



Computational Bounds for Doing Harmonic Analysis on Permutation Modules of Finite Groups

Michael Hansen¹ · Masanori Koyama² · Matthew B. A. McDermott³ · Michael E. Orrison⁴ · Sarah Wolff⁵

Received: 30 October 2019 / Revised: 5 August 2021 / Accepted: 11 August 2021 /
Published online: 3 September 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We develop an approach to finding upper bounds for the number of arithmetic operations necessary for doing harmonic analysis on permutation modules of finite groups. The approach takes advantage of the intrinsic orbital structure of permutation modules, and it uses the multiplicities of irreducible submodules within individual orbital spaces to express the resulting computational bounds. We conclude by showing that these bounds are surprisingly small when dealing with certain permutation modules arising from the action of the symmetric group on tabloids.

Keywords Harmonic analysis · Permutation modules · Finite groups

Communicated by Hans G. Feichtinger.

✉ Sarah Wolff
wolffs@denison.edu

Michael Hansen
mhansen@gmail.com

Masanori Koyama
masomatics@preferred.jp

Matthew B. A. McDermott
mmd@mit.edu

Michael E. Orrison
orrison@hmc.edu

¹ Rover, Seattle, WA, USA

² Preferred Networks, Tokyo, Japan

³ Computer Science & Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

⁴ Department of Mathematics, Harvey Mudd College, Claremont, CA, USA

⁵ Department of Mathematics and Computer Science, Denison University, Granville, OH, USA

Mathematics Subject Classification Primary 43A85 · 20C40 · 65T50

1 Introduction

Let G be a finite group acting on a finite set X . Let $\mathbb{C}G$ be the complex group algebra of G , and let $\mathbb{C}X$ be the vector space of complex-valued functions defined on X . The action of G on X turns $\mathbb{C}X$ into a $\mathbb{C}G$ -permutation module, where if $\alpha = \sum_{g \in G} \alpha(g)g \in \mathbb{C}G$ and $f \in \mathbb{C}X$, then

$$(\alpha \cdot f)(x) = \sum_{g \in G} \alpha(g) f(g^{-1}x)$$

for all $x \in X$. In this setting, we say that a basis \mathcal{B} for $\mathbb{C}X$ is a *harmonic basis* (with respect to $\mathbb{C}G$) if it can be partitioned into subsets such that each subset forms a basis for an irreducible $\mathbb{C}G$ -submodule of $\mathbb{C}X$. Note that a vector in a harmonic basis is called a *fundamental harmonic* in [14].

The problem addressed in this paper is the following: Given an arbitrary $f \in \mathbb{C}X$, how may we efficiently compute the coefficients necessary to express f in terms of a harmonic basis for $\mathbb{C}X$? In other words, if $X = \{x_1, \dots, x_m\}$, and we are given $f(x_1), \dots, f(x_m)$, then how may we efficiently find the complex coefficients β_1, \dots, β_m such that $f = \beta_1 b_1 + \dots + \beta_m b_m$ for some harmonic basis $\mathcal{B} = \{b_1, \dots, b_m\}$ of $\mathbb{C}X$? In this paper, we create and then analyze a change of basis algorithm for computing the coefficients β_1, \dots, β_m when \mathcal{B} is a harmonic basis that is also a *symmetry-adapted basis* associated to a chain of subgroups of G (see Sect. 2.3).

The problem of computing such coefficients arises when doing *harmonic analysis on finite groups*, especially when the function $f \in \mathbb{C}X$ corresponds to complex-valued data defined on a set X with an underlying symmetry group G (see, for example, [3,13,26]). It also arises in *generalized spectral analysis*, which was developed by Diaconis [9,10] and which extends the classical spectral analysis of time series to the analysis of functions in $\mathbb{C}X$.

For a well-known example, let $X = \{x_1, \dots, x_m\}$ and let G be the cyclic group $\mathbb{Z}/m\mathbb{Z}$ acting on X by cyclically permuting its elements. The elements of $\mathbb{C}X$ may be viewed as discretized periodic *signals* on m points, in which case the irreducible submodules of $\mathbb{C}X$ correspond to different *frequencies* in the usual signal processing sense. In this case, the coefficients β_1, \dots, β_m for $f \in \mathbb{C}X$ are then just fixed scalar multiples of the usual Fourier coefficients of f , which may be found by applying the usual discrete Fourier transform (DFT) to f . The classical fast Fourier transform (FFT) may therefore be used to efficiently compute the β_i using $O(m \log m)$ arithmetic operations (see, for example, [5,27]).

For another well-known example, suppose respondents in a survey are asked to choose their top k favorite items from a set of n items, where $k \leq n/2$. In this case, the set X is the set of all k -element subsets of the items, and the group G is the symmetric group S_n , whose natural action on the items induces an action on the set X . If $f \in \mathbb{C}X$ is the function defined by setting $f(x)$ to be the number of people who

choose the k -element set x , then the irreducible $\mathbb{C}S_n$ -modules in $\mathbb{C}X$ correspond to summary statistics about j -element subsets of the items where $0 \leq j \leq k$. As such, the associated β_j for such survey data can be used to uncover hidden relationships among the n items being ranked (see, for example, [7,9,10,17]).

The techniques and insights for the analysis of such “top k ” survey data extend to the situation in which survey respondents are asked to choose and rank their top k favorite items from the set of n items. The efficient harmonic analysis of such *partially ranked data* is addressed later in this paper (see Sect. 4). Along the way, we also develop useful results concerning the harmonic analysis of finite-dimensional permutation modules in general, and we recover a well-known result due to Clausen [4] concerning a bound for the number of arithmetic operations necessary to apply generalized discrete Fourier transforms of finite groups (see Sect. 3).

The overarching approach that we take in this paper is relatively simple. Given a chain $\{1\} = G_1 < \cdots < G_n = G$ of subgroups of G , we will associate to each subgroup G_j a special kind of harmonic basis \mathcal{B}_j of $\mathbb{C}X$, where \mathcal{B}_1 is the standard basis of $\mathbb{C}X$. In particular, these harmonic bases will be *symmetry-adapted* (see Sect. 2.3). We will then bound the number of arithmetic operations necessary to do a change of basis from \mathcal{B}_1 to \mathcal{B}_2 , then from \mathcal{B}_2 to \mathcal{B}_3 , and so on, until we reach the change of basis from \mathcal{B}_{n-1} to \mathcal{B}_n (see Sect. 3.2). Combining these results leads to an overall bound for the number of arithmetic operations necessary to do a change of basis from \mathcal{B}_1 to \mathcal{B}_n . As we will show, in some cases, this bound is surprisingly small.

Using symmetry-adapted bases to create fast algorithms for doing harmonic analysis is not new (see, for example, [2,5,18,19,21]). However, most books and papers on the subject focus on the regular representation of a group, i.e., when $X = G$, and the action of G on X is simply given by group multiplication. The resulting algorithms are then sometimes modified to deal with homogeneous spaces $X = G/H$ where $H = G_i$ for some subgroup G_i in the chain $\{1\} = G_1 < \cdots < G_n = G$.

In this paper, we show that it can also be fruitful to use symmetry-adapted bases when dealing with permutation representations in general. In particular, we do not require that X be a homogeneous space associated with a subgroup in the chain $\{1\} = G_1 < \cdots < G_n = G$. Instead, by using symmetry-adapted bases that respect the intrinsic orbital structure of a permutation representation, we are able to provide straightforward but nontrivial bounds expressed in terms of the multiplicities of the irreducible representations that arise when we restrict the action of G to subgroups in the chain $\{1\} = G_1 < \cdots < G_n = G$ (see Sect. 3).

Although we hope that our results appeal to a wide-ranging audience, for both convenience and space considerations, we will assume throughout the rest of the paper that the reader has a working knowledge of the basic representation theory of the symmetric group, and that the reader is familiar with the representation theory of finite groups in general. See, for example, the books by Sagan [23] and Serre [24].

2 Background and Lemmas

In this section, we explain some of the terminology and notation we will be using throughout the rest of the paper. We also introduce a couple of foundational lemmas

that will be used in the next section when we discuss general upper bounds. Good references for the ideas and terminology in this section are [3,5,13,26].

