An Eigenspace Approach to Decomposing Representations of Finite Groups

A Thesis

Submitted to the Faculty

in partial fulfillment of the requirements for the

degree of

Doctor of Philosophy

in

Mathematics

by

Michael E. Orrison Jr.

DARTMOUTH COLLEGE

Hanover, New Hampshire

June 1st, 2001

Examining Committee:

(chairman) Daniel Rockmore

Kenneth Bogart

Richard Foote

David Webb

Roger D. Sloboda
Dean of Graduate Students

Abstract

The first half of this thesis develops an eigenspace approach to computing the basis-independent isotypic decomposition of a vector in a representation of a finite group. The approach takes advantage of well-chosen diagonalizable linear transformations to compute isotypic projections through a series of eigenspace projections, and at its heart is an efficient eigenspace projection method built around a modified Gram-Schmidt algorithm known as the Lanczos iteration. In the second half, it is shown how this eigenspace approach gives rise to an efficient decomposition method for permutation representations of distance transitive graphs, the symmetric group, the hyperoctahedral group, the finite general linear group, and the finite symplectic group.

Acknowledgements

This thesis is dedicated to my wife, Jody—she is the music in my life and the reason I smile. The idea of computing isotypic projections with the aid of the Lanczos iteration is due to David Maslen—his thoughts on the subject have provided much of the impetus for the approach presented here. I would like to thank my advisor, Dan Rockmore, who has been patient with me as I have slowly learned how to be a mathematician, and Andrei Zelevinsky whose encouraging words and suggestions gave life to many of these results. I would also like to thank Ken Bogart and David Webb for all of their advice and for their willingness to listen to me as I worked through many of these ideas. Lastly, I would like to thank my parents for always believing in me and for helping me to capture a dream that started long ago.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The action of a finite group G on a finite set X gives rise to a permutation representation $\mathbb{C}[X]$ of G where $\mathbb{C}[X]$ is the vector space of complex-valued functions on X. Because $\mathbb{C}[X]$ is a representation of G, there is a canonical (i.e., basis-independent) decomposition

$$\mathbb{C}[X] = M_1 \oplus \cdots \oplus M_n$$

of $\mathbb{C}[X]$ into G-invariant subspaces known as *isotypic subspaces*. Given an arbitrary function f in $\mathbb{C}[X]$, one may therefore consider the problem of efficiently computing the projections of f onto each of the isotypic subspaces of $\mathbb{C}[X]$.

The problem of computing projections onto isotypic subspaces arises in *spectral analysis* which is a non-model based approach to the analysis of data that may be viewed as a complex-valued function f on a set X that has an underlying symmetry group G. Developed by Persi Diaconis in [14, 15], the subject extends the classical spectral analysis of time series and requires the computation of projections of f onto a subset of G-invariant subspaces of $\mathbb{C}[X]$.

As an example, let X be the set $\{x_0, \ldots, x_{n-1}\}$ and let G be the cyclic group

$\mathbb{Z}/n\mathbb{Z}$ acting on X by cyclicly permuting its elements. The elements of $\mathbb{C}[X]$ my be viewed as *signals* on n points and the isotypic subspaces of $\mathbb{C}[X]$ as corresponding to the different *frequencies* that make up these signals. The isotypic projections of $f \in \mathbb{C}[X]$ may be computed with the aid of the usual discrete Fourier transform (DFT). The now classical fast Fourier transform (FFT) may therefore be used to efficiently compute the projections of f onto the isotypic subspaces of $\mathbb{C}[X]$ (see, e.g., [33]).

As another example, suppose voters are asked to rank k candidates in order of preference. The set X is then the set of orderings of the k candidates and G is the symmetric group S_k whose natural action on the set of candidates induces an action on the set of orderings. If $f \in \mathbb{C}[X]$ is such that $f(x)$ is the number of voters choosing the ordering x, then there are natural statistics associated to f. For example, the *mean response* of f is the value $(1/|X|)\sum_{x\in X} f(x)$, whereas a *first order summary* of f counts the number of voters that ranked candidate i in position j. Similarly, there are *higher order summaries* associated to f. For example, we could compute the number of voters that ranked candidates i and j in positions k and l, either respectively or so that order does not matter. These summaries, however, contain redundant information. Removing this redundant information, or finding the *pure higher order effects* of f, is equivalent to computing the isotypic projections of f (see [15, 40]).

A naive approach (see, e.g., [41]) to computing the n isotypic projections of $f \in \mathbb{C}[X]$ requires $O(n|G||X|)$ operations where we count a complex multiplication followed by a complex addition as one *operation*. Diaconis and Rockmore [18] show that a careful reorganization of this approach reduces the number of necessary operations to $O(n|X|^2)$. The advantage of their approach is that it relies only on the knowledge of the characters of G. In terms of operation counts, however, the number

of operations required by a direct matrix multiplication approach is also $O(n|X|^2)$ which has prompted the search for other approaches to computing isotypic projections. For example, Driscoll, Healy and Rockmore [21] show that if X is a distance transitive graph, then fast discrete polynomial transforms may be used to compute the n isotypic projections of $f \in \mathbb{C}[X]$ with at most $O(|X|^2 + |X| n \log^2 n)$ operations. This bound, however, assumes the use of exact arithmetic. Stability issues arise when their algorithm is implemented using finite precision arithmetic.

In this thesis, we develop an approach to computing isotypic projections that uses diagonalizable linear transformations to decompose a representation through a series of eigenspace projections. The collections of diagonalizable transformations that we use are known as *separating sets* because they allow us to *separate* a module into its isotypic components. The approach may be seen as a generalization of the Gentleman-Sande, or *decimation in frequency*, fast Fourier transform in that we too will be iteratively computing projections of projections (see [27]).

As a simple example of how a separating set is used to compute isotypic projections, suppose that $\mathbb{C}[X]$ has three isotypic subspaces M_1, M_2, and M_3. Thus $\mathbb{C}[X] = M_1 \oplus M_2 \oplus M_3$, and each $f \in \mathbb{C}[X]$ may be written uniquely as $f = f_1 + f_2 + f_3$ where $f_i \in M_i$. Additionally, suppose that T and T' are diagonalizable linear transformations on $\mathbb{C}[X]$ such that the eigenspaces of T are $M_1 \oplus M_2$ and M_3, and the eigenspaces of T' are M_1 and $M_2 \oplus M_3$. In this case, $\{T, T'\}$ is a separating set for $\mathbb{C}[X]$. We may compute the f_i by first projecting f onto the eigenspaces of T to compute $f_1 + f_2$ and f_3, and then projecting both $f_1 + f_2$ and f_3 onto the eigenspaces of T' to compute f_1, f_2 and f_3. Note that each computation is done with respect to a fixed basis of $\mathbb{C}[X]$. This process of decomposing $\mathbb{C}[X] = M_1 \oplus M_2 \oplus M_3$ is illustrated in Figure 1.1.

The efficiency of this approach, and of its generalization that we present, depends

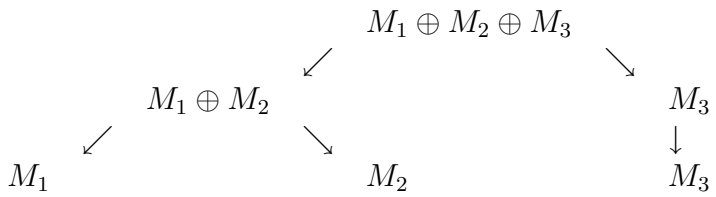

Figure 1.1: Decomposing $\mathbb{C}[X] = M_1 \oplus M_2 \oplus M_3$ using T and T'.

on an efficient eigenspace projection method. Since the separating sets we use consist of real symmetric matrices, we look to the *Lanczos iteration* for such a method. This is an algorithm that may be used to efficiently compute the eigenspace projections of a real symmetric matrix when it has relatively few eigenspaces and when it may be applied efficiently to arbitrary vectors, either directly or through a given subroutine (see [37]).

We proceed as follows. In Chapter 2 we review some useful facts about algebras and develop a recursive eigenspace approach to decomposing their modules. In Chapter 3 we turn our attention to group algebras and permutation modules. In Chapter 4 we show how an eigenspace approach to computing isotypic decompositions for abelian groups leads to the Gentleman-Sande FFT and we remark on the use of generalized FFTs in computing isotypic decompositions for nonabelian groups. Chapter 5 contains a review of the Lanczos iteration and its use in computing eigenspace projections. It also contains our main theorem, Theorem 5.4.1, which gives an upper bound on the number of operations required to compute isotypic projections for a module over a semisimple algebra. In Chapter 6 we apply the results of Chapter 5 to permutation modules associated to distance transitive graphs. In Chapters 7 and 8 we consider the symmetric and hyperoctahedral groups, respectively. Chapter 9 deals with the finite general linear group and Chapter 10 deals with the finite symplectic group. Finally, Chapter 11 contains a few thoughts on future directions.

Chapter 2

Representations of Algebras

In this chapter, we review some useful facts about algebras and begin to develop an eigenspace approach to decomposing their modules. A good reference is [22].

2.1 Algebras and Modules

A (complex) *algebra* is a vector space A over $\mathbb{C}$ with multiplication such that A is a ring with identity and $(\alpha a_1)a_2 = a_1(\alpha a_2) = \alpha(a_1 a_2)$ for all $a_1, a_2 \in A$ and $\alpha \in \mathbb{C}$. We will assume that all algebras are finite-dimensional. A subspace A' of A is a *subalgebra* of A if it is itself an algebra with respect to the operations in A, and if it has the same identity. The *center* of A is the subalgebra $Z(A) = \{z \in A \mid za = az \text{ for all } a \in A\}$.

An *algebra homomorphism* from an algebra A to an algebra B is a linear map $\phi : A \to B$ that preserves multiplication and the identity. If ϕ is bijective, then ϕ is an *algebra isomorphism*. We say that the algebras A and B are *isomorphic* if there exists an algebra isomorphism between them.

Let n be a positive integer. Let $M_n(\mathbb{C})$ be the algebra of $n \times n$ matrices with entries in $\mathbb{C}$ where addition and multiplication are the usual matrix addition and

multiplication. A *representation* of an algebra A is an algebra homomorphism

$$\rho : A \to M_n(\mathbb{C}).$$

The *degree* of the representation ρ is n. Its *character* is the function $\chi_\rho : A \to \mathbb{C}$ given by $\chi_\rho(a) = tr(\rho(a))$ where $tr(\rho(a))$ is the usual trace of the matrix $\rho(a)$, i.e., the sum of the diagonal entries of $\rho(a)$.

A *right module* over an algebra A, or a (right) A-*module*, is a vector space M over $\mathbb{C}$ whose elements can be multiplied on the right by elements of A. In other words, there is a map $M \times A \to M$ with $(m, a) \mapsto ma$ such that

$$\begin{aligned}
(m_1 + m_2)a &= m_1a + m_2a \\
m(a_1 + a_2) &= ma_1 + ma_2 \\
(\alpha m)a &= \alpha(ma) = m(\alpha a) \\
m(a_1a_2) &= (ma_1)a_2 \\
m1 &= m
\end{aligned}$$

for all $a, a_1, a_2 \in A$, $m, m_1, m_2 \in M$, and $\alpha \in \mathbb{C}$. *Left* A-*modules* are defined analogously. By M_A (or ${}_AM$) we understand that M is a right (or left) A-module. A particularly important example of an A-module is the *regular module* A_A where the multiplication is the same as that in A. We will assume that all modules are finite dimensional.

Let M and N be right A-modules. An A-*module homomorphism* from M to N is a linear map $\psi : M \to N$ such that $\psi(ma) = \psi(m)a$ for all $m \in M$ and $a \in A$. If ψ is bijective, then ψ is an A-*module isomorphism*. We say that M and N are *isomorphic* as A-modules if there exists an A-module isomorphism between them.

Note that $M \oplus N$ is also an A-module where we define $(m,n)a = (ma, na)$ for all $m \in M$, $n \in N$, and $a \in A$.

A *submodule* of an A-module M is a subspace N of M that is invariant under multiplication by A, i.e., $na \in N$ for all $n \in N$ and $a \in A$. A module M is *simple* if its only submodules are $\langle 0 \rangle$ and M, and is *semisimple* if it is isomorphic to a direct sum of simple modules. The algebra A is said to be *semisimple* if the regular module A_A is semisimple. Moreover, it can be shown that if A is semisimple, then every A-module is semisimple.

Let M be a right A-module with basis $\{b_1, \ldots, b_n\}$. By definition, each element $a \in A$ may be viewed as a linear transformation $a : M \to M$ (which we write on the right) where $(m)a = ma$ for all $m \in M$. Let $\rho : A \to M_n(\mathbb{C})$ be the map that assigns to each element a the matrix $\rho(a)$ corresponding to this linear transformation with respect to the basis $\{b_1, \ldots, b_n\}$ of M. In other words, $\rho(a)$ is the $n \times n$ matrix whose (i,j) entry is α_{ij} where

$$b_i a = \sum_{i=1}^{n} \alpha_{ij} b_j.$$

It is straightforward to show that the map $\rho : A \to M_n(\mathbb{C})$ is a representation of A. Conversely, given a representation of A, it is easy to construct an associated A-module. For this reason, we also refer to A-modules as *representations* of A. Furthermore, since the trace of a linear operator is independent of the choice of basis, the character of an A-module is uniquely defined.

Let M and N be right A-modules. Let $\mathrm{Hom}_A(M,N)$ denote the complex vector space of A-module homomorphisms between M and N where if $\psi, \psi' \in \mathrm{Hom}_A(M,N)$ and $\alpha \in \mathbb{C}$, then $\psi + \psi' : M \to N$ is defined by $(\psi + \psi')(m) = \psi(m) + \psi'(m)$ and $\alpha\psi : M \to N$ is defined by $(\alpha\psi)(m) = \alpha(\psi(m))$.

Elements of $\mathrm{Hom}_A(M,M)$ are known as *endomorphisms*. We denote the vector

space $\mathrm{Hom}_A(M, M)$ by $\mathrm{End}_A(M)$. The vector space $\mathrm{End}_A(M)$ is also an algebra, the *endomorphism algebra* (of M with respect to A), where multiplication is given by function composition. Furthermore, the right A-module M is also a left $\mathrm{End}_A(M)$-module where $\psi m = \psi(m)$. In fact, M is simultaneously a right A-module and a left $\mathrm{End}_A(M)$-module since, for all $\psi \in \mathrm{End}_A(M)$ and $a \in A$,

$$\psi(ma) = (\psi m)a.$$

Moreover, $\mathrm{End}_A(M)$ is semisimple if A is semisimple.

2.2 Using Separating Sets to Decompose Modules

Let A be a semisimple algebra and let M be a right A-module. The A-module M is semisimple and may be decomposed as

$$M = U_1 \oplus \cdots \oplus U_s$$

into a direct sum of simple submodules. Let $W_1, \ldots, W_n$ be a collection of pairwise non-isomorphic simple A-modules so that each U_i is isomorphic to some W_j. Let M_j be the direct sum of those U_i that are isomorphic to W_j. The decomposition

$$M = M_1 \oplus \cdots \oplus M_n$$

is known as the *isotypic decomposition* of M. Each M_i is an *isotypic subspace* of M. This decomposition is independent of the decomposition of M into simple submodules. An isotypic decomposition of left A-modules is defined analogously.

Given that M is a direct sum of the isotypic subspaces $M_1, \ldots, M_n$, each $m \in M$

may be written uniquely as

$$m = m_1 + \cdots + m_n$$

where $m_i \in M_i$. We say that m_i is the *isotypic projection* of m onto the isotypic subspace M_i.

Suppose T is a diagonalizable linear transformation on M whose eigenspaces are precisely the isotypic subspaces of M. We could then compute the isotypic projections of $m \in M$ by computing the projections of m onto the eigenspaces of T. More generally, suppose that $\{T_1, \ldots, T_k\}$ is a collection of diagonalizable linear transformations on M whose eigenspaces are direct sums of the isotypic subspaces of M. For each isotypic subspace M_i of M, let $c_i = (\mu_1^i, \ldots, \mu_k^i)$ be the k-tuple of eigenvalues where μ_j^i is the eigenvalue of T_j restricted to M_i. If $c_i \neq c_{i'}$ for all $M_i \neq M_{i'}$, then we say that $\{T_1, \ldots, T_k\}$ is a *separating set* for M.

