Eigenvalues for Equivariant Matrices

[1]Krister Åhlander and [2]Hans Munthe-Kaas

[1]*Department of Information Technology, Uppsala University*
[2]*Department of Informatics, University of Bergen*
emails: krister@tdb.uu.se, hans@ii.uib.no

Abstract

An equivariant matrix $\mathbf{A}$ commutes with a group of permutation matrices. Such matrices often arise in numerical applications where the computational domain exhibits geometrical symmetries, for instance triangles, cubes, or icosahedra.
The theory for block diagonalizing equivariant matrices via the Generalized Fourier Transform (GFT) is reviewed and applied to eigenvalue computations. For dense matrices which are equivariant under large symmetry groups, we give theoretical estimates that show a substantial performance gain. In case of cubic symmetry, the gain is about 800 times, which is verified by numerical results.
It is also shown how the multiplicity of the eigenvalues is determined by the symmetry, which thereby restricts the number of distinct eigenvalues. The inverse GFT is used to compute the corresponding eigenvectors. It is emphasized that the inverse transform in this case is very fast, due to the sparseness of the eigenvectors in the transformed space.

Key words: Generalized Fourier transform, symmetrical geometry, eigenvalue computation.

1 Introduction

An equivariant matrix $\mathbf{A}$ commutes with a group $\mathcal{G}$ of permutation matrices, thus $P\mathbf{A} = \mathbf{A}P$ for all P in $\mathcal{G}$. A well-known example is block-circulant matrices, which commute with shift matrices. Equivariant matrices may arise in applications where a symmetric domain is discretized, for instance when the boundary element method is used to solve a partial differential equation in a domain with the symmetry of a cube or of an icosahedron, see Figure 1. We refer to for instance [1] for several fine examples of applications where symmetric domains lead to equivariant matrices.

The equivariance property may be exploited to devise better numerical algorithms. The key is to use the generalized Fourier transform (GFT) in

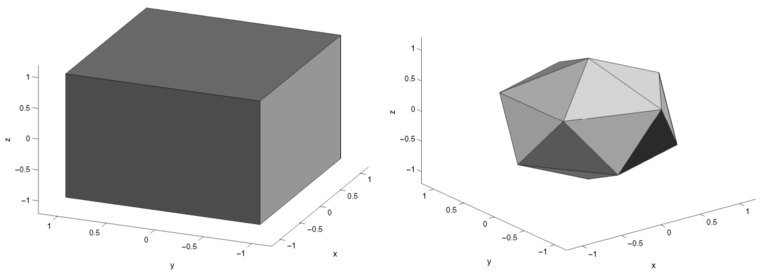

(a) (b)

Figure 1: Equivariant matrices arise for instance when partial differential equations are numerically solved in symmetrical domains, here exemplified by a cube and by an icosahedron.

order to block diagonalize $\mathbf{A}$. The GFT relies on representation theory for groups. The application of the GFT to solve a system of equations $\mathbf{A}x = b$ was recognized by Allgower and others [2, 3] and related to the symmetry utilizing methods described by [4, 5]. In [6], we survey the application of the GFT to numerical linear algebra, emphasizing the connection to the group algebra.

Symmetries and eigenvalue computations have been briefly discussed in earlier literature [2]. In this paper we present a more detailed study of the use of GFT in eigenvalue computations, and discuss the computational gains from this approach, both from a theoretical and a computational point of view.

In Section 2, we first introduce representation theory for groups. By using the relationship between equivariance and the regular representation, we then explain *why* it is possible to block-diagonalize an equivariant matrix, and finally we show *how* to achieve this via the GFT. In Section 3 we apply the theory to eigenvalue and eigenvector computations. In Section 4, we derive theoretical estimates on the performance gain, which are confirmed via numerical experiments.

2 Theory

The theory for the GFT relies on representation theory for groups. For readers unfamiliar with this theory, we recommend [7] and the references above for more details. Here, we basically follow the notation of [6].

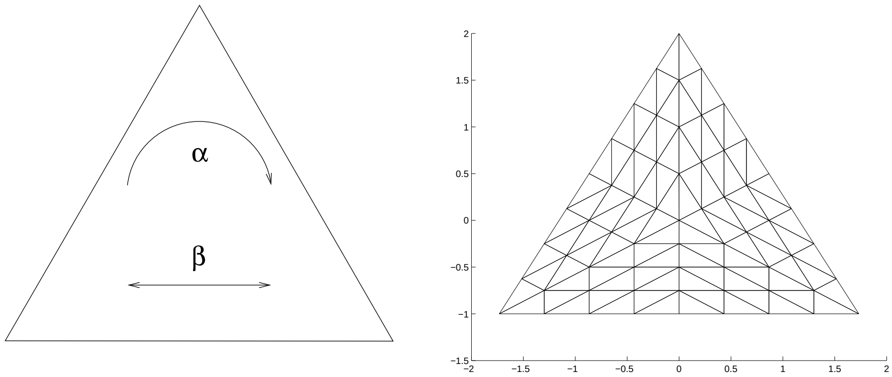

(a) (b)

