \in \mathcal{R}.$$

Thus all eigenvalues of \mathbf{A} are also found as eigenvalues for \hat{A} . In the case of non-free actions, we saw in (3.21) that A has an invariant eigenspace $W = V^\perp$ with eigenvalue 0, and hence \hat{A} has an invariant eigenspace \hat{W} with eigenvalue 0. This eigenspace does not correspond to eigenvectors of the original \mathbf{A} . It is, however, not difficult to remove the spurious 0 eigenvalues in \hat{A} , since these correspond exactly to the kernel of $\hat{\pi}$. Alternatively, we may work directly with a representation of \hat{A} in terms of a local basis on V . In this case the spurious 0 eigenvalues will not appear.

For computation of **matrix exponentials**, we note that

$$\exp(\hat{A}) \circ \text{GFT} \circ \mu = \text{GFT} \circ \mu \circ \exp(\mathbf{A}),$$

thus $\exp(\mathbf{A})$ is recovered from $\exp(\hat{A})$ by the IGFT and inverting ν .

For any matrix *iterative* algorithm for computing eigenvalues or solving linear systems, we remark that since $\hat{A} : \hat{V} \rightarrow \hat{V}$, the GFT based algorithms work fine both in the free and the non-free case. For example, in iterative solution of $\mathbf{A} \mathbf{x} = \mathbf{b}$, we have that $\hat{\mathbf{b}} \in \hat{V}$, thus an iterative algorithm produces $\hat{\mathbf{x}} \in \hat{V}$.

Almost equivariant matrices. In circulant matrix theory, it is well known that any matrix can be approximated by a circulant by averaging over the circulant diagonals. The approximation is optimal in the Frobenius norm. This has applications in circulant preconditioners. Similar results hold in the more general case. By averaging the original matrix over the \mathcal{G} -orbits, one obtains the best Frobenius norm approximation in the space of equivariant matrices, see [35].

Partially equivariant matrices. Bonnet [8] discusses applications where the domain consists of two disconnected components where only one of the components is symmetric. The resulting matrix will then be partially equivariant, and this part may be block-diagonalized.

5.4 Computational complexity.

Consider a free action of \mathcal{G} on \mathcal{I} , which partitions \mathcal{I} into m orbits. Let W_{direct} denote the computational work, in terms of floating point operations, for performing a given linear algebra algorithm on the original data \mathbf{A}, \mathbf{b} , etc., and let W_{fspace} be the cost of doing the same algorithm on the corresponding block diagonal GFT transformed data \hat{A}, \hat{b} etc.

First we consider algorithms of complexity $\mathcal{O}(n^3)$, such as solving linear systems and computing eigenvalues. Thus $W_{\text{direct}} = c(m|\mathcal{G}|)^3 = cm^3 (\sum_{R \in \mathcal{R}} d_R^2)^3$, $W_{\text{fspace}} = cm^3 \sum_{R \in \mathcal{R}} d_R^3$ and the ratio becomes

$$\mathcal{O}(n^3) : W_{\text{direct}}/W_{\text{fspace}} = \left(\sum_{R \in \mathcal{R}} d_R^2 \right)^3 / \sum_{R \in \mathcal{R}} d_R^3.$$

As an illustration, Table 5.1 tabulates this factor for the symmetries of the triangle, the tetrahedron, the 3D cube and the maximally symmetric discretization of a 3D sphere (icosahedral symmetry with reflections).

Table 5.1: Gain in computational complexity depending on algorithm type for a few typical symmetry groups.

Domain	\mathcal{G}	$ \mathcal{G} $	$\{d_R\}_{R \in \mathcal{R}}$	$W_{\text{direct}}/W_{\text{fspace}}$	
				$\mathcal{O}(n^2)$	$\mathcal{O}(n^3)$
triangle	\mathcal{D}_3	6	{1, 1, 2}	3.6	21.6
tetrahedron	\mathcal{S}_4	24	{1, 1, 2, 3, 3}	9	216
cube	$\mathcal{S}_4 \times \mathcal{C}_2$	48	{1, 1, 1, 1, 2, 2, 3, 3, 3, 3}	18	864
icosahedron	$\mathcal{A}_5 \times \mathcal{C}_2$	120	{1, 1, 3, 3, 3, 3, 4, 4, 5, 5}	29.5	3541

