SPARSE GENERALIZED FOURIER TRANSFORMS *

KRISTER ÅHLANDER[1]

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

Abstract.

Block-diagonalization of sparse equivariant discretization matrices is studied. Such matrices typically arise when partial differential equations that evolve in symmetric geometries are discretized via the finite element method or via finite differences.

By considering sparse equivariant matrices as equivariant graphs, we identify a condition for when block-diagonalization via a sparse variant of a generalized Fourier transform (GFT) becomes particularly simple and fast.

Characterizations for finite element triangulations of a symmetric domain are given, and formulas for assembling the block-diagonalized matrix directly are presented. It is emphasized that the GFT preserves symmetric (Hermitian) properties of an equivariant matrix.

By simulating the heat equation at the surface of a sphere discretized by an icosahedral grid, it is demonstrated that the block-diagonalization pays off. The gain is significant for a direct method, and modest for an iterative method.

A comparison with a block-diagonalization approach based upon the continuous formulation is made. It is argued that the sparse GFT method is an appropriate way to discretize the resulting continuous subsystems, since the spectrum and the symmetry are preserved.

Key words: Non commutative Fourier analysis, block-diagonalization, equivariant operators, sparse matrices, finite differences, finite elements.

1 Introduction.

Many applications display various kinds of symmetries, which scientists and engineers strive to exploit in order to develop faster algorithms that consume less memory; see [7, 10, 23] for a few recent examples. In this paper, we are concerned with linear partial differential equations (PDEs) that evolve in symmetrical domains, i.e., geometries that are invariant under a (finite) symmetry group $\mathcal{G}$. We assume that the linear operator of the PDE is *equivariant* with respect to the symmetry group, i.e., the PDE operator and the actions of the group commute.

Two decades ago, Bossavit treated this problem in the context of finite element method (FEM) applications [8]. By adopting a continuous perspective, he showed how to reformulate a problem equivariant under a symmetry group of size $|\mathcal{G}|$ into $|\mathcal{G}|$ smaller subproblems. Each of these subproblems are essentially $|\mathcal{G}|$ times smaller than the original, but several of the subproblems are coupled

*December 2005.

via boundary conditions. The method was refined in [9]. Independently, Allgower and others applied similar ideas to dense equivariant problems mainly arising in the context of boundary element method (BEM) applications [4, 6]. In the latter approach, a discrete point of view is adopted: by assuming a *symmetry respecting discretization* an equivariant PDE operator is discretized into an equivariant system matrix. An equivariant matrix $\bar{A}$ is a matrix for which

$$\bar{A}_{i,j} = \bar{A}_{ig,jg} \tag{1.1}$$

holds for all indices i, j and all $g \in \mathcal{G}$. The notation ig denotes the induced action of the symmetry group on the index set. This property obviously offers the possibility of memory reductions, but it is even more important that an equivariant matrix is block-diagonalized via a generalized Fourier transform (GFT). This leads to a significant speedup for algorithms with N^3 complexity such as matrix exponentiation, dense eigenvalue computations and solution of dense systems of equations, see e.g. [2].

The purpose of the present paper is to apply the GFT approach to *sparse* equivariant systems, including systems arising from FEM discretizations, and to compare this method with the approach of Bossavit. We pay explicit attention to symmetric systems, which are common in applications. By considering sparse equivariant matrices as *equivariant graphs*, we define *closed dependency* as a condition that guarantees a particularly simple and fast GFT of an equivariant graph. We discuss how to exploit geometrical symmetries when assembling a stiffness matrix and under what criteria a stiffness matrix that respects the closed dependency condition is assembled. We provide numerical examples for solving Poisson's equation in a domain with triangular symmetry as well as for the heat equation on a sphere discretized with an icosahedral grid.

The outline of this paper is as follows. The mathematical background is reviewed in Section 2. In Section 3, we propose a definition of symmetry that covers both equivariance and other symmetrical concepts such as Hermitian matrices, and we show that the Hermitian property is preserved when transforming an equivariant matrix. Section 4 focuses on the GFT of equivariant sparse matrices regarded as equivariant graphs. The assembly of equivariant stiffness matrices is studied in Section 5, where different triangulation properties are characterized and their effect on the assembled stiffness matrix is studied. Numerical examples are given in Section 6. The continuous approach is compared with the discrete approach in Section 7. Finally, Section 8 concludes the paper.

2 Background

As a convenience, we summarize briefly in Section 2.1 how representation theory can be used to block-diagonalize equivariant matrices, and in Section 2.2 we illustrate the notation and the approach by simple examples. For a thorough treatment on representation theory, [20] is a classical textbook. Our notation is based on [2].

2.1 Theory

Recall that a group $\mathcal{G}$ with identity e acts from the right on a set $\mathcal{X}$ if $(x, g) \mapsto xg$ in such a way that $xe = x$ and $x(gh) = (xg)h$ for all $x \in \mathcal{X}$ and all $g, h \in \mathcal{G}$. The isotropy subgroup at x is $\mathcal{G}_x = \{g \in \mathcal{G} \mid xg = x\}$. The orbit of x is $x\mathcal{G} = \{y \in \mathcal{X} \mid y = xg, g \in \mathcal{G}\}$. We will frequently refer to a set of orbit representatives as a fundamental set. If all isotropy subgroups are trivial, i.e., $\mathcal{G}_x = \{e\}$ for all x, the action is said to be *free*. We are mostly concerned with finite groups where $|\mathcal{G}| < \infty$. The relation $|\mathcal{G}_x||x\mathcal{G}| = |\mathcal{G}|$ shows that an alternative characterization of a free action is that all orbits have the size of the group.

Recall further the concept of cosets. Let $\mathcal{H}$ and $\mathcal{K}$ be two subgroups of a group $\mathcal{G}$. The *right cosets* $\mathcal{H}\backslash\mathcal{G}$ are the disjunct sets $\mathcal{H}s$ where $s \in \mathcal{G}$. The *left cosets* $\mathcal{G}/\mathcal{H}$ are the disjunct sets $s\mathcal{H}$ where $s \in \mathcal{G}$. The *double cosets* $\mathcal{H}\backslash\mathcal{G}/\mathcal{K}$ are the disjunct sets $\mathcal{H}s\mathcal{K}$ where $s \in \mathcal{G}$.

The basic tool for block-diagonalizing equivariant matrices is the GFT, which relies on representation theory. For a given group $\mathcal{G}$, a complex *representation* ρ of dimension $d = d_\rho$ is a homomorphism $\rho : \mathcal{G} \to \mathbb{C}^{d\times d}$, i.e., for all $g, h \in \mathcal{G}$, $\rho(gh) = \rho(g)\rho(h)$. Two representations ρ and σ are *equivalent* if there exists a shift of basis T such that $\rho(g) = T\sigma(g)T^{-1}$ for all $g \in \mathcal{G}$. A representation ρ is *reducible* if there exists a shift of basis T such that $T\rho(g)T^{-1}$ has a strictly block-diagonal structure for all $g \in \mathcal{G}$. A representation for which such a basis shift does not exist is *irreducible*. For every finite group $\mathcal{G}$, there exists a complete list $\mathcal{R}$ of nonequivalent irreducible representations for which $\sum_{\rho\in\mathcal{R}} d_\rho^2 = |\mathcal{G}|$. The list of nonequivalent irreducible representations is unique up to equivalence. For complex representations, it is always possible (see, e.g., [20, Section 2.3]) to make the list *unitary*, which means that $\rho(g) = \rho(g^{-1})^H$ for all $\rho \in \mathcal{R}$.

The GFT is an isomorphism between the group algebra $\mathbb{C}\mathcal{G}$ where the group elements are a basis and the corresponding Fourier algebra $\widehat{\mathbb{C}\mathcal{G}}$. The *group algebra* $\mathbb{C}\mathcal{G}$ is the vector space $\mathbb{C}^{|\mathcal{G}|}$ equipped with the convolution product

$$(x * y)(g) = \sum_{h\in\mathcal{G}} x(gh^{-1})y(h).$$

The Fourier algebra is the block-diagonal matrix algebra $\bigoplus_{\rho\in\mathcal{R}} \mathbb{C}^{d_\rho\times d_\rho}$. The different blocks of $\hat{x} \in \widehat{\mathbb{C}\mathcal{G}}$ are referred to as $\hat{x}(\rho),\ \rho \in \mathcal{R}$. The GFT is given by

$$\hat{x}(\rho) = \mathrm{gft}(x) = \sum_{g\in\mathcal{G}} x(g)\rho(g), \tag{2.1}$$

for $x \in \mathbb{C}\mathcal{G}$ and $\rho \in \mathcal{R}$, and the IGFT is given by

$$x(g) = \mathrm{igft}(\hat{x}) = \frac{1}{|\mathcal{G}|} \sum_{\rho\in\mathcal{R}} d_\rho \mathrm{trace}\left(\hat{x}(\rho)\rho(g^{-1})\right), \tag{2.2}$$

for $\hat{x} \in \widehat{\mathbb{C}\mathcal{G}}$ and $g \in \mathcal{G}$[1]. Computations of the above transforms can be carried

[1]We remark that we use capital letters for the abbreviation GFT, whereas the mapping gft is written in lowercase.

out more efficiently by using so called *fast GFTs*, see e.g. [17] for a survey. For block-diagonalization of discretization matrices, the discussions in e.g. [2] and the results in Section 6.2 show that the gain of the block-diagonalization itself often is much bigger than the cost for applying the transform. Therefore, fast GFTs are not studied further here, even though the topic is interesting.

Block versions of the (I)GFT are used for block-diagonalizing equivariant matrices. Let $\mathbb{C}^{m\times n}\mathcal{G} = \mathbb{C}^{m\times n}\otimes\mathbb{C}\mathcal{G}$, and let $\widehat{\mathbb{C}^{m\times n}\mathcal{G}} = \mathbb{C}^{m\times n}\otimes\widehat{\mathbb{C}\mathcal{G}}$. Lift the GFT and the IGFT component-wise, and define a block-convolution $*$: $\mathbb{C}^{m\times n}\mathcal{G}\times\mathbb{C}^{n\times p}\mathcal{G}\to\mathbb{C}^{m\times p}$ by

$$(A*B)_{i,j}(g) = \sum_k\sum_h A_{i,k}(gh^{-1})B_{k,j}(h).$$

The block-GFT then becomes an algebra isomorphism that preserves the convolution product.