Figure 2: (a) The symmetry group of the triangle is generated by the rotation α and the reflection β. (b) A symmetry respecting discretization of the triangle. Since every transformation apart from the identity transformation maps every triangle element to another element, the action is free on the triangle elements. The action on the nodes, however, is not free. The uppermost corner, for example, is fix under the reflection β.

Throughout our exposition, we will use the symmetry group of the triangle, see Figure 2, to exemplify the concepts. This example is chosen because it illustrates all interesting points of the theory, yet it is small enough to be discussed in some detail.

2.1 Introduction to groups and representations

A *group* is a set $\mathcal{G}$ with an associative binary relation, thus $(gh)t = g(ht)$, which has an identity e such that $eg = ge = g$, and where each element g has an inverse g^{-1} such that $gg^{-1} = g^{-1}g = e$. The order of the group is denoted $|\mathcal{G}|$. We focus on finite groups where $|\mathcal{G}| < \infty$.

Example 1 *All rotations and reflections which map the triangle onto itself form a group, see Figure 2a. This group is recognized as $\mathcal{D}_3$, the dihedral group with $|\mathcal{D}_3| = 6$ elements, generated by a rotation 120 degrees, α, and a reflection in the vertical axis, β. Its elements are $\{e, \alpha, \alpha^2, \beta, \alpha\beta, \alpha^2\beta\}$.*

In applications, groups are usually of importance because of their actions on relevant sets. A *group action* (from the right) of a group $\mathcal{G}$ on a set $\mathcal{I}$ is a relation $\mathcal{I} \times \mathcal{G} \to \mathcal{I}$ such that $ie = i$ and $i(gh) = (ig)h$ for all i in $\mathcal{I}$ and all g, h in $\mathcal{G}$. The *isotropy subgroup* $\mathcal{G}_i$ of an element i is the group of elements for which i is unaffected by the action; $\mathcal{G}_i = \{g \in \mathcal{G} \mid ig = i\}$. The *orbit* $i\mathcal{G}$ of an element i is the set of elements which are obtained when $\mathcal{G}$ acts on i. A group action is *free* if all isotropy subgroups are trivial, i.e. $\mathcal{G}_i = \{e\}$, for all i. If there is only one orbit of the action, it is said to be *transitive.*

Example 2 *Let $\mathcal{D}_3$ act on the set of triangles in the triangulization shown in Figure 2b. Since every triangle element is affected by an action other than the identity transform, we see that the action is free. On the other hand, when $\mathcal{D}_3$ acts on the nodes of the triangulization, we note that the isotropy subgroup of for instance the uppermost node is $\{e, b\}$, and the action is not free. If we only consider the action of $\mathcal{D}_3$ on the three corners of the triangle, the action is transitive but not free.*

A *representation* ρ of dimension d is a map $\rho : \mathcal{G} \to \mathbb{C}^{d\times d}$ for which $\rho(gh) = \rho(g)\rho(h)$ (a group homomorphism). We will frequently denote the dimensions of a representation σ by d_σ, and the dimension of a representation ρ_i by d_i. For every group, the 1-dimensional trivial representation $\tau : g \mapsto 1$, can be defined. Two representations ρ and σ are said to be *isomorphic* if there exists a nonsingular T such that $\sigma(g) = T\rho(g)T^{-1}$ for all g in $\mathcal{G}$. Thus, isomorphic representations differ only by the choice of basis. If there exists a basis in which a representation ρ is block diagonal, it is said to be *reducible* into subrepresentations, denoted $\rho = \rho_1 \oplus \rho_2$. Two key results in representation theory states that every representation is reducible into *irreducible* representations, and for every finite group there exists a finite list $\mathcal{R}$ of nonisomorphic irreducible representations, for which it holds that $\sum_{\rho\in\mathcal{R}} d_\rho^2 = |\mathcal{G}|$.

Example 3 *For $\mathcal{D}_3$, a 1-dimensional representation σ can be defined by $\sigma(\alpha) = 1$ and $\sigma(\beta) = -1$ and by using the homomorphism property. A 2-dimensional representation can be constructed by considering the action of $\mathcal{D}_3$ on the triangle in the plane. It may be realized by representing α and β by*

$$\rho(\alpha) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \text{ and } \rho(\beta) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix},$$

where θ is 120 degrees. Together with the trivial representation τ, these representations constitute a complete list $\mathcal{R} = \{\tau, \sigma, \rho\}$ of nonisomorphic irreducible representations of $\mathcal{D}_3$. Note that $d_\tau^2 + d_\sigma^2 + d_\rho^2 = 6 = |\mathcal{D}_3|$.