The computational complexity of the GFT is an interesting open problem. Let $W_{\text{GFT}}(\mathbb{C}\mathcal{G})$ denote the cost of the GFT. We say that \mathcal{G} allows a *fast transform* if there exists an algorithm such that $W_{\text{GFT}}(\mathbb{C}\mathcal{G}) = |\mathcal{G}| \log^c(|\mathcal{G}|)$ for some c . It is conjectured but not proven [32] that any group allows a fast transform. It is known that e.g. all the symmetric groups \mathcal{S}_n allow fast transforms. The history of fast GFT algorithms goes back to Karpovsky [23], and there has been a significant development of fast algorithms in recent years [7, 11, 12, 13, 15, 27, 28, 26, 30, 31]. A *slow transform* is a direct computation of the GFT (4.4). It is straightforward to verify that the cost of the slow transform (counting real flops) is:

$$(5.11) \quad W_{\text{GFT}}(\mathbb{C}\mathcal{G}) = \begin{cases} 8|\mathcal{G}|^2 & \text{if } \mathcal{R} \text{ is complex} \\ 2|\mathcal{G}|^2 & \text{if } \mathcal{R} \text{ is real.} \end{cases}$$

From (4.10) and Proposition 4.4 we find:

$$W_{\text{GFT}}(\mathbb{C}(\mathcal{G}_1 \times \mathcal{G}_2)) = W_{\text{GFT}}(\mathbb{C}\mathcal{G}_1)|\mathcal{G}_2| + |\mathcal{G}_1|W_{\text{GFT}}(\mathbb{C}\mathcal{G}_2),$$

$$W_{\text{GFT}}(\mathbb{C}^{m \times k}\mathcal{G}) = mkW_{\text{GFT}}(\mathbb{C}\mathcal{G}).$$

Thus, even if we use the slow transform, the cost of the transform on $\mathbb{C}^{m \times m}\mathcal{G}$ becomes $\mathcal{O}((m|\mathcal{G}|)^2)$, which is significantly less than W_{direct} .

For linear algebra algorithms of complexity $\mathcal{O}(n^2)$, such as e.g. iterative solution of $\mathbf{Ax} = \mathbf{b}$ in the case where A is a dense well conditioned matrix, the speed of the transform is more important. The cost of the direct matrix vector multiplication is $W_{\text{direct}} = c(m|\mathcal{G}|)^2$ whereas the same operation in the transformed space costs $W_{\text{fspace}} = c \sum_R (md_R)^2 d_R$, yielding the ratio

$$\mathcal{O}(n^2) : W_{\text{direct}}/W_{\text{fspace}} = \left(\sum_{R \in \mathcal{R}} d_R^2 \right)^2 / \sum_{R \in \mathcal{R}} d_R^3.$$

See Table 5.1 for typical values of this factor. The cost of a slow GFT on $\mathbb{C}^{m \times m}\mathcal{G}$ is of the same order as W_{direct} and the gain is more modest, which emphasizes the need for fast transforms in the $\mathcal{O}(n^2)$ case. Our numerical examples in Section 6.3 indicate that the issue of efficiently implementing the GFT is important for the total runtime of GFT based algorithms in the $\mathcal{O}(n^3)$ case as well.

6 Numerical examples.

We present examples of block-diagonalizing equivariant matrices under various groups. First, we discuss equivariance under the simplest non-abelian group, \mathcal{D}_3 , introduced in Section 2. We continue with equivariance under the group of the cube, and conclude with some numerical results.

6.1 Triangular symmetry.

The Triangle Example has been used throughout this text to exemplify the theory. We have now developed the required tools for summing up this example.

Recall that we want to solve a linear system of equations, $\mathbf{Ax} = \mathbf{b}$, where \mathbf{A} is *equivariant* w.r.t. \mathcal{D}_3 . Thus, $\mathbf{A}_{ig,jg} = \mathbf{A}_{i,j}$ for all $g \in \mathcal{D}_3$ acting on all $i \in \mathcal{I}$.