In order to treat equivariant matrices (1.1), the idea is to express the matrix-vector multiplication as a block-convolution. Therefore, let $\mathcal{E}_{\mathcal{G}}(\mathbb{C}^{\bar m})$ denote the space of equivariant $\mathbb{C}^{\bar m\times\bar m}$ matrices, and assume that the group action partitions the index set $\overline{\mathcal{I}}$ of size $\bar m$ into m different orbits, represented by the fundamental index set $\mathcal{I}$.[2] Define $\mu:\mathbb{C}^{\bar m}\to\mathbb{C}^m\mathcal{G}$, $\bar x\mapsto x$, and $\nu:\mathcal{E}_{\mathcal{G}}(\mathbb{C}^{\bar m})\to\mathbb{C}^{m\times m}\mathcal{G}$, $\bar A\mapsto A$ by

$$x_i(g) = \frac{1}{\sqrt{|\mathcal{G}_i|}}\bar x_{ig}, \tag{2.3}$$

$$A_{i,j}(g) = \frac{1}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}\bar A_{ig,j}. \tag{2.4}$$

By these mappings, $\bar A\bar x$ is mapped into $A*x$. The choice is not unique, as remarked in [2]. As shown in Section 3.2, the mappings above have the advantage that symmetry is preserved. The image $\mathcal{V}$ of μ equals $\mathbb{C}\mathcal{G}$ if the action is free, but in general it is a pure subset of $\mathbb{C}\mathcal{G}$. Theory for the general case was first developed in [6]. It can be explained by studying how the identity matrix is mapped. Let $\pi=\nu(I_{\bar m})$ and $\hat\pi=\mathrm{gft}(\pi)$. In the Fourier space, $\hat\pi$ is an orthogonal projection onto $\widehat{\mathcal{V}}=\mathrm{gft}(\mathcal{V})$. Factorize $\hat\pi=\varphi\varphi^T$ such that $\varphi^T\varphi=I_{\bar m}$ is an identity matrix. Let $\tilde x=\varphi^T\hat x\in\widetilde{\mathbb{C}^m\mathcal{G}}$ and let $\tilde A=\varphi^T\hat A\varphi\in\widetilde{\mathbb{C}^{m\times m}\mathcal{G}}$. By composing mappings, we obtain $\mathrm{gft}:\mathbb{C}^{\bar m}\to\widetilde{\mathbb{C}^m\mathcal{G}}$ and $\mathrm{gft}:\mathcal{E}_{\mathcal{G}}(\mathbb{C}^{\bar m})\to\widetilde{\mathbb{C}^{m\times m}\mathcal{G}}$. See (2.6)–(2.9) below for explicit formulas. Inverse mappings are defined accordingly. The situation is summarized in Figure 2.1 and by the following theorem.

THEOREM 2.1. *The vector spaces* $\mathbb{C}^{\bar m}$, $\mathcal{V}$, $\widehat{\mathcal{V}}$, *and* $\widetilde{\mathbb{C}^m\mathcal{G}}$ *are isomorphic. The GFT mappings preserve relevant structures; specifically is* $\mathrm{gft}(\bar A\bar x)=\mathrm{gft}(\bar A)\mathrm{gft}(\bar x)$ *and* $\mathrm{igft}(\mathrm{gft}(\bar x))=\bar x$.

We refer to [2] for more information. Before finishing our brief summary, we give some more details on the projections π and $\hat\pi$, the structure of $\widetilde{\mathbb{C}^m\mathcal{G}}$ and $\widetilde{\mathbb{C}^{m\times m}\mathcal{G}}$, and explicit formulas for the GFTs. The projection π written as an

[2] Our notational convention is that quantities with a "bar" refer to the original space.

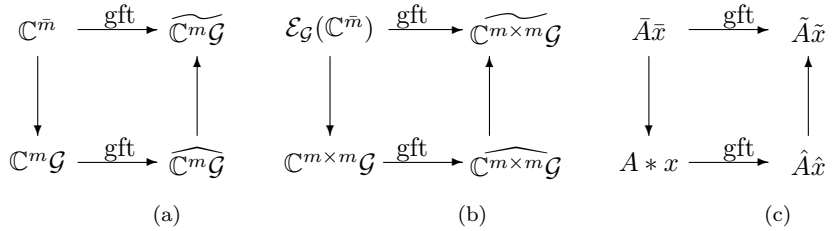

Figure 2.1: The relevant mappings for (a) vectors and (b) equivariant matrices preserve algebraic structures, specifically matrix vector product (c).

element in $\mathbb{C}^{m\times m}\mathcal{G}$ is $\pi_{i,j}(g) = \frac{1}{|\mathcal{G}_i|}$, if $i = j$ and $g \in \mathcal{G}_i$, and 0 else. Its Fourier transform $\hat{\pi}_{i,j}(\rho)$ is thus 0 for $i \neq j$ and for $i = j$ it becomes

$$\hat{\pi}_i(\rho) = \hat{\pi}_{i,i}(\rho) = \frac{1}{|\mathcal{G}_i|}\sum_{h\in\mathcal{G}_i}\rho(h). \tag{2.5}$$

Factorize $\hat{\pi}_i(\rho) = \varphi_i(\rho)\varphi_i(\rho)^T$ such that $\varphi_i(\rho)^T\varphi_i(\rho)$ is an identity matrix of rank $m_{i,\rho}$, where $m_{i,\rho}$ is the rank of $\hat{\pi}_i(\rho)$. Let $m_\rho = \sum_{i\in\mathcal{I}} m_{i,\rho}$. From these expressions we can describe the block-diagonal matrix spaces $\widetilde{\mathbb{C}^m\mathcal{G}}$ and $\widetilde{\mathbb{C}^{m\times m}\mathcal{G}}$:

$$\widetilde{\mathbb{C}^m\mathcal{G}} = \bigoplus_{\rho\in\mathcal{R}}\mathbb{C}^{m_\rho\times d_\rho},\ \widetilde{\mathbb{C}^{m\times m}\mathcal{G}} = \bigoplus_{\rho\in\mathcal{R}}\mathbb{C}^{m_\rho\times m_\rho}.$$

Explicit block formulas for the GFT follow from (2.1), (2.3), and (2.4).

$$\hat{x}_i(\rho) = \frac{1}{\sqrt{|\mathcal{G}_i|}}\sum_{g\in\mathcal{G}}\bar{x}_{ig}\rho(g), \tag{2.6}$$

$$\tilde{x}_i(\rho) = \frac{1}{\sqrt{|\mathcal{G}_i|}}\sum_{g\in\mathcal{G}}\bar{x}_{ig}\varphi_i(\rho)^T\rho(s), \tag{2.7}$$

$$\hat{A}_{i,j}(\rho) = \frac{1}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}\sum_{g\in\mathcal{G}}\bar{A}_{ig,j}\rho(g), \tag{2.8}$$

$$\tilde{A}_{i,j}(\rho) = \frac{1}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}\sum_{g\in\mathcal{G}}\bar{A}_{ig,j}\varphi_i(\rho)^T\rho(g)\varphi_j(\rho). \tag{2.9}$$

For later reference, we provide the following proposition.

PROPOSITION 2.2.

- *In $\mathbb{C}^m\mathcal{G}$, x_i is constant on right cosets of $\mathcal{G}_i$, i.e., if g and h are in the same coset $\mathcal{G}_i s$, $x_i(g) = x_i(h)$.*
- *In $\mathbb{C}^{m\times m}\mathcal{G}$, $A_{i,j}$ is constant on double cosets of $\mathcal{G}_i$ and $\mathcal{G}_j$, i.e., if g and h are in the same double coset $\mathcal{G}_i s\mathcal{G}_j$, $A_{i,j}(g) = A_{i,j}(h)$.*

PROOF. The first statement follows immediately from (2.3).The second follows from (2.4) and the equivariance (1.1). ∎

2.2 Illustration

To illustrate the method and the notation, we transform two matrices that are equivariant under the simplest nontrivial group. Let $\mathcal{G} = \mathcal{C}_2 = \langle c \mid c^2 = e\rangle = \{e, c\}$. The representations are $\mathcal{R} = \{\tau, \sigma\}$ where $\tau(c) = 1$ and $\sigma(c) = -1$. Note that $d_\tau = d_\sigma = 1$; representations are, in fact, always one-dimensional for abelian groups. Let $\overline{\mathcal{I}} = \{1, \dots, 5\}$. Define an action $\overline{\mathcal{I}} \times \mathcal{C}_2 \to \overline{\mathcal{I}}$ by $\{1, 2, 3, 4, 5\}c = \{5, 4, 3, 2, 1\}$.

The matrices

$$\bar{A} = \begin{pmatrix} -2 & 1 & & & \\ 1 & -2 & 1 & & \\ & 1 & -2 & 1 & \\ & & 1 & -2 & 1 \\ & & & 1 & -2 \end{pmatrix}, \quad \bar{B} = \begin{pmatrix} 1 & & & & 2 \\ 3 & 4 & 5 & & \\ & & 6 & & \\ & & 5 & 4 & 3 \\ 2 & & & & 1 \end{pmatrix} \tag{2.10}$$

are both equivariant with respect to this action[3]. A fundamental set may be chosen as $\mathcal{I} = \{1, 2, 3\}$. Isotropy subgroups are $\mathcal{G}_1 = \mathcal{G}_2 = \{e\}$ and $\mathcal{G}_3 = \{e, c\}$. The projection matrices are 1 except for $\hat{\pi}_3(\sigma)$ which is 0. The factorizations $\hat{\pi} = \varphi\varphi^T$ are thus trivial.