For every finite group, the standard basis of the vector space $\mathbb{C}^{|\mathcal{G}|}$ may be indexed by group elements. The *regular representation* is defined by the group action on this basis, $\{\hat{e}_h\}$, via the group operation. We distinguish between the *left* regular representation, $L(g) : \hat{e}_h \mapsto \hat{e}_{gh}$, and the *right* regular representation, $R(g) : \hat{e}_h \mapsto \hat{e}_{hg^{-1}}$.

We will base our current exposition on the following important property of the regular representations (see e.g. [7, Ch. 2.4]).

Theorem 1 *Every irreducible representation $\rho_i \in \mathcal{R}$ is contained in the (left or right) regular representation with multiplicity equal to its dimension d_i.*

Example 4 *For $\mathcal{D}_3$ of order 6, the theorem states that there exists a shift of basis F such that $\bar{L}(g) = FL(g)F^{-1}$ is block diagonal, containing two 1×1 blocks and two 2×2 blocks. The 1×1 blocks correspond to the irreducible representations τ and σ, respectively, whereas the 2×2 blocks are identical, corresponding to the 2-dimensional irreducible representation ρ.*

2.2 Equivariance

We are interested in applications where a symmetric domain is discretized, and we want to exploit the inherent symmetry of the underlying problem. For instance, if a spatial discretization operator commutes with the symmetry group of a symmetric domain, it is said to be *equivariant* with respect to the symmetry group. Provided a symmetry respecting discretization (see Figure 2b) is used, this property carries over to the discrete case, and an equivariant matrix which commutes with a group of permutation matrices is obtained. An equivalent definition of equivariance may be stated in terms of the symmetry group $\mathcal{G}$ acting on the indices $\mathcal{I}$. Thus, for an equivariant matrix,

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

In [6], we illustrate the connection between equivariance and the group algebra, and we briefly recapitulate this point of view in the next section. Here, we focus on the connection between Theorem 1 and the block diagonalization of equivariant matrices. It is then convenient to distinguish between different kinds of actions.

Free transitive actions If $\mathcal{G}$ acts on $\mathcal{I}$ transitively, $\mathcal{I}$ contains just one orbit. If the action is free, every isotropy subgroup is trivial, and we can identify the group action with a regular action on $\mathbb{C}^{|\mathcal{G}|}$. The connection between block diagonalization and equivariance is in this case revealed as follows.

Proposition 1 *Let* $\mathbf{A}$ *be equivariant with respect to a regular action from the right, i.e.* $\mathbf{A}_{gt,ht} = \mathbf{A}_{g,h}$. *Then* $\mathbf{A}$ *commutes with the right regular representation,* $\mathbf{A}R(g) = R(g)\mathbf{A}$ *for all* g *in* $\mathcal{G}$. *Moreover,* $\mathbf{A}$ *is a linear combination of the left regular representation,*

$$\mathbf{A} = \sum_{g\in\mathcal{G}} a(g)L(g), \tag{1}$$

where $a(g) = \mathbf{A}_{g,e}$.

The proof relies on the easily shown facts that $L(g)$ is equivariant and that $L(g)$ and $R(h)$ commute for all g, h in $\mathcal{G}$. Together with Theorem 1, an immediate conclusion is that if L is block diagonalized via a basis shift $\bar{L} = FLF^{-1}$, then will $\bar{\mathbf{A}} = F\mathbf{A}F^{-1}$ be block diagonal with the same block structure.

Free actions If $\mathcal{G}$ acts on $\mathcal{I}$ freely, we can identify $\mathcal{I}$ with $\mathcal{S}\times\mathcal{G}$, where $\mathcal{S}$ is a selection of orbit representatives, via the relation

$$(i, g) = ig \text{ for } i \in \mathcal{S},\ g \in \mathcal{G}. \tag{2}$$

Proposition 1 is readily generalized. If $\mathbf{A}$ is equivariant under a free action which partitions $\mathcal{I}$ into $m = |\mathcal{S}|$ orbits, $\mathbf{A}$ commutes with the representation given by $g \mapsto I_m \otimes R(g)$, i.e., the kronecker product between the $m\times m$ identity matrix and the right regular representation. Similarly,

$$\mathbf{A} = \sum_{g\in\mathcal{G}} A(g) \otimes L(g), \tag{3}$$

where $A(g)$ is an $m \times m$ matrix given by $A(g)_{i,j} = \mathbf{A}_{ig,j}$ for all $i, j \in \mathcal{S}$ and all $g \in \mathcal{G}$. We refer to A as the *essential* part of $\mathbf{A}$. The existence of a basis shift F which makes $\bar{\mathbf{A}} = F\mathbf{A}F^{-1}$ block diagonal is once again an immediate consequence of Theorem 1, but this time each block is $m \times m$ times bigger.