The case where the action is free and transitive corresponds to $\mathcal{I} = \{1, \dots, 6\}$ in Figure 2.1. In this case the equivariance of \mathbf{A} would yield a matrix with the structure illustrated to the left in Figure 6.1. In numerical applications, one would typically compute and store only the first column of \mathbf{A} , and represent \mathbf{A} as A according to (3.3). In order to exploit equivariance fully, we apply the GFT (4.4). Complete lists of irreducible representations \mathcal{R} for \mathcal{D}_n is given in Appendix A, stating that the three irreducible representations of \mathcal{D}_3 have dimensions 1, 1, and 2, respectively. Therefore, the transformed block matrix \hat{A} has the structure shown to the right in Figure 6.1. The transformed vectors have, in the case of free transitive actions, the same block structure. The original 6×6 system is thus transformed into three independent systems, where the first two

are trivial 1×1 “systems” and the third is a 2×2 system with two right hand sides.

$$\begin{pmatrix} a & b & e & f & c & d \\ b & a & f & e & d & c \\ c & d & a & b & e & f \\ d & c & b & a & f & e \\ e & f & c & d & a & b \\ f & e & d & c & b & a \end{pmatrix} \xrightarrow{\text{GFT}} \begin{pmatrix} \times & & & & & \\ & \times & & & & \\ & & \times & & & \\ & & & \times & \times & \\ & & & & \times & \times \end{pmatrix}$$

Figure 6.1: In the case when the action is free and transitive, the \mathcal{D}_3 -equivariance of \mathbf{A} is illustrated to the left and the structure of $\hat{\mathbf{A}} \in \widehat{\mathcal{CD}}_3$ is shown to the right.

For non transitive actions, the method is readily generalized via block GFTs. Assume a free action with m orbits, which implies that each orbit is of size $|\mathcal{D}_3| = 6$. The original system of size $6m \times 6m$ is transformed via the block GFT, yielding three independent systems of sizes $m \times m$, $m \times m$, and $2m \times 2m$, respectively. Again, the third of these systems has two right hand sides.

To illustrate general actions, we use the discretization in Figure 2.1, where $\mathcal{I} = \{1, \dots, 10\}$. In this case, we choose $\mathcal{S} = \{1, 7, 10\}$ and the three orbits are $\mathcal{O}_1 = \{1, \dots, 6\}$, $\mathcal{O}_7 = \{7, 8, 9\}$, and $\mathcal{O}_{10} = \{10\}$. The rank $r_{i,R}$ of the projection matrices $\hat{\pi}_i(R)$ is given in Table 6.1, for each $i \in \mathcal{S}$ and each $R = \rho_\ell \in \mathcal{R}$. The transformed equation system in \hat{V} will therefore have the structure shown in Figure 6.2.

Table 6.1: The rank $r_{i,R} = \text{rank}(\hat{\pi}_i(R))$, given the discretization in Figure 2.1.

R	$i \in \mathcal{S}$		
	1	7	10
ρ_0	1	1	1
ρ_1	1	0	0
ρ_2	2	1	0

$$\begin{pmatrix} \times & \times & \times & & & & & & & & \\ \times & \times & \times & & & & & & & & \\ \times & \times & \times & & & & & & & & \\ & & & \times & & & & & & & \\ & & & & \times & \times & \times & & & & \\ & & & & \times & \times & \times & & & & \\ & & & & \times & \times & \times & & & & \end{pmatrix} \begin{pmatrix} \times & & & & & & & & & & \\ \times & & & & & & & & & & \\ \times & & & & & & & & & & \\ & \times & & & & & & & & & \\ & & \times & \times & & & & & & & \\ & & \times & \times & & & & & & & \\ & & \times & \times & & & & & & & \end{pmatrix} = \begin{pmatrix} \times & & & & & & & & & & \\ \times & & & & & & & & & & \\ \times & & & & & & & & & & \\ & \times & & & & & & & & & \\ & & \times & \times & & & & & & & \\ & & \times & \times & & & & & & & \\ & & \times & \times & & & & & & & \end{pmatrix}$$

Figure 6.2: The structure of the transformed equation system $\hat{\mathbf{A}}''\hat{x}' = \hat{b}'$, given the discretization in Figure 2.1.

