

2-24-2004

Computing Isotypic Projections with the Lanczos Iteration

David K. Maslen
Susquehanna Partners

Michael E. Orrison Jr.
Harvey Mudd College

Daniel N. Rockmore
Dartmouth College

Recommended Citation

Maslen, David K., Orrison, Michael E., and Rockmore, Daniel N. "Computing isotypic projections with the Lanczos iteration." *SIAM Journal on Matrix Analysis and Applications* 25.3 (2004): 784-803.

This Article is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

COMPUTING ISOTYPIC PROJECTIONS WITH THE LANCZOS ITERATION*

DAVID K. MASLEN[†], MICHAEL E. ORRISON[‡], AND DANIEL N. ROCKMORE[§]

Abstract. When the isotypic subspaces of a representation are viewed as the eigenspaces of a symmetric linear transformation, isotypic projections may be achieved as eigenspace projections and computed using the Lanczos iteration. In this paper, we show how this approach gives rise to an efficient isotypic projection method for permutation representations of distance transitive graphs and the symmetric group.

Key words. isotypic projection, Lanczos iteration, separating set

AMS subject classifications. 20C40, 65F10

DOI. 10.1137/S0895479801399778

1. Introduction. Let G be a finite group acting on a finite set X . Let $L(X)$ be the vector space of complex valued functions on X . The action of G on X gives rise to a *permutation representation* ρ of G defined on $L(X)$ by

$$(\rho(g)(f))(x) = f(g^{-1}x)$$

for all $g \in G$, $f \in L(X)$, and $x \in X$. Because $L(X)$ is a representation of G , there is a basis independent decomposition

$$L(X) = V_1 \oplus \cdots \oplus V_n$$

of $L(X)$ into G -invariant subspaces known as *isotypic subspaces*. The problem addressed in this paper is the following: Given an arbitrary $f \in L(X)$, how may we efficiently compute the projection of f onto each isotypic subspace of $L(X)$?

The problem of computing projections onto isotypic subspaces arises in *spectral analysis* which is a nonmodel-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 Diaconis [5, 6], the subject extends the classical spectral analysis of time series and requires the computation of projections of f onto subsets of G -invariant subspaces of $L(X)$.

As an example, let X be the set $\{x_0, \dots, x_{n-1}\}$ and let G be the cyclic group $\mathbb{Z}/n\mathbb{Z}$ acting on X by cyclicly permuting its elements. The elements of $L(X)$ may be viewed as *signals* on n points and the isotypic subspaces of $L(X)$ as corresponding to the different *frequencies* that make up these signals. The isotypic projections of $f \in L(X)$ may be computed with the aid of the usual discrete Fourier transform (DFT). The classical fast Fourier transform (FFT) may therefore be used to efficiently compute the projections of f onto the isotypic subspaces of $L(X)$ (see, e.g., [13]).

*Received by the editors December 14, 2001; accepted for publication (in revised form) by S. A. Vavasis June 16, 2003; published electronically February 24, 2004.

<http://www.siam.org/journals/simax/25-3/39977.html>

[†]Susquehanna Partners GP, 401 City Line Ave., Bala Cynwyd, PA 19004 (david@maslen.net).

[‡]Department of Mathematics, Harvey Mudd College, Claremont, CA 91711 (orrison@hmc.edu). The research of this author was supported by AFOSR.

[§]Department of Mathematics, Dartmouth College, Hanover, NH 03755 (daniel.rockmore@dartmouth.edu). The research of this author was supported by NSF-DMS, AFOSR, and the Santa Fe Institute.

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 $f \in L(X)$ 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 higher order 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 [6, 17]).

A naive approach (see, e.g., [19]) to computing the n isotypic projections of $f \in L(X)$ requires $O(n|G||X|)$ operations where we count a complex multiplication followed by a complex addition as one *operation*. Diaconis and Rockmore [7] 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 [8] show that if X is a distance transitive graph, then fast discrete polynomial transforms may be used to compute the n isotypic projections of $f \in L(X)$ 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 paper, we develop an approach to computing isotypic projections that relies on a method for computing projections onto the eigenspaces of a collection of simultaneously diagonalizable linear transformations. We call the collections of transformations that we use *separating sets* because they allow us to *separate* a representation into its isotypic components. The approach may be seen as a generalization of the Gentleman–Sande, or *decimation in frequency*, FFT in that we too will be iteratively computing projections of projections (see [10]). Such collections have also been used in [3], for example, where certain separating sets are known as *complete sets of commuting operators*.

As a simple example of how a separating set is used to compute isotypic projections, suppose that $L(X)$ has three isotypic subspaces V_1 , V_2 , and V_3 . Thus $L(X) = V_1 \oplus V_2 \oplus V_3$ and each $f \in L(X)$ may be written uniquely as $f = f_1 + f_2 + f_3$, where $f_i \in V_i$. Additionally, suppose that T and T' are diagonalizable linear transformations on $L(X)$ such that the eigenspaces of T are $V_1 \oplus V_2$ and V_3 , and the eigenspaces of T' are V_1 and $V_2 \oplus V_3$. As we shall see, $\{T, T'\}$ is a separating set for $L(X)$. We may therefore 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 $L(X)$. This process of decomposing $L(X) = V_1 \oplus V_2 \oplus V_3$ is illustrated in Figure 1.