The existence of a separating set $\{T_1, \ldots, T_k\}$ for M means that the computation of the isotypic projections of $m \in M$ can be achieved through a series of eigenspace projections:

Stage 1: Compute the projections of m onto each of the eigenspaces of T_1.

Stage 2: Compute the projections of each of the previously computed projections onto each of the eigenspaces of T_2.

$\vdots$

Stage k: Compute the projections of each of the previously computed projections onto each of the eigenspaces of T_k.

Lemma 2.2.1. *The projections computed at Stage k are precisely the isotypic projections of the vector m.*

Proof. The projections at each stage are sums of the isotypic projections of m. If a projection at Stage k was the sum of two or more isotypic projections, then the corresponding isotypic subspaces must have been in the same eigenspace for each of the T_j. This, however, would contradict the fact that $\{T_1, \ldots, T_k\}$ is a separating set for M. □

A natural place to find separating sets for an A-module M is the center $Z(A)$ of the algebra A. This is because each $z \in Z(A)$, when viewed as a linear transformation on M, is diagonalizable with eigenspaces that are direct sums of isotypic subspaces. In fact, if χ_i is the character of the simple A-module corresponding to the isotypic subspace M_i, then the eigenvalue of $z \in Z(A)$ that is associated to M_i is $\chi_i(z)/\chi_i(1)$. A subset $\{z_1, \ldots, z_k\}$ of $Z(A)$ is therefore a separating set for M if $c_i = (\chi_i(z_1)/\chi_i(1), \ldots, \chi_i(z_k)/\chi_i(1))$ is distinct for each M_i.

Another natural place to find a separating set for M is the center $Z(\mathrm{End}_A(M))$ of the endomorphism algebra $\mathrm{End}_A(M)$. This is because the isotypic decomposition of M as a right A-module and as a left $\mathrm{End}_A(M)$-module coincide (see, e.g., Section 2.6 in [22]). Moreover, the endomorphism algebra $\mathrm{End}_A(M)$ and its center may be much easier to work with than the algebra A and its center.

Suppose now that A' is a semisimple subalgebra of A. Any A-module M is then naturally an A'-module. In particular, each isotypic subspace M_i of M is an A'-module and is therefore a direct sum

$$M_i = M_{i1} \oplus \cdots \oplus M_{in_i} \tag{2.1}$$

of isotypic subspaces with respect to A'. The A-module M may therefore be written

as

$$M = \bigoplus_{i=1}^{n} \bigoplus_{j=1}^{n_i} M_{ij}. \tag{2.2}$$

For convenience, we refer to the M_{ij} as *sub-isotypic spaces* (of M with respect to A'). By taking the direct sums of those sub-isotypic spaces whose corresponding simple A'-modules are isomorphic, we create the isotypic decomposition

$$M_{A'} = M'_1 \oplus \cdots \oplus M'_{n'}$$

of M when viewed as an A'-module.

Let N be a sub-isotypic space of M, U' be a simple A'-module, and U be a simple A-module. Suppose that the isotypic subspace of $M_{A'}$ that contains N corresponds to U' and that the isotypic subspace of M_A that contains N corresponds to U. We then say that N is of *type* (U', U).

Suppose now that $\{T_1, \ldots, T_k\}$ is a collection of diagonalizable linear transformations on M whose eigenspaces are direct sums of the sub-isotypic spaces. For each sub-isotypic space M_{ij}, let $c_{ij} = (\mu_1^{ij}, \ldots, \mu_k^{ij})$ be the k-tuple of eigenvalues where μ_l^{ij} is the eigenvalue of T_l restricted to M_{ij}. We then say that $\{T_1, \ldots, T_k\}$ is an *A'-separating set for M* if $c_{ij} \neq c_{i'j'}$ whenever the sub-isotypic spaces M_{ij} and $M_{i'j'}$ are distinct and contained in the same isotypic subspace of $M_{A'}$. In other words, for every simple submodule U' of $M_{A'}$, the transformations $T_1, \ldots, T_k$ must be able to distinguish between those sub-isotypic spaces of type (U', U) where U is a simple submodule of M_A.

The existence of an A'-separating set $\{T_1, \ldots, T_k\}$ for M means that the isotypic projections of $m \in M$ may be computed as follows:

Stage 0: Compute the isotypic projections of m with respect to A'.

Stage 1: Compute the projections of each of the previously computed projections onto each of the eigenspaces of T_1.

$$\vdots$$

Stage k: Compute the projections of each of the previously computed projections onto each of the eigenspaces of T_k.

Stage $k+1$: For each isotypic subspace M_i of M, compute the sum m_i of the previously computed projections that are contained in M_i.

Lemma 2.2.2. *The projections computed at Stage $k+1$ are precisely the isotypic projections of the vector m.*

Proof. By (2.2), each $m \in M$ may be written uniquely as

$$m = \sum_{i=1}^{n} \sum_{j=1}^{n_i} m_{ij}$$

where m_{ij} is contained in the sub-isotypic subspace M_{ij}. Arguing as we did in Lemma 2.2.1, we see that the projections computed at Stage k are precisely these m_{ij}. By (2.1), the isotypic projection $m_i \in M_i$ of m is then the sum

$$m_i = \sum_{j=1}^{n_i} m_{ij}$$

which is computed at Stage $k+1$. □

Lastly, note that if E is the endomorphism algebra $\mathrm{End}_A(M)$ and E' is a semisimple subalgebra of E, then we may also define E'-separating sets analogously to A'-separating sets. This fact will be used in Chapters 9 and 10.

Chapter 3

Representations of Finite Groups

We now turn our attention to algebras and representations that are associated to finite groups. Good references are [22, 41].

3.1 Group Algebras

Let G be a finite group and let $\mathbb{C}[G]$ be the vector space over $\mathbb{C}$ of formal linear combinations of elements of G. The vector space $\mathbb{C}[G]$ is easily given an algebra structure where multiplication in $\mathbb{C}[G]$ is a linear extension of the group multiplication in G. In other words, if $\sum_{h\in G} \alpha_h h$ and $\sum_{g\in G} \beta_g g$ are elements of $\mathbb{C}[G]$, then

$$(\sum_{h\in G} \alpha_h h)(\sum_{g\in G} \beta_g g) = \sum_{h,g\in G} \alpha_h \beta_g (hg).$$

The algebra $\mathbb{C}[G]$ is known as the *group algebra* of G and $\mathbb{C}[G]$-modules are *representations* of $\mathbb{C}[G]$. The algebra $\mathbb{C}[G]$ is semisimple, therefore every $\mathbb{C}[G]$-module M is semisimple.

Choosing a basis for an n-dimensional $\mathbb{C}[G]$-module creates an associated repre-

sentation

$$\rho : \mathbb{C}[G] \to M_n(\mathbb{C})$$

of the group algebra $\mathbb{C}[G]$. Let $GL(n, \mathbb{C})$ be the group of invertible $n \times n$ matrices with entries in $\mathbb{C}$. Since each $g \in G$ is an invertible element of $\mathbb{C}[G]$, the restriction of ρ to G induces a group homomorphism

$$\rho|_G : G \to GL(n, \mathbb{C}).$$

Any such homomorphism is a *representation* of G. Since any representation of the group G may be extended linearly to a representation of the group algebra $\mathbb{C}[G]$, $\mathbb{C}[G]$-modules are also known as *representations* of G.

Let $X = \{x_1, \ldots, x_n\}$ be a finite set and suppose that G acts on X on the right. In other words, suppose there is a map $X \times G \to X$ where $(x, g) \mapsto xg$ such that $x1 = x$ and $(xh)g = x(hg)$ for all $x \in X$ and $h, g \in G$. Let $\mathbb{C}[X]$ be the vector space over $\mathbb{C}$ of formal linear combinations of elements of X. The right action of G on X induces a right $\mathbb{C}[G]$-module structure on $\mathbb{C}[X]$ where

$$(\sum_{x \in X} \alpha_x x)(\sum_{g \in G} \beta_g g) = \sum_{x \in X, g \in G} \alpha_x \beta_g (xg).$$

Elements of $\mathbb{C}[X]$ may also be viewed as complex-valued functions on X where the value $f \in \mathbb{C}[X]$ at the element $x \in X$, which we denote by $f(x)$, is the coefficient of x in f. Thus, as an element of $\mathbb{C}[X]$, $x \in X$ corresponds to the function $\delta_x : X \to \mathbb{C}$ where

$$\delta_x(x') = \begin{cases} 1 & \text{if } x = x' \\ 0 & \text{otherwise.} \end{cases}$$

For this reason, we refer to the basis $\{x_1, \ldots, x_n\}$ of $\mathbb{C}[X]$ as the *delta basis* of $\mathbb{C}[X]$.

When viewed as a linear transformation on $\mathbb{C}[X]$ with respect to the delta basis, each $g \in G$ corresponds to an $n \times n$ matrix with a single 1 in each row and column and zeros elsewhere. Such a matrix is a *permutation matrix* and the associated representation of G is a *permutation representation*. We also say that $\mathbb{C}[X]$ is a $\mathbb{C}[G]$-*permutation module*.

As an example of a permutation module, let B be a subgroup of G and let G/B denote the set of right cosets of B in G. The group G acts naturally on G/B where $(Bh, g) \mapsto B(hg)$. The resulting permutation module is $\mathbb{C}[G/B]$.

If G acts on X, we may place an equivalence relation on X by saying that the elements x and x' are equivalent if there is a $g \in G$ such that $xg = x'$. If $X_1, \ldots, X_k$ are the disjoint equivalence classes of X under this equivalence relation, then each $\mathbb{C}[X_i]$ is a submodule of $\mathbb{C}[X]$ and

$$\mathbb{C}[X] = \mathbb{C}[X_1] \oplus \cdots \oplus \mathbb{C}[X_k]. \tag{3.1}$$

The X_i are said to be *orbits*. We say that the action of G on X is *transitive* if X has only one orbit. In this case, X is also called a *homogeneous space* for G.

Note that (3.1) reduces the study of $\mathbb{C}[G]$-permutation modules to the study of those $\mathbb{C}[G]$-permutation modules arising from the transitive action of G on some set X. In this case we may identify the set X with the set of right cosets G/B where B is the stabilizer of a particular element x in X. More precisely, if $x \in X$ and $B = \{b \in G \mid xb = b\}$ is the stabilizer of x, then there is a bijection $X \to G/B$ where $xg \mapsto Bg$. Moreover, this bijection induces a $\mathbb{C}[G]$-module isomorphism $\phi : \mathbb{C}[X] \to \mathbb{C}[G/B]$. We may therefore assume that all $\mathbb{C}[G]$-permutation modules are of the form $\mathbb{C}[G/B]$ for some subgroup B of G.

3.2 Decomposing Representations

Let G be a finite group, let M be a finite-dimensional $\mathbb{C}[G]$-module, and let

$$M = M_1 \oplus \cdots \oplus M_n$$

be the isotypic decomposition of M. Let U_i be the simple $\mathbb{C}[G]$-module corresponding to the isotypic subspace M_i, let χ_i be the character of U_i, and let d_i be the dimension of U_i.

If $f \in M$, then one approach to computing the isotypic projections of f is to directly apply the projection operator defined in the following classical theorem:

Theorem 3.2.1. *If $f \in M$ and $p_i = (d_i/|G|) \sum_{g \in G} \chi_i(g^{-1})g$, then the isotypic projection of f onto the isotypic subspace M_i is fp_i.*

Proof. See, e.g., Theorem 8 in [41]. □

By Theorem 3.2.1, the projection of f onto the isotypic subspace M_i may be computed by directly applying p_i to f. There are, however, drawbacks to this approach. First, the computation of fp_i requires the choice of a basis for M. If p_i is realized as a matrix with respect to this basis, then directly applying p_i to f requires $O(\dim(M)^2)$ operations. This could be prohibitive if $\dim M$ is large. Second, to construct p_i using the above formula requires a sum over the group G together with an explicit knowledge of the action of each element of G on M. This too could be prohibitive if G is large.

Consider now an eigenspace approach to computing isotypic projections for the $\mathbb{C}[G]$-module M. Let C be a conjugacy class of G and let $z(C) \in \mathbb{C}[G]$ be defined by

$$z(C) = \sum_{g \in C} g.$$

The element $z(C)$ is the *class sum* of the conjugacy class C. Each class sum is contained in the center $Z(\mathbb{C}[G])$ of $\mathbb{C}[G]$. In fact, if $C_1, \ldots, C_l$ are the conjugacy classes of G, then the set $\{z(C_1), \ldots, z(C_l)\}$ of class sums forms a basis for $Z(\mathbb{C}[G])$.

Lemma 3.2.2. *If M is a $\mathbb{C}[G]$-module, then $\{z(C_1), \ldots, z(C_l)\}$ is a separating set for M.*

Proof. This is Theorem 3.3 in [9] rewritten using the language of separating sets. □

We may, of course, be able to find subsets of the set of class sums that also form separating sets for a given $\mathbb{C}[G]$-module M. Such subsets are used extensively in Chen [9], for example, where separating sets are known as *complete sets of commuting operators* when M is the regular representation $\mathbb{C}[G]_{\mathbb{C}[G]}$. Lastly, note that if G' is a subgroup of G, then the group algebra $\mathbb{C}[G']$ is a semisimple subalgebra of $\mathbb{C}[G]$. As described in Section 2.2, we may therefore take advantage of $\mathbb{C}[G']$-separating sets for a $\mathbb{C}[G]$-module M to compute the isotypic projections of an element $m \in M$.

3.3 Endomorphism Algebras

To conclude this chapter, we review the structure of the endomorphism algebra of a permutation module arising from the transitive action of a group on a set of right cosets. The results in this section will be used in Chapters 9 and 10 when we turn our attention to computing isotypic projections for the finite general linear group and the finite symplectic group. We follow closely the treatment found in [26].

Let G be a finite group, let B be a subgroup of G, and let $\mathbb{C}[G/B]$ be the corresponding $\mathbb{C}[G]$-permutation module. The endomorphism algebra $\mathrm{End}_{\mathbb{C}[G]}(\mathbb{C}[G/B])$ is known as the *Hecke algebra* of G with respect to B. Recall that $\mathbb{C}[G]$ acts on the right of $\mathbb{C}[G/B]$ while endomorphisms are written on the left.

For each $\varphi \in \mathrm{End}_{\mathbb{C}[G]}(\mathbb{C}[G/B])$, define $\dot{\varphi} : G/B \times G/B$ by

$$\varphi(Bx') = \sum_{Bx \in G/B} \dot{\varphi}(Bx, Bx')Bx$$

for all $x' \in G$. Let $D(G,B)$ be a set of double coset representatives of B in G such that $1 \in D(G,B)$. For each $x \in D(G,B)$, define $\epsilon_x \in \mathrm{End}_{\mathbb{C}[G]}(\mathbb{C}[G/B])$ by the condition

$$\dot{\epsilon}_x(Bx_1, Bx_2) = \begin{cases} 1 & \text{if } x_1 x_2^{-1} \in BxB \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 3.3.1. (Schur) *The set* $\mathcal{B} = \{\epsilon_x \mid x \in D(G,B)\}$ *forms a basis* (the Schur basis) *for the endomorphism algebra* $End_{\mathbb{C}[G]}(\mathbb{C}[G/B])$. *Moreover, for any* $x, y \in D(G,B)$

$$\epsilon_x \epsilon_y = \sum_{z \in D(G,B)} a_{xyz} \epsilon_z$$

where

$$a_{xyz} = \left| \frac{ByB \cap Bx^{-1}Bz}{B} \right|.$$

Before continuing, we recall the definition of a Coxeter group. Let I be a finite set and for any $i, j \in I$, let m_{ij} be a positive integer with $m_{ii} = 1$ and $m_{ij} = m_{ji} \geq 2$ if $i \neq j$. Let M denote the set of m_{ij}. The *Coxeter group* of type M is the group W given by generators $S = \{s_i\}_{i \in I}$ and relations as

$$W = \langle s_i \mid s_i^2 = (s_i s_j)^{m_{ij}} = 1 \text{ for all } i, j \in I \rangle.$$

Each generator has order 2 and is therefore an involution. We refer to the pair (W, S) as a *Coxeter system*. To each Coxeter system (W, S) we also associate a length function $l : W \to \mathbb{Z}$ where $l(w)$ is the least number of generators required to form w.

Suppose now that G is a group with a so-called BN-pair, in which case we have the following:

1. There exist subgroups $B, N \leq G$ such that $B \cap N$ is normal in N, and the quotient $W = N/B \cap N$ is generated by a set S of involutions.