6.2 The cube.

The group of the cube consists of all linear transformations that maps a cube onto itself. This group is $\mathcal{K} = \mathcal{S}_4 \times \mathcal{C}_2$, which can be seen by the following

discussion given in [33, p. 43]. Consider a cube with vertices in $(\pm 1, \pm 1, \pm 1)$. Let $\mathcal{C}_2 = \langle \zeta \rangle$, where $\zeta(x_1, x_2, x_3) = -(x_1, x_2, x_3)$. Consider the tetrahedron with corners in $(1, 1, 1)$, $(1, -1, -1)$, $(-1, -1, 1)$, and $(-1, 1, -1)$. The group of the tetrahedron consists of all linear transformations that maps a tetrahedron onto itself. Since every such mapping can be expressed simply by stating how each of the 4 corners is mapped onto each others, we see that this group is isomorphic to \mathcal{S}_4 , the group of all permutations of 4 symbols. It is clear that $g \in \mathcal{S}_4$ as well as ζ maps the cube with its corners in $(\pm 1, \pm 1, \pm 1)$ onto itself. Moreover, we note that $g\zeta = \zeta g$ for each $g \in \mathcal{S}_4$. An alternative way to express \mathcal{K} is to combine the 6 transformations that permute coordinates (x_1, x_2, x_3) with the 8 reflections $(x_1, x_2, x_3) \mapsto (\pm x_1, \pm x_2, \pm x_3)$. We conclude that $|\mathcal{K}| = 48$ and we have $\mathcal{K} = \mathcal{S}_4 \times \mathcal{C}_2$.

In order to generate a symmetry respecting discretization of the cube surface, we note that it is enough to discretize a fundamental domain given by the triangle ABC in Figure 6.3, and then use the group \mathcal{K} to obtain a discretization of the whole cube boundary.

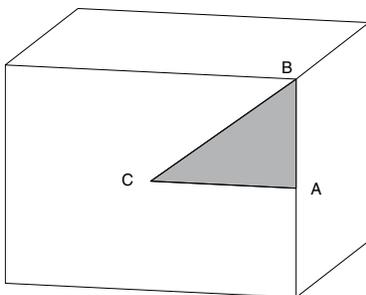


Figure 6.3: A symmetry respecting discretization of the cube boundary may be generated from discretization nodes in the triangle ABC .

If the discretization of ABC only contains points in the interior Ω , it implies that the action is free. In this case, the approach described in Section 5.1 is used. As $\mathcal{K} = \mathcal{S}_4 \times \mathcal{C}_2$, the 10 irreducible representations ρ_i are obtained from the irreducible representations of \mathcal{C}_2 and \mathcal{S}_4 , see Proposition 4.4. For $(g, \zeta^k) \in \mathcal{S}_4 \times \mathcal{C}_2$, let

$$\mathcal{R} = \{ \rho_\ell(g, \zeta^k) = \sigma_\ell(g), \rho_{\ell+5}(g, \zeta^k) = (-1)^k \sigma_\ell(g) \}_{\ell=0, \dots, 4}$$

be a list of irreducible representations for \mathcal{K} , where the irreducible representations σ_i for \mathcal{S}_4 are given in Appendix A.

If the discretization contains points on the boundary of the fundamental domain ABC , the action will no longer be free and we apply the general algorithm in Section 5.2. To illustrate how this affects the dimensions of the transformed quantities, Table 6.2 lists the ranks $r_{i,\ell} = \text{rank}(\hat{\pi}_i(\rho_\ell))$ for the the various domains of ABC . For the interior Ω of the triangle ABC , the action is free. In this case, $r_{i,\ell} = d_\ell$, the dimension of ρ_ℓ . Notice that it holds that $\mathcal{O}(x_i) = \sum_{\ell=0}^9 r_{i,\ell} d_\ell$.

Table 6.2 can be used for computing the size of $\hat{A}(\rho_\ell)$. If a discretization of the fundamental domain contains q_Δ points of each domain Δ in the table, $\hat{A}(\rho_\ell)$ has $\sum_\Delta q_\Delta r_{\Delta,\ell}$ rows and columns.

Table 6.2: The ranks $r_{i,\ell} = \text{rank}(\hat{\pi}_i(\rho_\ell))$. For a point x_i , this depends on the representation ρ_ℓ and the isotropy subgroup \mathcal{K}_i of x_i .