2.1 Functions on X

Let $X = \{x_1, \dots, x_m\}$ be a finite set. We denote the complex vector space of complex-valued functions defined on X by $\mathbb{C}X$. For convenience, we will identify $x \in X$ with the function that is 1 on x and 0 on all of the other elements. The x_i then form a basis for $\mathbb{C}X$, which we call the *standard basis* for $\mathbb{C}X$ and denote by \mathcal{B}_1 .

If $f = f(x_1)x_1 + \dots + f(x_m)x_m$, then the coordinate vector of f with respect to \mathcal{B}_1 is the column vector

$$[f]_{\mathcal{B}_1} = \begin{bmatrix} f(x_1) \\ \vdots \\ f(x_m) \end{bmatrix}.$$

More generally, if $\mathcal{B} = \{b_1, \dots, b_m\}$ is a basis for $\mathbb{C}X$, and $f = \beta_1 b_1 + \dots + \beta_m b_m$ for some complex coefficients β_1, \dots, β_m , then the coordinate vector of f with respect to \mathcal{B} is

$$[f]_{\mathcal{B}} = \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_m \end{bmatrix}.$$

2.2 Permutation Modules

If G is a finite group acting on the left on a finite set $X = \{x_1, \dots, x_m\}$, then G acts naturally on $\mathbb{C}X$, where if $g \in G$ and $f \in \mathbb{C}X$, then

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

for all $x \in X$. Linearly extending this action to $\mathbb{C}G$ then turns $\mathbb{C}X$ into a $\mathbb{C}G$ -permutation module. The action of G on $\mathbb{C}X$ gives rise to a *permutation representation*

$$\varphi : G \rightarrow GL_{|X|}(\mathbb{C})$$

where $\varphi(g)$ is the permutation matrix that encodes the action of $g \in G$ on X with respect to the standard basis \mathcal{B}_1 :

$$[\varphi(g)]_{ij} = \begin{cases} 1 & \text{if } gx_j = x_i \\ 0 & \text{otherwise.} \end{cases}$$

In particular, for all $g \in G$ and for all $f \in \mathbb{C}X$,

$$[g \cdot f]_{\mathcal{B}_1} = \varphi(g)[f]_{\mathcal{B}_1}.$$

In general, if \mathcal{B} is a basis for $\mathbb{C}X$, then we will denote the matrix encoding of the action of $g \in G$ with respect to \mathcal{B} by $[g]_{\mathcal{B}}$, which is the unique matrix with the property that for all $f \in \mathbb{C}X$,

$$[g]_{\mathcal{B}}[f]_{\mathcal{B}} = [g \cdot f]_{\mathcal{B}}.$$

In this case, note that $\varphi(g) = [g]_{\mathcal{B}_1}$, and that the map $g \mapsto [g]_{\mathcal{B}}$ creates a representation of G that is equivalent to φ but uses the basis \mathcal{B} instead of \mathcal{B}_1 .

2.3 Symmetry-Adapted Bases

Let G be a finite group, and let $\mathcal{R}(G)$ denote a fixed maximal set of pairwise inequivalent irreducible complex representations of G . Let M be a finite-dimensional $\mathbb{C}G$ -module with basis \mathcal{B} . Recall that \mathcal{B} is a *harmonic basis* if it can be partitioned into subsets such that each subset forms a basis for an irreducible $\mathbb{C}G$ -submodule of M . Furthermore, the basis \mathcal{B} is a *symmetry-adapted basis of M with respect to $\mathcal{R}(G)$* if for all $g \in G$, $[g]_{\mathcal{B}}$ is block diagonal, and each block of $[g]_{\mathcal{B}}$ is of the form $\rho(g)$ for some $\rho \in \mathcal{R}(G)$. A symmetry-adapted basis is therefore a special kind of harmonic basis that reflects the action of G as encoded in the representations in $\mathcal{R}(G)$. (See [13] for an algorithm for constructing symmetry-adapted bases, and see [1] for an algorithm for computing irreducible representations of finite groups.)

Suppose now that $\{1\} = G_1 < \dots < G_n = G$ is chain of subgroups of G . It is possible to find $\mathcal{R}(G_1), \dots, \mathcal{R}(G_n)$ and a basis \mathcal{B} of M such that for all $1 \leq i \leq n$, when M is viewed as a $\mathbb{C}G_i$ -module (by restricting the action of G to G_i), \mathcal{B} is simultaneously a symmetry-adapted basis with respect to $\mathcal{R}(G_1), \dots, \mathcal{R}(G_n)$ (see Theorem 13.45 in [2]). In this case, we say that the $\mathcal{R}(G_i)$ are *compatible* with respect to M , and that \mathcal{B} is a *symmetry-adapted basis with respect to the list $\mathcal{R}(G_1), \dots, \mathcal{R}(G_n)$* .

In what follows, we will make use of bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ where \mathcal{B}_j is a symmetry-adapted basis with respect to the list $\mathcal{R}(G_1), \dots, \mathcal{R}(G_j)$. The goal will be to show that the change of basis from \mathcal{B}_1 to \mathcal{B}_n can sometimes be computed efficiently by computing a change of basis from \mathcal{B}_1 to \mathcal{B}_2 , then from \mathcal{B}_2 to \mathcal{B}_3 , then from \mathcal{B}_3 to \mathcal{B}_4 , and so on until we reach \mathcal{B}_n .

Since we are dealing with permutation representations in this paper, we will take advantage of the orbit structure that is inherently present. To explain, suppose the action of G on X partitions X into orbits X_1, \dots, X_t . With a slight abuse of notation, we may view each $\mathbb{C}X_i$ as a submodule of $\mathbb{C}X$, in which case we may write $\mathbb{C}X$ as the direct sum

$$\mathbb{C}X = \mathbb{C}X_1 \oplus \dots \oplus \mathbb{C}X_t.$$

An *orbital harmonic basis* for $\mathbb{C}X$ is then a basis for $\mathbb{C}X$ that can be partitioned into subsets that form harmonic bases for the $\mathbb{C}X_i$. Note that the definition of orbital

harmonic basis depends only on the orbits of the action of G on X and does not require symmetry-adaptedness with respect to a list of compatible representations. We will specify when we are working with harmonic bases that are both orbital and symmetry-adapted with respect to a list of compatible representations. As we will see, working with such bases will sometimes lead to a fast algorithm for doing harmonic analysis on a permutation module.

As a simple but helpful example, consider the situation where $G = S_3$ acts on the set $X = \{1, 2, 3\}$ in the usual way. Expressing the vectors in $\mathbb{C}X$ using the coordinate vectors of the standard basis, we can write $\mathbb{C}X$ as a direct sum of irreducible $\mathbb{C}S_1$ -, $\mathbb{C}S_2$ -, and $\mathbb{C}S_3$ -modules, respectively, as

$$\begin{aligned} \mathbb{C}X &= \left\langle \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\rangle \oplus \left\langle \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\rangle \oplus \left\langle \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\rangle \\ &= \left\langle \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \right\rangle \oplus \left\langle \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} \right\rangle \oplus \left\langle \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \right\rangle \\ &= \left\langle \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\rangle \oplus \left\langle \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1/2 \\ 1/2 \\ -1 \end{bmatrix} \right\rangle. \end{aligned}$$

The vectors appearing in each line above form orbital harmonic bases $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{B}_3 of $\mathbb{C}X$ with respect to the compatible *seminormal representations* (see, for example, [16]) of S_1, S_2 , and S_3 , respectively. The change of basis matrix from \mathcal{B}_1 to \mathcal{B}_2 is

$$\begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/2 & -1/2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and the change of basis matrix from \mathcal{B}_2 to \mathcal{B}_3 is

$$\begin{bmatrix} 2/3 & 0 & 1/3 \\ 0 & 1 & 0 \\ 2/3 & 0 & -2/3 \end{bmatrix}.$$

The change of basis matrix from \mathcal{B}_1 to \mathcal{B}_3 is therefore the product

$$\begin{bmatrix} 2/3 & 0 & 1/3 \\ 0 & 1 & 0 \\ 2/3 & 0 & -2/3 \end{bmatrix} \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/2 & -1/2 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ 1/2 & -1/2 & 0 \\ 1/3 & 1/3 & -2/3 \end{bmatrix}.$$

2.4 Frequency Spaces

Let $\rho \in \mathcal{R}(G)$, and suppose ρ has degree d . For each i such that $1 \leq i \leq d$, there is a primitive idempotent $e = \sum_{g \in G} \epsilon(g)g \in \mathbb{C}G$ associated to ρ such that

- (1) $\sum_{g \in G} \epsilon(g)\rho(g)$ is a $d \times d$ matrix filled with zeros except for a 1 in the ii -th position, and
- (2) $\sum_{g \in G} \epsilon(g)\rho'(g)$ is the zero matrix for all $\rho' \in \mathcal{R}(G)$ such that $\rho' \neq \rho$.