2. We have a double coset decomposition $G = \coprod_{w \in W} B\dot{w}B$ where $\dot{w}$ denotes a representative of $w \in W$ in N.

3. The pair (W, S) is a Coxeter system, and if l is the corresponding length function, then
$$\dot{s}B\dot{w} = \begin{cases} B\dot{s}\dot{w} & \text{if } l(sw) > l(w) \\ B\dot{s}\dot{w} \cup B\dot{w}B & \text{if } l(sw) < l(w). \end{cases}$$

The group W is called the *Weyl group* of G. The double coset decomposition of G is called the *Bruhat decomposition* of G. Note that we may write Bw without specifying a representative of $w \in W$ in N since any two representatives differ only by an element of $B \cap N$.

Because of the Bruhat decomposition of G, we know that the elements $\{\dot{w} \mid w \in W\}$ form a set of representatives for the double cosets of B in G. For any $w \in W$, denote by T_w the corresponding Schur basis element of $\mathrm{End}_{\mathbb{C}[G]}(\mathbb{C}[G/B])$.

Theorem 3.3.2. (Iwahori, Matsumoto) *The Schur basis $\{T_w \mid w \in W\}$ of the Hecke algebra $End_{\mathbb{C}[G]}(\mathbb{C}[G/B])$ is such that*

$$T_s T_w = \begin{cases} T_{sw} & \text{if } l(sw) > l(w) \\ q_s T_{sw} + (q_s - 1)T_w & \text{if } l(sw) < l(w) \end{cases}$$

where $q_s = |B\dot{s}B/B|$ for all $s \in S$.

Corollary 3.3.3. *For each $Bg \in G/B$ and $s \in S$, $T_s(Bg)$ is the sum of q_s distinct right cosets of B.*

Proof. By definition, $T_s(B)$ is the sum of q_s distinct right cosets of B. Since $T_s(Bg) = (T_s(B))g$, the corollary follows. □

In his thesis, Hoefsmit [28] showed how the representation theory of the Hecke algebra $\mathrm{End}_{\mathbb{C}[G]}(\mathbb{C}[G/B])$ is related to that of the associated Weyl group W. In Chapters 9 and 10, we will make use of his results and of recent work by Ram [38] when we consider computing isotypic projections for permutation representations of the finite general linear group and the finite symplectic group. These are examples of groups with a BN-pair and their corresponding Schur bases are particularly useful when computing isotypic projections.

Chapter 4

Abelian Groups and FFTs

In this chapter we show how an eigenspace approach to computing isotypic projections for abelian groups leads to the Gentleman-Sande, or *decimation in frequency*, FFT (see [27]). We also remark on the use of generalized FFTs for computing isotypic projections for nonabelian groups.

4.1 The DFT and Isotypic Projections

Every abelian group is the direct product of cyclic groups. The problem of computing isotypic projections for abelian groups is therefore reduced to that of computing isotypic projections for cyclic groups (see, e.g., Section 3.2 of [41]).

Let G be the cyclic group $\mathbb{Z}/n\mathbb{Z}$ and let X be the set $\{x_0, \ldots, x_{n-1}\}$. Let ω be a primitive nth root of unity, let g be a generator for G, and let G act on X by setting $x_i g^j = x_{i+j}$ where all subscripts are taken modulo n. The resulting permutation representation $\mathbb{C}[X]$ has n isotypic subspaces $M_0, \ldots, M_{n-1}$ where each M_i is one-dimensional (and hence simple) with character χ_i defined by $\chi_i(g^j) = \omega^{ij}$.

Each element g^j of G forms a conjugacy class $C_j = \{g^j\}$. The eigenvalue of the

class sum $z(C_j)$ associated to the isotypic subspace M_i is therefore

$$\chi_i(z(C_j))/\chi_i(1) = \chi_i(g^j)/1 = \omega^{ij}.$$

Let $f \in \mathbb{C}[X]$ and let f_i be the isotypic projection of f onto the isotypic subspace M_i. Since ω is a primitive nth root of unity, the class sum $z(C_1)$ forms a separating set for $\mathbb{C}[X]$. The isotypic projection f_i may therefore be viewed as the projection of f onto the eigenspace of $z(C_1)$ with eigenvalue ω^i. By Theorem 3.2.1, this may be computed as

$$f_i = f(\frac{1}{n}\sum_{j=0}^{n-1}\omega^{-ij}g^j).$$

Note that $f_i(x_0) = \omega^{ik}f_i(x_k)$ and that f_i is therefore determined by

$$f_i(x_0) = f(x_0)(\frac{1}{n}\sum_{j=0}^{n-1}\omega^{-ij}g^j) = \frac{1}{n}\sum_{j=0}^{n-1}\omega^{ij}f(x_j),$$

which is the ith coefficient of the usual discrete Fourier transform (DFT) when applied to the vector f.

4.2 The Gentleman-Sande FFT

Suppose now that $n = pq$. Let G' be the subgroup of G of order q that is generated by g^p. Since the class sum $z(C_1)$ forms a separating set for $\mathbb{C}[X]$ it also forms a $\mathbb{C}[G']$-separating set. We could therefore compute the isotypic projections of f by first computing the isotypic projections of f with respect to $\mathbb{C}[G']$ and then projecting each of these projections onto the eigenspaces of $z(C_1)$.

The class sum $z(C_p)$ forms a separating set for $\mathbb{C}[X]$ with respect to $\mathbb{C}[G']$. The corresponding isotypic subspaces are $W_0, \ldots, W_{q-1}$ where the eigenvalue of $z(C_p)$ that

is associated to W_k is ω^{pk} and

$$W_k = M_k \oplus M_{k+q} \oplus \cdots \oplus M_{k+(p-1)q}.$$

The projection f'_k of f onto W_k is therefore

$$f_k + f_{k+q} + \cdots + f_{k+(p-1)q}. \tag{4.1}$$

By Theorem 3.2.1 we have that

$$f'_k = f(\frac{1}{q}\sum_{t=0}^{q-1} \omega^{-pkt} g^{pt}).$$

Note that $f'_k(x_s) = \omega^{pkt} f'_k(x_{s+pt})$ and that f'_k is therefore determined by the values $f'_k(x_0), \ldots, f'_k(x_{p-1})$. In this sense, since $f'_k(x_j)$ requires $O(q)$ operations to compute, f'_k requires $O(pq)$ operations to compute. Thus the isotypic projections $f'_0, \ldots, f'_{q-1}$ of f with respect to $\mathbb{C}[G']$ may be computed using $O(pq^2)$ operations.

Since $n = pq$, each $0 \leq i, j \leq n-1$ can be uniquely represented as $i = k + lq$ and $j = s + tp$ for some $0 \leq k, t \leq q-1$ and $0 \leq l, s \leq p-1$. Moreover, by (4.1), the isotypic projection $f_i = f_{k+lq}$ may be computed by projecting f'_k onto the eigenspace of $z(C_1)$ with eigenvalue $\omega^{(k+lq)}$. Recall that f_{k+lq} is determined by $f_{k+lq}(x_0)$ which

we may compute as

$$\begin{aligned}
f_{k+lq}(x_0) &= f'_k(x_0)\left(\frac{1}{n}\sum_{j=0}^{n-1}\omega^{-(k+lq)j}g^j\right) \\
&= \frac{1}{n}\sum_{j=0}^{n-1}\omega^{(k+lq)j}f'_k(x_j) \\
&= \frac{1}{pq}\sum_{s=0}^{p-1}\sum_{t=0}^{q-1}\omega^{(k+lq)(s+tp)}f'_k(x_{s+tp}) \\
&= \frac{1}{p}\sum_{s=0}^{p-1}\omega^{(k+lq)s}\frac{1}{q}\sum_{t=0}^{q-1}\omega^{(k+lq)tp}f'_k(x_{s+tp}) \\
&= \frac{1}{p}\sum_{s=0}^{p-1}\omega^{(k+lq)s}\frac{1}{q}\sum_{t=0}^{q-1}\omega^{pkt}f'_k(x_{s+tp}) \\
&= \frac{1}{p}\sum_{s=0}^{p-1}\omega^{(k+lq)s}\frac{1}{q}\sum_{t=0}^{q-1}f'_k(x_s) \\
&= \frac{1}{p}\sum_{s=0}^{p-1}\left(\omega^{ks}f'_k(x_s)\right)(\omega^q)^{ls}.
\end{aligned}$$

This is a DFT on p points applied to the function $\omega^{ks}f'_k$. Thus, if we have computed $f'_0, \ldots, f'_{q-1}$, we may compute the isotypic projection f_i using $O(p)$ operations. Since there are pq isotypic projections and the f'_k require $O(pq^2)$ operations to compute, we may compute the isotypic projections of $f \in \mathbb{C}[X]$ using $O(p^2q + pq^2) = O((p+q)pq)$ operations.

This particular fast Fourier transform is known as the Gentleman-Sande, or *decimation in frequency*, FFT (see [27]). Moreover, the approach to decomposing representations that is presented in this thesis may be viewed as a generalization of decimation in frequency since we too will be iteratively computing projections of projections. Another fast Fourier transform, the well-known Cooley-Tukey or *decimation in time* FFT (see [12]), has been generalized to many nonabelian groups. For exam-

ples and other references, see [4, 10, 17, 33]. As noted in the following remark, these generalized FFTs may also be used to compute isotypic projections.

Remark. Every semisimple algebra A is isomorphic to a direct sum $\oplus_{i=1}^{n} M_{N_i}(\mathbb{C})$ of matrix algebras (see, e.g., [22]). Moreover, the corresponding decomposition of A is the isotypic decomposition of the regular module A_A. Thus, if $\varphi : A \to \oplus_{i=1}^{n} M_{N_i}(\mathbb{C})$ is an isomorphism and $\psi_j : \oplus_{i=1}^{n} M_{N_i}(\mathbb{C}) \to \oplus_{i=1}^{n} M_{N_i}(\mathbb{C})$ is the projection of $\oplus_{i=1}^{n} M_{N_i}(\mathbb{C})$ onto its jth component $M_{N_j}(\mathbb{C})$, then the isotypic projection of $f \in A$ that corresponds to $M_{N_j}(\mathbb{C})$ is $(\varphi^{-1}\psi_j\varphi)f$.

The isomorphism $\varphi : A \to \oplus_{i=1}^{n} M_{N_i}(\mathbb{C})$ is often called a *discrete Fourier transform* for the algebra A, and an efficient algorithm for computing φf for arbitrary $f \in A$ is known as a *fast Fourier transform* or FFT. In particular, this terminology agrees with the use of the terms discrete and fast Fourier transform when A is the group algebra $\mathbb{C}[\mathbb{Z}/n\mathbb{Z}]$ (see, e.g., [10]). FFTs and their equally efficient inverses have been developed for groups such as the symmetric group, the hyperoctahedral group, and supersolvable groups. These FFTs and inverses require $O(|G|(\log|G|)^c)$ operations, where G is the group and c is a number that depends on G (see, e.g., [34]). Thus, by the comments above, n isotypic projections of a vector in $\mathbb{C}[G]$ may be computed using $O(n|G|(\log|G|)^c)$ operations.

The approach to computing isotypic projections presented in this thesis achieves similar bounds, but in a way that is straightforward to implement and that is easily applied to permutation representations of the form $\mathbb{C}[G/B]$. In particular, we need not worry about using *adapted representations* (see, e.g., [33]) which have been central to the development of generalized FFTs, but which limit the usefulness of FFTs when computing isotypic projections for arbitrary permutation modules.

Chapter 5

The Lanczos Iteration

In this chapter we review the Lanczos iteration and show how it is used to compute eigenspace projections for a real symmetric matrix. We also note how this gives rise to an efficient eigenspace projection method when the number of eigenspaces is relatively small and when the matrix may be applied efficiently. Good references for the Lanczos iteration are [13, 37, 44, 45].

5.1 Krylov Subspaces

Let $\mathbb{C}^N$ be the usual complex vector space of N-tuples with complex coefficients. We will view the elements of $\mathbb{C}^N$ as column matrices of size N. The matrices $M_N(\mathbb{C})$ may therefore be viewed as linear transformations of $\mathbb{C}^N$ with respect to the standard basis of $\mathbb{C}^N$.

Let $T \in M_N(\mathbb{C})$, let T^t denote the transpose of T, and let T^* denote the conjugate transpose of T. If $v, w \in \mathbb{C}^N$, then the usual inner product of v and w is v^*w. The norm of v is $||v|| = (v^*v)^{1/2}$. T is *symmetric* if $T = T^t$ and *hermitian* if $T = T^*$, in which case T is diagonalizable with real eigenvalues.

If $f \in \mathbb{C}^N$, then the *jth Krylov subspace* generated by T and f is the subspace $\mathcal{K}_j$ of $\mathbb{C}^N$ that is spanned by the vectors $f, Tf, \ldots, T^{j-1}f$. We write this as

$$\mathcal{K}_j = \langle f, Tf, \ldots, T^{j-1}f \rangle.$$

The T-invariant subspace $\mathcal{K} = \langle f, Tf, T^2f, \ldots \rangle$ is the *Krylov subspace* generated by T and f. Note that $\mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \mathcal{K}_3 \subseteq \cdots$ and that for some m, $\mathcal{K}_m = \mathcal{K}_{m+1} = \cdots = \mathcal{K}$.

Suppose now that $T \in M_N(\mathbb{C})$ is diagonalizable with n distinct eigenvalues. Then

$$\mathbb{C}^N = V_1 \oplus \cdots \oplus V_n$$

where the V_i are the n distinct eigenspaces of T. Each $f \in \mathbb{C}^N$ may therefore be written uniquely as $f = f_1 + \cdots + f_n$ where $f_i \in V_i$. We say that f_i is the *eigenspace projection of f onto the eigenspace V_i*. By the following lemma, we may restrict our attention to the Krylov subspace generated by T and f when computing these f_i.

Lemma 5.1.1. *If $T \in M_N(\mathbb{C})$ is diagonalizable and $f \in \mathbb{C}^N$, then the non-trivial projections of f onto the eigenspaces of T form a basis for the Krylov subspace generated by T and f.*

Proof. Suppose that T has n distinct eigenvalues $\mu_1, \ldots, \mu_n$ and that $f = f_1 + \cdots + f_n$, where f_i is the projection of f onto the eigenspace corresponding to the eigenvalue μ_i. We then have the following system of equations:

$$
\begin{aligned}
f &= f_1 + f_2 + \cdots + f_n \\
Tf &= \mu_1 f_1 + \mu_2 f_2 + \cdots + \mu_n f_n \\
T^2 f &= \mu_1^2 f_1 + \mu_2^2 f_2 + \cdots + \mu_n^2 f_n \\
&\vdots \\
T^{n-1} f &= \mu_1^{n-1} f_1 + \mu_2^{n-1} f_2 + \cdots + \mu_n^{n-1} f_n.
\end{aligned} \tag{5.1}
$$

The coefficients of the f_i in (5.1) form a Vandermonde matrix

$$
\begin{pmatrix}
1 & 1 & \cdots & 1 \\
\mu_1 & \mu_2 & \cdots & \mu_n \\
\mu_1^2 & \mu_2^2 & \cdots & \mu_n^2 \\
\vdots & \vdots & \ddots & \vdots \\
\mu_1^{n-1} & \mu_2^{n-1} & \cdots & \mu_n^{n-1}
\end{pmatrix}
$$

which is invertible since the μ_i are distinct (see, e.g., [23]). We may therefore solve the system for the f_i in terms of the $T^j f$. This shows that each f_i is contained in $\mathcal{K} = \langle f, Tf, T^2 f, \dots \rangle$. On the other hand, any power of T applied to f is a linear combination of the f_i. Thus $\mathcal{K}$ is spanned by the f_i. Since the non-trivial f_i are linearly independent, the lemma follows. □

Corollary 5.1.2. *The dimension of $\mathcal{K} = \langle f, Tf, T^2 f, \dots \rangle$ is equal to the number of non-trivial projections of f onto the eigenspaces of T.*

Corollary 5.1.3. *Eigenvectors of the restriction of T to $\mathcal{K}$ are scalar multiples of the eigenspace projections of f.*

Proof. This follows from the fact that each eigenspace of the restriction of T to $\mathcal{K}$ is

one-dimensional and is spanned by one of the non-trivial projections of f onto the eigenspaces of T. $\square$

If u is an eigenvector of the restriction of T to $\mathcal{K}$, then we may scale u into an eigenspace projection of f by Corollary 5.1.3. If the eigenspaces of the restriction of T to $\mathcal{K}$ are orthogonal, this may be computed as

$$\frac{u^* f}{u^* u} u. \tag{5.2}$$

Moreover, these computations may be done relative to a basis of $\mathcal{K}$ allowing us to gain efficiency if the dimension of $\mathcal{K}$ is small relative to N. For example, suppose $n = \dim \mathcal{K}$. Relative to a basis of $\mathcal{K}$, the computation in (5.2) requires $3n + 1$ operations. Relative to a basis of $\mathbb{C}^N$, however, this computation requires $3N + 1$ operations.