Applying the formulas to $\bar{B}$ yield

$$B(e) = \begin{pmatrix} 1 & & \\ 3 & 4 & \frac{5}{\sqrt{2}} \\ & & 3 \end{pmatrix}, \quad B(c) = \begin{pmatrix} 2 & & \\ & & \frac{5}{\sqrt{2}} \\ & & 3 \end{pmatrix},$$

$$\hat{B}(\tau) = \begin{pmatrix} 3 & & \\ 3 & 4 & \frac{10}{\sqrt{2}} \\ & & 6 \end{pmatrix}, \quad \hat{B}(\sigma) = \begin{pmatrix} -1 & & \\ 3 & 4 & \\ & & 0 \end{pmatrix}.$$

The blocks of $\tilde{B}$ are $\tilde{B}(\tau) = \hat{B}(\tau)$ and $\tilde{B}(\sigma)$, which is equal to $\hat{B}(\sigma)$ with the last row and column omitted. Similarly, transforming $\bar{A}$ gives

$$\tilde{A}(\tau) = \begin{pmatrix} -2 & 1 & \\ 1 & -2 & \sqrt{2} \\ & \sqrt{2} & -2 \end{pmatrix}, \quad \tilde{A}(\sigma) = \begin{pmatrix} -2 & 1 \\ 1 & -2 \end{pmatrix}.$$

Note that

$$\tilde{A} = \begin{pmatrix} \tilde{A}(\tau) & \\ & \tilde{A}(\sigma) \end{pmatrix} \tag{2.11}$$

is a block-diagonalized version of the original matrix $\bar{A}$. It is symmetric and has the same spectrum as $\bar{A}$, see e.g. [3].

[3] $\bar{A}$ displays, of course, other symmetries as well. The motivation behind this example is revealed in Secion 7.

3 Symmetry and equivariance

The equivariance property (1.1) describes a certain symmetry of a matrix. Minkwitz [18] generalized to the case where the action on row indices and column indices differ. This notion is also used e.g. in [5]. However, this generalization does not capture other familiar symmetries, for example Hermitian matrices. In Section 3.1 we propose a generalized symmetry definition that combines the definition of Minkwitz with ideas from [1]. Section 3.2 shows that the GFT transforms an equivariant Hermitian matrix into a block-diagonal Hermitian matrix.

3.1 Generalized symmetry

Let us consider equivariant matrices, symmetric matrices ($B = B^T$), and Hermitian matrices ($B = B^H$). We notice that both equivariance and symmetry could be described by the following definition of a symmetric table.

> A symmetric table $A : \mathcal{J} \to \mathcal{D}$ is a table that is constant on orbits of an action $\mathcal{J} \times \mathcal{G} \to \mathcal{J}$.

For square matrices, the index space is $\mathcal{J} = \mathcal{I} \times \mathcal{I}$. In case of symmetry, the group is $\mathcal{C}_2 = \{e, c\}$ and the action is given by $(i, j)c = (j, i)$. In case of equivariance, the action of $g \in \mathcal{G}$ is $(i, j) \mapsto (ig, jg)$.

However, this definition does not capture the case of a Hermitian matrix. Therefore, we generalize the idea and assume an action $* : \mathcal{G} \times \mathcal{D} \to \mathcal{D}$ on the matrix elements.

Definition 3.1. *A generally symmetric table $A : \mathcal{J} \to \mathcal{D}$ is a table where $A(jg) = A(j) * g$.*

Hermitian matrices are then recognized as generally symmetric tables under $\mathcal{C}_2$, where c acts by transposition on the indices and by conjugation $zc = z^*$ on the matrix elements.

We point out that symmetry is a concept of many different nuances. The purpose of this definition is to clarify that orbits may depend both on "traditional" symmetry and on equivariance, an observation that may be exploited by numerical algorithms. The following proposition is an immediate consequence.

Proposition 3.1. *Assume that $A : \mathcal{J} \to \mathcal{D}$ is a table that respects a generalized symmetry under $\mathcal{G}$. If two indices i and j belong to the same orbit when $\mathcal{G}$ acts on $\mathcal{J}$, then $A(i)$ and $A(j)$ belong to the same orbit when $\mathcal{G}$ acts on $\mathcal{D}$.*

3.2 Block-diagonalization of equivariant Hermitian matrices

In section 2.1, we reviewed how to block-diagonalize equivariant matrices. In this section, we study additional effects of traditional symmetry. We think it is useful to state these results explicitly, even though they are implicit in previous work, e.g. [6].

Proposition 3.2. *Let $\bar{A}$ be a Hermitian and equivariant matrix, and let $\bar{A} \mapsto A \mapsto \hat{A} \mapsto \tilde{A}$. Then, $A(g) = A(g^{-1})^H$. If the representations $\rho \in \mathcal{R}$ are unitary, $\hat{A}$ and $\tilde{A}$ are Hermitian.*

PROOF. Using (2.4), we obtain

$$A_{i,j}(g) = \frac{\bar{A}_{ig,j}}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}} = \frac{\bar{A}_{i,jg^{-1}}}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}} = \frac{\bar{A}^*_{jg^{-1},i}}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}} = A_{j,i}(g^{-1})^*.$$

The GFT and the unitary representations give

$$\hat{A}_{i,j}(\rho) = \sum_{g\in\mathcal{G}} A_{i,j}(g)\rho(g) = \sum_{g\in\mathcal{G}} A_{i,j}(g^{-1})\rho(g^{-1}) = \sum_{g\in\mathcal{G}} A_{j,i}(g)^*\rho(g)^H.$$

Thus $\hat{A} = \hat{A}^H$. Furthermore, we have $\tilde{A} = \varphi^T\hat{A}\varphi$. Complex conjugation and transposing yields $\tilde{A} = \tilde{A}^H$.

□

We remark that a real list of representations with the property $\rho(g) = \rho(g^{-1})^T$ for all $\rho \in \mathcal{R}$ would preserve the symmetry of a real matrix $\bar{A}$. For a given group $\mathcal{G}$, it is not certain that such a list of representations can be found, see e.g. [16]. Many important subgroups of $\mathcal{O}(d)$ do however admit such lists of representations, e.g., the dihedral groups and the invariance group of the icosahedron. See the examples in Section 6.

4 Sparse equivariant matrices

We now consider *sparse* equivariant matrices. Throughout this section, let $\bar{A} \in \mathbb{C}^{\bar{m}\times\bar{m}}$ be a sparse matrix that is equivariant under an action $\overline{\mathcal{I}} \times \mathcal{G} \to \overline{\mathcal{I}}$, and let $\mathcal{I}$ denote a fundamental index set.

Sparse matrices are often conveniently represented as graphs. Considering $\bar{A}$ as a graph, $\overline{\mathcal{I}}$ is its vertices and $\bar{A}_{i,j} \neq 0$ is the edge from vertex i to vertex j. For nonsymmetric matrices, this edge is directed, whereas it is undirected for symmetric matrices.

DEFINITION 4.1. *For $i, j \in \overline{\mathcal{I}}$, j is* dependent *on i, denoted $i \to j$, if $\bar{A}_{i,j} \neq 0$.*

Since we consider equivariant graphs, Proposition 2.2 shows that if g and h belong to the same double coset $\mathcal{G}_i s\mathcal{G}_j$, then is $ig \to j$ iff $ih \to j$. Consequently $i \to j$ implies $ig \to j$ for all $g \in \mathcal{G}_i\mathcal{G}_j = \mathcal{G}_i e\mathcal{G}_j$. The following definition characterizes indices that are dependent "at a distance".

DEFINITION 4.2. *For $i, j \in \overline{\mathcal{I}}$, j is* distantly dependent *on i, denoted $i \overset{\text{dist}}{\to} j$, if $ig \to j$ for some $g \notin \mathcal{G}_i e\mathcal{G}_j$.*

We will derive an efficient formula for how to compute the GFT $\hat{A}_{i,j}$ of edges $\bar{A}_{i,j}$ where $i \overset{\text{djst}}{\not\to} j$. First, we study the set $\mathcal{G}_i\mathcal{G}_j$ in more detail.

LEMMA 4.1. *Each element $g \in \mathcal{G}_i\mathcal{G}_j$ can be written uniquely as $g = sht$, where $h \in \mathcal{H} = \mathcal{G}_i \cap \mathcal{G}_j$, s is a representative for the left cosets $\mathcal{G}_i/\mathcal{H}$, and t is a representative for the right cosets $\mathcal{H}\backslash\mathcal{G}_j$.*

PROOF. Any element $g \in \mathcal{G}_i\mathcal{G}_j$ is a product g_1g_2 for $g_1 \in \mathcal{G}_i$ and $g_2 \in \mathcal{G}_j$. It is easy to show that $\mathcal{H}$ is a subgroup of both $\mathcal{G}_i$ and $\mathcal{G}_j$. The construction of the cosets implies that $g_1 \in \mathcal{G}_i$ can be written uniquely as sh_1 for some coset representative s and some $h_1 \in \mathcal{H}$. By the same argument, $g_2 \in \mathcal{G}_j$ can be

written uniquely as $h_2 t$ for some $h_2 \in \mathcal{H}$ and some coset representative t. Thus, $g = sht$ for a unique $h = h_1 h_2 \in \mathcal{H}$. □

Summation of representations over $\mathcal{G}_i \mathcal{G}_j$ is now readily derived.

LEMMA 4.2. *Let ρ be a representation of $\mathcal{G}$. Let $\mathcal{G}_i$ and $\mathcal{G}_j$ be subgroups of $\mathcal{G}$. Then is*

$$\sum_{g \in \mathcal{G}_i \mathcal{G}_j} \rho(g) = \frac{|\mathcal{G}_i||\mathcal{G}_j|}{|\mathcal{G}_i \cap \mathcal{G}_j|} \hat{\pi}_i(\rho) \hat{\pi}_j(\rho),$$

where $\hat{\pi}_i$ and $\hat{\pi}_j$ are projections given by (2.5).

PROOF. Let $\mathcal{H} = \mathcal{G}_i \cap \mathcal{G}_j = \{h\}$ and let $\{s\}$ and $\{t\}$ be coset representatives for $\mathcal{G}_i/\mathcal{H}$ and $\mathcal{H}\backslash\mathcal{G}_j$, respectively. By the previous lemma,

$$\sum_{g \in \mathcal{G}_i \mathcal{G}_j} \rho(g) = \sum_s \sum_h \sum_t \rho(sht) = |\mathcal{H}| \sum_s \rho(s) \left(\frac{1}{|\mathcal{H}|} \sum_h \rho(h) \right) \sum_t \rho(t).$$

The expression in parenthesis is recognized—see (2.5)—as a projection $\hat{\pi}_{\mathcal{H}}$. By using $\hat{\pi}_{\mathcal{H}} = \hat{\pi}_{\mathcal{H}} \hat{\pi}_{\mathcal{H}}$ we obtain

$$\begin{aligned} \sum_{g \in \mathcal{G}_i \mathcal{G}_j} \rho(g) &= \frac{1}{|\mathcal{H}|} \sum_s \rho(s) \sum_h \rho(h) \sum_{h'} \rho(h') \sum_t \rho(t) \\ &= \frac{1}{|\mathcal{H}|} \sum_{s,h} \rho(sh) \sum_{h',t} \rho(h't) \;=\; \frac{1}{|\mathcal{H}|} \sum_{g \in \mathcal{G}_i} \rho(g) \sum_{g' \in \mathcal{G}_j} \rho(g'). \end{aligned}$$

By using (2.5) again, the lemma follows. □

The main result of this section is how to compute the GFT of edges that are not distantly dependent.