It follows that if \mathcal{B} is a symmetry-adapted basis with respect to $\mathcal{R}(G)$ for the $\mathbb{C}G$ -module M , then each basis vector in \mathcal{B} is an eigenvector for e with eigenvalue 1 or 0.

We call the subspace of M spanned by the vectors in \mathcal{B} that are eigenvectors for e with eigenvalue 1 the *frequency space* corresponding to e . It is the subspace

$$eM = \{em \mid m \in M\}.$$

If we let $\mathcal{E}(G)$ denote the set of primitive idempotents in $\mathbb{C}(G)$ corresponding to $\mathcal{R}(G)$, then we may write M as a direct sum

$$M = \bigoplus_{e \in \mathcal{E}(G)} eM$$

of frequency spaces. (This is known as a *Pierce decomposition* of the module M . See, for example, [12].) In this case, we will say that eM is a *frequency space of M with respect to $\mathcal{R}(G)$* .

Lemma 1 *Let G be a finite group, and let M be a finite-dimensional $\mathbb{C}G$ -module. Suppose $\mathcal{R}(G) = \{\rho_1, \dots, \rho_h\}$ and that U_j is an irreducible $\mathbb{C}G$ -module corresponding to ρ_j . If e is a primitive idempotent associated to ρ_j as described above, and $M \cong \alpha_1 U_1 \oplus \dots \oplus \alpha_h U_h$, then the frequency space eM has dimension α_j .*

Proof Let \mathcal{B} be a symmetry-adapted basis of M with respect to $\mathcal{R}(G)$. If $e = \sum_{g \in G} \epsilon(g)g$, then by construction, the matrix $[e]_{\mathcal{B}} = \sum_{g \in G} \epsilon(g)[g]_{\mathcal{B}}$ will have zeros everywhere except for α_j 1's on its diagonal, one for each copy of U_j in M . It follows that $[e]_{\mathcal{B}}$ has rank α_j , and thus the dimension of eM is α_j . \square

Suppose now that H is a subgroup of G , and that \mathcal{B} is a symmetry-adapted basis with respect to $\mathcal{R}(H)$ and $\mathcal{R}(G)$. The following lemma shows that the frequency spaces for H and G are nicely related.

Lemma 2 *Let G be a finite group, let H be a subgroup of G , and let M be a finite-dimensional $\mathbb{C}G$ -module. If $\mathcal{R}(H)$ and $\mathcal{R}(G)$ are compatible with respect to M , then each frequency space of M with respect to $\mathcal{R}(H)$ is a direct sum of frequency spaces of M with respect to $\mathcal{R}(G)$.*

Proof Without loss of generality, we may assume that $\mathcal{R}(H)$ and $\mathcal{R}(G)$ are also compatible with respect to the regular $\mathbb{C}G$ -module. The result then follows from the fact that, by our definition of compatibility, each primitive idempotent in $\mathcal{E}(H)$ is a sum of primitive idempotents in $\mathcal{E}(G)$. \square

Finally, we introduce some notation and one more term that will be helpful in the next section.

Suppose G acts transitively on a finite set X . Let H be a subgroup of G , let N_1, \dots, N_s be a complete collection of irreducible $\mathbb{C}H$ -modules, and suppose that as a $\mathbb{C}H$ -module, we have that

$$\mathbb{C}X \cong \kappa_1 N_1 \oplus \dots \oplus \kappa_s N_s.$$

We will let $K_X(G, H)$ denote $\max\{\kappa_1, \dots, \kappa_s\}$. Also, when we consider the frequency spaces of $\mathbb{C}X$ as a $\mathbb{C}H$ -module, we will let $\Phi_X(G, H)$ denote the sum of the squares of the dimensions of these frequency spaces. By Lemma 1, it follows that

$$\Phi_X(G, H) = \kappa_1^2 \dim(N_1) + \dots + \kappa_s^2 \dim(N_s).$$

If the action of G on X is not transitive, and the orbits are X_1, \dots, X_t , then we will use $\Phi_X(G, H)$ to denote the sum

$$\Phi_{X_1}(G, H) + \dots + \Phi_{X_t}(G, H)$$

where we apply the previous definition of $\Phi_{X_i}(G, H)$ to each of the X_i . Similarly, we will use $K_X(G, H)$ to instead denote $\max\{K_{X_1}(G, H), \dots, K_{X_t}(G, H)\}$. Finally, we will refer to the frequency spaces of the $\mathbb{C}X_i$ with respect to $\mathcal{R}(H)$ as (G, H) -orbital frequency spaces.

3 General Upper Bounds

Let G be a finite group acting on a finite set X , let $\mathbb{C}X$ be the associated $\mathbb{C}G$ -permutation module, and let $f \in \mathbb{C}X$. In this section, we will assume we are given the subgroup chain

$$\{1\} = G_1 < \dots < G_n = G$$

and collections $\mathcal{R}(G_1), \dots, \mathcal{R}(G_n)$ of irreducible representations that are compatible with respect to $\mathbb{C}X$. We will also assume that we are given harmonic bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ of $\mathbb{C}X$, where

- (1) \mathcal{B}_j is an orbital symmetry-adapted basis with respect to $\mathcal{R}(G_1), \dots, \mathcal{R}(G_j)$, and
- (2) \mathcal{B}_1 is the standard basis of $\mathbb{C}X$.

In this section, we find bounds for the number of arithmetic operations necessary to compute the coordinate vector $[f]_{\mathcal{B}_n}$ when given the coordinate vector $[f]_{\mathcal{B}_1}$.

3.1 Nonzero Entries in Matrices

If A is a nonzero $m \times m$ matrix with complex entries, then we will denote the number of nonzero entries in A by $\nu(A)$. We will view $\nu(A)$ as a measure of how difficult it is to multiply the matrix A and an arbitrary vector in \mathbb{C}^m . In particular, note that such a

product will require no more than $\nu(A)$ multiplications and strictly fewer than $\nu(A)$ additions.

If we let $\omega(A)$ denote the total number of arithmetic operations required to compute the product of a nonzero $m \times m$ matrix A and an arbitrary vector in \mathbb{C}^m , note that we then have that $\omega(A) < 2\nu(A)$. Furthermore, if A can be factored as a product $A = A_1 \cdots A_l$, then $\omega(A) < 2\nu(A_1) + \cdots + 2\nu(A_l)$.

Suppose now that V is an m -dimensional vector space, and that \mathcal{B} and \mathcal{B}' are bases for V . Let $C(\mathcal{B}, \mathcal{B}')$ be the change of basis matrix from the basis \mathcal{B}' to the basis \mathcal{B} . In other words, $C(\mathcal{B}, \mathcal{B}')$ is the unique $m \times m$ matrix such that

$$C(\mathcal{B}, \mathcal{B}')[v]_{\mathcal{B}'} = [v]_{\mathcal{B}}$$

for all $v \in V$. In this case, note that $\nu(C(\mathcal{B}, \mathcal{B}')) \leq m^2$. Also, if \mathcal{B}'' is another basis for V , then $C(\mathcal{B}, \mathcal{B}'') = C(\mathcal{B}, \mathcal{B}')C(\mathcal{B}', \mathcal{B}'')$ and thus

$$\omega(C(\mathcal{B}, \mathcal{B}'')) < 2\nu(C(\mathcal{B}, \mathcal{B}')) + 2\nu(C(\mathcal{B}', \mathcal{B}'')).$$

It follows that if we are given harmonic bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ of $\mathbb{C}X$ corresponding to the chain

$$\{1\} = G_1 < \cdots < G_n = G$$

then we may compute $[f]_{\mathcal{B}_n}$ by iteratively computing the product $C(\mathcal{B}_j, \mathcal{B}_{j-1})[f]_{\mathcal{B}_{j-1}} = [f]_{\mathcal{B}_j}$ as j goes from 2 to n . Furthermore, this approach will be more efficient than the usual naive approach of simply computing $C(\mathcal{B}_n, \mathcal{B}_1)[f]_{\mathcal{B}_1} = [f]_{\mathcal{B}_n}$ whenever

$$\nu(C(\mathcal{B}_n, \mathcal{B}_{n-1})) + \cdots + \nu(C(\mathcal{B}_2, \mathcal{B}_1))$$

is small relative to $(\dim \mathbb{C}X)^2 = |X|^2$. We would therefore like to find upper bounds for each of the $\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1}))$ that are much smaller than $|X|^2$.