The efficiency of this approach depends 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




FIG. 1. Decomposing $L(X) = V_1 \oplus V_2 \oplus V_3$ using T and T' .

be used to efficiently compute the eigenspace projections of a real symmetric matrix when, as in all of our examples, it has relatively few eigenspaces and when it may be applied efficiently to arbitrary vectors, either directly or through a given subroutine (see, e.g., [16]).

We proceed as follows. In section 2, we describe the isotypic decomposition of a representation and introduce the idea of a separating set of diagonalizable linear transformations. In section 3, we show how an eigenspace approach to computing isotypic projections for cyclic groups leads to the Gentleman–Sande FFT. In section 4, we review how the Lanczos iteration may be used to compute the projections of a vector onto the eigenspaces of a real symmetric matrix. We then use the results of section 2 to create an isotypic projection method. This method is then shown to be efficient for certain permutation representations of distance transitive graphs in section 5 and the symmetric group in section 6.

2. Isotypic subspaces. In this section, we describe the isotypic decomposition of a representation and we introduce the idea of a separating set of diagonalizable linear transformations. We then show how these separating sets may be used to compute isotypic projections. A good reference for representations of finite groups is [19].

2.1. Complex representations. Let G be a finite group, let V be a finite dimensional vector space over \mathbb{C} , and let $GL(V)$ be the group of automorphisms of V . A *representation* of G is a homomorphism $\rho : G \rightarrow GL(V)$. If the homomorphism ρ is understood, then we also say that V is a representation of G . The *character* of ρ is the function $\chi : G \rightarrow \mathbb{C}$, where $\chi(g)$ is the usual trace of $\rho(g)$. Note that the character of a representation of G is constant on the conjugacy classes of G .

A subspace W of V is *invariant* if $\rho(g)(w) \in W$ for all $g \in G$, $w \in W$. A representation is said to be *simple* if it contains no nontrivial invariant subspaces. If C_1, \dots, C_h are the distinct conjugacy classes of G , then there are h distinct (up to isomorphism) simple representations W_1, \dots, W_h of G . Let d_i be the dimension of W_i , let χ_i be the character of W_i , and let $\chi_i(C_j)$ be the value of χ_i on C_j .

2.2. The isotypic decomposition. Every representation of G is a direct sum of simple representations. Thus, V is a direct sum of simple representations, say, U_1, \dots, U_l . Denote by V_i the direct sum of those U_1, \dots, U_l that are isomorphic to W_i . Removing the trivial V_i (and renumbering if necessary) creates the *isotypic decomposition*

$$V = V_1 \oplus \cdots \oplus V_n,$$

where each V_i is then an *isotypic subspace* of V . Each $v \in V$ may therefore be written uniquely as

$$v = v_1 + \cdots + v_n,$$

where $v_i \in V_i$ is called the *isotypic projection* of v onto the isotypic subspace V_i . The isotypic decomposition of V is independent of the choice of U_j .

THEOREM 2.1. *The projection p_i of V onto V_i along $\oplus_{j \neq i} V_j$ is given by the formula*

$$p_i = \frac{d_i}{|G|} \sum_{g \in G} \chi_i(g)^* \rho(g).$$

Proof. See, for example, Theorem 8 in [19]. □

By Theorem 2.1, the isotypic projection v_i may be computed by directly applying p_i to v . There are, however, drawbacks to this approach. First, directly applying p_i to an arbitrary vector in V requires $O(\dim(V)^2)$ operations which may be prohibitive if $\dim(V)$ is large. Second, to construct p_i using the above formula requires a sum over the group G as well as an explicit knowledge of the representations of each element of G . This too may be prohibitive if G is large.

2.3. Separating sets. Suppose now that $\{T_1, \dots, T_k\}$ is a collection of diagonalizable linear transformations on V whose eigenspaces are direct sums of the isotypic subspaces of V . For each isotypic subspace V_i , let $c_i = (\mu_{i1}, \dots, \mu_{ik})$ be the k -tuple of eigenvalues where, for $1 \leq j \leq k$, μ_{ij} is the eigenvalue of T_j associated to V_i . If $c_i \neq c_{i'}$ whenever $V_i \neq V_{i'}$, then we say that $\{T_1, \dots, T_k\}$ is a *separating set* for V .

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

Stage 1. Compute the projections of v 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 .

⋮

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

LEMMA 2.2. *The computed projections at Stage k are precisely the isotypic projections of the vector v .*

Proof. The projections at each stage are sums of the isotypic projections of v . 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 $\{T_1, \dots, T_k\}$ is a separating set for V . □

We may easily find separating sets for V by looking to the conjugacy classes C_1, \dots, C_h of G . In particular, if $T_j = \sum_{c \in C_j} \rho(c)$ is the *class sum* of C_j (with respect to ρ) and $\mu_{ij} = |C_j| \chi_i(C_j) / d_i$, then we have the following lemma.