THEOREM 4.3. *Let $\bar{A}$ be an equivariant matrix. For $i, j \in \mathcal{I}$ where $i \overset{\text{dist}}{\not\mapsto} j$, we have*

$$\hat{A}_{i,j}(\rho) = \frac{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}{|\mathcal{G}_i \cap \mathcal{G}_j|} \bar{A}_{i,j} \hat{\pi}_i(\rho) \hat{\pi}_j(\rho) \tag{4.1}$$

$$\tilde{A}_{i,j}(\rho) = \frac{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}{|\mathcal{G}_i \cap \mathcal{G}_j|} \bar{A}_{i,j} \varphi_i(\rho)^T \varphi_j(\rho) \tag{4.2}$$

PROOF. Since $i \overset{\text{dist}}{\not\mapsto} j$, $\bar{A}_{ig,j} = 0$ for $g \neq \mathcal{G}_i \mathcal{G}_j$, and because of Proposition 2.2 we have $\bar{A}_{ig,j} = \bar{A}_{i,j}$ for $g \in \mathcal{G}_i \mathcal{G}_j$. Thus,

$$\hat{A}_{i,j}(\rho) = \sum_{g \in \mathcal{G}_i \mathcal{G}_j} A_{i,j}(g) \rho(g) = \frac{\bar{A}_{i,j}}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}} \sum_{g \in \mathcal{G}_i \mathcal{G}_j} \rho(g).$$

To see (4.1), we use Lemma 4.2 and obtain

$$\frac{\bar{A}_{i,j}}{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}} \frac{|\mathcal{G}_i||\mathcal{G}_j|}{|\mathcal{G}_i \cap \mathcal{G}_j|} \hat{\pi}_i(\rho) \hat{\pi}_j(\rho) = \frac{\sqrt{|\mathcal{G}_i||\mathcal{G}_j|}}{|\mathcal{G}_i \cap \mathcal{G}_j|} \bar{A}_{i,j} \hat{\pi}_i(\rho) \hat{\pi}_j(\rho).$$

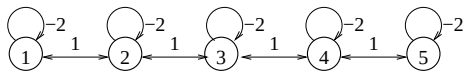

(a)

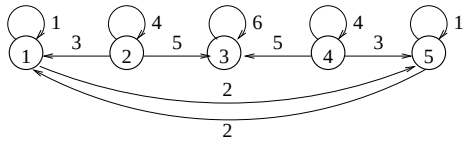

(b)

Figure 4.1: The equivariant matrices $\bar{A}$ and $\bar{B}$ of (2.10) represented as graphs. The graph in (b) is not closely dependent because $1 \overset{\text{dist}}{\rightarrow} 5$. In (a), $\mathcal{I} = \{1, 2, 3\}$ constitute a closely dependent fundamental set, whereas $\{2, 3, 5\}$ is a fundamental set that is not closely dependent.

Since $\varphi_i(\rho)^T \hat{\pi}_i(\rho) = \hat{\pi}_i(\rho)$ and $\hat{\pi}_j(\rho)\varphi_j(\rho) = \varphi_i(\rho)$, (4.2) follows. □

When computing the GFT of a sparse equivariant matrix, it is obviously desirable to have many distantly independent fundamental indices. This motivates the following definition.

DEFINITION 4.3. *If $i \overset{\text{dist}}{\not\mapsto} j$ for all $i, j \in \mathcal{I}$, $\mathcal{I}$ is a* closely dependent fundamental set. *If $\bar{A}$ admits a closely dependent fundamental set, $\bar{A}$ is* closely dependent.

Note that even if $\bar{A}$ is closely dependent, it does not imply that all choices of fundamental index sets $\mathcal{I}$ are closely dependent. Figure 4.1 illustrates these concepts.

5 Finite element assembly of equivariant matrices

Let us now consider the assembly of a linear operator that is equivariant w.r.t. reflections and rotations in a symmetric domain. First, we review the assembly process; see any FEM textbook for more details, e.g. [15]. Next, we discuss implications of equivariance and how to exploit these in the assembly process, similar to the considerations in [9]. Continuing—and this is novel—we characterize different kinds of symmetry respecting discretizations. Particularly, we identify triangulations that assert that the assembled matrix becomes closely dependent. We remark on connections between continuous symmetry and discrete symmetry. Finally, we emphasize that the GFT of an equivariant matrix is very fast.

5.1 Finite element assembly

To establish the notation, we consider the assembly of $-\nabla^2$ in a domain $\overline{\Omega}$ with closure $\mathrm{cl}(\overline{\Omega})$, discretized into a triangulation $\overline{\mathcal{T}}$. Each element $t \in \overline{\mathcal{T}}$ has a set $\mathrm{ixs}(t)$ of corner indices, $\mathrm{ixs}(t) \subset \overline{\mathcal{I}}$. In case of 2D triangles, $\mathrm{ixs}(t)$ has three elements. Each index $i \in \overline{\mathcal{I}}$ corresponds uniquely to a point $p_i \in \overline{\mathcal{P}}$. In total, we have $|\overline{\mathcal{T}}|$ elements and $\bar{m} = |\overline{\mathcal{I}}| = |\overline{\mathcal{P}}|$ indices. Figure 5.1 illustrates various 2D triangulations. We assume a set of basis functions ϕ_i where $\phi_i(p_j) = \delta_{i,j}$ for all $i, j \in \overline{\mathcal{I}}$, and we assume that ϕ_i has support only in elements t for which $i \in \mathrm{ixs}(t)$. The elements of the stiffness matrix $\bar{A} \in \mathbb{C}^{\bar{m}\times\bar{m}}$ are defined by $\bar{A}_{i,j} = a(\phi_i, \phi_j)$, for all $i, j \in \overline{\mathcal{I}}$, where $a(u, v)$ is given by the integral $\int_{\overline{\Omega}} \nabla u \cdot \nabla v^* \, dx$. The standard assembly process computes the integrals over each element. Denote the element stiffness matrix for element t by

$$a^t_{i,j} = \int_t \nabla\phi_i \cdot \nabla\phi^*_j \, dx \text{ for all } i, j \in \overline{\mathcal{I}}, \tag{5.1}$$

and note that $a^t_{i,j}$ is zero unless i and j are both members of $\mathrm{ixs}(t)$. By assembling contributions from all element stiffness matrices, we obtain the total stiffness matrix $\bar{A} = \mathrm{assemble}(\overline{\mathcal{T}})$ by:

$$\bar{A}_{i,j} = \sum_{t\in\overline{\mathcal{T}}} a^t_{i,j} = \sum_{t\in\overline{\mathcal{T}} \;|\; t\ni i,j} a^t_{i,j}.$$

The stiffness matrix is sparse and symmetric. If the geometry is symmetric, it can be triangulated such that $\bar{A}$ becomes equivariant. Exploitation of equivariance, symmetry, and sparsity is studied next.

5.2 Assembly over a symmetric region

Consider the assembly of $\bar{A}$ over a symmetric domain $\overline{\Omega}$ that admits a group action $\overline{\Omega} \times \mathcal{G} \to \overline{\Omega}$. Throughout this section, we assume that $\mathcal{G}$ is a finite subgroup of $\mathcal{O}(d)$ for d dimensions. Thus, the group actions preserve distance. In order to exploit the symmetries of the domain, it is important to have a symmetry respecting triangulation.

DEFINITION 5.1. *A triangulation $\overline{\mathcal{T}}$ of a symmetric domain $\overline{\Omega}$ is* symmetry respecting *if the group action on the domain induces group actions $\overline{\mathcal{T}} \times \mathcal{G} \to \overline{\mathcal{T}}$, $\overline{\mathcal{P}} \times \mathcal{G} \to \overline{\mathcal{P}}$, $\overline{\mathcal{I}} \times \mathcal{G} \to \overline{\mathcal{I}}$. If the action $\overline{\mathcal{T}} \times \mathcal{G} \to \overline{\mathcal{T}}$ is free, we say that the triangulation is* freely symmetry respecting.

PROPOSITION 5.1. *The stiffness matrix $\bar{A} = \mathrm{assemble}(\overline{\mathcal{T}})$ becomes equivariant and symmetric if $\overline{\mathcal{T}}$ is symmetry respecting.*

PROOF. Symmetry follows trivially from the symmetry of the element stiffness matrices, see (5.1). Equivariance is shown e.g. in [9]; it follows from

$$\bar{A}_{i,j} = \sum_{t\in\overline{\mathcal{T}}} a^t_{i,j} = \sum_{t\in\overline{\mathcal{T}}} a^{tg}_{ig,jg} = \sum_{s\in\overline{\mathcal{T}}} a^s_{ig,jg} = \bar{A}_{ig,jg}.$$

Note that $a^{tg}_{ig,jg} = a^t_{i,j}$ holds because group actions are assumed to preserve distance. ∎

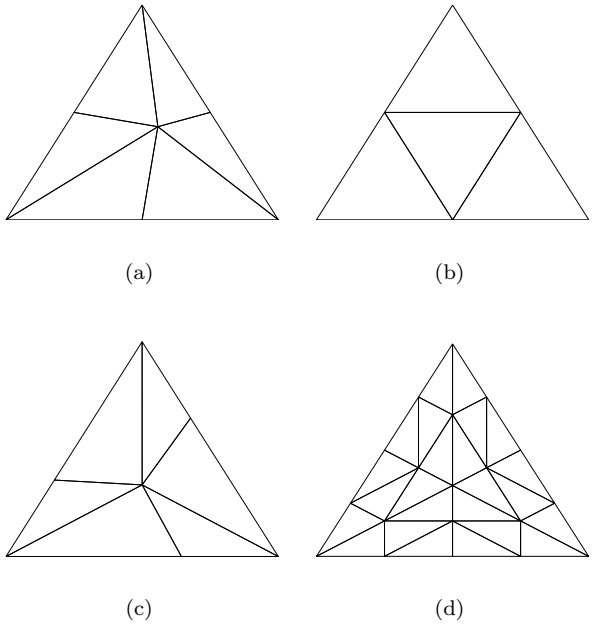

Figure 5.1: Various triangulations of an equilateral triangle that is symmetric under $\mathcal{C}_3$ as well as under $\mathcal{D}_3$. In (a) the triangulation breaks symmetry. In (b) the triangulation respects symmetry, but not freely. In (c), the symmetry with respect to $\mathcal{D}_3$ is broken, but the symmetry with respect to $\mathcal{C}_3$ remains. The triangulation is not closely adjacent symmetry respecting, though. The triangulation in (d) is a closely adjacent symmetry respecting triangulation with respect to $\mathcal{D}_3$.