3.2 Bounds Based on the Dimensions of Frequency Spaces

The following theorem is our main theorem. It provides a bound on $\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1}))$ in terms of the dimensions of the frequency spaces of $\mathbb{C}X$ when viewed as a $\mathbb{C}G_{j-1}$ -module.

Theorem 3 *Let $2 \leq j \leq n$. The number of nonzero entries in the change of basis matrix $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ is bounded above by $\Phi_X(G_j, G_{j-1})$. In other words,*

$$\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})) \leq \Phi_X(G_j, G_{j-1}).$$

Furthermore, each column of $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ has at most $K_X(G_j, G_{j-1})$ nonzero entries, and thus we also have that

$$\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})) \leq K_X(G_j, G_{j-1})|X|.$$

Proof This follows from Lemma 1, Lemma 2, and the fact that \mathcal{B}_{j-1} and \mathcal{B}_j are orbital symmetry-adapted bases. In particular, suppose the action of G_j on X partitions X into orbits X_1, \dots, X_t . By Lemma 2, the change of basis from \mathcal{B}_{j-1} to \mathcal{B}_j only involves doing a change of basis within every frequency space, when viewed as a $\mathbb{C}G_{j-1}$ -module, of each submodule $\mathbb{C}X_i$ of $\mathbb{C}X$. Thus, by Lemma 1, each column of $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ will have at most $K_X(G_j, G_{j-1})$ nonzero entries. Also, since a change of basis matrix for a d -dimensional vector space has at most d^2 nonzero entries, the number of nonzero entries in the change of basis matrix $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ is bounded above by $\Phi_X(G_j, G_{j-1})$, which is simply the sum of the squares of the dimensions of the (G_j, G_{j-1}) -orbital frequency spaces for $\mathbb{C}G_{j-1}$ (in each $\mathbb{C}X_i$). \square

Corollary 4 *If $f \in \mathbb{C}X$, then the number $\omega(C(\mathcal{B}_n, \mathcal{B}_1))$ of arithmetic operations necessary to compute $[f]_{\mathcal{B}_n}$ when given $[f]_{\mathcal{B}_1}$ is strictly bounded above by*

$$2\Phi_X(G_n, G_{n-1}) + \dots + 2\Phi_X(G_2, G_1).$$

Proof Recall that $C(\mathcal{B}_n, \mathcal{B}_1)[f]_{\mathcal{B}_1} = [f]_{\mathcal{B}_n}$. The result then follows directly from the fact that

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) \leq \omega(C(\mathcal{B}_n, \mathcal{B}_{n-1})) + \dots + \omega(C(\mathcal{B}_2, \mathcal{B}_1))$$

and because by Theorem 3 we have that

$$\omega(C(\mathcal{B}_j, \mathcal{B}_{j-1})) < 2\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})) \leq 2\Phi_X(G_j, G_{j-1})$$

for all j such that $2 \leq j \leq n$. \square

3.3 Regular Representations

Before we apply the results above to the harmonic analysis of partially ranked data in the next section, we finish this section by considering the case where the permutation module in question is the regular representation of G .

Suppose H is a finite group, and that d_1, \dots, d_s are the dimensions of the irreducible representations of H . Define $d^3(H) = d_1^3 + \dots + d_s^3$, and let $[G : H]$ denote the index of H in G .

Lemma 5 *Let G be a finite group acting on $X = G$ by left multiplication so that $\mathbb{C}X = \mathbb{C}G$ is the regular $\mathbb{C}G$ -module. If H is a subgroup of G , then $\Phi_G(G, H) = [G : H]^2 d^3(H)$.*

Proof Let N_1, \dots, N_s be a complete collection of irreducible $\mathbb{C}H$ -modules. The orbits of G under the action of H are the right cosets of H in G . Therefore, as a $\mathbb{C}H$ -module, $\mathbb{C}G \cong \bigoplus_{j=1}^{[G:H]} \mathbb{C}H$. As a $\mathbb{C}H$ -module, $\mathbb{C}H \cong \bigoplus_{i=1}^s \dim(N_i)N_i$, so we therefore have that

$$\mathbb{C}G \cong [G : H] \dim(N_1)N_1 \oplus \dots \oplus [G : H] \dim(N_s)N_s.$$

The result now follows from Lemma 1 and the discussion after Lemma 2. □

By recursively applying Lemma 5 to a chain of subgroups, we get the following theorem, which is essentially Theorem 1.1 of [4], and which provides a general upper bound for applying a discrete Fourier transform to a finite group and therefore a bound for doing harmonic analysis on the regular representation of a finite group.

Theorem 6 *Let G be a finite group acting on $X = G$ by left multiplication so that $\mathbb{C}X = \mathbb{C}G$ is the regular $\mathbb{C}G$ -module. Let*

$$\{1\} = G_1 < \dots < G_n = G$$

be a chain of subgroups of G , and suppose $\mathcal{R}(G_1), \dots, \mathcal{R}(G_n)$ are compatible with respect to $\mathbb{C}X$. Let $\mathcal{B}_1, \dots, \mathcal{B}_n$ be harmonic bases of $\mathbb{C}X$, where \mathcal{B}_j is an orbital symmetry-adapted basis with respect to $\mathcal{R}(G_1), \dots, \mathcal{R}(G_j)$, and \mathcal{B}_1 is the standard basis of $\mathbb{C}X$. If $q_j = [G_j : G_{j-1}]$, then

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < 2 \sum_{j=2}^n (q_j^2 q_{j+1} \dots q_n d^3(G_{j-1})).$$

Proof By Theorem 3 and Lemma 5, the number of nonzero entries in $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ is bounded above by

$$\Phi_G(G_j, G_{j-1}) = [G : G_j][G_j : G_{j-1}]^2 d^3(G_{j-1}) = q_n \dots q_{j+1} q_j^2 d^3(G_{j-1}).$$

The result now follows from Corollary 4. □

4 The Symmetric Group Acting on Tabloids

We now shift our attention to the permutation modules that arise when the symmetric group acts on *tabloids* (see below for the definition). These objects have been used to index both fully and partially ranked data (see, for example, [7,8]). The results in this section may therefore be viewed as statements about the efficient analysis of such data. Good references for many of the ideas and much of the notation and terminology found in this section are [16,23,25].

4.1 Tabloids

Let n be a positive integer. A *weak composition* of n is a sequence $\alpha = (\alpha_1, \dots, \alpha_k)$ of nonnegative integers satisfying $\sum \alpha_i = n$. If each of the summands $\alpha_1, \dots, \alpha_k$ is a positive integer, then α is a *composition* of n . A *partition* of n is then a composition $\lambda = (\lambda_1, \dots, \lambda_k)$ of n such that $\lambda_1 \geq \dots \geq \lambda_k$. If $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition of n then we write $\lambda \vdash n$ and say that λ has k parts.