LEMMA 2.3. *The class sum T_j is a diagonalizable linear transformation on V whose eigenspaces are direct sums of isotypic subspaces, and μ_{ij} is the eigenvalue of T_j that is associated to the isotypic subspace V_i .*

Proof. This is a variation of Proposition 6 in [19]. □

The complete collection of class sums forms a separating set of V . In fact, by Theorem 2.1, every separating set for V is composed of linear combinations of class sums. We may, however, be able to find much smaller separating sets than the complete collection of class sums.

2.4. Permutation representations. Suppose now that G acts on a finite set X . Let $L(X)$ be the vector space of complex valued functions on X . The action of G on X induces a *permutation representation* $\rho : G \rightarrow GL(L(X))$ defined by

$$(\rho(g)(f))(x) = f(g^{-1}x)$$

for all $g \in G$, $f \in L(X)$, and $x \in X$. The vector space $L(X)$ has a natural basis $\{\delta_x\}_{x \in X}$, where

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

We will refer to $\{\delta_x\}_{x \in X}$ as the *delta basis* of $L(X)$. Note that $\dim(L(X)) = |X|$.

By choosing a basis for $L(X)$, we may identify each linear transformation on $L(X)$ with an $|X| \times |X|$ matrix. Thus, we will assume that each linear transformation on $L(X)$ is written as a matrix with respect to the delta basis of $L(X)$. In particular, if $g \in G$, then $\rho(g)$ corresponds to an $|X| \times |X|$ matrix with one 1 in each row and column, and zeros elsewhere.

3. Cyclic groups. In this section, we show how using separating sets to compute isotypic projections for cyclic groups leads to the Gentleman–Sande, or *decimation in frequency*, FFT (see [10]).

3.1. The DFT and isotypic projections. Let G be the cyclic group $\mathbb{Z}/n\mathbb{Z}$ and let X be the set $\{x_0, \dots, x_{n-1}\}$. Let ω be a primitive n th root of unity, let g be a generator for G , and let G act on X by setting $g^j x_i = x_{i+j}$, where all subscripts are taken modulo n . The resulting permutation representation

$$\rho : \mathbb{Z}/n\mathbb{Z} \rightarrow GL(L(X))$$

has n isotypic subspaces V_0, \dots, V_{n-1} , where each V_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 T_j of C_j associated to the isotypic subspace V_i is therefore $\chi_i(C_j)/d_i = \chi_i(g^j)/1 = \omega^{ij}$.

Let $f \in L(X)$ and let f_i be the isotypic projection of f onto the isotypic subspace V_i . Since ω is a primitive n th root of unity, the class sum T_1 forms a separating set for $L(X)$. The isotypic projection f_i may therefore be viewed as the projection of f onto the eigenspace of T_1 with eigenvalue ω^i . By Theorem 2.1, this may be computed as

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

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

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

This, however, is the i th coefficient of the usual DFT applied to f . An FFT on n points may therefore be thought of as an efficient algorithm for computing isotypic projections of vectors in $L(X)$.

3.2. The Gentleman–Sande FFT. Suppose now that $n = pq$. Since $\{T_1\}$ is a separating set for $L(X)$, so is $\{T_1, T_p\}$. We could therefore compute the isotypic projections of f by first computing the projections of f onto the eigenspaces of T_p and then projecting each of these projections onto the eigenspaces of T_1 .

The eigenspaces of T_p are W_0, \dots, W_{q-1} , where the eigenvalue of T_p that is associated to W_k is ω^{pk} and

$$W_k = V_k \oplus V_{k+q} \oplus \dots \oplus V_{k+(p-1)q}.$$

The projection f'_k of f onto W_k is therefore

$$(3.1) \quad f_k + f_{k+q} + \dots + f_{k+(p-1)q}.$$

In fact, the W_k are the isotypic subspaces of $L(X)$ with respect to the action on X of the subgroup of G that is generated by g^p . This subgroup is cyclic with order q . Thus, by Theorem 2.1,

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

Note that $f'_k(x_s) = \omega^{pkt} f'_k(x_{s+pt})$ and that f'_k is therefore determined by the values $f'_k(x_0), \dots, f'_k(x_{p-1})$. In this sense, since $f'_k(x_j)$ requires $O(q)$ operations to compute, f'_k requires $O(pq)$ operations to compute. The projections f'_0, \dots, f'_{q-1} may therefore 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 (3.1), the isotypic projection $f_i = f_{k+lq}$ may be computed by projecting f'_k onto the eigenspace of T_1 with eigenvalue $\omega^{(k+lq)}$. Recall that f_{k+lq} is determined by $f_{k+lq}(x_0)$, which we may compute as

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

This is a DFT on p points applied to the function $\omega^{ks} f'_k$. Thus, if we have computed f'_0, \dots, 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 $f \in L(X)$ using $O(p^2q + pq^2) = O((p+q)pq)$ operations.