General actions In general, we may have nontrivial isotropy subgroups, which implies that $N = |\mathcal{S}\times\mathcal{G}| > |\mathcal{I}| = n$. This case is discussed in some detail in [6], using the GFT. The exposition here is instead based on matrix formulation, cf. [3, 8].

Assume that $\mathbf{A}$ is equivariant under an action which partitions the set $\mathcal{I}$ of n indices into m orbits represented by $\mathcal{S}$. Assume that $N = m|\mathcal{G}| > n$, which implies that the action is not free. The idea is to use $\mathbf{A} : \mathbb{C}^n \to \mathbb{C}^n$ to construct $\mathbf{A}^+ : \mathbb{V} \to \mathbb{V} \subset \mathbb{C}^N$, which is equivariant under a *free* action $(\mathcal{S}\times\mathcal{G})\times\mathcal{G} \to \mathcal{S}\times\mathcal{G}$. This is achieved by defining the $N \times n$ matrix $V_{(i,g),j} =$

$\delta_{ig,j}$, for $(i,g) \in \mathcal{S} \times \mathcal{G}$, $j \in \mathcal{I}$, using the kronecker symbol δ. It is easily shown that $V^T V = I_n$, and that $\Pi = VV^T$ is a projection matrix onto the subspace $\mathbb{V}$ spanned by V. Thus, $V^T : \mathbb{V} \to \mathbb{C}^n$ is the inverse of V, and $\mathbb{V}$ is isomorphic to $\mathbb{C}^n$. Let $\mathbf{A}^+ = V\mathbf{A}V^T$, which is equivariant with respect to the free induced action on $\mathcal{S} \times \mathcal{G}$. Also note that Π is equivariant. By applying the theory for free actions, we obtain block diagonal $\bar{\mathbf{A}}^+ = F\mathbf{A}^+F^{-1}$ and $\bar{\Pi} = F\Pi F^{-1}$. Let r_j be the rank of $\bar{\Pi}^{(j)}$. By factorizing $\bar{\Pi} = QP$, we draw the following conclusion.

Theorem 2 *A matrix* $\mathbf{A}$ *equivariant under a general action can be block diagonalized via a basis shift* $\bar{\mathbf{A}} = PFV\mathbf{A}V^TF^{-1}Q$*, where each block* $\bar{\mathbf{A}}^{(j)}$ *has dimension* $r_j \times r_j$*.*

2.3 The Generalized Fourier Transform

The previous section showed that it was possible to find a basis shift which makes an equivariant matrix $\mathbf{A}$ block diagonal. Here, we proceed and describe *how* to find this basis shift. The essential tool which is used for block diagonalizing $\mathbf{A}$ is the generalized Fourier transform (GFT). Formally, it is an algebra isomorphism between the group algebra $\mathbb{C}\mathcal{G}$ and its Fourier space, $\widehat{\mathbb{C}\mathcal{G}}$. The group algebra $\mathbb{C}\mathcal{G}$ is the vector space $\mathbb{C}^{\mathcal{G}}$ equipped with the natural product,

$$x * y = \left(\sum_{g\in\mathcal{G}} x(g)\hat{e}_g\right) * \left(\sum_{h\in\mathcal{G}} y(h)\hat{e}_h\right) = \sum_{g,h\in\mathcal{G}} x(g)y(h)\hat{e}_{gh}.$$

The Fourier space is a block diagonal matrix algebra, with block dimensions according to the irreducible representations, $\mathcal{R}$, of the group: $\widehat{\mathbb{C}\mathcal{G}} = \bigoplus_{\rho\in\mathcal{R}} \mathbb{C}^{d_\rho \times d_\rho}$. The GFT $\hat{x} = \mathrm{gft}(x)$ of $x \in \mathbb{C}\mathcal{G}$ is given by

$$\hat{x}(\rho) = \sum_{g\in\mathcal{G}} x(g)\rho(g), \text{ for all } \rho \in \mathcal{R}.$$

Recall that the algebra isomorphism property implies that products are preserved, $\mathrm{gft}(x * y) = \mathrm{gft}(x)\mathrm{gft}(y)$, where the product in the Fourier space is matrix multiplication, and that the inverse transform exists, $x = \mathrm{igft}(\mathrm{gft}(x))$. We remark that the computations may be organized to exploit structure further, see [9]. It is beyond the scope of this paper to discuss so called fast GFTs in detail, and for eigenvalue computations, the GFT should not be a time critical operation.