Fig. 1 Two equivalent tableaux of shape $(2, 3, 1, 3)$ and their tabloid

| | | |
|---|---|---|
| 2 | 7 | |
| 3 | 9 | 1 |
| 5 | | |
| 6 | 8 | 4 |

| | | |
|---|---|---|
| 7 | 2 | |
| 9 | 1 | 3 |
| 5 | | |
| 6 | 4 | 8 |

| | | |
|---|---|---|
| 2 | 7 | |
| 1 | 3 | 9 |
| 5 | | |
| 4 | 6 | 8 |

Fig. 2 The action of $\sigma = (147)(56)$ on a tabloid of shape $(2, 3, 1, 3)$

$$(147)(56) \cdot \begin{array}{|c|c|} \hline 2 & 7 \\ \hline 1 & 3 & 9 \\ \hline 5 & & \\ \hline 4 & 6 & 8 \\ \hline \end{array} = \begin{array}{|c|c|} \hline 2 & 1 \\ \hline 4 & 3 & 9 \\ \hline 6 & & \\ \hline 7 & 5 & 8 \\ \hline \end{array}$$

Note that a weak composition α can be identified with a unique partition $\bar{\alpha}$ by writing the positive summands in weakly decreasing order. For example, if $\alpha = (3, 4, 0, 1, 0, 2)$, then

$$\bar{\alpha} = (4, 3, 2, 1).$$

We will make use of this fact below.

Given a weak composition $\alpha = (\alpha_1, \dots, \alpha_k)$ of n , the *Young diagram of shape α* is the left-justified array of boxes with k rows and α_i boxes in the i th row. Filling these boxes with the numbers $1, \dots, n$, without repetition, creates a *Young tableau of shape α* . Two Young tableaux of shape α are then *row-equivalent* if they have the same set of numbers in each row. Each equivalence class of tableaux of shape α under this relation is a *tabloid of shape α* . We will use X^α to denote the set of all tabloids of shape α .

It is common to denote a tabloid by first forming a representative tableau and then removing the vertical dividers within each row (see Fig. 1).

4.2 The Action of S_n

We may view $\mathbb{C}X^\alpha$ as a $\mathbb{C}S_n$ -permutation module under the natural action of the symmetric group S_n on X^α , where if $\sigma \in S_n$ and $T \in X^\alpha$, then $\sigma \cdot T$ is the tabloid obtained from T by applying the permutation σ to each entry of T (see Fig. 2).

Note that the action of S_n on X^α does not depend on the order of the rows of α . In particular, as $\mathbb{C}S_n$ -modules, we have that $\mathbb{C}X^\alpha \cong \mathbb{C}X^{\bar{\alpha}}$. For convenience, we will therefore always assume that we are starting with a permutation module $\mathbb{C}X^\lambda$ where λ is a partition.

Fortunately, there are well-studied collections $\mathcal{R}(S_1), \dots, \mathcal{R}(S_n)$ of irreducible representations that are compatible with respect to $\mathbb{C}X^\lambda$. For example, two possible such collections are Young’s seminormal representations and Young’s orthogonal representations. In fact, these are the representations most often used when doing harmonic analysis on tabloids, and these are the ones we suggest using if the reader wishes to implement any of the ideas that follow.

4.3 Irreducible Representations

There is a well-known parametrization of the irreducible representations of S_n by the partitions of n . As is typically done, we use S^μ to denote the irreducible $\mathbb{C}S_n$ -module corresponding to $\mu \vdash n$. (S^μ is called a *Specht module*.) The multiplicity of S^μ in the decomposition of $\mathbb{C}X^\lambda$ into irreducible $\mathbb{C}S_n$ -modules is given by the well-studied but elusive *Kostka numbers*, which are defined below.

Definition 7 Given two partitions $\lambda = (\lambda_1, \dots, \lambda_k), \mu = (\mu_1, \dots, \mu_{k'})$ of n , we say that μ *dominates* λ , and write $\lambda \trianglelefteq \mu$, if

$$\sum_{i=1}^j \lambda_i \leq \sum_{i=1}^j \mu_i$$

for all j such that $1 \leq j \leq \max(k, k')$. If $\lambda \trianglelefteq \mu$, then a *Kostka filling of μ by λ* is any filling of a tableau of shape μ using exactly λ_i i 's such that the entries in each row, when read from left to right, are non-decreasing and the entries of each column, when read from top to bottom, are strictly increasing.

Example 8 If $\lambda = (3, 1, 1)$ and $\mu = (4, 1)$, then $\lambda \trianglelefteq \mu$ and there are two Kostka fillings of μ by λ , namely

$$\begin{array}{|c|c|c|c|} \hline 1 & 1 & 1 & 2 \\ \hline 3 & & & \\ \hline \end{array} \quad \text{and} \quad \begin{array}{|c|c|c|c|} \hline 1 & 1 & 1 & 3 \\ \hline 2 & & & \\ \hline \end{array}$$

Let $\kappa_{\mu,\lambda}$ denote the number of Kostka fillings of μ by λ . These are the *Kostka numbers*, and Young's rule (see, for example, Theorem 2.11.2 in [23]) states that as a $\mathbb{C}S_n$ -module,

$$\mathbb{C}X^\lambda \cong \bigoplus_{\mu: \lambda \trianglelefteq \mu} \kappa_{\mu,\lambda} S^\mu. \tag{1}$$

In order to make use of Theorem 3 when dealing with $\mathbb{C}X^\lambda$, we need to find suitable bounds for $K_{X^\lambda}(S_j, S_{j-1})$ for each $2 \leq j \leq n$. Recall that if the orbits of X^λ under the action of S_j are X_1, \dots, X_t , then $K_{X^\lambda}(S_j, S_{j-1})$ is the maximum over all of the $\mathbb{C}X_i$ of the dimensions of their (S_j, S_{j-1}) -orbital frequency spaces. Combining an understanding of the orbits of X_i under the action of S_{j-1} , and decomposing the resulting spaces into irreducible $\mathbb{C}S_{j-1}$ -modules according to Young's rule will therefore enable us to determine the maximum frequency space dimension $K_{X^\lambda}(S_j, S_{j-1})$. We begin by considering the case where $j = n$, and then proceed recursively.

4.4 Orbits Under the Action of $\mathbb{C}S_{n-1}$

Suppose $\lambda \vdash n$. The orbits of S_{n-1} acting on X^λ are given by the ways to obtain a weak composition of $n - 1$ from λ by subtracting 1 from a nonzero part λ_i of λ . To describe this more formally, we introduce the following notation.

Definition 9 If $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition, and $\alpha = (\lambda_1, \dots, \lambda_{i-1}, \lambda_i - 1, \lambda_{i+1}, \dots, \lambda_k)$, then we will define λ^i to be the partition $\bar{\alpha}$.

Example 10 Let $\lambda = (4, 2, 1, 1)$. Then $\lambda^1 = (3, 2, 1, 1)$, $\lambda^2 = (4, 1, 1, 1)$, and $\lambda^3 = \lambda^4 = (4, 2, 1)$.

Theorem 11 If $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition of n , then as a $\mathbb{C}S_{n-1}$ -module,

$$\mathbb{C}X^\lambda \cong \bigoplus_{i=1}^k \mathbb{C}X^{\lambda^i}.$$

Proof Two tabloids $T, T' \in X^\lambda$ are in the same orbit under the action of S_{n-1} if and only if they contain n in the same row. Furthermore, the action of S_{n-1} on the orbit of a tabloid T that contains n in its i th row corresponds exactly to the action of S_{n-1} on X^{λ^i} . It follows that, as a $\mathbb{C}S_{n-1}$ -module, $\mathbb{C}X^\lambda \cong \bigoplus_{i=1}^k \mathbb{C}X^{\lambda^i}$. \square

4.5 The Case $\lambda = (n - k, k)$

The action of the symmetric group on two-rowed tabloids gives rise to a permutation module with a particularly simple decomposition into irreducible modules using Young’s rule. As noted in Sect. 1, these tabloids arise when dealing with survey data for which respondents have been asked simply to choose their top k items from a set of n items, where $k \leq n/2$.

Lemma 12 If $\lambda = (n - k, k)$ is a partition of n , then as a $\mathbb{C}S_n$ -module,

$$\mathbb{C}X^\lambda \cong \bigoplus_{0 \leq l \leq k} S^{(n-l, l)}.$$

In particular, this is a multiplicity-free decomposition.