Fundamental sets of indices, points, and elements may be chosen in various ways, c.f. the example in Figure 4.1. We find it convenient to make the following assumption.

ASSUMPTION 5.1. *Let a symmetric domain $\overline{\Omega}$ be discretizaed by a symmetry respecting triangulation $\overline{\mathcal{T}}$, with points $\overline{\mathcal{P}}$, indexed by $\overline{\mathcal{I}}$. We assume that the corresponding fundamental sets considered as sets in $\mathbb{R}^d$ are related via $\mathcal{P} \subseteq \mathrm{cl}(\Omega) \subseteq \mathcal{T}$. We also assume that $\mathcal{I} = \mathcal{P}$.*

In order to study the relationship between the triangulation and the sparse graph of an assembled stiffness matrix, we commence with a few definitions. First, we collect definitions in immediate correspondance with those in Section 4; next, we introduce some concepts that are useful in order to describe triangulation properties.

DEFINITION 5.2. *Let $i, j \in \overline{\mathcal{I}}$. Then, i and j are* adjacent*, denoted $i \bowtie j$, if there exists a element t such that* $\mathrm{ixs}(t) \ni i, j$.[4] *The indices i and j are* distantly

[4]More generally, i and j are adjacent if the the support of basis function ϕ_i overlaps the support of basis function ϕ_j.

adjacent, *denoted* $i \overset{\text{dist}}{\bowtie} j$, *if* $ig \bowtie j$ *for some* $g \notin \mathcal{G}_i e \mathcal{G}_j$. *If* $i \overset{\text{dist}}{\not\bowtie} j$ *for all* $i, j \in \mathcal{I}$, $\mathcal{I}$ *is a* closely adjacent fundamental set. *If* $\overline{\mathcal{T}}$ *admits a closely adjacent fundamental set* $\mathcal{I}$, $\overline{\mathcal{T}}$ *is a* closely adjacent *symmetry respecting triangulation.*

Definition 5.3. *For* $i \in \overline{\mathcal{I}}$, $\text{adj}(i) = \{k \in \overline{\mathcal{I}} \mid i \bowtie k\}$ *is the* adjacency set *of* i. *An index* $i \in \overline{\mathcal{I}}$ *is* internal *if* $\text{adj}(i) \in \mathcal{I}$. *An index* $i \in \overline{\mathcal{I}}$ *is* isotropically internal *if* $(\text{adj}(i) \cap \mathcal{I})\mathcal{G}_i = \text{adj}(i)$.

Consider edge dependencies in $\bar{A} = \text{assemble}(\overline{\mathcal{T}})$. Equation (5.1) shows that $a^t_{i,j} = 0$ if $i \not\bowtie j$. For $i, j \in \overline{\mathcal{I}}$, $i \overset{\text{dist}}{\not\bowtie} j$ thus implies $i \overset{\text{dist}}{\not\mapsto} j$. Hence, a closely adjacent symmetry respecting triangulation leads to a closely dependent stiffness matrix. We state this as a proposition, and we give a criterion for when the triangulation is closely adjacent.

Proposition 5.2. *If* $\overline{\mathcal{T}}$ *is closely adjacent,* $\bar{A} = \text{assemble}(\overline{\mathcal{T}})$ *is closely dependent.*

Proposition 5.3. *The triangulation* $\overline{\mathcal{T}}$ *is closely adjacent iff* i *is isotropically internal for all* $i \in \mathcal{I}$.

Proof.

$\Rightarrow$ Assume that $\overline{\mathcal{T}}$ is closely adjacent. Suppose there exists $i \in \mathcal{I}$ which is not isotropically internal, hence there exists $k \in \text{adj}(i) \setminus (\text{adj}(i) \cap \mathcal{I})\mathcal{G}_i$. Obviously, $k \notin \mathcal{I}$ so $k = jg$ for some $j \in \mathcal{I}$, some $g \in \mathcal{G}$. Since $\overline{\mathcal{T}}$ is closely adjacent, $g \in \mathcal{G}_j \mathcal{G}_i$, i.e., $g = h_1 h_2$ for $h_1 \in \mathcal{G}_j$ and $h_2 \in \mathcal{G}_i$. Thus is $i \bowtie k = jh_2$ for some $h_2 \in \mathcal{G}_i$. But then is $i = ih_2^{-1} \bowtie j$, which means that $k \in (\mathcal{I} \cap adj(i))\mathcal{G}_i$. This is a contradiction and we conclude that i is isotropically internal.

$\Leftarrow$ Assume i is isotropically internal for all $i \in \mathcal{I}$ but suppose there exist $i, j \in \mathcal{I}$ such that $i \overset{\text{dist}}{\bowtie} j$, i.e., $ig = j$ for some $g \notin \mathcal{G}_i \mathcal{G}_j$. Since $g \neq \mathcal{G}_i$ and $i \in \mathcal{I}$, we conclude that $ig \notin \mathcal{I}$. But since j is isotropically internal, we see that $ig = kh$ for some $k \in (\text{adj}(j) \cap \mathcal{I})$ and $h \in \mathcal{G}_j$. Then $k \in i\mathcal{G}$ and $i, k \in \mathcal{I}$ implies $i = k$ and $g \in \mathcal{G}_i h^{-1}$ which contradicts $i \overset{\text{dist}}{\bowtie} j$. Thus, $i \overset{\text{dist}}{\not\bowtie} j$ for all $i, j \in \mathcal{I}$ and the triangulation is closely adjacent.

□

Let us now characterize different kinds of symmetry respecting triangulations, illustrated in Figure 5.1. We have symmetry respecting triangulations, freely symmetry respecting triangulations, and closely adjacent symmetry respecting triangulations. Of these criteria, just "symmetry respecting" is the weakest by definition. The following proposition clarifies the relationship between the remaining criteria.

Proposition 5.4. *A closely adjacent symmetry respecting triangulation is freely symmetry respecting.*

Proof. We prove the opposite implication. Consider a triangulation that is not freely symmetry respecting. Then, there exists an element t in the fundamental domain $\mathcal{T}$ such that $tg = t$ for some $g \neq e$ in $\mathcal{G}$. A consequence is that the corners of t are permuted by g. At least two of these indices therefore belong to the same orbit, say $i \in \mathcal{I}$ and $ig = k$. We note that $k \notin (\text{adj}(i) \cap \mathcal{I})\mathcal{G}_i$.

Since, $k \in \text{adj}(i)$, we have shown that i is not isotropically internal. According to Proposition 5.3, the triangulation is not closely adjacent. □

We proceed with the question of assembling the stiffness matrix $\bar{A}$ efficiently. This is accomplished by assembling just one entry per nonzero orbit, given the generalized action of both symmetry and equivariance, c.f. Proposition 3.1. The following proposition shows that the assembly of $\bar{A}$ is trivial for entries corresponding to internal indices.

PROPOSITION 5.5. *If i or j are internal indices,*

$$\bar{A}_{i,j} = F_{i,j}.$$

PROOF. Assume i is internal. Every element t with $i \in \text{ixs}(t)$ must then have $\text{ixs}(t) \subset \mathcal{I}$. Hence, $t \in \mathcal{T}$ according to Assumption 5.1, and

$$\bar{A}_{i,j} = \sum_{t \in \overline{\mathcal{T}}} a^t_{i,j} = \sum_{t \in \mathcal{T}} a^t_{i,j} = F_{i,j}.$$

□

In order to assemble all of $\bar{A}$, the triangulation characteristics should be considered. We introduce the following notation.

DEFINITION 5.4.

- *Let $\mathcal{I}^+$ be the set of all indices that are corners to some $t \in \mathcal{T}$.*
- *Let $\mathcal{T}^+$ be the set of all elements that has some corner in $\mathcal{I}^+$.*
- *Let $F = \text{assemble}(\mathcal{T})$.*
- *Let $F^+ = \text{assemble}(\mathcal{T}^+)$.*

We remark that $\mathcal{I} \subseteq \mathcal{I}^+$ according to Assumption 5.1. In next section, it is noted that equality holds if the triangulation is closely adjacent.

THEOREM 5.6. *The set $\mathcal{J} = \mathcal{I}^+ \times \mathcal{I}$ contains a set of representatives of the nonzero orbits of $\bar{A}$ under the equivariance and symmetry actions. For $(i,j) \in \mathcal{J}$, $\bar{A}_{i,j}$ may be deduced as follows:*

1. *If the triangulation is symmetry respecting,*

$$\bar{A}_{i,j} = F^+_{i,j}.$$

2. *If the triangulation is freely symmetry respecting, there exists $\mathcal{K} \subseteq \mathcal{I}^+ \times \mathcal{I}^+$ such that*

$$\bar{A}_{i,j} = \sum_{(k,l) \in \mathcal{K}} F_{k,l}.$$

3. *If the triangulation is closely adjacent symmetry respecting,*

$$\bar{A}_{i,j} = |\mathcal{G}_i \cap \mathcal{G}_j| F_{i,j}.$$

PROOF. Equivariance implies that representations of the index set may be chosen as $\overline{\mathcal{I}} \times \mathcal{I}$. Considering symmetry and the fact that contributions only come from adjacent indices, nonzero orbits may be represented by a subset of $\mathcal{I}^+ \times \mathcal{I}$.