This particular FFT is known as the Gentleman–Sande, or *decimation in frequency*, FFT (see [10]). The approach to decomposing representations that is presented in this paper may be viewed as a generalization of decimation in frequency since we too will be iteratively computing projections of projections.

4. The Lanczos iteration. Given a separating set, isotypic projections become eigenspace projections. In this section, we show how the Lanczos iteration gives rise to an efficient isotypic projection method when the number of isotypic subspaces is relatively small and the linear transformations in a separating set are real symmetric matrices that can be applied efficiently. Good references for the Lanczos iteration are [4, 16, 22, 23].

4.1. Krylov subspaces. Let \mathbb{C}^N be the usual complex vector space of N -tuples with complex coefficients. Let $M_N(\mathbb{C})$ be the set of $N \times N$ matrices 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 j th *Krylov subspace* generated by T and f is the subspace \mathcal{K}_j of \mathbb{C}^N that is spanned by the vectors $f, Tf, \dots, T^{j-1}f$. We write this as

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

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

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

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

where the V_i are the n distinct eigenspaces of T . Each $f \in \mathbb{C}^N$ may therefore be written uniquely as $f = f_1 + \dots + f_n$, 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 4.1. *If $T \in M_N(\mathbb{C})$ is diagonalizable and $f \in \mathbb{C}^N$, then the nontrivial 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 μ_1, \dots, μ_n and that $f = f_1 + \dots + f_n$, 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:

$$\begin{aligned}
 f &= f_1 + f_2 + \cdots + f_n, \\
 Tf &= \mu_1 f_1 + \mu_2 f_2 + \cdots + \mu_n f_n, \\
 T^2 f &= \mu_1^2 f_1 + \mu_2^2 f_2 + \cdots + \mu_n^2 f_n, \\
 &\vdots \\
 T^{n-1} f &= \mu_1^{n-1} f_1 + \mu_2^{n-1} f_2 + \cdots + \mu_n^{n-1} f_n.
 \end{aligned}
 \tag{4.1}$$

The coefficients of the f_i in (4.1) 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., [9]). 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 nontrivial f_i are linearly independent, the lemma follows. \square

COROLLARY 4.2. *The dimension of $\mathcal{K} = \langle f, Tf, T^2 f, \dots \rangle$ is equal to the number of nontrivial projections of f onto the eigenspaces of T .*

COROLLARY 4.3. *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 nontrivial 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 4.3. If the eigenspaces of the restriction of T to \mathcal{K} are orthogonal, this may be computed as

$$\frac{u^* f}{u^* u} u.
 \tag{4.2}$$

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 (4.2) requires $3n + 1$ operations. Relative to a basis of \mathbb{C}^N , however, this computation requires $3N + 1$ operations.

4.2. 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 j th *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  $i = 1, 2, 3, \dots$ 
   $v = Tq_i$ 
   $\alpha_i = q_i^* v$ 
   $v = v - \beta_{i-1}q_{i-1} - \alpha_i q_i$ 
   $\beta_i = \|v\|$ 
  if  $\beta_i \neq 0$ 
     $q_{i+1} = v/\beta_i$ 
  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 4.4. *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., [22]). For this reason, some form of reorthogonalization is usually introduced. For example, the *Lanczos iteration with complete reorthogonalization*, as described in [16], reorthogonalizes 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 Reorthogonalization
(assuming finite precision arithmetic)