Proof By Young’s rule (1), we are only concerned with S^μ such that $\lambda \trianglelefteq \mu$. Since $\lambda = (n - k, k)$, we see that μ must have the form $(n - l, l)$ for $0 \leq l \leq k$. The multiplicity of S^μ is given by $\kappa_{\mu, \lambda}$, which is the number of Kostka fillings of $\mu = (n - l, l)$ with $(n - k)$ 1’s and k 2’s. To be a Kostka filling, all $(n - k)$ 1’s must be in the first row of μ , completely determining the filling. Thus $\kappa_{\mu, \lambda} = 1$. \square

Together with Theorems 3 and 11, the simple decomposition of $\mathbb{C}X^{(n-k, k)}$ into irreducible modules in Lemma 12 gives rise to an efficient approach to doing a change of basis from the standard basis \mathcal{B}_1 to a harmonic basis \mathcal{B}_n , especially when compared to the naive bound of $\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < 2\binom{n}{k}^2$.

Theorem 13 Suppose $\lambda = (n - k, k)$ is a partition of n , and that $\mathcal{R}(S_1), \dots, \mathcal{R}(S_n)$ are compatible with respect to $\mathbb{C}X^\lambda$. Let $\mathcal{B}_1, \dots, \mathcal{B}_n$ be harmonic bases of $\mathbb{C}X^\lambda$ where \mathcal{B}_j is an orbital symmetry-adapted basis with respect to $\mathcal{R}(S_1), \dots, \mathcal{R}(S_j)$, and \mathcal{B}_1

is the standard basis of $\mathbb{C}X^\lambda$. Then

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < 4(n - 1) \binom{n}{k}.$$

Proof Let $2 \leq j \leq n$. The action of S_j on X^λ partitions X^λ into orbits $X^{\lambda'}$ that correspond to tabloids with at most two rows and with entries from $\{1, \dots, j\}$. With the use of Theorem 11 and Lemma 12, we see that restricting the action further to S_{j-1} on $X^{\lambda'}$ gives rise to a decomposition of the $\mathbb{C}X^{\lambda'}$ into irreducible $\mathbb{C}S_{j-1}$ -modules each with multiplicity no more than two. It follows that

$$K_{X^\lambda}(S_j, S_{j-1}) \leq 2.$$

By Theorem 3, we have that

$$\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})) \leq K_{X^\lambda}(S_j, S_{j-1})|X^\lambda| \leq 2 \binom{n}{k}.$$

The theorem statement now follows from the fact that

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < \sum_{j=2}^n 2\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})).$$

□


Remark Theorem 13 is essentially Theorem 2 in [15]. In that paper, the authors use Gelfand–Tsetlin bases, which are a special case of what we are calling harmonic bases, and the bound they provide is $2(n - 1) \binom{n}{k}$. It differs from our bound of $4(n - 1) \binom{n}{k}$ because the computational model they use counts a single complex multiplication and addition as one operation.

4.6 The Case $\lambda = (n - k, 1, \dots, 1)$

We now consider the situation where we are dealing with partitions of n of the form $\lambda = (n - k, 1, \dots, 1)$. As noted in Sect. 1, tabloids of this shape arise when dealing with survey data for which respondents have been asked to rank their top k items from a set of n items. Such data is called *hook data* in [22], in which a detailed analysis of such a data set is given, together with a discussion of the underlying representation theory.

As we saw when $\lambda = (n - k, k)$, the two main steps in determining bounds on the number of nonzero entries in a full factorization of the change of basis matrix $C(\mathcal{B}_n, \mathcal{B}_1)$ involve decomposing X^λ into orbits $X^{\lambda'}$ under the action of S_j , and then decomposing the resulting permutation modules $\mathbb{C}X^{\lambda'}$ into irreducible $\mathbb{C}S_{j-1}$ -submodules. Bounds then come from the maximum dimensions of the

Fig. 3 Two standard skew tableaux of shape μ/λ where $\mu = (5, 4, 2)$ and $\lambda = (3, 1)$



(S_j, S_{j-1}) -orbital frequency spaces, which amounts to summing Kostka numbers across certain S_{j-1} -orbits.

Though the Kostka numbers are well-studied, Stanley remarks in [25] that it is unlikely that a general formula for $\kappa_{\mu,\lambda}$ exists. Fortunately, when $\lambda = (n - k, k)$, the Kostka number $\kappa_{\mu,\lambda}$ is trivially 1 for each μ dominating λ . When $\lambda = (n - k, 1, \dots, 1)$, this is no longer the case. However, the orbits in this case still exhibit enough structure for us to find bounds for the multiplicities. We use the language of *skew tableaux* to describe the orbits (see, for example, [16]).

Let $\lambda = (\lambda_1, \dots, \lambda_k)$ and $\mu = (\mu_1, \dots, \mu_\ell)$ be partitions. The partition μ contains λ if $k \leq \ell$ and $\lambda_i \leq \mu_i$ for all i such that $1 \leq i \leq k$. If μ contains λ , then the *skew diagram of shape μ/λ* is the set of boxes in the Young diagram of shape μ that are not in the Young diagram of shape λ . If n is the number of boxes in the skew diagram of shape μ/λ , then a *standard skew tableau of shape μ/λ* is a filling of the boxes of the skew diagram with the numbers $1, \dots, n$, without repetition, so that the entries in each row and each column are strictly increasing (see Fig. 3).

Let $d_{\mu/\lambda}$ denote the number of standard skew tableaux of shape μ/λ . For convenience, if μ does not contain λ , then we will set $d_{\mu/\lambda} = 0$. For positive integers k and n such that $k \leq n$, let $(n)_k$ denote the falling factorial

$$(n)_k = n(n - 1)(n - 2) \cdots (n - (k - 1)).$$

Note that if $\lambda = (n - k, 1, \dots, 1)$ is a partition of n , then $\dim \mathbb{C}X^\lambda = (n)_k$.

For nonnegative integers ℓ and r such that $\ell \geq r$, let $D(\ell, r) = \{\mu \mid \mu \vdash \ell, \mu_1 \geq r\}$. Let

$$M(n, k) = \max_{\mu \in D(n-1, n-k-1)} \{d_{\mu/(n-k-1)} + kd_{\mu/(n-k)}\}.$$

Note that if μ has first part $\mu_1 = n - k - 1$, the sum reduces to $d_{\mu/(n-k-1)}$.

As $\mu/(n - k - 1)$ represents the skew diagram resulting from removing $n - k - 1$ of the boxes in the first row of μ , and similarly for $\mu/(n - k)$, for large enough n , the value of $M(n, k)$ depends only on k . For example, $M(7, 6) = 112$, $M(8, 6) = 155$, and $M(n, 6) = 160$ for all $n > 8$. Let $N(k)$ denote the maximum value of $M(n, k)$ over all positive integers n such that $n > k$.

Theorem 14 *Let $\lambda = (n - k, 1, \dots, 1) \vdash n$ and let $\mathcal{R}(S_1), \dots, \mathcal{R}(S_n)$ be compatible with respect to $\mathbb{C}X^\lambda$. Let $\mathcal{B}_1, \dots, \mathcal{B}_n$ be harmonic bases of $\mathbb{C}X^\lambda$ where \mathcal{B}_j is an orbital symmetry-adapted basis with respect to $\mathcal{R}(S_1), \dots, \mathcal{R}(S_j)$, and \mathcal{B}_1 is the standard basis of $\mathbb{C}X^\lambda$. Then*

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < 2N(k)(n - 1)(n)_k$$

as opposed to the naive bound of $2((n)_k)^2$.

In Table 1, we compare the bound of Theorem 14 to the naive bound.

Proof Let $\lambda = (n - k, 1, \dots, 1) \vdash n$. By Theorem 11, as a $\mathbb{C}S_{n-1}$ -module,

$$\mathbb{C}X^\lambda \cong \bigoplus_{i=1}^{k+1} \mathbb{C}X^{\lambda^i} \cong \mathbb{C}X^{\lambda^1} \oplus k\mathbb{C}X^{\lambda^2}$$

where $\lambda^1 = (n - k - 1, 1, \dots, 1) \vdash (n - 1)$ and $\lambda^2 = \lambda^3 = \dots = \lambda^{k+1} = (n - k, 1, \dots, 1) \vdash (n - 1)$. Using Young’s rule (1) to decompose $\mathbb{C}X^{\lambda^1}$ as a $\mathbb{C}S_{n-1}$ -module we first note that the set of partitions of $n - 1$ that dominate λ^1 is given by $D(n - 1, n - k - 1)$.