1. For a symmetry respecting triangulation, $i \in \mathcal{I}^+$, and $j \in \mathcal{I}$, we obtain $\bar{A}_{i,j} = F^+_{i,j}$ since only elements in $\mathcal{T}^+$ contribute, c.f. the proof of Proposition 5.5.

2. A freely symmetry respecting triangulation implies that

$$\bar{A}_{i,j} = \sum_{g \in \mathcal{G}} \sum_{t \in \mathcal{T}} a^{tg}_{i,j}$$

for all $i, j \in \overline{\mathcal{I}}$. Since $a^{tg}_{i,j} = a^t_{k,\ell}$ for $k = ig^{-1}, \ell = jg^{-1}$, the statement follows.

3. Consider an element $t \in \mathcal{T}$ and two indices $i, j \in \mathrm{ixs}(t)$. For any $g \in \mathcal{G}$, element tg has corners ig, jg. Suppose tg have corners i and j. The requirement that i and j are isotropically internal then assures us that $g \in \mathcal{G}_i \cap \mathcal{G}_j$. By Proposition 5.4, the action $\overline{\mathcal{T}} \times \mathcal{G} \to \overline{\mathcal{T}}$ is free, and

$$\begin{aligned}\bar{A}_{i,j} &= \sum_{t \in \mathcal{T}} \sum_{g \in \mathcal{G}} a^{tg}_{i,j} = \sum_{t \in \mathcal{T}} \sum_{g \in (\mathcal{G}_i \cap \mathcal{G}_j)} a^{tg}_{i,j} = \sum_{t \in \mathcal{T}} \sum_{g \in (\mathcal{G}_i \cap \mathcal{G}_j)} a^t_{i,j} \\ &= |\mathcal{G}_i \cap \mathcal{G}_j| \sum_{t \in \mathcal{T}} a^t_{i,j} = |\mathcal{G}_i \cap \mathcal{G}_j| F_{i,j}.\end{aligned}$$

□

As discussed above, only representatives of nonzero orbits need to be assembled. These are found by considering nonzeros of F^+ and F, respectively. In Section 5.4 we combine the assembly process with the GFT, but first we make some observations regarding the symmetries of the continuous domain and its discrete counterpart.

5.3 Continuous considerations

As already mentioned, a continuous approach was proposed by Bossavit [8]. This method is further discussed in Section 7. Here, we focus on properties of the discrete and the continuous fundamental domains.

Bossavit states the following definitions for a domain $\overline{\Omega}$ symmetric under a finite group $\mathcal{G}$.

DEFINITION 5.5. *A domain $\overline{\Omega}$ has a* regular boundary Γ *if Γ is the union of a finite number of differentiable closed surfaces and if every $x \in \Gamma$ is adherent both to $\overline{\Omega}$ and to $\mathbb{R}^d \setminus \mathrm{cl}(\overline{\Omega})$. A symmetry cell C (a fundamental domain) relative to $(\overline{\Omega}, \mathcal{G})$ is any subdomain with regular boundary such that*

$$\overline{\Omega} \subset \mathrm{cl}\Big(\bigcup_{g \in \mathcal{G}} Cg\Big)$$

and

$$Cg \cap Ch = \emptyset, \text{ if } g \neq h. \tag{5.2}$$

The set $\Sigma = \partial C \setminus \Gamma$ *is the "new boundary" of the symmetry cell* C. Bossavit uses these definitions to prove the following.

PROPOSITION 5.7. *If* $x \in \Sigma$, *there is at least one* $g \in \mathcal{G}$, $g \neq e$, *such that* $xg \in \Sigma$. *For almost all* $x \in \Sigma$ *this* g *is unique.*

The proof is very elegant. He first shows that if $y \in C \cup \Sigma$ and $yg \notin C \cup \Sigma$ for all $g \neq e$, then is $y \in C$. Next he studies a regular point x at the new boundary Σ and assumes that xg and xh are also at Σ for $g \neq e \neq h$. By following the three normals, he concludes from (5.2) that at most two of the points x, xg, xh differ.

When subgroups of $\mathcal{O}(d)$ are considered, it is evident that reflections enforce "reflection boundaries" as new regular boundaries. We note that regular points x at a reflection boundary have isotropy subgroups of size 2. However, if $\overline{\Omega}$ is symmetric under a subgroup $\mathcal{H}$ of $\mathcal{SO}(d)$, no segment of the new boundary will be a reflection boundary. We conclude that all regular points x at the new boundary have trivial isotropy subgroups $\mathcal{H}_x = \{e\}$.

Regarding the discrete case, it is obvious that a discretization of the computational cell as defined by Bossavit yields a freely symmetry respecting discretization, since all basis functions ϕ_i have support only in elements adjacent to i. Is it always possible to obtain a closely adjacent symmetry respecting triangulation? No: domains that are symmetric under groups that only contain rotations can not be discretized by a closely adjacent triangulation, as the following argument shows.

A discrete point at a reflection boundary corresponds to an isotropically internal index. A discrete point at a new boundary which is not a reflection boundary will—apart from pathological cases with extremely coarse discretizations—be distantly adjacent to some internal point. A closely adjacent symmetry respecting triangulation therefore requires all new boundaries to be reflection boundaries. We note that a consequence is that $\mathcal{I}$ and $\mathcal{I}^+$ are equal, which leads to a simple way of recognizing closely adjacent triangulations.

5.4 The GFT of an equivariant stiffness matrix

It is time to sum up how the GFT of an equivariant stiffness matrix is computed. We focus on symmetric domains $\overline{\Omega}$ where all regular points of the new boundary Σ of Ω are reflection points.

1. Discretize $\overline{\Omega}$ by a closely adjacent symmetry respecting triangulation.
2. Compute $F = \text{assemble}(\mathcal{T})$. Since F is symmetric, only the overtriangular part is needed.
3. Theorems 4.3 and 5.6 and Proposition 5.3 show that $\hat{A}$ and $\tilde{A}$ are given by

$$\hat{A}_{i,j(\rho)} = \sqrt{|\mathcal{G}_i||\mathcal{G}_j|} F_{i,j} \hat{\pi}_i(\rho) \hat{\pi}_j(\rho) \text{ for } i, j \in \mathcal{I}, \tag{5.3}$$

$$\tilde{A}_{i,j}(\rho) = \sqrt{|\mathcal{G}_i||\mathcal{G}_j|} F_{i,j} \varphi_i(\rho)^T \varphi_j(\rho) \text{ for } i,j \in \mathcal{I}. \tag{5.4}$$

By Proposition 3.2, $\hat{A}$ and $\tilde{A}$ are Hermitian, and only the overtriangular parts are required.

If the triangulation is not closely adjacent, for instance because the symmetry group does not contain reflections, it is straightforward to adjust for points at the new boundary. Note that the new boundary is relatively a small fraction of the computational domain. The isotropy subgroup is trivial for interior indices, which means that the interior part of $\tilde{A}(\rho)$ equals $I \otimes F'$, where I is a $d_\rho \times d_\rho$ identity matrix and where F' denotes the interior part of F.

Thus we conclude that, for most entries of $\tilde{A}(\rho)$, the GFT can be carried out with *no* arithmetical operations. The complexity can therefore be considered constant as the number of discretization points grow, but it requires some bookkeeping in the Fourier space. Still, even if the sparse matrices $\tilde{A}(\rho)$ are explicitly formed, the complexity is linear in the number of nonzeros of F. Consequently, the sparse GFT of an equivariant stiffness matrix is fast.

6 Numerical experiments

We present in detail how Poisson's equation is solved in a domain with triangular symmetry. This example is illustrative because it considers symmetry under the smallest nonabelian group. Since the group is small, the gain is modest even for dense systems; see [2] for some numerical results. To demonstrate performance speedup of symmetry exploiting, we proceed with a more challenging example with a larger symmetry group: simulation of the heat equation at the surface of a sphere discretized with an icosahedral grid. This problem is also discussed in [19, 22]. Here, finite differences are used to obtain a sparse equivariant matrix. This matrix is block-diagonalized via a GFT but without exploiting the sparse GFT. The time differentiation is then carried out by a Lie group integrator, which uses matrix exponentiation.

6.1 Poisson's equation in a geometry with triangular symmetry

Consider Poisson's equation $-\nabla^2 u = g$ in the domain $\overline{\Omega}$ of Figure 6.1(a) with Dirichlet boundary conditions $u = f$ at $\partial\overline{\Omega}$. Figure 6.1(b) shows a computed solution where $g = -\nabla^2 f$ and $f(x,y) = \cos(x-1) + \sin(y+1)$ was chosen. The solution is obviously $u = f$. This problem is useful to verify the order of the numerical method, and we have verified that the GFT approach and the direct approach give the same answer up to numerical accuracy. Note that the domain is symmetric under $\mathcal{D}_3$ whereas the boundary conditions and the forcing function are not.

The symmetry group $\mathcal{D}_3$ has 6 elements generated by a reflection b in the y-axis and a rotation a of 120°. Since the symmetry group contains reflections, we discretize the domain with a closely adjacent symmetry respecting triangulation. FEM with standard "hat" basis functions is used, leading to a second order approximation [15]. We use the Dirichlet boundary conditions and remove the

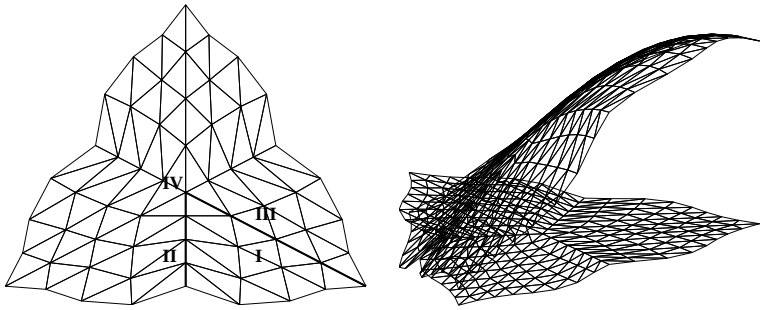

(a) (b)

Figure 6.1: A discretization of a domain with a triangular symmetry (a) and a simulated solution to a nonsymmetric problem (b).

unknowns corresponding to outer boundary points, leaving a symmetric and equivariant stiffness matrix $\bar{A}$ corresponding to the interior nodes of $\overline{\Omega}$. The right-hand side $\bar{g}$ is assembled and adjusted accordingly.