	$ \mathcal{O}(x_i) $	$ \mathcal{K}_i $	$r_{i,0}$	$r_{i,1}$	$r_{i,2}$	$r_{i,3}$	$r_{i,4}$	$r_{i,5}$	$r_{i,6}$	$r_{i,7}$	$r_{i,8}$	$r_{i,9}$
Ω	48	1	1	1	2	3	3	1	1	2	3	3
AB	24	2	1	0	1	2	1	1	0	1	2	1
AC	24	2	1	1	2	1	1	0	0	0	2	2
BC	24	2	1	0	1	2	1	1	0	1	2	1
A	12	4	1	0	1	1	0	0	0	0	1	1
B	8	6	1	0	0	1	0	1	0	0	1	0
C	6	8	1	0	1	0	0	0	0	0	1	0

6.3 Results.

We demonstrate performance results for solving equation systems via a direct method. The experiments are done in Matlab and executed on a Sun Enterprise 420R, which has 4 UltraSPARC-II with 450 MHz. The purpose is to illustrate how the GFT based algorithms perform for a few different groups. We consider equivariance under free actions of \mathcal{D}_3 , \mathcal{S}_4 and \mathcal{K} – i.e., the group of the triangle, the tetrahedron, and the cube, respectively. In Figure 6.4, we show the total time used for solving systems with n unknowns. Figures 6.5 and 6.6 show the time spent for the Fourier transforms and for the solves in the Fourier spaces, respectively. Tables 6.3–6.5 summarize the results for each group. For comparison, we list execution times for solving systems using direct solves in Table 6.6. We use real data and real representations, see Appendix A.

We stress the following points.

- Exploiting symmetry via the GFT leads to a considerable gain. For $n = 5760$, the gain is a factor 10 for \mathcal{S}_4 and a factor 16 for \mathcal{K} . The gain increases with n .
- For a small group such as \mathcal{D}_3 , the time for performing the GFT is negligible when n grows. For larger groups, and moderate sized matrices, the big gain comes from the fact that the diagonal blocks become small. Computing $t_{\text{direct}}/t_{\text{solve}}$ for the different groups, we find that the ratios for $n = 5760$ are 16, 119, and 340, for the three symmetry groups, respectively. This is significant, even if it is less than what was predicted in Section 5.4. The discrepancy is explained by the fact that the cn^3 model for the solution times is not adequate in the current range, see data in Table 6.6.
- When computing the GFT, it is important to exploit structure. In our implementation, each GFT for the cube exploits two GFTs for \mathcal{S}_4 . For example, for $m = 120$ orbits, t_{GFT} for the cube is about twice t_{GFT} for \mathcal{S}_4 , as it should be. For $m = 180$, the factor has increased, which we believe is due to cache size limitations and other memory effects.

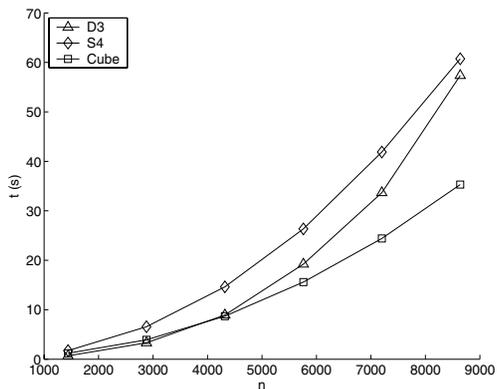


Figure 6.4: Total time.

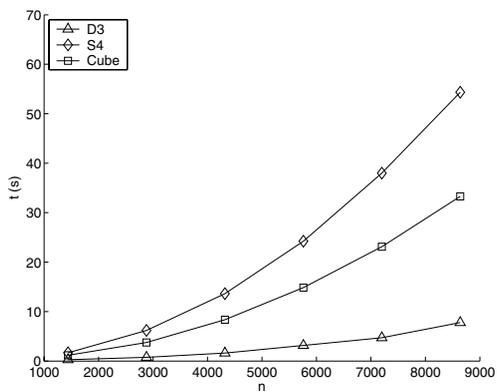


Figure 6.5: GFT times.

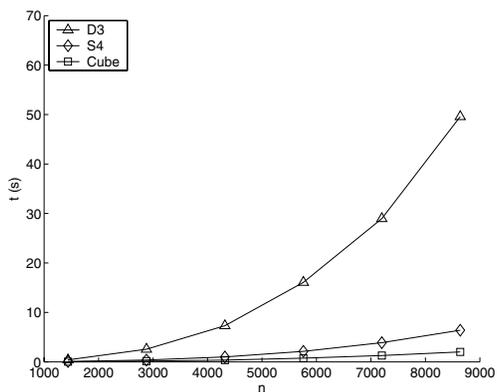


Figure 6.6: Fourier space solve times.