Suppose $\mu \in D(n - 1, n - k - 1)$. To determine the multiplicity of S^μ in the decomposition of $\mathbb{C}X^{\lambda^1}$ as a $\mathbb{C}S_{n-1}$ -module, we need to determine κ_{μ, λ^1} . This is the number of Kostka fillings of μ with $(n - k - 1)$ 1’s, along with the entries $2, \dots, k + 1$. In order to be a valid Kostka filling, this means that the first $n - k - 1$ entries of the first row of μ must be a 1, leaving exactly one of each number $2, \dots, k + 1$ to fill the remaining k boxes of μ . This is exactly the number of standard skew tableaux of shape $\mu / (n - k - 1)$. Thus, as a $\mathbb{C}S_{n-1}$ -module,

$$\mathbb{C}X^{\lambda^1} \cong \bigoplus_{\mu \in D(n-1, n-k-1)} d_{\mu / (n-k-1)} S^\mu.$$

As $\lambda^2 = (n - k, 1, \dots, 1) \vdash (n - 1)$, we see by the same argument that each of the k copies of $\mathbb{C}X^{\lambda^2}$ decomposes as a $\mathbb{C}S_{n-1}$ -module as

$$\mathbb{C}X^{\lambda^2} \cong \bigoplus_{\mu \in D(n-1, n-k)} d_{\mu / (n-k)} S^\mu.$$

Then in the complete decomposition of $\mathbb{C}X^\lambda$ into irreducible $\mathbb{C}S_{n-1}$ -modules, only those S^μ with $\mu \in D(n - 1, n - k - 1)$ will appear, and they appear with multiplicity

$$d_{\mu / (n-k-1)} + kd_{\mu / (n-k)}.$$

Note that if $\mu_1 = (n - k - 1)$, this sum reduces to $d_{\mu / (n-k-1)}$, as $d_{\mu / (n-k)} = 0$. By Theorem 3, the number of nonzero entries in each column of the change of basis matrix $C(\mathcal{B}_n, \mathcal{B}_{n-1})$ is bounded by $K_{X^\lambda}(S_n, S_{n-1}) = M(n, k)$.

Now consider the change of basis matrix $\mathcal{C}(\mathcal{B}_j, \mathcal{B}_{j-1})$. By repeated applications of Theorem 11, as a $\mathbb{C}S_j$ -module,

$$\mathbb{C}X^\lambda \cong \bigoplus \mathbb{C}X^{\lambda'}$$

where the λ' are a collection of (possibly repeated) partitions of j , each of which has form $\lambda' = (j - m, 1, \dots, 1)$ for some nonnegative integer m , where $k - (n - j) \leq m \leq k$.

By Theorem 3, we need only determine

$$K_{X^\lambda}(S_j, S_{j-1}) = \max_{\lambda'} K_{X^{\lambda'}}(S_j, S_{j-1}).$$

Applying the same arguments as above to decompose $\mathbb{C}X^{\lambda'}$ into irreducible $\mathbb{C}S_{j-1}$ -submodules using Theorem 11 and Young’s rule (1), we see that the number of nonzero entries in each column of the change of basis matrix $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ is bounded by $M(j, m)$, where

$$M(j, m) = \max_{\mu \in D(j-1, j-m-1)} \{d_{\mu/(j-m-1)} + md_{\mu/(j-m)}\}.$$

Suppose $\mu = (\mu_1, \dots, \mu_s) \in D(j-1, j-m-1)$. Then $\mu \vdash j-1$ and $\mu_1 = j-m-1+t$ for some t such that $0 \leq t \leq m$. The skew diagram $\mu/(j-m-1)$ has the same shape as μ for all but its first row and has t boxes in its first row.

Let $\mu' = (\mu_1 + n - j, \mu_2, \dots, \mu_s)$. Then $\mu' \vdash n-1$ and $\mu'_1 = n-m-1+t \geq n-k-1+t$, so $\mu' \in D(n-1, n-k-1)$. The skew diagram $\mu'/(n-k-1)$ has the same shape as μ for all but its first row and has $t+k-m \geq t$ boxes its first row. Moreover, as $n-k-1 \geq j-m-1$, the first box in this row is either in the same position or to the right of the first box in the first row of $\mu/(j-m-1)$.

A standard skew tableau of shape $\mu/(j-m-1)$ yields a standard skew tableau of shape $\mu'/(n-k-1)$ as follows: fill the boxes in rows μ_2, \dots, μ_s the same way as in $\mu/(j-m-1)$ and fill the first t boxes of the first row of $\mu'/(n-k-1)$ as they are filled in the first row of $\mu/(j-m-1)$. If boxes remain in the first row of $\mu'/(n-k-1)$, fill them with the numbers that remain, in ascending order from left to right.

This implies that there are at least as many standard skew tableau of shape $\mu'/(n-k-1)$ as there are of shape $\mu/(j-m-1)$. Thus, $d_{\mu/(j-m-1)} \leq d_{\mu'/(n-k-1)}$. Similarly, $d_{\mu/(j-m)} \leq d_{\mu'/(n-k)}$ and so

$$d_{\mu/(j-m-1)} + md_{\mu/(j-m)} \leq d_{\mu'/(n-k-1)} + kd_{\mu'/(n-k)}$$

implying that $M(j, m) \leq M(n, k) \leq N(k)$. By Theorem 3, we have that

$$\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})) \leq K_{X^\lambda}(S_j, S_{j-1})|X^\lambda| \leq N(k)(n)_k.$$

The theorem statement now follows from the fact that

$$\omega(C(\mathcal{B}_n, \mathcal{B}_1)) < \sum_{j=2}^n 2\nu(C(\mathcal{B}_j, \mathcal{B}_{j-1})).$$

□

5 Conclusion and Open Questions

In this paper, we have developed a framework for computing the coefficients of $f \in \mathbb{C}X$ in terms of a harmonic basis \mathcal{B} of $\mathbb{C}X$ using intermediate bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ and

Table 1 Comparison of bound given by Theorem 14 with naive bound

| k | $N(k)(n-1)(n)_k$ | |
|-----|-----------------------|----------------|
| 2 | $4(n-1)(n)_2$ | $((n)_2)^2$ |
| 3 | $9(n-1)(n)_3$ | $((n)_3)^2$ |
| 4 | $18(n-1)(n)_4$ | $((n)_4)^2$ |
| 5 | $60(n-1)(n)_5$ | $((n)_5)^2$ |
| 6 | $160(n-1)(n)_6$ | $((n)_6)^2$ |
| 7 | $420(n-1)(n)_7$ | $((n)_7)^2$ |
| 8 | $1344(n-1)(n)_8$ | $((n)_8)^2$ |
| 9 | $5376(n-1)(n)_9$ | $((n)_9)^2$ |
| 10 | $16800(n-1)(n)_{10}$ | $((n)_{10})^2$ |
| 11 | $59400(n-1)(n)_{11}$ | $((n)_{11})^2$ |
| 12 | $222750(n-1)(n)_{12}$ | $((n)_{12})^2$ |
| 13 | $878592(n-1)(n)_{13}$ | $((n)_{13})^2$ |

iteratively computing a change of basis from \mathcal{B}_{j-1} to \mathcal{B}_j . In Sect. 3, we saw that when \mathcal{B}_1 is the standard basis and \mathcal{B}_j is an orbital symmetry-adapted basis with respect to compatible representations $\mathcal{R}(G_1), \dots, \mathcal{R}(G_j)$, we can bound the number of nonzero entries in the change of basis matrix $C(\mathcal{B}_j, \mathcal{B}_{j-1})$ in terms of the multiplicities of the irreducible $\mathbb{C}G_{j-1}$ -submodules in the orbital decomposition of $\mathbb{C}X$ under the action of G_j .