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

```

for  $i = 1, 2, 3, \dots$ 
   $v = Tq_i$ 
   $\alpha_i = q_i^* v$ 
   $v = v - \beta_{i-1}q_{i-1} - \alpha_i q_i$ 
  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 reorthogonalization is much more stable than the Lanczos iteration without reorthogonalization. In fact, the numerical stability of the Lanczos iteration with reorthogonalization is comparable to that of

the Givens and Householder algorithms, which, like the Lanczos iteration, reduce a matrix to tridiagonal form (see Chapter 6, section 41 of [23]).

To get a sense of how much work it takes to compute the Lanczos iteration with complete reorthogonalization, let T^{op} be the number of operations needed to 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 4.5. *If T is an $N \times N$ real symmetric matrix and $f \in \mathbb{C}^N$, then*

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

operations are required to compute n iterations of the Lanczos iteration with complete reorthogonalization for T and f .

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

4.3. 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 4.6. *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 require $O(nT^{\text{op}} + n^2N)$ operations.*

Proof. The claim follows directly from the discussion in [16] of the Rayleigh–Ritz procedure applied to the sequence of Krylov subspaces $\mathcal{K}_1, \mathcal{K}_2, \dots$ generated by T and f . The method is important, however, so we include the details.

Suppose that f has m nonzero projections f_1, \dots, 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 m th Lanczos matrix generated during the Lanczos iteration with respect to T and f . Let $\{q_1, \dots, q_m\}$ be the corresponding orthonormal basis of the Krylov subspace \mathcal{K} generated by T and f .

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

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 $\|\tilde{u}_i\| = 1$. Since L_m is a real symmetric matrix, $\{\tilde{u}_1, \dots, \tilde{u}_m\}$ is an orthonormal basis for \mathcal{K} .

Since $q_1 = \|f\|^{-1}f$, $\tilde{f} = (\|f\|, 0, \dots, 0)^t$. 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 $\{q_1, \dots, q_m\}$. Thus, if Q_m is the $N \times m$ matrix whose i th 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 μ_1, \dots, μ_m and corresponding eigenvectors $\tilde{u}_1, \dots, \tilde{u}_m$ 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 $O(mT^{\text{op}} + m^2N)$ operations and Stage 2 requires $O(m^3)$ operations due to the tridiagonal form of T_m (see [23]). Stage 3 requires $O(m^2)$ operations and Stage 4 requires $O(m^2N)$ operations. Since $m \leq n \leq N$, the theorem follows. \square

Remark. The coefficient implied by $O(nT^{\text{op}} + n^2N)$ in Theorem 4.6 is independent of n , T^{op} , and N . We will implicitly make use of this fact throughout the rest of the paper.

We will refer to the projection method outlined in Theorem 4.6 as the *Lanczos eigenspace projection method* or LEPM.

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.

4.4. The Lanczos isotypic projection method. In this section, we combine the results of sections 2.3 and 4.3 to create an isotypic projection method that relies on the use of separating sets of real symmetric matrices.

Let G be a finite group, let V be a finite dimensional representation of G , and let $\{T_1, \dots, T_k\}$ be a separating set of real symmetric matrices for V . By Lemma 2.2, we may compute the isotypic projections of a vector $v \in V$ as follows.

Stage 1. Using the LEPM, compute the projections of v onto each of the eigenspaces of T_1 .

Stage 2. Using the LEPM, compute the projections of each of the previously computed projections onto each of the eigenspaces of T_2 .

⋮

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

We will refer to this approach to computing isotypic projections as the *Lanczos isotypic projection method* or LIPM.

Let $\iota(V)$ be the least number of operations needed to compute the isotypic projections of an arbitrary vector in V . We may now state our main theorem.

MAIN THEOREM 4.7. *Let G be a finite group acting on a finite set X . Let $L(X)$ be the resulting permutation representation. If $L(X) = V_1 \oplus \dots \oplus V_n$ is the isotypic decomposition of $L(X)$ and $\{T_1, \dots, T_k\}$ is an isotypic separating set of real symmetric matrices for $L(X)$, then*

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

Proof. The number of operations needed at the i th stage of the LIPM is never more than $O(nT_i^{\text{op}} + n^2|X|)$. The theorem follows immediately. \square

5. 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 2). A good reference for distance transitive graphs is [2].




FIG. 2. A distance transitive graph.

Let X be a distance transitive graph, let G be a distance transitive group of automorphisms of X , and let $L(X)$ 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 $A : L(X) \rightarrow L(X)$, where

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