Table 6.3: Results (s) for \mathcal{D}_3 .

n	1440	2880	4320	5760	7200	8640
m	240	480	720	960	1200	1440
t_{solve}	0.4000	2.5600	7.3200	16.0800	28.9600	49.5900
t_{GFT}	0.2800	0.7600	1.6200	3.1500	4.7100	7.7700
t_{tot}	0.6800	3.3200	8.9400	19.2300	33.6700	57.3600

Table 6.4: Results (s) for \mathcal{S}_4 .

n	1440	2880	4320	5760	7200	8640
m	60	120	180	240	300	360
t_{solve}	0.0700	0.3800	1.0000	2.1300	3.8900	6.4000
t_{GFT}	1.6600	6.1800	13.6300	24.2200	37.9900	54.3400
t_{tot}	1.7300	6.5600	14.6300	26.3500	41.8800	60.7400

Table 6.5: Results (s) for the cube.

n	1440	2880	4320	5760	7200	8640
m	30	60	90	120	150	180
t_{solve}	0.0400	0.1500	0.3500	0.7500	1.2900	2.0200
t_{GFT}	1.1600	3.7600	8.3600	14.8500	23.1100	33.2900
t_{tot}	1.2000	3.9100	8.7100	15.6000	24.4000	35.3100

Table 6.6: Results (s) for direct solves.

n	540	1440	2880	4320	5760
t_{direct}	0.3900	5.5800	36.7700	111.6800	254.8200

7 Summary.

We have summarized how noncommutative Fourier analysis can be exploited in numerical analysis, based on the GFT. In particular, we have pointed out the role of the convolutions in the group algebra and the relation to standard Fourier analysis. For the case with fixed points, we have extended previous work and developed formulas which exploit invariance under isotropy subgroups.

We have carried out numerical experiments for solving linear systems of equations, equivariant under \mathcal{D}_3 , \mathcal{S}_4 , and \mathcal{K} , the symmetry groups of the triangle, the tetrahedron and the group, respectively. For the cube, we exploit that \mathcal{K} is a direct product $\mathcal{S}_4 \times \mathcal{C}_2$, which allows the use of a fast transform. The numerical results are supported by complexity analysis.

In our exposition, we have strived to use a coherent notation for the various mathematical entities. We find that this is a very important issue, not only for understanding and conveying the ideas, but also when considering implemen-

tation. In [24], we have started to explore how to use modern programming techniques such as generic programming in order to obtain both general and efficient software. We also want to consider parallel aspects.

For practical 3D applications, we finally notice that the number of finite groups which we need to consider is actually limited, see, e.g., the discussion in [2]. Any finite subgroup of $SO(3)$ (3×3 orthogonal matrices with determinant 1, corresponding to groups of rotations) is isomorphic to either a cyclic group, a dihedral group, the alternating group \mathcal{A}_4 on 4 symbols, \mathcal{S}_4 , or \mathcal{A}_5 . Any finite symmetry group in 3D is either one of these, or a direct product with \mathcal{C}_2 and one of these. The symmetry group of the icosahedron, for example, is $\mathcal{C}_2 \times \mathcal{A}_5$. Our goal is to provide an efficient library for all of these symmetries. Particularly, we believe that an efficient implementation of the GFT for the icosahedron may be useful as an approximation for the sphere.

Acknowledgements.

We thank Gene Allgower, Henrik Brandén, Craig Douglas, Malin Ljungberg and Will Wright for valuable inputs.

A Appendix: Representations for some common groups.

A.1 Representations for dihedral groups.

The dihedral groups \mathcal{D}_n correspond to the symmetries of regular n -sided polygons, and are generated by

$$\mathcal{D}_n = \langle \alpha, \beta \mid \alpha^n = \beta^2 = e, \alpha\beta = \beta\alpha^{-1} \rangle.$$

The complete set of irreducible representations depends on whether n is odd ($2q+1$) or even ($2q+2$). Table A.1 lists possible sets of irreducible representations for \mathcal{D}_n . If real representations are desired, it is possible to change the basis of a two-dimensional representation via the transform

$$\rho_{1+k} \mapsto T\rho_{1+k}T^{-1}, \quad T = \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix}.$$

For example, we obtain for \mathcal{D}_3 that

$$\rho_2(\alpha) = \begin{pmatrix} -1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{pmatrix}, \quad \rho_2(\beta) = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

A.2 Representations for \mathcal{S}_4 .