5.2 Restricting Real Symmetric Matrices to Krylov Subspaces

Let T be an $N \times N$ real symmetric matrix. For $f \in \mathbb{C}^N$, define the jth *Lanczos matrix* L_j to be the symmetric tridiagonal matrix

$$L_j = \begin{pmatrix} \alpha_1 & \beta_1 & & \\ \beta_1 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_{j-1} \\ & & \beta_{j-1} & \alpha_j \end{pmatrix}$$

whose entries are defined recursively using the *Lanczos iteration*:

The Lanczos Iteration
(assuming exact arithmetic)

$\beta_0 = 0$, $q_0 = 0$, $q_1 = f/||f||$

```
for i = 1, 2, 3, ...
    v = T q_i
    α_i = q_i^* v
    v = v − β_{i−1} q_{i−1} − α_i q_i
    β_i = ||v||
    if β_i ≠ 0
        q_{i+1} = v/β_i
    else
        q_{i+1} = 0
```

The Lanczos iteration is a modified version of the classical Gram-Schmidt orthogonalization process. At its heart is an efficient three-term recurrence which arises because the matrix T is real and symmetric. The usefulness of the Lanczos matrices, together with the q_i that are generated during the Lanczos iteration, is revealed in the following lemma:

Lemma 5.2.1. *If the dimension of the Krylov subspace $\mathcal{K} = \langle f, Tf, T^2f, \dots \rangle$ is m, then $\{q_1, \dots, q_m\}$ is an orthonormal basis for $\mathcal{K}$ and L_m is the restriction of T to $\mathcal{K}$ with respect to this basis.*

Although the Lanczos iteration is easily implemented, in finite precision arithmetic the q_i quickly lose their property of being orthogonal. They may even become linearly dependent (see, e.g., [44]). For this reason, some form of re-orthogonalization is usually introduced. For example, the *Lanczos iteration with complete re-orthogonalization*, as described in Parlett [37], re-orthogonalizes v against *all* of the previous $q_1, \dots, q_i$ after computing α_i and $v = \beta_i q_{i+1}$:

The Lanczos Iteration
with Complete Re-orthogonalization
(assuming finite precision arithmetic)

$\beta_0 = 0$, $q_0 = 0$, $q_1 = f/||f||$, $\epsilon =$ tolerance

`for` $i = 1, 2, 3, \ldots$
    $v = Tq_i$
    $\alpha_i = q_i^* v$
    $v = v - \beta_{i-1} q_{i-1} - \alpha_i q_i$
    `for` $j = 1$ `to` i
        $\gamma = q_{i-j+1}^* v$
        $v = v - \gamma q_{i-j+1}$
    $\beta_i = ||v||$
    `if` $\beta_i > \epsilon$
        $q_{i+1} = v/\beta_i$
    `else`
        $q_{i+1} = 0$

Remark. The Lanczos iteration with complete re-orthogonalization is much more stable than the Lanczos iteration without re-orthogonalization. In fact, the numerical stability of the Lanczos iteration with re-orthogonalization is comparable to that of the Givens and Householder algorithms that, like the Lanczos iteration, reduce a matrix to tridiagonal form (see Chapter 6, Section 41 of [45]).

To get a sense of how much work it takes to compute the Lanczos iteration with complete re-orthogonalization, let T^{op} be the number of operations needed apply the matrix T to an arbitrary vector, either directly or through a given subroutine. Note that T^{op} is never more than the number of nonzero entries of T.

Lemma 5.2.2. *If T is an $N \times N$ real symmetric matrix and $f \in \mathbb{C}^N$, then*

$$O(nT^{\text{op}} + n^2 N)$$

operations are required to compute n iterations of the Lanczos iteration with complete

re-orthogonalization for T and f.

Proof. It is easy to see that the Lanczos iteration without re-orthogonalization requires $O(nT^{\text{op}} + nN)$ operations. Since complete re-orthogonalization requires an additional $O(n^2N)$ operations, the lemma follows. □

5.3 The Lanczos Eigenspace Projection Method

We may now state the following theorem. Its proof outlines a method for computing projections onto the eigenspaces of a real symmetric matrix.

Theorem 5.3.1. *If T is an $N \times N$ real symmetric matrix with n distinct eigenvalues and f is a nonzero vector in $\mathbb{C}^N$, then the projections of f onto the eigenspaces of T requires $O(nT^{\text{op}} + n^2N)$ operations.*

Proof. The claim follows directly from Parlett's discussion in [37] of the Rayleigh-Ritz procedure applied to the sequence of Krylov subspaces $\mathcal{K}_1, \mathcal{K}_2, \ldots$ generated by T and f. The method is important, however, so we include the details.

Suppose that f has $m \leq n$ nontrivial projections $f_1, \ldots, f_m$ onto the eigenspaces of T. Let μ_i be the eigenvalue corresponding to the eigenspace containing f_i. Let L_m be the mth Lanczos matrix generated during the Lanczos iteration with respect to T and f. Let $\{q_1, \ldots, q_m\}$ be the corresponding orthonormal basis of the Krylov subspace $\mathcal{K}$ generated by T and f.

It will be useful to express the elements of $\mathcal{K}$ with respect to the basis $\{q_1, \ldots, q_m\}$. Thus, if $v \in \mathcal{K}$, let $\tilde{v}$ denote v with respect to $\{q_1, \ldots, q_m\}$. In other words, if $v = \sum_{i=1}^{m} \alpha_i q_i$, then $\tilde{v} = (\alpha_1, \ldots, \alpha_m)^t$.

Since $\mathcal{K}$ is spanned by the f_i, $\mathcal{K} = \mathcal{K}_m$ and each μ_i is an eigenvalue of L_m. Let $\tilde{u}_i$ be an eigenvector of L_m with eigenvalue μ_i such that $||\tilde{u}_i|| = 1$. Note that since L_m

is a real symmetric matrix, $\{\tilde{u}_1, \ldots, \tilde{u}_m\}$ is an orthonormal basis for $\mathcal{K}$.

Since $q_1 = ||f||^{-1} f$, $\tilde{f} = (||f||, 0, \ldots, 0)^t$. It follows that $\tilde{f}_i = (\tilde{u}_i^* \tilde{f})\tilde{u}_i$ is the eigenspace projection f_i with respect to the basis $\{q_1, \ldots, q_m\}$. Thus, if Q_m is the $N \times m$ matrix whose ith column is the vector q_i, then $f_i = Q_m \tilde{f}_i$. We may therefore compute the eigenspace projections of f as follows:

Stage 1: Generate L_m and Q_m by using the Lanczos iteration with complete reorthogonalization with T and f until a zero vector appears.

Stage 2: Compute the m eigenvalues $\mu_1, \ldots, \mu_m$ and corresponding eigenvectors $\tilde{u}_1, \ldots, \tilde{u}_m$ of L_m.

Stage 3: For $1 \leq i \leq m$, compute $\tilde{f}_i = (\tilde{u}_i^* \tilde{f})\tilde{u}_i$.

Stage 4: For $1 \leq i \leq m$, compute $f_i = Q_m \tilde{f}_i$.

Stage 1 requires $O(mT^{\text{op}} + m^2 N)$ operations and Stage 2 requires $O(m^3)$ operations due to the tridiagonal form of T_m (see [45]). Stage 3 requires $O(m^2)$ operations and Stage 4 requires $O(m^2 N)$ operations. Since $m \leq n \leq N$, the theorem follows. □

Remark. Note that the coefficient implied by $O(nT^{\text{op}} + n^2 N)$ in Theorem 5.3.1 is independent of n, T^{op}, and N. We will implicitly make use of this fact throughout the rest of the thesis.

We will refer to the projection method outlined in Theorem 5.3.1 as the *Lanczos Eigenspace Projection Method*, or LEPM. As an example of how the LEPM works,

suppose we want to compute the projections of the vector

$$f = \begin{pmatrix} 4 \\ 3 \\ 9 \\ 1 \\ 7 \\ 3 \end{pmatrix}$$

onto the eigenspaces of the real symmetric matrix

$$T = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}. \tag{5.3}$$

In Stage 1 of the LEPM, the Lanczos iteration with complete re-orthogonalization, when applied to T and f, creates the matrices

$$L_3 = \begin{pmatrix} 2.9091 & 1.8417 & 0 \\ 1.8417 & 0.7751 & 0.8113 \\ 0 & 0.8113 & -1.6842 \end{pmatrix}$$

and

$$Q_3 = \begin{pmatrix} 0.3114 & 0.3535 & -0.5331 \\ 0.2335 & 0.3497 & 0.4146 \\ 0.7006 & -0.3881 & 0.0592 \\ 0.0778 & 0.5956 & 0.5331 \\ 0.5449 & -0.1422 & 0.1777 \\ 0.2335 & 0.4765 & -0.4739 \end{pmatrix}.$$

Recall that the columns of Q_3 form an orthonormal basis of $\mathcal{K} = \langle f, Tf, T^2f, \dots \rangle$ and that L_3 is the restriction of T to $\mathcal{K}$ with respect to this basis. In Stage 2, we compute the eigenvalues $\mu_1 = 4.0000$, $\mu_2 = 0.0000$, $\mu_3 = 2.0000$ of L_3 and their corresponding eigenvectors

$$\tilde{u}_1 = \begin{pmatrix} 0.8581 \\ 0.5083 \\ 0.0725 \end{pmatrix}, \tilde{u}_2 = \begin{pmatrix} 0.4954 \\ -0.7826 \\ -0.3770 \end{pmatrix}, \tilde{u}_3 = \begin{pmatrix} 0.1348 \\ -0.3594 \\ 0.9234 \end{pmatrix}.$$

In Stage 3, we scale the $\tilde{u}_i$ into the eigenspace projections of f by computing $(\tilde{f}^* \tilde{u}_i)\tilde{u}_i$ (recall that $\tilde{f} = (||f||, 0, 0)^t$):

$$\tilde{f}_1 = \begin{pmatrix} 9.4588 \\ 5.6029 \\ 0.7997 \end{pmatrix}, \tilde{f}_2 = \begin{pmatrix} 3.1529 \\ -4.9803 \\ -2.3990 \end{pmatrix}, \tilde{f}_3 = \begin{pmatrix} 0.2335 \\ -0.6225 \\ 1.5993 \end{pmatrix}.$$

Finally, in Stage 4, we apply Q_3 to each of the $\tilde{f}_i$ to express the eigenspace projections

of f in the original basis:

$$f_1 = \begin{pmatrix} 4.5000 \\ 4.5000 \\ 4.5000 \\ 4.5000 \\ 4.5000 \\ 4.5000 \end{pmatrix}, f_2 = \begin{pmatrix} 0.5000 \\ -2.0000 \\ 4.0000 \\ -4.0000 \\ 2.0000 \\ -0.5000 \end{pmatrix}, f_3 = \begin{pmatrix} -1.0000 \\ 0.5000 \\ 0.5000 \\ 0.5000 \\ 0.5000 \\ -1.0000 \end{pmatrix}.$$

Remark. The LEPM is a sensible way of computing eigenspace projections only if n is much less than N and T^{op} is much less than N^2. After all, a naive algorithm that uses matrix multiplication to directly compute the f_i requires $O(nN^2)$ operations. Thus, for our method to be efficient, we must have an efficient algorithm for applying the real symmetric matrix T and the number of distinct eigenvalues of T must be small relative to the dimension of the space upon which T acts.

5.4 The Lanczos Isotypic Projection Method

In this section, we combine the results of Section 2.2 and Section 5.3 to create an isotypic projection method that relies on the use of separating sets of real symmetric matrices.

Let A be a semisimple algebra, let A' be a semisimple subalgebra of A, and let M be an A-module. Let $\mathcal{B}$ be a basis of M and assume that all vectors and linear transformations of M are viewed as matrices with respect to this basis.

Let $\{T_1, \ldots, T_k\}$ be an A'-separating set for M. If $\{T_1, \ldots, T_k\}$ forms a collection of real symmetric matrices with respect to $\mathcal{B}$, then by Lemma 2.2.2 we may compute the isotypic projections of a vector $m \in M$ as follows:

Stage 0: Compute the isotypic projections of m with respect to the subalgebra A'.

Stage 1: Using the LEPM, compute the projections of each of the previously computed projections onto each of the eigenspaces of T_1.

$\vdots$

Stage k: Using the LEPM, compute the projections of each of the previously computed projections onto each of the eigenspaces of T_k.

Stage $k+1$: For each isotypic subspace M_i of M, compute the sum m_i of the previously computed projections that are contained in M_i.

We will refer to this approach to computing isotypic projections as the *Lanczos Isotypic Projection Method* or LIPM.

Let $\iota_{\mathcal{B}}(M)$ be the least number of operations needed to compute the isotypic projections of an arbitrary vector in M with respect to the basis $\mathcal{B}$ of M. We will also denote $\iota_{\mathcal{B}}(M)$ by $\iota(M)$ if the basis $\mathcal{B}$ is understood. We may now state our main theorem.

Main Theorem 5.4.1. *Let A be a semisimple algebra, let A' be a semisimple subalgebra of A, and let M be an N-dimensional A-module with basis $\mathcal{B}$. Suppose that $M_{A'}$ has n' isotypic subspaces $M'_1, \ldots, M'_{n'}$ and that M'_j contains τ_j sub-isotypic spaces. If $\{T_1, \ldots, T_k\}$ is an A'-separating set of real symmetric matrices for M with respect to $\mathcal{B}$, then*

$$\iota_{\mathcal{B}}(M) = \iota_{\mathcal{B}}(M_{A'}) + O\left(\sum_{j=1}^{n'} \sum_{i=1}^{k} \left(\tau_j T_i^{\mathrm{op}} + \tau_j^2 N\right)\right).$$

Proof. We may compute the isotypic projections of a vector $m \in M$ by using the LIPM with the A'-separating set $\{T_1, \ldots, T_k\}$. Stage 0 of the LIPM requires $\iota_{\mathcal{B}}(M_{A'})$

operations and Stage $k+1$ requires no more than $(\sum_{j=1}^{n'} \tau_j)N$ operations. By Theorem 5.3.1, for each isotypic subspace M'_j of $M_{A'}$, Stage i requires $O\left(\tau_j T_i^{\text{op}} + \tau_j^2 N\right)$ operations for $1 \leq i \leq k$. Thus Stage 1 through Stage $k+1$ of the LIPM requires $O(\sum_{j=1}^{n'} \sum_{i=1}^{k} (\tau_j T_i^{\text{op}} + \tau_j^2 N))$ operations. The theorem follows immediately. □

Remark. Note that Theorem 5.4.1 also holds for left A-modules.

Corollary 5.4.2. *Let A be a semisimple algebra and let M be an N-dimensional A-module with n isotypic subspaces. If $\{T_1, \ldots, T_k\}$ is a separating set of real symmetric matrices for M, then*

$$\iota(M) = O\left(\sum_{i=1}^{k} (nT_i^{\text{op}} + n^2 N)\right).$$

Proof. Let A' be the trivial subalgebra of A containing only scalar multiples of the identity. The separating set $\{T_1, \ldots, T_k\}$ is then also an A'-separating set. The corresponding sub-isotypic spaces are precisely the isotypic subspaces of M, and Stage 0 and Stage $k+1$ of the LIPM with $\{T_1, \ldots, T_k\}$ are trivial. □

Corollary 5.4.3. *Let A be a semisimple algebra and let M be an N-dimensional A-module with n isotypic subspaces. If $\{T_1, \ldots, T_k\}$ is a separating set of real symmetric matrices for M and $l = (1/knN) \sum_{i=1}^{k} T_i^{\text{op}}$, then*

$$\iota(M) = O\left((l+1)kn^2 N\right).$$

Proof. This follows directly by substitution. □

Chapter 6

Distance Transitive Graphs

Let X be a connected graph and denote the distance function of X by d. Let k be the *diameter* of X which is the maximum distance between any two vertices of X. A group G of automorphisms of X is said to be *distance transitive* on X if G is transitive on each of the sets $\{(x, x') \mid x, x' \in X \text{ and } d(x, x') = i\}$ for $0 \leq i \leq k$. A graph is said to be *distance transitive* if it is connected and has a distance transitive group of automorphisms. For example, the 2-element subsets of a 4-element set form a distance transitive graph where two 2-element subsets are adjacent if their intersection has size 1 (see Figure 6.1). A good reference for distance transitive graphs is [7].