Example 5 *For* $x \in \mathbb{C}\mathcal{D}_3$*,* $\hat{x} \in \widehat{\mathbb{C}\mathcal{D}_3}$ *is a block diagonal matrix with two 1-dimensional blocks,* $\hat{x}(\tau)$ *and* $\hat{x}(\sigma)$ *and one 2-dimensional block,* $\hat{x}(\rho)$*. Explicitly, we have* $\hat{x}(\tau) = \sum_{g\in\mathcal{D}_3} x(g)$*,* $\hat{x}(\sigma) = \sum_{g\in\mathcal{D}_3} x(g)\sigma(g)$ *and* $\hat{x}(\rho) = \sum_{g\in\mathcal{D}_3} x(g)\rho(g)$*, where* σ *and* ρ *are given in Example 3.*

In order to explain how the GFT is used to block diagonalize a linear system of equations, $\mathbf{A}x = b$, we emphasize the connection to the group algebra. Again, it is convenient to distinguish between different kinds of action.

Transitive free actions If the action is free and transitive, it is easy to see that $\mathbf{A}x = b$ corresponds to an equation $a * x = b$ in the group algebra, where $a(g) = \mathbf{A}_{e,g}$, cf. Section 2.2. We define the GFT of the matrix $\hat{A} = \mathrm{gft}(\mathbf{A})$ as the GFT of the corresponding element a in the group algebra.

Example 6 *Consider $\mathcal{D}_3$ and let $\mathbf{A} = L(\alpha)$, the left regular representation of α. The corresponding element in the group algebra is $\hat{e}_\alpha$. We obtain $\hat{a}(\tau) = \hat{a}(\sigma) = 1$, and $\hat{a}(\rho) = \rho(\alpha)$, cf. Example 3. Note that $a * a = \hat{e}_{\alpha^2}$ is the corresponding component to $L(\alpha^2)$, which is a direct consequence of the isomorphism between $\mathbb{C}\mathcal{G}$ and equivariant matrices under the regular action.*

Free actions If the action is free with m orbits, the identification between $\mathcal{I}$ and $\mathcal{S} \times \mathcal{G}$ allows us to generalize the previous case. $\mathbf{A}x = b$ now corresponds to an equation $A * x = b$, where $A(g) \in \mathbb{C}^{m\times m}$, and $x(g), b(g) \in \mathbb{C}^m$, and where the product $*$ is both a convolution product and a matrix multiplication. The GFT is computed elementwise; we define $\hat{A} = \mathrm{gft}(\mathbf{A})$ and $\hat{x} = \mathrm{gft}(x)$ by

$$\hat{A}_{i,j} = \mathrm{gft}(A_{i,j}) \text{ and } \hat{x}_i = \mathrm{gft}(x_i)$$

for matrices and vectors, respectively. Notice that each block of $\hat{A}(\rho_i)$ has dimension $md_i \times md_i$, whereas each block $\hat{x}(\rho_i)$ of a transformed vector has dimension $md_i \times d_i$, where d_i is the dimension of ρ_i.

General actions For general actions, we proceed in the same way as in Section 2.2. If the action partitions $\mathcal{I}$ into m orbits where $n = |\mathcal{I}| < m|\mathcal{G}| = N$, it implies that the action is not free. Let x and b in $\mathbb{C}^n$ be induced into x^+ and b^+ in the isomorphic subspace $\mathbb{V} \subset \mathbb{C}^N$. Similarly, the $n \times n$ matrix $\mathbf{A}$ is induced into the $N \times N$ matrix $\mathbf{A}^+$, and we obtain a free equivariant system. Let $\hat{x}^+ = \mathrm{gft}^+(x) = \mathrm{gft}(x^+)$ denote the GFT of the induced vector and let $\hat{x} = \mathrm{gft}(x)$ be the projection of $\hat{x}^+$ onto the space $\widehat{\mathbb{V}}$ spanned by $\hat{\Pi}$. Similarly $\hat{A}^+ = \mathrm{gft}(\mathbf{A}^+)$ is projected to $\hat{A} = \mathrm{gft}(\mathbf{A})$.

Notice that each block of $\hat{A}^+(\rho_i)$ has dimension $md_i \times md_i$. However, if we let r_i be the rank of $\hat{\Pi}(\rho_i)$, we see that each block $\hat{A}(\rho_i)$ of the projected matrix is $r_i \times r_i$. Similarly, for the transformed vectors, we have that $\hat{x}^+(\rho_i)$ is an $md_i \times d_i$ block whereas the projection $\hat{x}(\rho_i)$ is $r_i \times d_i$.