In Sect. 4, we applied these results to permutation modules that arise when the symmetric group acts on a set of tabloids X^λ . In particular, for both $\lambda = (n-k, k) \vdash n$ and $\lambda = (n-k, 1, \dots, 1) \vdash n$, we provided a coarse bound based only on bounding the largest number of nonzero entries per column in $C(\mathcal{B}_n, \mathcal{B}_{n-1})$. Despite the coarseness of these bounds, they greatly improve on the naive bound, which is $2(\dim(\mathbb{C}X^\lambda))^2$ computations (see Table 1).

In truth, both results could be refined further by bounding the number of nonzero entries in each column of each change of basis matrix, rather than focusing on the maximum possible number of nonzero entries in a column over all change of basis matrices $C(\mathcal{B}_j, \mathcal{B}_{j-1})$. Indeed, the proof of Theorem 14 leads to a more refined bound than the theorem statement itself. The proof provides a bound for the number of nonzero entries in each individual column of $C(\mathcal{B}_n, \mathcal{B}_{n-1})$. Summing over the columns leads to a bound of at most

$$\sum_{\mu \in D(n-1, n-k-1)} (d_{\mu/(n-k-1)} + kd_{\mu/(n-k)}) \dim(S^\mu)$$

nonzero entries in $C(\mathcal{B}_n, \mathcal{B}_{n-1})$. Continuing in this manner yields similar looking bounds for the number of nonzero entries in each change of basis matrix $C(\mathcal{B}_j, \mathcal{B}_{j-1})$. Preliminary numerical results suggest that these more refined bounds are worth investigating. Indeed, for small values of n and k they differ by only about a factor of 2 from the bound given in the recent algorithm of Clausen and Hühne [6], which uses

detailed knowledge about the appearance of repeated entries in the corresponding representation matrices.

It is also worth noting that the results of Sect. 3 apply for any set of harmonic bases $\mathcal{B}_1, \dots, \mathcal{B}_n$ where \mathcal{B}_1 is the standard basis and \mathcal{B}_j is an orbital symmetry-adapted basis with respect to a compatible collection $\mathcal{R}(G_1), \dots, \mathcal{R}(G_j)$ of irreducible representations. Is there a particular choice for these bases that can lead to a more efficient change of basis computation? Algorithms exist for constructing symmetry-adapted bases (see, for example, [13]) but at several key steps in these algorithms there is a degree of choice. As some choices could lead to more efficient bounds than others, it would be interesting to refine our bounds further by streamlining these basis constructions from the perspective of creating sparse change of basis matrices.

Finally, it would be interesting to explore the extent to which the change of basis algorithm presented in this paper might be helpful when dealing with other computational problems associated with permutation modules (see, for example, [11,20]).

References

1. Babai, L., Rónyai, L.: Computing irreducible representations of finite groups. *Math. Comput.* **55**(192), 705–722 (1990)
2. Bürgisser, P., Clausen, M., Shokrollahi, M.A.: Algebraic Complexity Theory. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 315. Springer, Berlin (1997). With the collaboration of T. Lickteig
3. Ceccherini-Silberstein, T., Scarabotti, F., Tolli, F.: Harmonic Analysis on Finite Groups. *Cambridge Studies in Advanced Mathematics*, vol. 108. Cambridge University Press, Cambridge (2008). Representation theory, Gelfand pairs and Markov chains
4. Clausen, M.: Fast generalized Fourier transforms. *Theor. Comput. Sci.* **67**(1), 55–63 (1989)
5. Clausen, M., Baum, U.: Fast Fourier Transforms. *Bibliographisches Institut, Mannheim* (1993)
6. Clausen, M., Hühne, P.: Linear time Fourier transforms of S_{n-k} -invariant functions on the symmetric group S_n . In: *ISSAC'17—Proceedings of the 2017 ACM International Symposium on Symbolic and Algebraic Computation*, pp. 101–108. ACM, New York (2017)
7. Crisman, K.-D., Orrison, M.E.: Representation theory of the symmetric group in voting theory and game theory. In: *Algebraic and Geometric Methods in Discrete Mathematics Contemporary Mathematics*, vol. 685, pp. 97–115. American Mathematical Society, Providence (2017)
8. Daugherty, Z., Eustis, A.K., Minton, G., Orrison, M.E.: Voting, the symmetric group, and representation theory. *Am. Math. Mon.* **116**(8), 667–687 (2009)
9. Diaconis, P.: Group Representations in Probability and Statistics. *Institute of Mathematical Statistics Lecture Notes—Monograph Series*, vol. 11. Institute of Mathematical Statistics, Hayward (1988)
10. Diaconis, P.: A generalization of spectral analysis with application to ranked data. *Ann. Stat.* **17**(3), 949–979 (1989)
11. Diaconis, P., Rockmore, D.: Efficient computation of isotypic projections for the symmetric group. In: *Groups and Computation* (New Brunswick, NJ, 1991). DIMACS: Series in Discrete Mathematics and Theoretical Computer Science, vol. 11, pp. 87–104. American Mathematical Society, Providence (1993)
12. Drozd, Y.A., Kirichenko, V.V.: *Finite-Dimensional Algebras*. Springer, Berlin (1994). Translated from the 1980 Russian original and with an appendix by V. Dlab
13. Fässler, A., Stiefel, E.: *Group Theoretical Methods and Their Applications*. Birkhäuser Boston, Boston (1992). Translated from the German by B.D. Wong
14. Gross, K.I.: On the evolution of noncommutative harmonic analysis. *Am. Math. Mon.* **85**(7), 525–548 (1978)
15. Iglesias, R., Natale, M.: Complexity of the Fourier transform on the Johnson graph. arXiv e-prints (2017). [arXiv:1704.06299](https://arxiv.org/abs/1704.06299)

16. James, G., Kerber, A.: The Representation Theory of the Symmetric Group. *Encyclopedia of Mathematics and Its Applications*, vol. 16. Addison-Wesley Publishing Co., Reading (1981). With a foreword by P.M. Cohn, With an introduction by G.B. Robinson
17. Marden, J.I.: *Analyzing and Modeling Rank Data*. Monographs on Statistics and Applied Probability, vol. 64. Chapman & Hall, London (1995)
18. Maslen, D.K.: The efficient computation of Fourier transforms on the symmetric group. *Math. Comput.* **67**(223), 1121–1147 (1998)
19. Maslen, D.K., Rockmore, D.N.: Separation of variables and the computation of Fourier transforms on finite groups. I. *J. Am. Math. Soc.* **10**(1), 169–214 (1997)
20. Maslen, D.K., Orrison, M.E., Rockmore, D.N.: Computing isotypic projections with the Lanczos iteration. *SIAM J. Matrix Anal. Appl.* **25**(3), 784–803 (2003)
21. Maslen, D., Rockmore, D.N., Wolff, S.: Separation of variables and the computation of Fourier transforms on finite groups. II. *J. Fourier Anal. Appl.* **24**(1), 226–284 (2018)
22. Rockmore, D.N.: Some applications of generalized FFTs. In: *Groups and Computation, II* (New Brunswick, NJ, 1995). DIMACS: Series in Discrete Mathematics and Theoretical Computer Science, vol. 28, pp. 329–369. American Mathematical Society, Providence (1997)
23. Sagan, B.E.: *The Symmetric Group*. Graduate Texts in Mathematics, vol. 203, 2nd edn. Springer, New York (2001). Representations, combinatorial algorithms, and symmetric functions
24. Serre, J.-P.: *Linear Representations of Finite Groups*. Springer, New York (1977). Translated from the 2nd French edition by L.L. Scott, Graduate Texts in Mathematics, vol. 42
25. Stanley, R.P.: *Enumerative Combinatorics*, vol. 2. Cambridge Studies in Advanced Mathematics, vol. 62. Cambridge University Press, Cambridge (1999). With a foreword by G.-C. Rota and Appendix 1 by S. Fomin
26. Terras, A.: *Fourier Analysis on Finite Groups and Applications*. London Mathematical Society Student Texts, vol. 43. Cambridge University Press, Cambridge (1999)
27. Van Loan, C.: *Computational Frameworks for the Fast Fourier Transform*, vol. 10. *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia (1992)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.