The group of all permutations of 4 symbols can be generated by the cyclic permutations

$$\alpha = (1\ 2), \beta = (1\ 2\ 3), \text{ and } \gamma = (1\ 2\ 3\ 4).$$

Table A.1: Complete sets of irreducible representations for \mathcal{D}_n . Here, $\omega = \exp(2\pi i/n)$. Note that the last two representations only apply when $n = 2q + 2$. In this case, the number of representations is $q + 4 = n/2 + 3$. If $n = 2q + 1$, the number of representations is $q + 2 = (n + 3)/2$.

	α	β	
ρ_0	1	1	
ρ_1	1	-1	
ρ_{1+k}	$\begin{pmatrix} \omega^k & 0 \\ 0 & \omega^{-k} \end{pmatrix}$	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	$k = 1, \dots, q$
ρ_{2+q}	-1	1	if n is even
ρ_{3+q}	-1	-1	if n is even

Every $g \in \mathcal{S}_4$ can be written as $\alpha^i \beta^j \gamma^k$ where $i \in \{0, 1\}$, $j \in \{0, 1, 2\}$, and $k \in \{0, 1, 2, 3\}$. The size of the group is $|\mathcal{S}_4| = 24$. A complete set of irreducible representations contains 5 representations of dimensions 1, 1, 2, 3, 3, see Table A.2.

Table A.2: A possible choice of real irreducible representations for \mathcal{S}_4 .

	α	β	γ
σ_0	1	1	1
σ_1	-1	1	-1
σ_2	$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$	$\begin{pmatrix} -1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{pmatrix}$	$\begin{pmatrix} -1/2 & \sqrt{3}/2 \\ \sqrt{3}/2 & 1/2 \end{pmatrix}$
σ_3	$\begin{pmatrix} -1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} -1 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} -1 & -1 & -1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$
σ_4	$\begin{pmatrix} 1 & 0 & 0 \\ -1 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}$	$\begin{pmatrix} -1 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{pmatrix}$

REFERENCES

1. K. Åhlander and H. Munthe-Kaas, *Eigenvalues for equivariant matrices*, J. Comput. Appl. Math. (2005). Available online via ScienceDirect.
2. E. L. Allgower, K. Böhmer, K. Georg and R. Miranda, *Exploiting symmetry in boundary element methods*, SIAM J. Numer. Anal., 29 (1992), pp. 534–552.
3. E. L. Allgower and A. F. Fässler, *Blockstructure and equivalence of matrices*, in Aspects of Complex Analysis, Differential Geometry, Mathematical Physics and Applications, pp. 19–34, World Scientific, 1999.
4. E. L. Allgower and K. Georg, *Exploiting symmetry in numerical solving*, in Proceedings of the Seventh Workshop on Differential Equations and its Applications, C.-S. Chien, ed., Taichung, Taiwan, 1999.