Example 7 *Consider the action of $\mathcal{D}_3$ on $n = 9$ indices $\mathcal{I}$, which partitions $\mathcal{I}$ into one full orbit with 6 elements and one orbit with 3 elements. Assume* *an ordering of* $\mathcal{I}$ *and of* $\mathcal{S} \times \mathcal{G}$ *in which the induction mapping* V *is given by*

$$V = \begin{pmatrix} I_6 & \\ & W \end{pmatrix} \text{ where } W = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & & & 1 & & \\ & 1 & & & & 1 \\ & & 1 & & 1 & \end{pmatrix}^T .$$

We may confirm that $V^T V = I_9$ *and that* $\Pi^2 = \Pi = VV^T$ *is indeed a projection. By applying the GFT to* Π*, we obtain*

$$\hat{\Pi}(\tau) = I_2, \quad \hat{\Pi}(\sigma) = \begin{pmatrix} 1 & \\ & 0 \end{pmatrix}, \text{ and } \hat{\Pi}(\sigma) = \begin{pmatrix} I_3 & \\ & 0 \end{pmatrix}.$$

Obviosly, the ranks for the different blocks of $\hat{\Pi}$ *are 2, 1, and 3, respectively. Consequently, the blocks of* $\hat{A}$ *are* 2×2*,* 1×1*, and* 3×3*, whereas the blocks of* $\hat{x}$ *are* 2×1*,* 1×1*, and* 3×2*.*

3 Application to eigenvalue computations

We will now discuss how to compute eigenvectors and eigenvalues of an equivariant matrix. First, we establish the connection between block diagonalization via a basis shift, Section 2.2, and the GFT, Section 2.3. By considering the GFT of a vector x as a linear mapping between two vector spaces of dimension n, we can obtain $\bar{x}$ from $\hat{x}$ simply by stacking each column of each block $\hat{x}(\rho_i)$ above one another, obtaining $\bar{x} = Fx$. Thus, the entries of F are recognized as the coefficients in the irreducible representations $\rho \in \mathcal{R}$.

Example 8 *The irreducible representations of* $\mathcal{D}_3$ *were defined in Example 3. Collecting their entries, we obtain*

$$F = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & -1 & -1 & -1 \\ 1 & c & c & -1 & -c & -c \\ 0 & -s & s & 0 & s & -s \\ 0 & s & -s & 0 & s & -s \\ 1 & c & c & 1 & c & c \end{pmatrix},$$

where $c = \cos(120°) = -0.5$ *and* $s = \sin(120°) = \sqrt{3}/2$*.*

The interpretation of the GFT as a shift of basis, implies that we may compute $\bar{\mathbf{A}}$ as $\bar{\mathbf{A}} = F\mathbf{A}F^{-1}$. We stress, however, that this is not efficient. Instead, we obtain $\bar{\mathbf{A}}$ from $\hat{A}$, simply by repeating each block $\hat{A}(\rho_i)$ d_i times. Of course, in practical computations we need only form $\hat{A}$. Still, the establishment of this connection makes it simple to deduce the following facts regarding the eigenvalues of $\mathbf{A}$.

Proposition 2 *Let $\mathbf{A}$ be an equivariant matrix under a general action of $\mathcal{G}$.*

1. *Each eigenvalue of $\mathbf{A}$ is also an eigenvalue of $\hat{A}$.*
2. *Each eigenvalue of $\hat{A}(\rho_i)$ is also an eigenvalue of $\mathbf{A}$, repeated d_i times.*
3. *$\hat{A}^+$ has $N-n$ spurious zero eigenvalues which are not eigenvalues of $\mathbf{A}$.*
4. *Each nonzero eigenvalue of $\hat{A}^+(\rho_i)$ is also an eigenvalue of $\mathbf{A}$, repeated d_i times.*

Proof Since $\bar{\mathbf{A}}$ is obtained from $\mathbf{A}$ via a basis shift, their sets of eigenvalues are identical. Statements 1 and 2 follows from the fact that each block in $\hat{A}$ is contained d_i times in $\bar{\mathbf{A}}$. The last two statements are consequences of the isomorphism between $\mathbb{V}$ and $\mathbb{C}^n$; consequently will $\mathbf{A}^+$ and $\mathbf{A}$ have the same rank. Finally, since the GFT is an algebra isomorphism, $\hat{A}^+$ has the same eigenvalues as $\mathbf{A}^+$.

The above proposition immediately yields a bound on the number of distinct eigenvalues.