We divide the fundamental discretization domain into different regions according to their isotropy subgroup, see Figure 6.1(a). In the interior, region **I**, all points have trivial isotropy subgroups. Along the regular boundary **II**, the isotropy subgroup is $\{e, b\}$. Along the regular boundary **III**, the isotropy subgroup is $\{e, a^2b\}$. Finally, for the nonregular point in the origin, region **IV**, the isotropy subgroup is all of $\mathcal{D}_3$.

There are three nonequivalent irreducible representations of $\mathcal{D}_3$. We denote them $\mathcal{R} = \{\tau, \sigma, \rho\}$, see Table 6.1. For each region I and for each representation R we precompute and factorize $\hat{\pi}_I(R) = \varphi_I(R)\varphi_I(R)^T$, see Table 6.2.

The stiffness matrix $F = \text{assemble}(\mathcal{T})$ is computed and $\tilde{A}$ is obtained via the sparse GFT (5.4). The right-hand side $\tilde{g}$ is given by (2.7). Three independent systems are obtained:

$$\tilde{A}(\tau)\tilde{u}(\tau) = \tilde{g}(\tau),\ \tilde{A}(\sigma)\tilde{u}(\sigma) = \tilde{g}(\sigma),\ \tilde{A}(\rho)\tilde{u}(\rho) = \tilde{g}(\rho). \tag{6.1}$$

Recall that each system has d_R columns; in our case $\tilde{A}(\rho)\tilde{u}(\rho) = \tilde{g}(\rho)$ has two right-hand sides. Also note that $\tilde{A}(\rho)$ is almost block-diagonal since $\tilde{A}_{i,j}(\rho)_{\mu,\nu} = \delta_{\mu,\nu}$ for interior indices i, j and $\mu, \nu \in [1, 2]$. After solving for $\tilde{u}$, we obtain $\bar{u}$ by the inverse GFT.

The potential symmetry gain is quite modest in this example, since the symmetry group is small. Next, we study equivariance under a larger symmetry group.

6.2 Simulating the heat equation on an icosahedral grid

Consider the heat equation $u_t = c\nabla^2 u$ at the surface of a sphere. We discretize the surface by an icosahedral grid and apply FEM to assemble $\bar{B}$ as a

Table 6.1: Generators for a list of irreducible nonequivalent representations of $\mathcal{D}_3$.

	a	b
τ	1	1
σ	1	-1
ρ	$\frac{1}{2}\begin{pmatrix} -1 & -\sqrt{3} \\ \sqrt{3} & -1 \end{pmatrix}$	$\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$

Table 6.2: Factorizations $\varphi_I(R)$ for different regions and different representations.

	I	**II**	**III**	**IV**
τ	1	1	1	1
σ	1	0	0	0
ρ	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$	$\frac{1}{2}\begin{pmatrix} \sqrt{3} \\ -1 \end{pmatrix}$	0

discrete approximation of the Laplacian ∇^2. Applying the trapezoidal rule for the time integration gives

$$\left(I - \frac{1}{2}hc\bar{B}\right)\bar{u}^{k+1} = \left(I + \frac{1}{2}hc\bar{B}\right)\bar{u}^{k}. \tag{6.2}$$

All matrices are equivariant with respect to the full symmetry group of the icosahedron. The rotational group of the icosahedron is isomorphic to $\mathcal{A}_5$, the group of even permutations on five symbols. In addition, the reflection matrix $-I$ commutes with any of these 60 rotations, and the full symmetry group with 120 elements is thus $\mathcal{C}_2 \times \mathcal{A}_5$. A list of nonequivalent irreducible representations of $\mathcal{A}_5$ is found in e.g. [16] and the dimensions are $1, 3, 3, 4, 5$. An irreducible representation for the full symmetry group is obtained as the tensor product of one of these with one of the onedimensional irreducible representations of $\mathcal{C}_2$ (see Section 2.2). The full symmetry group thus has 10 irreducible representations. Consequently, a block-diagonalization of $\bar{B}$ has 10 blocks; the two smallest are approximately 120 times smaller than the original and the two largest are about 24 times smaller.

The fundamental domain Ω is one sixth of one of the 20 equilateral triangles of an icosahedron, projected upon the surface of the sphere. It is partitioned into seven regions depending upon their isotropy subgroups: the interior, the three sides, and the three corners. In total, there are 120 factorization matrices $\varphi_I(R)$ to be precomputed. However, in the interior where the isotropy subgroup is trivial, the projection matrices and the factorizations are just identity matrices which need no consideration, and for the onedimensional representations we find that the projections are either 1 or 0, trivial to factorize. For the remaining projection matrices given by (2.5), we use SVD to obtain orthogonal

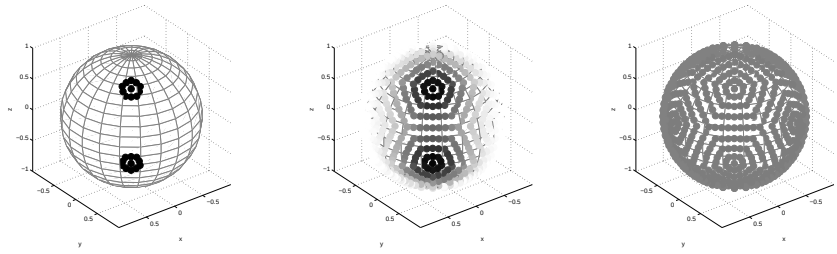

(a) t=0 (b) t=5 (c) t=150

Figure 6.2: Simulation of the heat equation at the surface of a sphere. Dark color means hot. Number of unknowns $n = 962$. Initially, in (a), two "hot spots" are located. The heat is gradually spread out (b) and, eventually, evenly distributed across the surface (c).

factorization matrices [12]. We stress that the SVD factorization is only needed once per symmetry group, and can therefore be neglected.

Applying the GFT yields the following equation in $\widetilde{\mathbb{C}^m\mathcal{G}}$ space:

$$\left(I - \frac{1}{2}hc\tilde{B}\right)\tilde{u}^{k+1} = \left(I + \frac{1}{2}hc\tilde{B}\right)\tilde{u}^{k}. \tag{6.3}$$

It is also possible to work in $\widehat{\mathbb{C}^m\mathcal{G}}$ space and solve the singular system

$$\left(\hat{\pi} - \frac{1}{2}hc\hat{B}\right)\hat{u}^{k+1} = \left(\hat{\pi} + \frac{1}{2}hc\hat{B}\right)\hat{u}^{k}. \tag{6.4}$$

by an iterative method.

Henriksson has implemented this example in his Master thesis; an illustration of his results is shown in Figure 6.2. We refer to [14] for details and summarize a few observations here. Table 6.3 shows elapsed time for solving the heat equation in different spaces and with different methods, using Matlab. We confirm that solution of (6.2) with a direct method is slow whereas an iterative method is much better. The iterative method used in these results is conjugate gradients (CG). Similar results were also obtained when minimum residual was studied.

Solving the block-diagonalized version (6.3) is faster, particularly for the direct method. Since the matrix is sparse, the gain is less than it would have been for a dense method. Still, the direct method outperforms the iterative method for system sizes in this range.

The singular system (6.4) is solved iteratively about as fast as the projected system (6.3). This is an interesting observation since it is somewhat simpler to transform to $\widehat{\mathbb{C}^m\mathcal{G}}$ space.

As mentioned, these results were measured in Matlab, and we believe that one should be careful not to draw too strong conclusions from these experiments.

We think that the asymptotic behaviour of the iterative method is surprisingly good for the block-diagonal system whereas it is less so for the system in the original space. A Fortran 90 version of this application has also been developed. This code uses inimum residual for solving the equations. Compared to Matlab the results were about 10 to 50 times faster. The time complexity, both for solving (6.2), (6.3) and (6.4), is approximately linear in the number of unknowns, see Table 6.4. This table also shows elapsed time for the GFTs. In our example, the cost for the GFTs can be amortized over the number of time steps, and is therefore relatively less important. However, if e.g. Poisson's equation is solved, it becomes relevant to optimize the GFTs, for example by using fast GFTs [17].

Table 6.3: Solution times (s), $hc = 0.1$. Matlab version.

n	$\bar{B}$ CG	$\bar{B}$ Direct	$\tilde{B}$ CG	$\tilde{B}$ Direct	$\hat{B}$ CG
962	0.44	0.66	0.08	0.01	0.11
1502	1.15	2.28	0.10	0.01	0.13
2162	2.42	6.10	0.11	0.02	0.17
2942	4.19	14.79	0.14	0.03	0.18
3842	7.10	30.91	0.20	0.05	0.21

Table 6.4: Time measurements (ms) for solving a single time step, either directly or indirectly. Fortran 90 Minimum residual version is used.

n	Total $\tilde{B}$	$F \to \tilde{B}$	gft($\bar{b}$)	igft($\hat{x}$)	Solve $\tilde{B}$	Total $\bar{B}$	Solve $\bar{B}$
542	9.3	2.6	2.3	1.3	3.1	5.4	3.7
962	13.3	2.9	3.1	2.2	5.1	10.5	7.5
1502	18.2	3.6	4.0	3.1	7.5	16.3	11.6
2162	24.9	4.5	5.9	4.2	10.3	22.7	16.1
2942	31.9	5.4	7.0	5.5	14.0	30.9	21.6
3842	41.1	6.7	9.2	7.0	18.2	43.5	31.2
4862	50.0	8.3	10.2	8.8	22.7	58.1	41.7
6002	62.0	10.3	12.0	10.8	28.9	76.6	56.0

7 Discussion

In this paper, we pursue the sparse GFT approach for stiffness matrices assembled on symmetrical geometries. Using continuous considerations, Bossavit developed essentially the same method [8, 9]. In order to compare the approaches, we first consider a trivial example.