for all $x \in X$. The operator A has $k + 1$ distinct eigenvalues which are also the zeros of certain polynomials associated with the graph X (see, e.g., [2]). For example, the adjacency operator of the graph in Figure 2, relative to its delta basis (as defined in section 2.4), is

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

It has three distinct eigenvalues.

LEMMA 5.1. *The isotypic subspaces of $L(X)$ are precisely the eigenspaces of A .*

Proof. This follows from section 2 of Stanton [21]. \square

THEOREM 5.2. *Let X be a distance transitive graph with diameter k , let G be a distance transitive group of automorphisms of X , and let $L(X)$ be the associated permutation representation of G . If A is the adjacency operator of X , then*

$$\iota(L(X)) = O(kA^{\text{op}} + k^2|X|).$$

Proof. Relative to the delta basis of $L(X)$, the adjacency operator A is a real symmetric matrix. Thus, the result follows from Theorem 4.7 and Lemma 5.1. \square

A direct matrix multiplication approach to computing isotypic projections for $L(X)$ requires $O(k|X|^2)$ operations. Although $O(kA^{\text{op}} + k^2|X|)$ 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.

5.1. Radon transforms. Let G be a finite group acting on finite sets X and Y and giving permutation representations $L(X)$ and $L(Y)$, 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* $R : L(X) \rightarrow L(Y)$ is then defined by

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

for all $x \in X$ (see [1]). The adjoint $R^* : L(Y) \rightarrow L(X)$ of R is defined by

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

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 $R : L(X) \rightarrow L(Y)$ 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 $L(X)$ and $L(Y)$, 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 5.3. *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. \square

LEMMA 5.4. *If $A : L(X) \rightarrow L(X)$ is the adjacency operator of X and $I : L(X) \rightarrow L(X)$ is the identity, then $A = (1/s)(R^*R - rI)$.*

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

$$\begin{aligned} (R^*Rf)(x) &= \sum_{y:x\sim y} \sum_{x':x'\sim y} f(x') = rf(x) + s \left(\sum_{x':d(x,x')=1} f(x') \right) \\ &= ((rI + sA)f)(x). \quad \square \end{aligned}$$

LEMMA 5.5. *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 $L(X)$ and $(R^*R)^{\text{op}} \leq 2r|X|$.*

Proof. Let A be the adjacency operator of X . The product R^*R and the adjacency operator A have the same eigenspaces by Lemma 5.4; therefore $\{(R^*R)\}$ is a separating set since $\{A\}$ is a separating set by Lemma 5.1.

We may apply R^*R to a vector $f \in L(X)$ by first computing Rf and then $R^*(Rf)$. Furthermore, when regarded as a matrix with respect to the delta bases

of $L(X)$ and $L(Y)$, both R and R^* contain $r|X|$ nonzero entries. It follows that $(R^*R)^{\text{op}} \leq R^{*\text{op}} + R^{\text{op}} \leq r|X| + r|X| = 2r|X|$. \square

By Theorem 4.7 and Lemma 5.5, we have the following theorem.

THEOREM 5.6. *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*

$$\iota(L(X)) = O(kr|X| + k^2|X|).$$

Remark. 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 5.6 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 $L(X)$. We illustrate this in the next two sections.

5.2. The Johnson graph. Let $n \geq 2$ and let $k \leq n/2$. The k -element subsets $X^{(n-k,k)}$ of $\{1, \dots, n\}$ 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 5.2, we therefore have that

$$(5.1) \quad \iota\left(L\left(X^{(n-k,k)}\right)\right) = O\left(k^2(n-k)\binom{n}{k}\right).$$

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 5.6, we have the following improvement to (5.1).

THEOREM 5.7. *If $n \geq 2$, $k \leq n/2$, and $L(X^{(n-k,k)})$ is the permutation representation of S_n associated to the Johnson graph $J(n, k)$, then*

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

We summarize the results of this section in Table 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 in [8].

TABLE 1
Upper bounds on $\iota(L(X^{(n-k,k)}))$.

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

5.3. 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 [2] for details concerning the Grassmann graph.

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

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

Although not obvious, this is a polynomial in q known as a *Gaussian polynomial* (see, e.g., [20]).

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 5.2, we have that

$$(5.2) \quad \iota(\mathbb{C}[X_{(n-k,k)}]) = O\left(kq[k][n - k] \binom{n}{k}_q\right).$$

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 5.6, we therefore have the following improvement to (5.2).

THEOREM 5.8. *Let $n \geq 2$ and $k \leq n/2$. Let $L(X_{(n-k,k)})$ be the permutation representation of $GL(n, q)$ associated to the Grassmann graph $G(n, k, q)$. Then*

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

We summarize the results of this section in Table 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 in [8].

TABLE 2
Upper bounds on $\iota(L(X_{(n-k,k)}))$.

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