Proposition 3 *Assume that $\mathbf{A}$ is equivariant under an action with m orbits.*

1. *If the action is free, $\mathbf{A}$ has at most $M = \sum_i m d_i$ distinct eigenvalues.*
2. *If the action is not free, $\mathbf{A}$ has at most $\sum_i r_i < M$ distinct eigenvalues, where r_i is the rank of $\hat{\Pi}(\rho_i)$.*

Eigenvectors In order to compute eigenvectors of an equivariant matrix, we note that an eigenvalue λ of $\hat{A}(\rho_i)$ with eigenvector v is a repeated eigenvalue of $\mathbf{A}$ with multiplicity d_i. To find the d_i eigenvectors which span the corresponding eigenspace, we construct d_i matrices $\hat{x}^{(j)} = v\hat{e}_j^T$ of dimension $r_i \times d_i$ for $j = 1 \ldots d_i$, i.e. matrices whose jth column is equal to v. Since $\hat{A}(\rho_i)\hat{x}^{(j)} = \lambda\hat{x}^{(j)}$, eigenvectors to $\mathbf{A}$ are found by applying the inverse GFT, $x^{(j)} = \text{igft}(\hat{x}^{(j)})$. Note that the $x^{(j)}$ vectors are independent, since $\hat{x}^{(j)}$ obviously are independent. Also notice that the inverse transform is very efficient, since we only need to consider a few columns of F^{-1} in order to compute the inverse transform of an eigenmatrix.

Example 9 *Consider the eigenvalues and eigenvectors of $\mathbf{A} = L(a)$, cf. Example 6. The eigenvalues of the 1×1 blocks $\hat{A}(\tau)$ and $\hat{A}(\sigma)$ are 1, let the corresponding eigenmatrices be $x(\tau) = x(\sigma) = 1$. Thus, the corresponding eigenvectors to $\mathbf{A}$ are simply the first two columns G_1 and G_2 of $G = F^{-1}$.*

$\hat{A}(\rho)$ has $(-1+i\sqrt{3})/2$ as a complex eigenvalue with multiplicity 2, let the eigenvectors be $v = (1,\ \pm i)^T$. Corresponding eigenmatrices to $\hat{A}(\rho)$ are

$$\hat{x}^{(1,2)}(\rho) = \begin{pmatrix} 1 & 0 \\ \pm i & 0 \end{pmatrix} \text{ and } \hat{x}^{(3,4)}(\rho) = \begin{pmatrix} 0 & 1 \\ 0 & \pm i \end{pmatrix}.$$

Eigenvectors to $\mathbf{A}$ are given by multiplying $(G_3,\ G_4)(1,\ \pm i)^T$ and $(G_5,\ G_6)(1,\ \pm i)^T$, again denoting by G_j the jth column of $G = F^{-1}$.

Algorithm summary We summarize the algorithm for finding eigenvalues and eigenvectors of an equivariant matrix.

1. Identify the essential part A of $\mathbf{A}$.
2. Compute $\hat{A}$, the GFT of $\mathbf{A}$.
3. For each block $\hat{A}(\rho_i)$, solve the eigenvalue problem $\hat{A}(\rho_i)v = \lambda v$. Each λ is an eigenvalue to $\mathbf{A}$ of multiplicity d_i.
4. Construct the d_i eigenmatrices $\hat{x}^{(j)}$ and apply the inverse GFT to obtain d_i independent eigenvectors of $\mathbf{A}$.

We note that it is simpler to efficiently implement the GFT for a free action, since the projection matrices need not be considered. Therefore, we point out that if we are only interested in nonzero eigenvalues of $\mathbf{A}$, we should consider the following algorithm.

1. Identify the essential part A of $\mathbf{A}$.
2. Compute $\hat{A}^+$, the GFT of $\mathbf{A}^+$.
3. For each block $\hat{A}^+(\rho_i)$, solve the eigenvalue problem $\hat{A}^+(\rho_i)v = \lambda v$. Each nonzero λ is an eigenvalue to $\mathbf{A}$ of multiplicity d_i.

If we are interested in eigenvectors as well, we may proceed as in the previous case, but we need to project the eigenvectors obtained.

4 Numerical Experiments

We have implemented the GFT for a few different groups, for instance the dihedral groups, the group of the tetrahedron and the group of the cube. Computing eigenvalues by exploiting the GFT yields a substantial gain, particularly for larger groups. Assuming that the free part of the action dominates, we derive the following theoretical estimates.

m	10	20	30	40	50	100	150
n	480	960	1440	1920	2400	4800	7200
t	0.14	0.53	1.21	2.22	3.57	16.38	44.24
t_{gft}	0.11	0.39	0.90	1.60	2.51	10.00	22.51
t_{eig}	0.03	0.14	0.31	0.62	1.06	6.38	21.73
t_{direct}	5.50	64.26	221.88				

Table 1: Timings for computing eigenvalues for an equivariant matrix under the group of the cube, for different number of orbits m, corresponding to an $n \times n$ matrix. The total time for the GFT approach is t (seconds), the time for doing the GFT is t_{gft} and the time for computing the eigenvalues of each block is t_{eig}. The last row illustrates, for small matrices, how much more time it takes to compute eigenvalues without exploiting the equivariance.