Let X be a distance transitive graph, let G be a distance transitive group of automorphisms of X, and let $\mathbb{C}[X]$ be the permutation representation of G induced by the action of G on the vertices of X. The *adjacency operator* of X is the linear transformation $A : \mathbb{C}[X] \to \mathbb{C}[X]$ where

$$A(x) = \sum_{x' : d(x,x')=1} x'$$

for all $x \in X$. The operator A has $k + 1$ distinct eigenvalues which are also the zeros

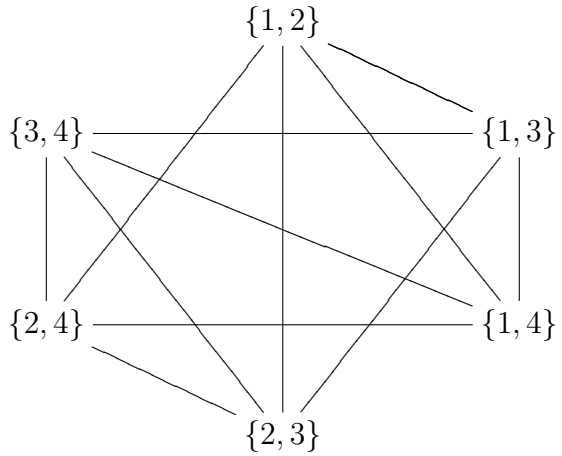

Figure 6.1: A Distance Transitive Graph

of certain polynomials associated with the graph X (see, e.g., [7]). For example, the adjacency operator of the graph in Figure 6.1, relative to its delta basis, is

$$A = \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}$$

which is also the matrix found in (5.3). It has 3 distinct eigenvalues.

Lemma 6.0.4. *The isotypic subspaces of $\mathbb{C}[X]$ are precisely the eigenspaces of A.*

Proof. This follows from Section 2 of Stanton [43]. □

Thus, by Corollary 5.4.2, we have the following theorem:

Theorem 6.0.5. *Let X be a distance transitive graph with diameter k, let G be a distance transitive group of automorphisms of X, and let $\mathbb{C}[X]$ be the associated*

permutation representation of G. If A is the adjacency operator of X, then

$$\iota(\mathbb{C}[X]) = O(kA^{\mathrm{op}} + k^2|X|).$$

A direct matrix multiplication approach to computing isotypic projections for $\mathbb{C}[X]$ requires $O(k|X|^2)$ operations. Although $O(kA^{\mathrm{op}} + k^2|X|)$ may yield a better upper bound, we may be able to gain even more efficiency by taking advantage of the graph structure of X. For this, the notion of a Radon transform is helpful.

6.1 Radon Transforms

Let G be a finite group acting on finite sets X and Y and giving permutation representations $\mathbb{C}[X]$ and $\mathbb{C}[Y]$, respectively. In addition, suppose there is an incidence relation between X and Y where we write $x \sim y$ if $x \in X$ is incident to $y \in Y$. The *Radon transform* $R : \mathbb{C}[X] \to \mathbb{C}[Y]$ is then defined by

$$R(x) = \sum_{y : x \sim y} y$$

for all $x \in X$ (see [6]). The adjoint $R^* : \mathbb{C}[Y] \to \mathbb{C}[X]$ of R is defined by

$$R^*(y) = \sum_{x : x \sim y} x$$

for all $y \in Y$.

Suppose now that X is a distance transitive graph with respect to G and let X' be a complete subgraph of X that contains at least two vertices. Recall that a graph is said to be *complete* if every pair of distinct vertices is adjacent. Let Y be the collection of distinct images of X' under the action of G on X and say that $x \in X$ is

incident to $y \in Y$ if x is a vertex of y. Let $R : \mathbb{C}[X] \to \mathbb{C}[Y]$ be the associated Radon transform. For convenience, we say that Y is a *complete covering of X with Radon transform R*. Note that, with respect to the delta bases of $\mathbb{C}[X]$ and $\mathbb{C}[Y]$, R^*R is a matrix with integer coefficients, $R^* = R^t$, and $(R^tR)^t = R^tR^{tt} = R^tR$. Thus R^*R is a real symmetric matrix.

We will make use of the integers r and s that are defined in the following lemma:

Lemma 6.1.1. *There are integers r and s such that*

$$|\{y \in Y \mid x \sim y\}| = r$$

for every vertex x of X and

$$|\{y \in Y \mid x \sim y \text{ and } x' \sim y\}| = s$$

for every edge $\{x, x'\}$ of X.

Proof. This follows from the fact that X is a distance transitive graph. □

Lemma 6.1.2. *If $A : \mathbb{C}[X] \to \mathbb{C}[X]$ is the adjacency operator of X and $I : \mathbb{C}[X] \to \mathbb{C}[X]$ is the identity, then $A = (1/s)(R^*R - rI)$.*

Proof. This follows from the fact that, for each $x \in X$,

$$(R^*R)(x) = \sum_{y : x \sim y} \sum_{x' : x' \sim y} x' = rx + s \left(\sum_{x' : d(x,x') = 1} x' \right) = (rI + sA)(x).$$

□

Lemma 6.1.3. *If X is a distance transitive graph and Y is a complete covering of X with Radon transform R, then $\{(R^*R)\}$ is a separating set for $\mathbb{C}[X]$ and $(R^*R)^{\mathrm{op}} \leq 2r|X|$.*

Proof. Let A be the adjacency operator of X. The product R^*R and the adjacency operator A have the same eigenspaces by Lemma 6.1.2, therefore $\{(R^*R)\}$ is a separating set since $\{A\}$ is a separating set by Lemma 6.0.4.

We may apply R^*R to a vector $f \in \mathbb{C}[X]$ by first computing Rf and then $R^*(Rf)$. Furthermore, when regarded as a matrix with respect to the delta bases of $\mathbb{C}[X]$ and $\mathbb{C}[Y]$, both R and R^* contain $r|X|$ nonzero entries. It follows that $(R^*R)^{\mathrm{op}} \leq R^{*\mathrm{op}} + R^{\mathrm{op}} \leq r|X| + r|X| = 2r|X|$. □

By Corollary 5.4.2 and Lemma 6.1.3, we therefore have the following theorem:

Theorem 6.1.4. *Let X be a distance transitive graph and let Y be a complete covering of X. If X has diameter k and $|\{y \in Y \mid x \sim y\}| = r$ for every vertex x of X, then*

$$\iota(\mathbb{C}[X]) = O\left(kr|X| + k^2|X|\right).$$

Since X is a distance transitive graph, there is an integer a such that, for every vertex x of X,

$$|\{x' \in X \mid d(x, x') = 1\}| = a.$$

Applying the adjacency operator of X directly therefore requires no more than $a|X|$ operations. Thus, if r is noticeably less than a, then by Theorem 6.1.4 we may want to use the associated Radon transform and its adjoint in the LIPM rather than the adjacency operator to compute the isotypic projections of a vector in $\mathbb{C}[X]$. We illustrate this in the next two sections.

6.2 The Johnson Graph

Let $n \geq 2$ and let $k \leq n/2$. The k-element subsets $X^{(n-k,k)}$ of $\{1, \ldots, n\}$ form a distance transitive graph with automorphism group S_n by defining two k-element subsets to be adjacent if their intersection has size $k-1$. The resulting graph is known as the *Johnson graph*. It has diameter k and is sometimes denoted by $J(n,k)$.

Each vertex of $J(n,k)$ is adjacent to $k(n-k)$ other vertices and $|X^{(n-k,k)}| = \binom{n}{k}$. The number of operations required to directly apply the adjacency operator A is therefore $k(n-k)\binom{n}{k}$. By Theorem 6.0.5, we therefore have that

$$\iota(\mathbb{C}[X^{(n-k,k)}]) = O\left(k^2(n-k)\binom{n}{k}\right). \tag{6.1}$$

For each $(k-1)$-element subset $y \in X^{(n-(k-1),k-1)}$ there is a corresponding complete subgraph of $J(n,k)$ consisting of those $x \in X^{(n-k,k)}$ that contain y. The collection Y of these subgraphs forms a complete cover of $J(n,k)$ and each vertex of $J(n,k)$ is contained in k such subgraphs. Thus, by Theorem 6.1.4, we have the following improvement to (6.1):

Theorem 6.2.1. *If $n \geq 2$, $k \leq n/2$, and $\mathbb{C}[X^{(n-k,k)}]$ is the $\mathbb{C}[S_n]$-permutation module associated to the Johnson graph $J(n,k)$, then*

$$\iota\left(\mathbb{C}[X^{(n-k,k)}]\right) = O\left(k^2\binom{n}{k}\right).$$

We summarize the results of this section in Table 6.1. Note that the bounds involving the LIPM compare favorably to the upper bound of

$$O\left(\binom{n}{k}^2 + \binom{n}{k} k \log^2 k\right)$$

given by [21].

LIPM with R^*R	LIPM with A	Direct Matrix Multiplication
$O\left(k^2\binom{n}{k}\right)$	$O\left(k^2(n-k)\binom{n}{k}\right)$	$O\left(k\binom{n}{k}^2\right)$

Table 6.1: Upper bounds on $\iota\left(\mathbb{C}[X^{(n-k,k)}]\right)$.

6.3 The Grassmann Graph

Let $n \geq 2$, let $k \leq n/2$, and let V be an n-dimensional vector space over the finite field $\mathbb{F}_q$ of q elements. Let $GL(n,q)$ be the group of automorphisms of V. The k-dimensional subspaces $X_{(n-k,k)}$ of V form a distance transitive graph with respect to $GL(n,q)$ by defining two k-dimensional subspaces to be adjacent if their intersection is a $(k-1)$-dimensional subspace of V. The resulting graph is known as the *Grassmann graph*. It has diameter k and is analogous to the Johnson graph $J(n,k)$. We will denote it by $G(n,k,q)$. See [7] for details concerning the Grassmann graph.

For each non-negative integer m, let $[m] = 1 + q + q^2 + \cdots + q^{m-1}$, let $[m]! = [m][m-1]\cdots[1]$ if $m > 0$, and let $[0]! = 1$. Note that $[m] = (q^m - 1)/(q-1)$, $[0] = 0$, and $[1] = 1$. Define

$$\binom{m}{l}_q = \begin{cases} [m]!/([l]![m-l]!) & \text{if } m \geq l \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Each vertex of $G(n,k,q)$ is adjacent to $q[k][n-k]$ other vertices and $|X_{(n-k,k)}| =$

$\binom{n}{k}_q$. Direct multiplication of the adjacency operator A of $G(n,k,q)$ therefore requires $q[k][n-k]\binom{n}{k}_q$ operations. By Theorem 6.0.5, we have that

$$\iota\left(\mathbb{C}[X_{(n-k,k)}]\right) = O\left(kq[k][n-k]\binom{n}{k}_q\right). \tag{6.2}$$

Each $(k-1)$-dimensional subspace $y \in X_{(n-(k-1),k-1)}$, in analogy with the Johnson graph, corresponds to a complete subgraph of $G(n,k,q)$ consisting of those $x \in X_{(n-k,k)}$ that contain y. The collection Y of such subgraphs forms a complete cover of $G(n,k,q)$ and each vertex of $G(n,k,q)$ is contained in $[k]$ such subgraphs. By Theorem 6.1.4, we therefore have the following improvement to (6.2):

Theorem 6.3.1. *Let $n \geq 2$ and $k \leq n/2$. Let $\mathbb{C}[X_{(n-k,k)}]$ be the $\mathbb{C}[GL(n,q)]$-permutation module associated to the Grassmann graph $G(n,k,q)$. Then*

$$\iota\left(\mathbb{C}[X_{(n-k,k)}]\right) = O\left(k[k]\binom{n}{k}_q\right).$$

We summarize the results of this section in Table 6.2. As with the Johnson Graph, note that the bounds involving the LIPM compare favorably to the upper bound of

$$O\left(\binom{n}{k}_q^2 + \binom{n}{k}_q k\log^2 k\right)$$

given by [21].

LIPM with R^*R	LIPM with A	Direct Matrix Multiplication
$O\left(k[k]\binom{n}{k}_q\right)$	$O\left(kq[k][n-k]\binom{n}{k}_q\right)$	$O\left(k\binom{n}{k}_q^2\right)$

Table 6.2: Upper bounds on $\iota\left(\mathbb{C}[X_{(n-k,k)}]\right)$.

Chapter 7

The Symmetric Group

Spectral analysis for nonabelian groups has found its most success with the analysis of ranked data (see [14, 15, 40]). Ranked data arises when respondents are given a list of n items which they are asked to rank in terms of preference. We say that such a ranking is *full* if the respondents are asked to rank each element of the list. On the other hand, we say that a ranking is a *partial ranking of shape* λ if for some sequence $\lambda = (\lambda_1, \ldots, \lambda_m)$ of positive integers whose sum is n, the respondents are asked to choose their top λ_1 items, then their next top λ_2 items, and so on, with no internal ordering. Note that a full ranking is a partial ranking of shape $(1, \ldots, 1)$.

If X^λ is the set of possible partial rankings of shape λ, the *partially ranked data of shape* λ is the function $f \in C[X^\lambda]$ where, for each $x \in X^\lambda$, $f(x)$ is the number of respondents choosing the partial ranking x. For an example of partially ranked data, consider a lottery in which participants are asked to choose five numbers from the set $\{1, \ldots, 39\}$. Each lottery ticket corresponds to a partial ranking of shape $(5, 34)$, and the relevant ranked data is then the function that assigns to each such ranking the number of tickets corresponding to that ranking that were sold.

For another example of ranked data, consider the partially ranked data that arises

when a film society asks its members to choose, from a list of ten movies, their three favorite movies, then their next three favorite movies. Their choices correspond to partial rankings of shape $(3, 3, 4)$, and the relevant partially ranked data is the function that assigns to each such ranking the number of members choosing that ranking.

The natural action of the symmetric group S_n on the n items in the list induces a $\mathbb{C}[S_n]$-permutation module structure on $\mathbb{C}[X^\lambda]$. Moreover, as noted in Chapter 1, the isotypic subspaces of $\mathbb{C}[X^\lambda]$ correspond to certain *pure higher order effects* associated to the ranked data $f \in \mathbb{C}[X^\lambda]$ (see [15, 40]). Computing the isotypic projections of f can therefore lead to some insight into how the respondents went about choosing their rankings.

In this chapter, we show how the LIPM gives rise to an efficient isotypic projection method for representations of the symmetric group that are associated to ranked data.

7.1 Representation Theory

We begin by reviewing the relevant representation theory of the symmetric group. See, for example, [29] for proofs and further details.

Let n be a positive integer and let S_n be the symmetric group of permutations of the set $\{1, \ldots, n\}$. The transpositions

$$s_i = (i-1, i) \text{ for } 2 \leq i \leq n$$

generate S_n and satisfy the relations

$$s_i s_j = s_j s_i, \text{ for } |i-j| > 1$$

$$s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}, \text{ for } 2 \leq i \leq n-1$$

$$s_i^2 = 1, \text{ for } 2 \leq i \leq n.$$

In particular, S_n is a Coxeter group of order $n!$ (see Section 3.3). For $1 \leq k \leq n$, we identify the subgroup of S_n that is generated by $s_2, \ldots, s_k$ with the group S_k.

A *composition* of n is a sequence $\lambda = (\lambda_1, \ldots, \lambda_m)$ of positive integers whose sum is n. If $\lambda_1 \geq \cdots \geq \lambda_m$, then λ is a *partition* of n. Note that to each composition λ, there corresponds a partition $\overline{\lambda}$ obtained by arranging the parts of λ in non-increasing order. The partitions of n form a partially ordered set under the *dominance order* where if λ and λ' are partitions of n, then we say that λ *dominates* λ' if $\lambda_1 + \cdots + \lambda_i \geq \lambda'_1 + \cdots + \lambda'_i$ for all $i \geq 1$. If λ dominates λ', then we write $\lambda \trianglerighteq \lambda'$.