6. The symmetric group. Spectral analysis for nonabelian groups has found its greatest success with the analysis of ranked data (see [5, 6, 17]). 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 $\lambda = (\lambda_1, \dots, \lambda_m)$ 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 $(1, \dots, 1)$.

If X^λ is the set of possible partial rankings of shape λ , the *partially ranked data of shape λ* is the function $f \in L(X^\lambda)$, 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 $\{1, \dots, 39\}$. 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 and 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 gives rise to an action of S_n on X^λ . Moreover, as noted in section 1, the isotypic subspaces of the resulting permutation representation $L(X^\lambda)$ correspond to certain *pure higher order effects* associated to the ranked data $f \in L(X^\lambda)$ (see [6, 17]). Computing the isotypic projections of f can therefore lead to some insight into how the respondents went about choosing their rankings.

6.1. Representation theory. Let n be a positive integer. A *composition* of n is a sequence $\lambda = (\lambda_1, \dots, \lambda_m)$ of positive integers whose sum is n . If $\lambda_1 \geq \dots \geq \lambda_m$, then λ is a *partition* of n . To each composition λ , there corresponds a partition $\bar{\lambda}$ obtained by arranging the parts of λ in nonincreasing 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 $\lambda_1 + \dots + \lambda_i \geq \lambda'_1 + \dots + \lambda'_i$ for all $i \geq 1$. If λ dominates λ' , then we write $\lambda \supseteq \lambda'$.

As is often the case, we identify the composition $\lambda = (\lambda_1, \dots, \lambda_m)$ of n with the *Ferrers diagram* of shape λ , which is the left-justified array of dots with λ_i dots in the i th row (see Figure 3). If the dots of a Ferrers diagram of shape λ are replaced by the numbers $1, \dots, n$ without repetition, then we create a *Young tableau* of shape




FIG. 3. The Ferrers diagram of shape $(2, 3, 1, 3)$.

λ . Two Young tableaux are said to be equivalent if they differ only by a permutation of the entries within the rows of the tableaux. An equivalence class of tableaux is a *tabloid*. A tabloid is denoted by forming a representative tableau and then drawing lines between the rows (see Figure 4).




FIG. 4. Two equivalent tableaux and their tabloid.

Let X^λ be the set of tabloids of shape λ . 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 nonincreasing order. We may therefore assume that λ is a partition of n .

Let λ be a partition of n . The action of S_n on $\{1, \dots, n\}$ induces an action of S_n on X^λ . For example, if $\sigma = (135)(27)$ and

$$t = \frac{\frac{5 \quad 2 \quad 3}{4 \quad 1 \quad 6}}{7},$$

then

$$\sigma t = \frac{\frac{\sigma(5) \quad \sigma(2) \quad \sigma(3)}{\sigma(4) \quad \sigma(1) \quad \sigma(6)}}{\sigma(7)} = \frac{\frac{1 \quad 7 \quad 5}{4 \quad 3 \quad 6}}{2}.$$

We denote the resulting permutation representation $L(X^\lambda)$ by M^λ .

For every partition μ of n , there is a simple representation W^μ of S_n . These representations form a complete (up to isomorphism) collection of simple representations of S_n . The representation M^λ is isomorphic to a direct sum of simple representations

$$M^\lambda \cong \bigoplus_{\mu \geq \lambda} \kappa_{\mu\lambda} W^\mu,$$

where the numbers $\kappa_{\mu\lambda}$ are *Kostka numbers* and denote the multiplicity of W^μ in M^λ . (Kostka numbers also count objects known as *semistandard tableaux*. See, e.g., [18].)

Note that the subspace of M^λ that is isomorphic to $\kappa_{\mu\lambda}W^\mu$ is the isotypic subspace of M^λ that corresponds to the simple representation W^μ .

6.2. Separating sets. Let C_i be the conjugacy class of i -cycles in S_n and let T_i be the corresponding class sum with respect to M^λ . For example, if $n = 4$, $i = 3$, and $\lambda = (2, 2)$, then

$$C_3 = \{(123), (132), (124), (142), (134), (143), (234), (243)\}$$

and, under a particular order of the delta basis of $M^{(2,2)}$,

$$T_3 = \begin{pmatrix} 0 & 2 & 2 & 2 & 2 & 0 \\ 2 & 0 & 2 & 2 & 0 & 2 \\ 2 & 2 & 0 & 0 & 2 & 2 \\ 2 & 2 & 0 & 0 & 2 & 2 \\ 2 & 0 & 2 & 2 & 0 & 2 \\ 0 & 2 & 2 & 2 & 2 & 0 \end{pmatrix}.$$

THEOREM 6.1 (Katriel). *If $\lambda = (\lambda_1, \dots, \lambda_m)$ is a partition of n , then $\{T_2, \dots, T_m\}$ is a separating set for M^λ .*

Proof. This is Theorem 3 in Katriel [11] rewritten using the language of separating sets. \square

Moreover, the number of T_i that are actually needed to form a separating set for any representation of S_n seems to be small relative to n . Katriel made this observation after calculations revealed that $\{T_2, \dots, T_{k+1}\}$ is a separating set for any representation of the symmetric group on $\varphi(k)$ or less symbols, where $\varphi(k)$ is much larger than k . For example, $\{T_2\}$ is a separating set for S_2, S_3, S_4 , and S_5 but not S_6 . Thus $\varphi(1) = 5$. Similarly, calculations have shown that $\varphi(2) = 14$, $\varphi(3) = 23$, $\varphi(4) = 41$, and $\varphi(5) \geq 72$ (see [11]). We therefore have the following theorem.