EXAMPLE 7.1. *Solve the onedimensional problem* $\partial^2 u/\partial x^2 = f$ *for* $-1 < x < 1$ *with Dirichlet boundary conditions, see Figure 7.1(a).*

It is clear that problem symmetries can be exploited just by treating the even and the odd parts separately. Let

$$f^+(x) = f(x) + f(-x) \text{ and } f^-(x) = f(x) - f(-x)$$

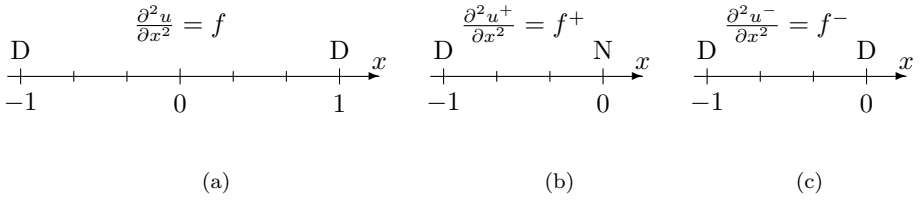

(a) (b) (c)

Figure 7.1: A trivial 1d problem (a) is partitioned by considering its even (b) and odd (c) parts. Dirichlet and Neumann boundary conditions are indicated by "D" and "N", respectively.

for $-1 < x < 0$. This gives two half-sized problems, and the boundary conditions at the new boundary $x = 0$ are homogeneous Neumann and Dirichlet conditions for the even and the odd problem, respectively; see Figure 7.1(b) and (c). When assembling via FEM or by using standard finite differences, the scaled discretization matrices become

$$A^+ = \begin{pmatrix} -2 & 1 & \\ 1 & -2 & 1 \\ & -1 & 1 \end{pmatrix} \text{ and } A^- = \begin{pmatrix} -2 & 1 \\ 1 & -2 \end{pmatrix}.$$

Referring back to Section 2.2, it is clear that $\bar{A}$ in (2.10) corresponds to a scaled discretization matrix for this problem when symmetries are not exploited. Why then is A^+ and A^- different from $\tilde{A}(\tau)$ and $\tilde{A}(\sigma)$? The reason is that the straightforward approach does not consider the isotropy of $x = 0$. The sparse GFT approach guarantees preservation of the symmetry as well as of the spectrum; note that $\tilde{A}$ in (2.11) has the same eigenvalues as the original matrix $\bar{A}$, whereas A^+ and A^- contain eigenvalues not present in the spectrum of $\bar{A}$.

The continuous approach may be applied to less trivial examples such as the ones discussed in Section 6. For example, Bossavit explains [8] how to reformulate the triangular example $-\nabla^2 u = g$ into six subproblems. Two of the subproblems correspond to the representations of dimension 1, and the remaining subproblems correspond to the third representation of $\mathcal{D}_3$. These problems are pair-wise coupled via Dirichlet boundary conditions. To illustrate the connection with the approach outlined in Section 6.1, let $v_1 = \tilde{u}(\tau)$, $v_2 = \tilde{u}(\sigma)$, $v^{ij} = \tilde{u}(\rho)_{ij}$, and denote the GFT of the right-hand side similarly. Thus, there are four independent systems.

$$\begin{aligned} -\nabla^2 v_1 &= g_1 \quad + \quad \text{Neumann BCs at } \Sigma \\ -\nabla^2 v_2 &= g_2 \quad + \quad \text{Dirichlet BCs at } \Sigma \\ \begin{cases} -\nabla^2 v^{11} = g^{11} \\ -\nabla^2 v^{21} = g^{21} \end{cases} &\quad + \quad \text{coupled Dirichlet BCs at } \Sigma \\ \begin{cases} -\nabla^2 v^{12} = g^{12} \\ -\nabla^2 v^{22} = g^{22} \end{cases} &\quad + \quad \text{coupled Dirichlet BCs at } \Sigma. \end{aligned}$$

It is clear that (6.1) represent a discrete form of the above. We find that the

continuous approach is a very elegant way to explain the symmetry exploitation. However, a straightforward discretization of this method may change the spectrum of the transformed system matrix, and the coupled Dirichlet boundary conditions may cause a loss of symmetry. Therefore we find that the sparse GFT method is the appropriate way to discretize the continuous approach.

Regarding symmetry respecting discretizations that are not free, see Figure 5.1(b), this is also addressed by the refined theory in [9]. This is perhaps less important for FEM discretizations, because it should always be possible to construct a freely symmetry respecting triangulation. However, we find this point interesting because similar situations arise when discretization matrices are obtained by other methods, e.g. higher order finite differences. For example, apply a stencil of width 5 to (a finer discretization of) the problem of Example 7.1. The resulting discretization matrix is equivariant. The sparse GFT will automatically yield discretizations for the Neumann and Dirichlet boundary conditions at $x = 0$. These boundary discretizations are guaranteed to yield the same answer as if symmetry would not have been exploited.

8 Conclusions

We have used the GFT approach to block-diagonalize sparse equivariant matrices, typically arising when FEM is used in geometrically symmetrical domains. The main theoretical contributions of this paper are the following:

- Our generalized symmetry definition relates both to equivariance and "traditional" symmetry. We use this concept to prescribe the smallest set of entries that need to be assembled if a stiffness matrix is both symmetrical and equivariant.
- Equivariant graphs are introduced to discuss equivariant sparse matrices. The notion of a closely dependent graph describes when the GFT becomes particularly simple and fast.
- The classification of different symmetry respecting triangulations is useful when a symmetrical domain is triangulated. We state when it is not possible to obtain a closely adjacent triangulation, and we discuss how to assemble and transform a stiffness matrix efficiently in those cases as well.

Our numerical examples show the viability of the approach. The simulation of the heat equation makes some comparisons between iterative methods and direct methods for solving linear systems of equations where the system matrix is sparse, symmetric, and equivariant under icosahedral symmetry. The results demonstrate that direct methods outperforms iterative methods, at least for systems of these sizes. We believe however, that iterative methods can be improved further, and we consider a thorough study of this question to be an issue for future work. In future, we also aim for harder problems, such as the shallow-water equations on the sphere. Compare with [11] where an icosahedral grid is used but without symmetry exploitation. Another important future question is adaptivity. One possibility is to require adaptivity in space to respect

symmetry, but this might be unnecessarily restrictive. Adaptive solution of the independent systems in Fourier space seems interesting, though.

Regarding the continuous formulation and the discrete GFT method, we find that the approaches complement each other. The former gives a good understanding of the mechanisms behind symmetry exploitation, but we find that the GFT approach is appropriate because it is guaranteed to preserve Hermitian symmetry as well as the spectrum. Another advantage with the sparse GFT approach is its ability to handle sparse equivariant matrices obtained by other numerical methods. We also remark that the GFT formulation more directly can benefit from the research area of fast GFTs.

Finally, we acknowledge that not all problems have the kind of symmetries that we address here. However, many problems may be nearly symmetric or partially symmetric. Exploitation of the symmetric part can in these cases be used to improve the numerical method, see e.g. [7, 21, 13].

Acknowledgements

I thank Master thesis student Daniel Henriksson for his work on the icosahedral example. Niklas Fors, Olof Dolfe, and Sverker Holmgren have also contributed to this example. I am also grateful to Auden Trønnes and Hans Munthe-Kaas for their cooperation.

REFERENCES

1. K. Åhlander, *Supporting tensor symmetries in EinSum*, Computers and Mathematics with Applications, 45 (2003), pp. 789–803. Also available as Technical Report 212 from Dept. of Informatics, University of Bergen.
2. K. Åhlander and H. Munthe-Kaas, *Applications of the Generalized Fourier Transform in Numerical Linear Algebra*, BIT Numerical Mathematics, (2005), pp. 819–850. Also available as Technical Report 2004-029, Department of Information Technology, Uppsala University.
3. ——, *Eigenvalues for Equivariant Matrices*, Journal of Computational and Applied Mathematics, (2005). Available online via ScienceDirect.
4. 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.
5. E. L. Allgower and K. Georg, *Numerical exploitation of symmetric structures in BEM*, in Boundary Integral Methods - Numerical and Mathematical Aspects, M. Golberg, ed., Computational Mechanics Publications, Southampton, 1998, pp. 289–306.
6. 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, Providence, RI, 1993, pp. 23–36.

7. M. Bonnet, *Exploiting partial or complete geometrical symmetry in 3D symmetric Galerkin indirect BEM formulations*, Intl. J. Numer. Meth. Engng, 57 (2003), pp. 1053–1083.
8. 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 (1986), pp. 167–215.
9. ——, *Boundary value problems with symmetry and their approximation by finite elements*, SIAM J. Appl. Math., 53 (1993), pp. 1352–1380.
10. S. Egner and M. Puschel, *Symmetry-based matrix factorization*, J. Symb. Comput., 37 (2004), pp. 157–186.
11. F. Giraldo and T. Warburton, *A nodal triangle-based spectral element method for the shallow water equations on the sphere*, Journal of Computational Physics, (2005), pp. 129–150.
12. G. Golub and C. van Loan, *Matrix computations*, The John Hopkins Univeristy Press, 1984.
13. A. Hall, *Pricing financial derivatives using radial basis functions and the generalized Fourier transform*, Master's thesis, Uppsala University, Uppsala, Sweden, 2005.
14. D. Henriksson, *Exploiting symmetries when solving the heat equation on a spherical icosahedral grid*, Master's thesis, Uppsala University, Uppsala, Sweden, 2006.
15. C. Johnsson, *Numerical solution of partial differential equations by the finite element method*, Studentlitteratur, 1987.
16. J. S. Lomont, *Applications of Finite Groups*, Academic Press, New York, 1959.
17. 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.
18. T. Minkwitz, *Algorithms explained by symmetry*, in Lecture Notes on Computer Science, vol. 900, 1995, pp. 157–167.
19. H. Munthe-Kaas, *On group Fourier analysis and symmetry preserving discretizations of PDEs*, Physica A, (2005). To appear.
20. J. P. Serre, *Linear Representations of Finite Groups*, Springer, 1977. ISBN 0387901906.
21. J. Tausch, *Equivariant preconditioners for boundary element methods*, SIAM Sci. Comp., 17 (1996), pp. 90–99.
22. A. Trønnes, *Symmetries and generalized Fourier transforms applied to computing the matrix exponential*, Master's thesis, University of Bergen, Bergen, Norway, 2005.

23. J. Turski, *Geometric Fourier analysis of the conformal camera for active vision*, SIAM Review, 46 (2004), pp. 230–255.