As is often the case, we will identify the composition $\lambda = (\lambda_1, \ldots, \lambda_m)$ of n with the *Ferrers diagram* of shape λ, which is the left-justified array of dots with λ_i dots in the ith row (see Figure 6.1). If the dots of a Ferrers diagram of shape λ are replaced

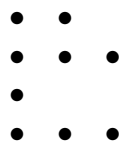

Figure 7.1: The Ferrers diagram of shape $(2, 3, 1, 3)$.

by the numbers $1, \ldots, n$ without repetition, then we create a *Young tableau* of shape λ. Two Young tableaux are said to be equivalent if they differ only by a permutation of the entries within the rows of each tableau. An equivalence class of tableaux is a *tabloid*. A tabloid is denoted by forming a representative tableau and drawing lines

between the rows (see Figure 6.2).

$$\begin{array}{llll} 4 & 9 & & \\ 5 & 2 & 3 & \\ 7 & & & \\ 8 & 1 & 6 & \end{array} \qquad \begin{array}{lll} 9 & 4 & \\ 2 & 5 & 3 \\ 7 & & \\ 6 & 8 & 1 \end{array}$$

$$\begin{array}{|lll|} \hline 4 & 9 & \\ \hline 2 & 3 & 5 \\ \hline 7 & & \\ \hline 1 & 6 & 8 \\ \hline \end{array}$$

Figure 7.2: Two equivalent tableaux and their tabloid.

Let X^λ be the set of tabloids of shape λ. Note that the set X^λ naturally corresponds to the set of rankings of shape λ since each row of a tabloid may be viewed as a ranked subset of an n-element set. Moreover, we may rearrange the subsets in each ranking so that their sizes are in non-increasing order. Thus, we may assume that λ is a partition of n.

The simple modules of $\mathbb{C}[S_n]$ are indexed by the partitions of n. Thus, for every partition μ of n, there is a simple $\mathbb{C}[S_n]$-module U^μ with character χ_μ. These modules form a complete (up to isomorphism) collection of simple $\mathbb{C}[S_n]$-modules.

Let λ be a partition of n. The action of S_n on $\{1, \ldots, n\}$ induces an action of S_n on X^λ and we denote the resulting $\mathbb{C}[S_n]$-permutation module $\mathbb{C}[X^\lambda]$ by M^λ. It is well known that the $\mathbb{C}[S_n]$-module M^λ is the direct sum of isotypic subspaces

$$M^\lambda = \bigoplus_{\mu \trianglerighteq \lambda} M^{\mu\lambda}$$

where $M^{\mu\lambda}$ is the isotypic subspace containing those simple submodules of M^λ that are isomorphic to U^μ.

As a $\mathbb{C}[S_{n-1}]$-module, the simple $\mathbb{C}[S_n]$-module U^μ is a direct sum of simple mod-

ules

$$U^{\mu} = \bigoplus_{\eta \in \mu^-} U^{\eta} \tag{7.1}$$

where the sum is over those partitions η of $n-1$ that are obtained from μ by removing one dot. Thus, as a $\mathbb{C}[S_{n-1}]$-module, the isotypic subspace $M^{\mu\lambda}$ of M^{λ} is a direct sum

$$M^{\mu\lambda} = \bigoplus_{\eta \in \mu^-} M^{\eta\mu\lambda}$$

of sub-isotypic spaces where $M^{\eta\mu\lambda}$ is the isotypic subspace of $M^{\mu\lambda}$ corresponding to the simple $\mathbb{C}[S_{n-1}]$-module U^{η}. We therefore have the decomposition

$$M^{\lambda} = \bigoplus_{\mu \trianglerighteq \lambda} \bigoplus_{\eta \in \mu^-} M^{\eta\mu\lambda} \tag{7.2}$$

of M^{λ} into sub-isotypic spaces with respect to $\mathbb{C}[S_{n-1}]$.

7.2 Separating Sets

Let $z^{(n)} \in C[S_n]$ be the class sum of transpositions. Thus

$$z^{(n)} = \sum_{1 \le i < j \le n} (i, j).$$

Let μ be a partition of n. The *content* of a dot d in μ is given by

$$\mathrm{ct}_{\mu}(d) = j - i$$

if d is in position (i, j).

Proposition 7.2.1. *If n is a positive integer and μ is a partition of n, then*

$$\chi_\mu(z^{(n)})/\chi_\mu(1) = \sum_{d\in\mu} \mathrm{ct}_\mu(d). \tag{7.3}$$

Proof. As noted in [38], this follows from **I**, §7, Ex. 7, and **I**, §1, Ex. 3, of [31]. □

Thus, by the comments of Sections 2.2 and 3.2, when $z^{(n)}$ is viewed as a linear transformation, the simple $\mathbb{C}[S_n]$-module U^μ is an eigenspace of $z^{(n)}$ with eigenvalue $\sum_{d\in\mu} \mathrm{ct}_\mu(d)$. Let $T^{(n)} = z^{(n)} - z^{(n-1)} = \sum_{i=1}^{n-1}(i,n)$.

Lemma 7.2.2. *$\{T^{(n)}\}$ is a $\mathbb{C}[S_{n-1}]$-separating set for every $\mathbb{C}[S_n]$-module M.*

Proof. By (7.1), every sub-isotypic space N of M is of type (U^η, U^μ) where η is some partition of $n-1$ and μ is a partition of n obtained from η by adding a dot e. Thus, by Proposition 7.2.1, the eigenvalue of $T^{(n)} = z^{(n)} - z^{(n-1)}$ restricted to N is

$$\sum_{d\in\mu} \mathrm{ct}_\mu(d) - \sum_{d\in\eta} \mathrm{ct}_\eta(d) = \mathrm{ct}_\mu(e).$$

Suppose now that N' is a sub-isotypic space of M of type $(U^\eta, U^{\mu'})$ where μ' is a partition of n obtained from η by adding a dot e'. The eigenvalue of $T^{(n)}$ restricted to N' is then $\mathrm{ct}_{\mu'}(e')$. Since the content of a dot depends only on the diagonal on which the dot lies, we have that $\mathrm{ct}_\mu(e) = \mathrm{ct}_{\mu'}(e')$ if and only if $\mu = \mu'$. Thus, $\{T^{(n)}\}$ is a $\mathbb{C}[S_{n-1}]$-separating set for M. □

Remark. The elements $T^{(2)}, \ldots, T^{(n)}$ are known as *Jucys-Murphy elements* of S_n (see, e.g., [16, 35, 38]). Analogues of these elements for other Weyl groups and Hecke algebras have been developed in [38]. In the subsequent chapters, we will make use of the fact that these analogues also form separating sets for their corresponding representations.

7.3 Upper Bounds

Now that we have a $\mathbb{C}[S_{n-1}]$-separating set for $\mathbb{C}[S_n]$-modules, we may turn our attention to computing an upper bound for $\iota(M^\lambda)$. We begin with a lemma.

Lemma 7.3.1. *Let n be a positive integer, λ be a partition of n, and $1 \leq k \leq n$. If $\zeta_\lambda^{(k)}$ is the number of isotypic subspaces of $M^\lambda_{\mathbb{C}[S_k]}$, then*

$$\iota\left(M^\lambda_{\mathbb{C}[S_k]}\right) = \iota\left(M^\lambda_{\mathbb{C}[S_{k-1}]}\right) + O\left(\zeta_\lambda^{(k-1)} k^2 \dim(M^\lambda)\right).$$

Proof. Each isotypic subspace of $M^\lambda_{\mathbb{C}[S_{k-1}]}$ contains less than k sub-isotypic spaces with respect to $\mathbb{C}[S_k]$. By Lemma 7.2.2, $T^{(k)} = \sum_{i=1}^{k-1}(i,k)$ is a $\mathbb{C}[S_{k-1}]$-separating set for $M^\lambda_{\mathbb{C}[S_k]}$. As a linear transformation on M^λ with respect to the usual basis of M^λ, $T^{(k)}$ is a real symmetric matrix, and ${T^{(k)}}^{\mathrm{op}} \leq (k-1)\dim(M^\lambda)$ since $T^{(k)}$ is the sum of $k-1$ permutation matrices. Thus, by Theorem 5.4.1,

$$\begin{aligned}\iota\left(M^\lambda_{\mathbb{C}[S_k]}\right) &= \iota\left(M^\lambda_{\mathbb{C}[S_{k-1}]}\right) + O\left(\sum_{j=1}^{\zeta_\lambda^{(k-1)}} \left(k(k-1)\dim(M^\lambda) + k^2\dim(M^\lambda)\right)\right) \\ &= \iota\left(M^\lambda_{\mathbb{C}[S_{k-1}]}\right) + O\left(\zeta_\lambda^{(k-1)} k^2 \dim(M^\lambda)\right).\end{aligned}$$

□

Lemma 7.3.2. *Let n be a positive integer and λ be a partition of n. For $1 \leq k \leq n$, let $\zeta_\lambda^{(k)}$ be the number of isotypic subspaces of $M^\lambda_{\mathbb{C}[S_k]}$. Then $\zeta_\lambda^{(k)} \leq \zeta_\lambda^{(l)}$ for $1 \leq k \leq l \leq n$.*

Proof. For $1 \leq k \leq n$, let P_k^λ be the set of partitions of k that correspond to the isotypic subspaces of $M^\lambda_{\mathbb{C}[S_k]}$. Thus $|P_k^\lambda| = \zeta_\lambda^{(k)}$. Adding $n-k$ dots to the first row of each partition in P_k^λ creates an injective map from P_k^λ into P_n^λ. Successively removing these dots then creates injective maps from P_k^λ into P_l^λ for $k \leq l \leq n-1$. Thus $\zeta_\lambda^{(k)} \leq \zeta_\lambda^{(l)}$ for $1 \leq k \leq l \leq n$. □

We may now give an upper bound for $\iota(M^\lambda)$.

Theorem 7.3.3. *Let n be a positive integer, let λ be a partition of n, and let M^λ be the corresponding $\mathbb{C}[S_n]$-permutation module. If M^λ has ζ_λ isotypic subspaces, then*

$$\iota\left(M^\lambda\right) = O\left(\zeta_\lambda n^3 \dim(M^\lambda)\right).$$

Proof. By Lemma 7.3.2, $1 = \zeta_\lambda^{(1)} \leq \zeta_\lambda^{(2)} \leq \cdots \leq \zeta_\lambda^{(n-1)} \leq \zeta_\lambda^{(n)} = \zeta_\lambda$. Thus, by Lemma 7.3.1 and induction, we have that

$$\iota\left(M^\lambda\right) = O\left(\sum_{k=2}^{n} \zeta_\lambda^{(k-1)} k^2 \dim(M^\lambda)\right).$$

It follows that $\iota\left(M^\lambda\right) = O\left(\zeta_\lambda n^3 \dim(M^\lambda)\right)$. □

Corollary 7.3.4. *If n is a positive integer, then*

$$\iota\left(\mathbb{C}[S_n]_{\mathbb{C}[S_n]}\right) = O\left(\frac{1}{4\sqrt{3n}} e^{\pi\sqrt{2n/3}} n^3 n!\right).$$

Proof. The regular representation of $\mathbb{C}[S_n]$ is isomorphic to M^λ when $\lambda = (1, \ldots, 1)$. The number $\zeta_{(1,\ldots,1)}$ of isotypic subspaces of $\mathbb{C}[S_n]$ is therefore equal to the number $p(n)$ of partions of n. Since

$$p(n) \sim \frac{1}{4\sqrt{3n}} e^{\pi\sqrt{2n/3}}$$

as $n \to \infty$ (see, e.g., [2]) and $\dim(M^{(1,\ldots,1)}) = n!$, the result follows immediately. □

Remark. An FFT and inverse for the symmetric group, both requiring $O(n^2 n!)$ operations, was constructed in [32]. Thus, the isotypic projections of a vector in $\mathbb{C}[S_n]$ may be computed using $O(p(n)n^2 n!)$ operations. This should be compared to

the bound of $O\left(p(n)n^3 n!\right)$ in Corollary 7.3.4.

Remark. Note that when $n \geq 2$ and $k \leq n/2$, we were able to find a bound for $\iota\left(M^{(n-k,k)}\right)$ in Section 6.2 by viewing the elements of $X^{(n-k,k)}$ as the vertices of a distance transitive graph. Moreover, the upper bound of

$$O\left(k^2\binom{n}{k}\right)$$

in Section 6.2 is much better than the upper bound of

$$O\left(kn^3\binom{n}{k}\right)$$

given by Theorem 7.3.3. It is also better than the upper bound of

$$O\left(k^2(n-k)\binom{n}{k}\right)$$

operations suggested by a modified FFT in [32].

Chapter 8

The Hyperoctahedral Group

In this chapter, we show how the LIPM gives rise to an efficient isotypic decomposition method for the regular representation of the hyperoctahedral group H_n. The hyperoctahedral group is a wreath product of the symmetric group and the cyclic group $\mathbb{Z}/2\mathbb{Z}$. Being able to efficiently decompose data on these and other wreath products is of interest because such groups are the symmetry groups of certain nested designs, and because of their recent use in signal and image processing (see, e.g., [24, 39]). See also [34] for an FFT for the hyperoctahedral group.

8.1 Representation Theory

As with the symmetric group, we begin with the necessary representation theory. See, for example, [26] for proofs and further details.

Let n be a positive integer. The *hyperoctahedral group* is the group H_n of *signed permutations* σ of $\{-n, \ldots, -1, 1, \ldots, n\}$ where $\sigma(-k) = -\sigma(k)$ for all $1 \leq k \leq n$.

The elements

$$s_1 = (1, -1) \text{ and } s_i = (i-1, i)(-(i-1), -i), \text{ for } 2 \leq i \leq n,$$

generate H_n and satisfy the relations

$$\begin{aligned} s_i s_j &= s_j s_i, \text{ for } |i-j| > 1 \\ s_i s_{i+1} s_i &= s_{i+1} s_i s_{i+1}, \text{ for } 2 \leq i \leq n-1 \\ s_1 s_2 s_1 s_2 &= s_2 s_1 s_2 s_1 \\ s_i^2 &= 1, \text{ for } 1 \leq i \leq n. \end{aligned}$$

In particular, H_n is a Coxeter group of order $2^n n!$ (see Section 3.3). For $1 \leq k \leq n$, we identify the subgroup of H_n that is generated by $s_1, \ldots, s_k$ with the group H_k.

A *double partition* of n is an ordered pair of partitions (λ, μ) such that if λ is a partition of n' and μ is a partition n'', then $n' + n'' = n$. Either λ or μ may be the empty partition. A double partition is identified with the corresponding pair of Ferrers diagrams.

The irreducible representations of H_n are indexed by the double partitions of n. Thus, for every double partition (λ, μ) of n, there is a simple $\mathbb{C}[H_n]$-module $U^{(\lambda,\mu)}$ with character $\chi_{(\lambda,\mu)}$. These modules form a complete collection (up to isomorphism) of simple $\mathbb{C}[H_n]$-modules. The regular representation of $\mathbb{C}[H_n]$ is therefore a direct sum

$$\mathbb{C}[H_n] = \bigoplus_{(\lambda,\mu)} M^{(\lambda,\mu)}$$

of isotypic subspaces where the sum is over the double partitions of n and $M^{(\lambda,\mu)}$ contains those simple submodules of $\mathbb{C}[H_n]$ that are isomorphic to $U^{(\lambda,\mu)}$.

As a $\mathbb{C}[H_{n-1}]$-module, the simple $\mathbb{C}[H_n]$-module $U^{(\lambda,\mu)}$ is a direct sum of simple modules

$$U^{(\lambda,\mu)} = \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} U^{(\eta,\nu)} \tag{8.1}$$

where the sum is over those double partitions (η,ν) of $n-1$ that are obtained from (λ,μ) by removing one dot from either Ferrers diagram. Thus, as a $\mathbb{C}[H_{n-1}]$-module, the isotypic subspace $M^{(\lambda,\mu)}$ of $\mathbb{C}[H_n]$ is a direct sum

$$M^{(\lambda,\mu)} = \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} M^{(\eta,\nu)(\lambda,\mu)}$$

of sub-isotypic spaces, where $M^{(\eta,\nu)(\lambda,\mu)}$ is the isotypic subspace of $M^{(\lambda,\mu)}$ corresponding to the simple $\mathbb{C}[H_{n-1}]$-module $U^{(\eta,\nu)}$. We therefore have the decomposition

$$\mathbb{C}[H_n] = \bigoplus_{(\lambda,\mu)} \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} M^{(\eta,\nu)(\lambda,\mu)}$$

of $\mathbb{C}[H_n]$ into sub-isotypic spaces with respect to $\mathbb{C}[H_{n-1}]$.