The cost of the direct algorithm for computing eigenvalues of $\mathbf{A}$ is $w_{\text{direct}} = k(m \sum_i d_i^2)^3$, while computing the eigenvalues for $\hat{A}$ costs $w_{\text{eig}} = k \sum_i (md_i)^3$. Thus

$$\frac{w_{direct}}{w_{\text{eig}}} = \frac{(\sum_i d_i^2)^3}{\sum_i d_i^3}.$$

For example, consider the triangle, the cube and the maximally symmetric discretization of the sphere (icosahedral symmetry with reflections). We have:

Symmetry	d_i	$\frac{w_{\text{direct}}}{w_{\text{eig}}}$
triangle (6 symm.)	$\{1, 1, 2\}$	22
cube (48 symm.)	$\{1, 1, 1, 1, 2, 2, 3, 3, 3, 3\}$	864
icosahedron (120 symm.)	$\{1, 1, 3, 3, 3, 3, 4, 4, 5, 5\}$	3541

To confirm these estimations, we present numerical experiment for eigenvalue computations of a matrix equivariant under the symmetry of the cube, see Table 1. The experiments was carried out using Matlab. For $m = 30$, we have $t_{\text{direct}}/t_{\text{eig}} = 716$ which confirms the estimations. The discrepancy is mostly due to the fact that the n^3 estimation of the performance of the eigenvalue computations is not valid in this regime of matrix sizes. Due to an inefficient memory access pattern, the GFT in our current Matlab implementaton is relatively slow. We address this issue by rewriting the implementation in C, emphasizing on an efficient memory layout of the data.

5 Conclusions

Equivariant matrices may be block diagonalized via the GFT. We have reviewed this technique and applied it to eigenvalue computations. We have shown how the equivariance restricts the number of distinct eigenvalues, and we have devised an algorithm for using the GFT to compute the eigenvalues much more efficiently than with a direct approach. This is confirmed by numerical computations.

We have also derived algorithms for computing eigenvectors of equivariant matrices. Here, we stress that the application of the inverse GFT is very efficient, since the eigenmatrices in the Fourier space are very sparse.

We are currently investigating how to exploit fast GFT algorithms in parallel, and we are addressing implementation issues [10]. We are also studying the generation of symmetry respecting discretizations [11].

Acknowledgments

This work has been partially supported by a grant from the Göran Gustafsson foundation.

References

[1] M. Bonnet. Exploiting partial or complete geometrical symmetry in 3D symmetric Galerkin indirect BEM formulations. *Intl. J. Numer. Meth. Engng*, 57:1053–1083, 2003.

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

[3] E. L. Allgower and K. Georg. Numerical exploitation of symmetric structures in BEM. In M. Golberg, editor, *Boundary Integral Methods - Numerical and Mathematical Aspects*, pages 289–306. Computational Mechanics Publications, Southampton, 1998.

[4] 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.

[5] A. Bossavit. Boundary value problems with symmetry and their approximation by finite elements. *SIAM J. Appl. Math.*, 53:1352–1380, 1993.

[6] Åhlander, K. and Munthe-Kaas, H. Generalized Fourier Transforms applied to Equivariant Systems. Technical Report 2004-029, Department of Information Technology, Uppsala University, July 2004.

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

[8] Georg K. and Tausch J. A generalized Fourier transform for boundary element methods with symmetries. Technical report, Colorado State University, Ft. Collins, Colorado, 1994.

[9] D. Maslen and D. Rockmore. Generalized FFTs - a survey of some recent results. Technical Report PCS-TR96-281, Dartmouth College, Department of Computer Science, Hanover, NH, 1996.

[10] M. Ljungberg and K. Åhlander. Generic Programming Aspects of symmetry exploiting numerical software. In P. Neittaanmäki et al., editors, *Proceedings of European Congress on Compoutational Methods in Applied Sciences and Engineering, ECCOMAS 2004, Jyväskylä, 24–28 July 2004*, 2004. Also available as Technical Report 2004-020 from the Department of Information Technology, Uppsala University.

[11] K. Åhlander. Mesh generation for symmetrical geometries. Accepted for presentation at CGA'05 and for publication in the conference proceedings of ICCSA'05 (LNCS).