THEOREM 6.2. *Let n and k be positive integers such that $n \leq \varphi(k)$. If λ is a partition of n , and ζ_λ is the number of isotypic subspaces of M^λ , then*

$$\iota(M^\lambda) = O\left(\sum_{i=2}^{k+1} \left(\zeta_\lambda(i-1)! \binom{n}{i} |X^\lambda| + \zeta_\lambda^2 |X^\lambda|\right)\right).$$

Proof. The collection $\{T_2, \dots, T_{k+1}\}$ is a separating set for M^λ since $n \leq \varphi(k)$. It is easy to show that each T_i is a real symmetric matrix with respect to the delta basis of M^λ . Thus, by Theorem 4.7,

$$\iota(M^\lambda) = O\left(\sum_{i=2}^{k+1} (\zeta_\lambda T_i^{\text{op}} + \zeta_\lambda^2 |X^\lambda|)\right).$$

Recall that T_i^{op} is no more than the number of nonzero entries in T_i , which is at most $|C_i||X^\lambda|$. Since $|C_i| = (i-1)! \binom{n}{i}$, the theorem follows. \square

We summarize the results of this section, and include some particular examples, in Table 3.

Remarks. Note that when $n \geq 2$ and $k \leq n/2$, we were able to find a bound for $\iota(M^{(n-k,k)})$ in section 5.2 by viewing the elements of $X^{(n-k,k)}$ as the vertices of a distance transitive graph. Moreover, the upper bound in section 5.2 is much better than the upper bound given by Theorem 6.2.

TABLE 3
Upper bounds on $\iota(M^\lambda)$.

λ	LIPM	Direct matrix multiplication
$(n-k, k)$	$O\left(k^2(n-k)\binom{n}{k}\right)$	$O\left(k\binom{n}{k}^2\right)$
$(n-k, k-1, 1)$	$O\left(k^3(n-k)\binom{n}{k}\right)$	$O\left(k^3\binom{n}{k}^2\right)$
$(\lambda_1, \dots, \lambda_m)$ $n \leq \varphi(k)$	$O\left(\sum_{i=2}^{k+1} (\zeta_\lambda(i-1)! \binom{n}{i} X^\lambda + \zeta_\lambda^2 X^\lambda)\right)$	$O(\zeta_\lambda X^\lambda ^2)$

Additionally, an FFT and inverse for the symmetric group, both requiring $O(n^2n!)$ operations, were constructed in [12]. Thus if $p(n)$ is the number of partitions of n , then the isotypic projections of a vector in $M^{(1, \dots, 1)}$ may be computed using $O(p(n)n^2n!)$ operations. See [14] for an FFT for the homogeneous space $M^{(n-k, k)}$ and [15] for some generalizations of the results presented in this paper.

REFERENCES

- [1] E. BOLKER, *The finite Radon transform*, in Integral Geometry (Brunswick, Maine, 1984), AMS, Providence, RI, 1987, pp. 27–50.
- [2] A. BROUWER, A. COHEN, AND A. NEUMAIER, *Distance-regular Graphs*, Springer-Verlag, Berlin, 1989.
- [3] J. CHEN, *Group Representation Theory for Physicists*, 2nd ed., World Scientific, Teaneck, NJ, 1989.
- [4] J. CULLUM AND R. WILLOUGHBY, *Lánczos Algorithms for Large Symmetric Eigenvalue Computations, Vol. I, Theory*, Birkhäuser Boston, Boston, MA, 1985.
- [5] P. DIACONIS, *Group Representations in Probability and Statistics*, Institute of Mathematical Statistics, Hayward, CA, 1988.
- [6] P. DIACONIS, *A generalization of spectral analysis with application to ranked data*, Ann. Statist., 17 (1989), pp. 949–979.
- [7] P. DIACONIS AND D. ROCKMORE, *Efficient computation of isotypic projections for the symmetric group*, in Groups and Computation (New Brunswick, NJ, 1991), AMS, Providence, RI, 1993, pp. 87–104.
- [8] 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), pp. 1066–1099.
- [9] D. DUMMIT AND R. FOOTE, *Abstract Algebra*, 2nd ed., John Wiley, New York, 1999.
- [10] W. GENTLEMAN AND G. SANDE, *Fast Fourier transforms for fun and profit*, in Proc. AFIPS, Fall Joint Computer Conference, Vol. 29, Spartan Books, NY, 1966, pp. 563–578.
- [11] J. KATRIEL, *Some useful results concerning the representation theory of the symmetric group*, J. Phys. A, 24 (1991), pp. 5227–5234.
- [12] D. MASLEN, *The efficient computation of Fourier transforms on the symmetric group*, Math. Comp., 67 (1998), pp. 1121–1147.
- [13] D. MASLEN AND D. ROCKMORE, *Generalized FFTs—a survey of some recent results*, in Groups and Computation, II, (New Brunswick, NJ, 1995), AMS, Providence, RI, 1997, pp. 183–237.
- [14] D. MASLEN AND D. ROCKMORE, *Separation of variables and the computation of Fourier transforms on finite groups. I*, J. Amer. Math. Soc., 10 (1997), pp. 169–214.
- [15] M. ORRISON, *An Eigenspace Approach to Decomposing Representations of Finite Groups*, Ph.D. thesis, Dartmouth College, Hanover, NH, 2001.

- [16] B. PARLETT, *The Symmetric Eigenvalue Problem*, Prentice-Hall Inc., Englewood Cliffs, NJ, 1980.
- [17] D. ROCKMORE, *Some applications of generalized FFTs*, in *Groups and Computation, II* (New Brunswick, NJ, 1995), AMS, Providence, RI, 1997, pp. 329–369.
- [18] B. SAGAN, *The Symmetric Group. Representations, Combinatorial Algorithms, and Symmetric Functions*, The Wadsworth & Brooks/Cole Mathematics Series, Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA, 1991.
- [19] J.-P. SERRE, *Linear Representations of Finite Groups*, Springer-Verlag, New York, 1977.
- [20] R. STANLEY, *Enumerative Combinatorics. Vol. 1*, Cambridge University Press, Cambridge, UK, 1997.
- [21] D. STANTON, *Orthogonal polynomials and Chevalley groups*, in *Special Functions: Group Theoretical Aspects and Applications*, Reidel, Dordrecht, The Netherlands, 1984, pp. 87–128.
- [22] L. TREFETHEN AND D. BAU, III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [23] J. WILKINSON, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, UK, 1965.