8.2 Separating Sets

Let $z_s^{(n)} \in \mathbb{C}[H_n]$ be the class sum

$$z_s^{(n)} = \sum_{i=1}^{n} (i,-i)$$

and let $z_l^{(n)} \in \mathbb{C}[H_n]$ be the class sum

$$z_l^{(n)} = \sum_{1\le i<j\le n} \left((i,j)(-i,-j) + (i,-j)(-i,j)\right).$$

Let (λ, μ) be a double partition of n. The *content* of a dot d in (λ, μ) is given by

$$\mathrm{ct}_{(\lambda,\mu)}(d) = j - i$$

if d is in position (i, j) in either λ or μ. The *sign* of d in (λ, μ) is given by

$$\mathrm{sgn}_{(\lambda,\mu)}(d) = \begin{cases} 1 & \text{if } d \in \lambda \\ -1 & \text{if } d \in \mu. \end{cases}$$

Proposition 8.2.1. (Ram) *If n is a positive integer and (λ, μ) is a double partition of n, then*

$$\chi_{(\lambda,\mu)}(z_s^{(n)})/\chi_{(\lambda,\mu)}(1) = \sum_{d \in (\lambda,\mu)} \mathrm{sgn}_{(\lambda,\mu)}(d)$$

and

$$\chi_{(\lambda,\mu)}(z_l^{(n)})/\chi_{(\lambda,\mu)}(1) = 2 \sum_{d \in (\lambda,\mu)} \mathrm{ct}_{(\lambda,\mu)}(d).$$

Proof. This is Proposition 4.8 in [38]. □

Thus, by the comments of Sections 2.2 and 3.2, when $z_s^{(n)}$ and $z_l^{(n)}$ are viewed as linear transformations, the simple $\mathbb{C}[H_n]$-module $U^{(\lambda,\mu)}$ is an eigenspace of $z_s^{(n)}$ with eigenvalue $\sum_{d \in (\lambda,\mu)} \mathrm{sgn}_{(\lambda,\mu)}(d)$ and an eigenspace of $z_l^{(n)}$ with eigenvalue $2 \sum_{d \in (\lambda,\mu)} \mathrm{ct}_{(\lambda,\mu)}(d)$. Let

$$T_s^{(n)} = z_s^{(n)} - z_s^{(n-1)} = (n, -n)$$

and let

$$T_l^{(n)} = z_l^{(n)} - z_l^{(n-1)} = \sum_{i=1}^{n-1} \left((i, n)(-i, -n) + (i, -n)(-i, n) \right).$$

Remark. The elements $T_s^{(1)}, \ldots, T_s^{(n)}$ and $T_l^{(2)}, \ldots, T_l^{(n)}$ of $\mathbb{C}[H_n]$ are analogues of the Jucys-Murphy elements for the symmetric group discussed in Chapter 6. See [38]

for more details.

Lemma 8.2.2. *$\{T_s^{(n)}, T_l^{(n)}\}$ is a $\mathbb{C}[H_{n-1}]$-separating set for every $\mathbb{C}[H_n]$-module M.*

Proof. By (8.1), every sub-isotypic space N of M is of type $(U^{(\eta,\nu)}, U^{(\lambda,\mu)})$ where (η,ν) is some double partition of $n-1$ and (λ,μ) is a double partition of n obtained from (η,ν) by adding a dot e. Thus, by Proposition 8.2.1, the eigenvalue of $T_s^{(n)} = z_s^{(n)} - z_s^{(n-1)}$ restricted to N is

$$\sum_{d\in(\lambda,\mu)} \mathrm{sgn}_{(\lambda,\mu)}(d) - \sum_{d\in(\eta,\nu)} \mathrm{sgn}_{(\eta,\nu)}(d) = \mathrm{sgn}_{(\lambda,\mu)}(e)$$

and the eigenvalue of $T_l^{(n)} = z_l^{(n)} - z_l^{(n-1)}$ restricted to N is

$$2\sum_{d\in(\lambda,\mu)} \mathrm{ct}_{(\lambda,\mu)}(d) - 2\sum_{d\in(\eta,\nu)} \mathrm{ct}_{(\eta,\nu)}(d) = 2\mathrm{ct}_{(\lambda,\mu)}(e).$$

Suppose now that N' is a sub-isotypic space of M of type $(U^{(\eta,\nu)}, U^{(\lambda',\mu')})$. It follows that N and N' are in the same eigenspace of both $T_l^{(n)}$ and $T_s^{(n)}$ if and only if $(\lambda,\mu) = (\lambda',\mu')$. Thus, $\{T_s^{(n)}, T_l^{(n)}\}$ is a $\mathbb{C}[H_{n-1}]$-separating set for M. □

8.3 Upper Bounds

Now that we have a $\mathbb{C}[H_{n-1}]$-separating set for $\mathbb{C}[H_n]$-modules, we may compute an upper bound for $\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_n]}\right)$. We begin with a lemma. Let $d(k)$ be the number of double partitions of a positive integer k.

Lemma 8.3.1. *If n is a positive integer and $1 \le k \le n$, then*

$$\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_k]}\right) = \iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_{k-1}]}\right) + O\left(d(k-1)k^2 2^n n!\right).$$

Proof. Each isotypic subspace of $\mathbb{C}[H_n]_{\mathbb{C}[H_{k-1}]}$ contains less than k sub-isotypic spaces with respect to $\mathbb{C}[H_n]_{\mathbb{C}[H_k]}$. We also know that the number of isotypic subspaces of $\mathbb{C}[H_n]_{\mathbb{C}[H_{k-1}]}$ is $d(k-1)$.

By Lemma 8.2.2, $\{T_s^{(k)}, T_l^{(k)}\}$ is a $\mathbb{C}[H_{k-1}]$-separating set for $\mathbb{C}[H_n]_{\mathbb{C}[H_k]}$. As linear transformations on $\mathbb{C}[H_n]$ with respect to the usual basis of $\mathbb{C}[H_n]$, both $T_s^{(k)}$ and $T_l^{(k)}$ are real symmetric matrices. Since the dimension of $\mathbb{C}[H_n]$ is $2^n n!$, ${T_s^{(k)}}^{\text{op}} \leq 2^n n!$ and ${T_l^{(k)}}^{\text{op}} \leq 2(k-1)2^n n!$. Thus, by Theorem 5.4.1,

$$\begin{aligned}\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_k]}\right) &= \iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_{k-1}]}\right) + O\left(\sum_{j=1}^{d(k-1)} (k^2 2^n n!)\right) \\ &= \iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_{k-1}]}\right) + O\left(d(k-1)k^2 2^n n!\right).\end{aligned}$$

□

We may now give an upper bound for $\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_n]}\right)$.

Theorem 8.3.2. *If n is a positive integer and $d(n)$ is the number of double partitions of n, then*

$$\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_n]}\right) = O\left(d(n)n^3 2^n n!\right).$$

Proof. Since $1 = d(0) \leq d(1) \leq \cdots \leq d(n)$, by induction and Lemma 8.3.1,

$$\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_n]}\right) = O\left(\sum_{k=1}^{n} d(k-1)k^2 2^n n!\right).$$

It follows that $\iota\left(\mathbb{C}[H_n]_{\mathbb{C}[H_n]}\right) = O\left(d(n)n^3 2^n n!\right)$. □

Remark. The bound given in Theorem 8.3.2 is also the bound suggested by the FFT for H_n that is found in [34].

Chapter 9

The Finite General Linear Group

The action of the finite general linear group on flags of subspaces is analogous to the action of the symmetric group on tabloids. These flags may be viewed as generalizations of point-block pairs in balanced incomplete block designs whose blocks are the k-dimensional subspaces of an n-dimensional vector space V, and whose points are the 1-dimensional subspaces of V (see, e.g., [5]). Such designs provide combinatorial structures used to measure the effects of several combinations (blocks) of certain factors (points) on the outcome of an experiment. Classically, data on such designs is analyzed using analysis of variance (ANOVA), however this analysis may sometimes be subsumed within a spectral analysis approach to the data (see [40]).

9.1 Representation Theory

As usual, we begin with the relevant representation theory. See [26, 28] for proofs and related material.

Let $\mathbb{F}_q$ be the field with q elements, let n be a positive integer, and let V be an n-dimensional vector space over $\mathbb{F}_q$. By definition, the automorphism group of V is

the finite general linear group $G_n = GL(n, q)$. Let $\{b_1, \ldots, b_n\}$ be a basis of V. For convenience we identify G_n with the group of invertible $n \times n$ matrices over $\mathbb{F}_q$ with respect to this basis. Let $V_{n-1}, \ldots, V_1$ be a collection of subspaces of V such that $\dim(V_i) = i$ and

$$V_{n-1} \supset \cdots \supset V_1.$$

We say that $V_{n-1}, \ldots, V_1$ forms a *maximal flag* of subspaces of V. We denote the set of maximal flags of V by X_n.

The action of G_n on V induces a transitive action of G_n on X_n. Thus, if $x \in X_n$ is the flag

$$\langle b_2, \ldots, b_n \rangle \supset \langle b_3, \ldots, b_n \rangle \supset \cdots \langle b_{n-1}, b_n \rangle \supset \langle b_n \rangle$$

and B_n is the stabilizer of x in G_n, then we may identify X_n with the set of right cosets G_n/B_n, and the resulting permutation representation $\mathbb{C}[X_n]$ with $\mathbb{C}[G_n/B_n]$.

Note that B_n corresponds to the subgroup of upper triangular matrices in G_n. Furthermore, if N_n is the subgroup in G_n of monomial matrices (i.e., matrices with only one nonzero entry in each row and column), then B_n and N_n form a BN-pair for G_n with Weyl group S_n (see, e.g., [8]). The associated Schur basis $\{T_w \mid w \in S_n\}$ of $\mathrm{End}_{\mathbb{C}[G_n]}\mathbb{C}[G_n/B_n]$ is such that

$$T_s T_w = \begin{cases} T_{sw} & \text{if } l(sw) > l(w) \\ qT_{sw} + (q-1)T_w & \text{if } l(sw) < l(w) \end{cases}$$

where s is a generator of S_n as defined in Section 7.1 and $w \in S_n$ (see, e.g., [26]).

Let $E_n = \mathrm{End}_{\mathbb{C}[G_n]}\mathbb{C}[G_n/B_n]$. For $1 \leq k \leq n$, let E_k be the subalgebra of E_n that is generated by $\{T_w \mid w \in S_k\}$. As an E_n-module, the isotypic decomposition of $\mathbb{C}[G_n/B_n]$ is analogous to the isotypic decomposition of the regular representation

of $\mathbb{C}[S_n]$. The simple E_n-modules are indexed by the partitions of n. Thus, for each partition μ of n, there is a simple E_n-module U_μ. The isotypic decomposition of $\mathbb{C}[G_n/B_n]$ as an E_n-module is

$$\mathbb{C}[G_n/B_n] = \bigoplus_{\mu} M_\mu \tag{9.1}$$

where the sum is over all partitions of n and M_μ is the isotypic subspace of $\mathbb{C}[G_n/B_n]$ corresponding to the simple E_n-module U_μ. Recall that (9.1) is also the isotypic decomposition of $\mathbb{C}[G_n/B_n]$ as a $\mathbb{C}[G_n]$-module (see Section 2.2).

As an E_{n-1} module, the simple E_n-module U_μ is a direct sum of simple modules

$$U_\mu = \bigoplus_{\eta \in \mu^-} U_\eta$$

where the sum is over those partitions η of $n-1$ that are obtained from μ by removing one dot. Thus, as an E_{n-1}-module, the isotypic subspace M_μ is a direct sum

$$M_\mu = \bigoplus_{\eta \in \mu^-} M_{\eta\mu}$$

of sub-isotypic spaces, where $M_{\eta\mu}$ is the isotypic subspace of M_μ corresponding to the simple E_{n-1}-module U_η. We therefore have the decomposition

$$\mathbb{C}[G_n/B_n] = \bigoplus_{\mu} \bigoplus_{\eta \in \mu^-} M_{\eta\mu}$$

of $\mathbb{C}[G_n/B_n]$ into sub-isotypic spaces with respect to E_{n-1}.

9.2 Separating Sets

For $2 \leq k \leq n$, define

$$J_k = q^{-(k-1)} T_{s_k} \cdots T_{s_3} T_{s_2} T_{s_2} T_{s_3} \cdots T_{s_k}. \tag{9.2}$$

Using the relation $T_{s_i}^2 = q + (q-1)T_{s_i}$, it may be shown that

$$\frac{J_k - 1}{q-1} = \sum_{i=1}^{k-1} q^{-(k-i)} T_{(i,k)}$$

which gives a q-analogue of the Jucys-Murphy elements associated to the symmetric group (see [38]).

Lemma 9.2.1. *$\{J_n\}$ is an E_{n-1}-separating set for every left E_n-module M.*

Proof. Let N be a sub-isotypic space of M of type (U_η, U_μ) where η is a partition of $n-1$ and μ is a partition of n obtained from η by adding a dot e. By Section 3 of [38], the eigenvalue of the restriction of J_n to N is $q^{\mathrm{ct}_\mu(e)}$. Arguing as we did in Lemma 7.2.2, it follows that $\{J_n\}$ is an E_{n-1}-separating set for M. □

9.3 Upper Bounds

Now that we have separating sets for $\mathbb{C}[G_n/B_n]$, we may give an upper bound for $\iota({}_{E_n}\mathbb{C}[G_n/B_n])$.

Theorem 9.3.1. *If n is a positive integer, then*

$$\iota\left({}_{E_n}\mathbb{C}[G_n/B_n]\right) = O\left(\frac{1}{4\sqrt{3}n} e^{\pi\sqrt{2n/3}} n^3 q \dim \mathbb{C}[G_n/B_n]\right).$$

Proof. The proof is essentially that of Corollary 7.3.4 with the appropriate modifications. We will therefore be brief.

By Lemma 9.2.1, J_k is an E_{k-1}-separating set for $\mathbb{C}[G_n/B_n]$ when viewed as a left E_k-module. As a linear transformation on $\mathbb{C}[G_n/B_n]$ with respect to the delta basis, T_{s_i} is a real symmetric matrix for $2 \leq i \leq n$. Thus, by (9.2), J_k is a real symmetric matrix with respect to the delta basis. By Corollary 3.3.3,

$$T_{s_i}^{\mathrm{op}} \leq q \dim \mathbb{C}[G_n/B_n]$$

for $2 \leq i \leq n$. Thus

$$\begin{aligned} J_k^{\mathrm{op}} &\leq \left(q^{-(k-1)} T_{s_k} \cdots T_{s_3} T_{s_2} T_{s_2} T_{s_3} \cdots T_{s_k}\right)^{\mathrm{op}} \\ &\leq 2(k-1) q \dim \mathbb{C}[G_n/B_n]. \end{aligned}$$

Let $p(k)$ be the number of partitions of a positive integer k. By Theorem 5.4.1, it follows that

$$\begin{aligned} \iota\left({}_{E_k}\mathbb{C}[G_n/B_n]\right) &= \iota\left({}_{E_{k-1}}\mathbb{C}[G_n/B_n]\right) + O\left(\sum_{j=1}^{p(k-1)} \left(k J_k^{\mathrm{op}} + k^2 \dim \mathbb{C}[G_n/B_n]\right)\right) \\ &= \iota\left({}_{E_{k-1}}\mathbb{C}[G_n/B_n]\right) + O\left(p(k-1) k^2 q \dim \mathbb{C}[G_n/B_n]\right) \end{aligned}$$

for $2 \leq k \leq n$. The theorem now follows by induction and the fact that

$$p(n) \sim \frac{1}{4\sqrt{3n}} e^{\pi\sqrt{2n/3}}$$

as $n \to \infty$ (see, e.g., [2]). □

Chapter 10

The Finite Symplectic Group

We now consider the question of efficiently computing isotypic projections for the finite symplectic group. Just as the finite general linear group is related to the symmetric group, the finite symplectic group is related to the hyperoctahedral group. As with the finite general linear group, we take advantage of elements in the corresponding endomorphism algebra that are analogues of Jucys-Murphy elements.