5. E. L. Allgower, K. Georg and R. Miranda, *Exploiting permutation symmetry with fixed points in linear equations*, in Lectures in Applied Mathematics, E. L. Allgower, K. Georg and R. Miranda, eds., vol. 29, American Mathematical Society, pp. 23–36, Providence, RI, 1993.
6. E. L. Allgower, K. Georg, R. Miranda and J. Tausch, *Numerical exploitation of equivariance*, *Z. Angew. Math. Mech.*, 78 (1998), pp. 185–201.
7. T. Beth, *Generalized Fourier Transforms*, in Trends in Computer Algebra, Lect. Notes Comput. Sci., vol. 296, pp. 92–118, Springer, Berlin, 1988.
8. M. Bonnet, *Exploiting partial or complete geometrical symmetry in 3D symmetric Galerkin indirect BEM formulations*, *Int. J. Numer. Methods Eng.*, 57 (2003), pp. 1053–1083.
9. A. Bossavit, *Symmetry, groups, and boundary value problems. a progressive introduction to noncommutative harmonic analysis of partial differential equations in domains with geometrical symmetry*, *Comput. Methods Appl. Mech. Eng.*, 56 (1986), pp. 167–215.
10. A. Bossavit, *Boundary value problems with symmetry and their approximation by finite elements*, *SIAM J. Appl. Math.*, 53 (1993), pp. 1352–1380.
11. A. Clausen, *Fast Fourier transforms for metabelian groups*, *SIAM J. Comput.*, 18 (1989), pp. 55–63.
12. A. Clausen and U. Baum, *Fast Fourier Transforms*, Wissenschaftsverlag, Mannheim, 1993.
13. A. Clausen and M. Müller, *Generating fast Fourier transforms of solvable groups*, *J. Symb. Comput.*, 37 (2004), pp. 137–156.
14. P. Diaconis, *Group representations in probability and statistics*, in Lecture Notes – Monograph Series, Institute of Mathematical Statistics, Hayward, CA, 1988.
15. P. Diaconis and D. N. Rockmore, *Efficient computation of the Fourier transform on finite groups*, *J. Am. Math. Soc.*, 3 (1990), pp. 297–332.
16. C. C. Douglas and J. Mandel, *Abstract theory for the domain reduction method*, *Computing*, 48 (1992), pp. 73–96.
17. S. Egner and M. Puschel, *Symmetry-based matrix factorization*, *J. Symb. Comput.*, 37 (2004), pp. 157–186.
18. A. F. Fässler and E. Stiefel, *Group Theoretical Methods and Their Applications*, Birkhäuser, Boston, 1992.
19. K. Georg and R. Miranda, *Exploiting symmetry in solving linear equations*, in Bifurcation and Symmetry, E. L. Allgower, K. Böhmer and M. Golubisky, eds., ISNM, vol. 104, pp. 157–168, Birkhäuser, Basel, 1992.
20. K. Georg and R. Miranda, *Symmetry aspects in numerical linear algebra with applications to boundary element methods*, in Lectures in Applied Mathematics, E. L. Allgower, K. Georg and R. Miranda, eds., vol. 29, pp. 213–228, American Mathematical Society, Providence, RI, 1993.
21. K. Georg and J. Tausch, *A generalized Fourier transform for boundary element methods with symmetries*, Technical report, Colorado State University, Ft. Collins, Colorado, 1994.
22. G. James and M. Liebeck, *Representations and Characters of Groups*, Cambridge University Press, 2nd edn., 2001. ISBN 052100392X.
23. M. Karpovsky, *Fast Fourier transforms on finite non-abelian groups*, *IEEE Transact. Comput.*, 10 (1977), pp. 1028–1030.
24. M. Ljungberg and K. Åhlander, *Generic Programming Aspects of symmetry exploiting numerical software*, in Proceedings of European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS 2004, Jyväskylä, 24–28 July 2004, P. Neittaanmäki et al., eds., 2004. Also available as Technical Report 2004-020 from the Department of Information Technology, Uppsala University.
25. J. S. Lomont, *Applications of Finite Groups*, Academic Press, New York, 1959.
26. D. K. Maslen, *The efficient computation of Fourier transforms on the symmetric group*, *Math. Comput.*, 67 (1998), pp. 1121–1147.

27. D. K. Maslen and D. N. Rockmore, *Generalized FFTs – a survey of some recent results*, in Proceedings of the 1995 DIMACS Workshop on Groups and Computation, L. Finkelstein and W. Kantor, eds., June 1997, pp. 183–237.
28. D. K. Maslen and D. N. Rockmore, *Separation of variables and the computation of Fourier transforms on finite groups*, J. Am. Math. Soc., 10 (1997), pp. 169–214.
29. 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>.
30. D. N. Rockmore, *Fast Fourier analysis for abelian group extensions*, Adv. Appl. Math., 11 (1990), pp. 164–204.
31. D. N. Rockmore, *Fast Fourier-transforms for wreath-products*, Appl. Comput. Harmonic Anal., 2 (1995), pp. 279–292.
32. D. N. Rockmore, *Some applications of generalized FFTs*, in Proceedings of the 1995 DIMACS Workshop on Groups and Computation, L. Finkelstein and W. Kantor, eds., June 1997, pp. 329–369.
33. J. P. Serre, *Linear Representations of Finite Groups*, Springer, 1977. ISBN 0387901906.
34. L. Stiller, *Exploiting symmetry on parallel architectures*, PhD thesis, Johns Hopkins University, 1995.
35. J. Tausch, *Equivariant preconditioners for boundary element methods*, SIAM Sci. Comput., 17 (1996), pp. 90–99.