10.1 Representation Theory

Once again, we begin by reviewing the underlying representation theory. See [26, 28] for details and proofs.

Let $\mathbb{F}_q$ be the field with q elements and let n be a positive integer. Let V be a $2n$-dimensional vector space over $\mathbb{F}_q$ with a non-degenerate, alternating, bilinear form $\langle \cdot, \cdot \rangle$. We may define such a form on V by defining it on a basis $\{b_1, \ldots, b_n, b'_1, \ldots, b'_n\}$

of V as

$$(b_i, b'_j) = \delta_{ij} = -(b'_j, b_i)$$
$$(b_i, b_j) = 0 = (b'_i, b'_j)$$

for $1 \leq i, j \leq n$. The *symplectic group* $Sp_n = Sp(n, q)$ is then the group of automorphisms of V that preserve $\langle \cdot, \cdot \rangle$. For convenience, we view the elements of Sp_n as matrices with respect to the basis $\{b_1, \ldots, b_n, b'_1, \ldots, b'_n\}$ and we call these matrices *symplectic matrices*. The symplectic group is, however, independent (up to isomorphism) of the specific choice of $\langle \cdot, \cdot \rangle$ (see, e.g., [3]).

Let B_n be the subgroup of Sp_n of upper triangular symplectic matrices and let N_n be the subgroup of monomial symplectic matrices. The subgroups B_n and N_n form a BN-pair for Sp_n with Weyl group the hyperoctahedral group H_n (see, e.g., [8]). The corresponding Schur basis $\{T_w \mid w \in H_n\}$ of $\mathrm{End}_{\mathbb{C}[Sp_n]}\mathbb{C}[Sp_n/B_n]$ is such that

$$T_s T_w = \begin{cases} T_{sw} & \text{if } l(sw) > l(w) \\ qT_{sw} + (q-1)T_w & \text{if } l(sw) < l(w) \end{cases}$$

where s is a generator of H_n as defined in Section 8.1 and $w \in H_n$ (see, e.g., [26]).

Let $E_n = \mathrm{End}_{\mathbb{C}[Sp_n]}\mathbb{C}[Sp_n/B_n]$. For $1 \leq k \leq n$, let E_k be the subalgebra of E_n that is generated by $\{T_w \mid w \in H_k\}$. Denote by E_0 the subalgebra of E_n that is generated by T_1, i.e., the subalgebra of scalar multiples of the identity.

As an E_n-module, the isotypic decomposition of $\mathbb{C}[Sp_n/B_n]$ is analogous to the isotypic decomposition of the regular representation of $\mathbb{C}[H_n]$. The simple E_n-modules are indexed by the double partitions of n. Thus, for each double partition (λ, μ) of n, there is a simple E_n-module $U_{(\lambda,\mu)}$. The isotypic decomposition of $\mathbb{C}[Sp_n/B_n]$ as

an E_n-module is

$$\mathbb{C}[Sp_n/B_n] = \bigoplus_{(\lambda,\mu)} M_{(\lambda,\mu)} \tag{10.1}$$

where the sum is over all double partitions of n and $M_{(\lambda,\mu)}$ is the isotypic subspace of $\mathbb{C}[Sp_n/B_n]$ corresponding to the simple E_n-module $U_{(\lambda,\mu)}$. This decomposition of $\mathbb{C}[Sp_n/B_n]$ is also the isotypic decomposition of $\mathbb{C}[Sp_n/B_n]$ as a $\mathbb{C}[Sp_n]$-module.

As an E_{n-1} module, the simple E_n-module $U_{(\lambda,\mu)}$ is a direct sum of simple modules

$$U_{(\lambda,\mu)} = \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} U_{(\eta,\nu)}$$

where the sum is over those double partitions (η, ν) of $n-1$ that are obtained from (λ, μ) by removing one dot from either λ or μ. Thus, as a E_{n-1}-module, the isotypic subspace $M_{(\lambda,\mu)}$ of $\mathbb{C}[Sp_n/B_n]$ is a direct sum

$$M_{(\lambda,\mu)} = \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} M_{(\eta,\nu)(\lambda,\mu)}$$

of sub-isotypic spaces, where $M_{(\eta,\nu)(\lambda,\mu)}$ is the isotypic subspace of $M_{(\lambda,\mu)}$ corresponding to the simple E_{n-1}-module $U_{(\eta,\nu)}$. We therefore have the decomposition

$$\mathbb{C}[Sp_n/B_n] = \bigoplus_{(\lambda,\mu)} \bigoplus_{(\eta,\nu)\in(\lambda,\mu)^-} M_{(\eta,\nu)(\lambda,\mu)}$$

of $\mathbb{C}[Sp_n/B_n]$ into sub-isotypic spaces with respect to E_{n-1}.

10.2 Separating Sets

For $1 \leq k \leq n$, define

$$L_k = q^{-(k-\frac{1}{2})} T_{s_k} \cdots T_{s_2} T_{s_1} T_{s_2} \cdots T_{s_k}. \tag{10.2}$$

Lemma 10.2.1. *$\{L_n\}$ is an E_{n-1}-separating set for every left E_n-module M.*

Proof. Let N be a sub-isotypic space of M of type $(U_{(\eta,\nu)}, U_{(\lambda,\mu)})$ where (η, ν) is a double partition of $n-1$ and (λ, μ) is a double partition of n obtained from (η, ν) by adding a dot e. By Section 4 of [38], the eigenvalue of the restriction of L_n to N is $\mathrm{sgn}_{(\lambda,\mu)}(e) q^{\frac{1}{2}\mathrm{sgn}_{(\lambda,\mu)}(e)} q^{\mathrm{ct}_{(\lambda,\mu)}(e)}$. Arguing as we did in Lemma 8.2.2, it follows that $\{J_n\}$ is an E_{n-1}-separating set for M. □

10.3 Upper Bounds

Now that we have separating sets for $\mathbb{C}[Sp_n/B_n]$, we may give an upper bound for $\iota({}_{E_n}\mathbb{C}[Sp_n/B_n])$.

Theorem 10.3.1. *If n is a positive integer and $d(n)$ is the number of double partitions of n, then*

$$\iota\left({}_{E_n}\mathbb{C}[Sp_n/B_n]\right) = O\left(d(n) n^3 q \dim \mathbb{C}[Sp_n/B_n]\right).$$

Proof. We will be brief since the proof is essentially that of Theorem 8.3.2 with only a few changes.

By Lemma 10.2.1, L_k is an E_{k-1}-separating set for $\mathbb{C}[Sp_n/B_n]$ when viewed as a left E_k-module . As a linear transformation on $\mathbb{C}[Sp_n/B_n]$ with respect to the delta basis, T_{s_i} is a real symmetric matrix for $1 \leq i \leq n$. Thus, by (10.2), L_k is a real

symmetric matrix with respect to the delta basis of $\mathbb{C}[Sp_n/B_n]$. By Corollary 3.3.3,

$$T_{s_i}^{\mathrm{op}} \leq q \dim \mathbb{C}[Sp_n/B_n]$$

for $1 \leq i \leq n$. Thus

$$\begin{aligned} L_k^{\mathrm{op}} &\leq \left(q^{-(k-\frac{1}{2})} T_{s_k} \cdots T_{s_2} T_{s_1} T_{s_2} \cdots T_{s_k}\right)^{\mathrm{op}} \\ &\leq (2k-1) q \dim \mathbb{C}[Sp_n/B_n]. \end{aligned}$$

By Theorem 5.4.1, it follows that

$$\begin{aligned} \iota\left({}_{E_k}\mathbb{C}[Sp_n/B_n]\right) &= \iota\left({}_{E_{k-1}}\mathbb{C}[Sp_n/B_n]\right) + O\left(\sum_{j=1}^{d(k-1)} (kL_k^{\mathrm{op}} + k^2 \dim \mathbb{C}[Sp_n/B_n])\right) \\ &= \iota\left({}_{E_{k-1}}\mathbb{C}[Sp_n/B_n]\right) + O\left(d(k-1)k^2 q \dim \mathbb{C}[Sp_n/B_n]\right) \end{aligned}$$

for $1 \leq k \leq n$. The theorem now follows by induction. □

Chapter 11

Future Directions

There are many refinements that could be made to the approach to decomposing representations that we have presented. For example, the number of sub-isotypic spaces has often been overestimated. Counts that are more accurate may therefore yield better bounds. We have also done all of our calculations with respect to one basis, namely the delta basis of the underlying permutation representation. Intermediate changes of bases (e.g., to reflect the orbits in permutation representations when restricting the action of a group to a subgroup) may also lead to improved bounds. At the very least, preliminary investigations suggest that these changes of bases lead to worthwhile simplifications in implementation.

In Chapters 7-10, we have taken advantage of the fact that Jucys-Murphy elements and their analogues form separating sets that are particularly easy to apply. It would be interesting to find and use analogous elements for groups that are not associated to Weyl groups. This idea has been touched upon, for example, in [16].

Lastly, the Jucys-Murphy elements for the symmetric group may be seen as linear combinations of certain idempotents which correspond to Young's seminormal bases for simple $\mathbb{C}[S_n]$-modules (see [36]). This suggests an eigenspace approach to comput-

ing the associated discrete Fourier transform by allowing the Jucys-Murphy elements to act on both the right and left of $\mathbb{C}[S_n]$ (see, e.g., Chapter 1 of [22]). The techniques presented in this thesis may, therefore, lead to an interesting eigenspace-based fast Fourier transform for the symmetric group and other Weyl groups.

Bibliography

[1] W. Adkins and S. Weintraub, *Algebra*, Springer-Verlag, New York, 1992, An approach via module theory.

[2] G. Andrews, *The theory of partitions*, Addison-Wesley Publishing Co., Reading, Mass.-London-Amsterdam, 1976, Encyclopedia of Mathematics and its Applications, Vol. 2.

[3] E. Artin, *Geometric algebra*, Interscience Publishers, Inc., New York-London, 1957.

[4] T. Beth, *Verfahren der schnellen Fourier-Transformation*, B. G. Teubner, Stuttgart, 1984.

[5] T. Beth, D. Jungnickel, and H. Lenz, *Design theory*, Cambridge University Press, Cambridge, 1986.

[6] E. Bolker, *The finite Radon transform*, Integral geometry (Brunswick, Maine, 1984), Amer. Math. Soc., Providence, RI, 1987, pp. 27–50.

[7] A. Brouwer, A. Cohen, and A. Neumaier, *Distance-regular graphs*, Springer-Verlag, Berlin, 1989.

[8] K. Brown, *Buildings*, Springer-Verlag, New York, 1989.

[9] J. Chen, *Group representation theory for physicists*, World Scientific Publishing Co. Inc., Teaneck, NJ, 1989.

[10] M. Clausen and U. Baum, *Fast Fourier transforms*, Bibliographisches Institut, Mannheim, 1993.

[11] W. Cochran et al., *What is the fast Fourier transform?*, IEEE Transactions on Audio and Electroacoustics **AU-15** (1967), no. 2, 45–55.

[12] J. Cooley and J. Tukey, *An algorithm for the machine calculation of complex Fourier series*, Math. Comp. **19** (1965), 297–301.

[13] J. Cullum and R. Willoughby, *Lánczos algorithms for large symmetric eigenvalue computations. Vol. I*, Birkhäuser Boston Inc., Boston, MA, 1985, Theory.

[14] P. Diaconis, *Group representations in probability and statistics*, Institute of Mathematical Statistics, Hayward, CA, 1988.

[15] ———, *A generalization of spectral analysis with application to ranked data*, Ann. Statist. **17** (1989), no. 3, 949–979.

[16] P. Diaconis and C. Greene, *Applications of Murphy's elements*, Tech. report, Department of Statistics, Stanford University, 1989.

[17] P. Diaconis and D. Rockmore, *Efficient computation of the Fourier transform on finite groups*, J. Amer. Math. Soc. **3** (1990), no. 2, 297–332.

[18] ———, *Efficient computation of isotypic projections for the symmetric group*, Groups and computation (New Brunswick, NJ, 1991), Amer. Math. Soc., Providence, RI, 1993, pp. 87–104.

[19] R. Dipper and G. James, *Blocks and idempotents of Hecke algebras of general linear groups*, Proc. London Math. Soc. (3) **54** (1987), no. 1, 57–82.

[20] R. Dipper, G. James, and E. Murphy, *Hecke algebras of type B_n at roots of unity*, Proc. London Math. Soc. (3) **70** (1995), no. 3, 505–528.

[21] J. Driscoll, D. Healy, and D. Rockmore, *Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs*, SIAM J. Comput. **26** (1997), no. 4, 1066–1099.

[22] Y. Drozd and V. Kirichenko, *Finite-dimensional algebras*, Springer-Verlag, Berlin, 1994.

[23] D. Dummit and R. Foote, *Abstract algebra, second edition*, John Wiley & Sons, Inc., New York, 1999.

[24] R. Foote, D. Healy, G. Mirchandani, T. Olson, and D. Rockmore, *A wreath product group approach to signal and image processing. I. Multiresolution analysis*, IEEE Trans. Signal Process. **48** (2000), no. 1, 102–132.

[25] ———, *A wreath product group approach to signal and image processing. II. Convolution, correlation, and applications*, IEEE Trans. Signal Process. **48** (2000), no. 3, 749–767.

[26] M. Geck and G. Pfeiffer, *Characters of finite Coxeter groups and Iwahori-Hecke algebras*, The Clarendon Press Oxford University Press, New York, 2000.

[27] W. Gentleman and G. Sande, *Fast Fourier transform for fun and profit*, Proc. AFIPS, Joint Computer Conference **29** (1966), 563–578.

[28] P. Hoefsmit, *Representations of Hecke algebras of finite groups with BN-pairs of classical type*, Ph.D. thesis, The University of British Columbia, 1974.

[29] G. James, *The representation theory of the symmetric groups*, Springer, Berlin, 1978.

[30] ———, *Representations of general linear groups*, Cambridge University Press, Cambridge, 1984.

[31] I. Macdonald, *Symmetric functions and Hall polynomials*, The Clarendon Press Oxford University Press, New York, 1979, Oxford Mathematical Monographs.

[32] D. Maslen, *The efficient computation of Fourier transforms on the symmetric group*, Math. Comp. **67** (1998), no. 223, 1121–1147.

[33] D. Maslen and D. Rockmore, *Generalized FFTs—a survey of some recent results*, Groups and computation, II (New Brunswick, NJ, 1995), Amer. Math. Soc., Providence, RI, 1997, pp. 183–237.

[34] ———, *Separation of variables and the computation of Fourier transforms on finite groups. I*, J. Amer. Math. Soc. **10** (1997), no. 1, 169–214.

[35] G. Murphy, *A new construction of Young's seminormal representation of the symmetric groups*, J. Algebra **69** (1981), no. 2, 287–297.

[36] ———, *The idempotents of the symmetric group and Nakayama's conjecture*, J. Algebra **81** (1983), no. 1, 258–265.

[37] B. Parlett, *The symmetric eigenvalue problem*, Prentice-Hall Inc., Englewood Cliffs, N.J., 1980.

[38] A. Ram, *Seminormal representations of Weyl groups and Iwahori-Hecke algebras*, Proc. London Math. Soc. (3) **75** (1997), no. 1, 99–133.

[39] D. Rockmore, *Fast Fourier transforms for wreath products*, Appl. Comput. Harmon. Anal. **2** (1995), no. 3, 279–292.

[40] ———, *Some applications of generalized FFTs*, Groups and computation, II (New Brunswick, NJ, 1995), Amer. Math. Soc., Providence, RI, 1997, pp. 329–369.

[41] J.-P. Serre, *Linear representations of finite groups*, Springer-Verlag, New York, 1977.

[42] R. Stanley, *Enumerative combinatorics. Vol. 1*, Cambridge University Press, Cambridge, 1997.

[43] D. Stanton, *Orthogonal polynomials and Chevalley groups*, Special functions: group theoretical aspects and applications, Reidel, Dordrecht, 1984, pp. 87–128.

[44] L. Trefethen and D. Bau III, *Numerical linear algebra*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997.

[45] J. Wilkinson, *The algebraic eigenvalue problem*, Clarendon Press, Oxford, 1965.

Index