Fast Fourier Transforms for Wreath Products *

Daniel N. Rockmore
Department of Mathematics and Computer Science
Dartmouth College
Hanover, NH 03755
rockmore@gilligans.dartmouth.edu

March 14, 1994

Abstract

In this paper fast Fourier transform algorithms (FFT's) are constructed for wreath products of the form $G[S_n]$. Examples of interest include the hyperoctahedral groups B_n (the symmetry groups of the n-cube) as well as more generally $A[S_n]$ for any abelian groups A and also the full wreath products $S_m[S_n]$ and their iterates, often identified as automorphism groups of spherically homogeneous rooted trees.

In general, direct computation of the discrete Fourier transform for any finite group H requires $|H|^2$ operations. The algorithms presented in this paper provide substantial speed-ups for general wreath products $G[S_n]$, which for the particular examples mentioned above reduce the $|H|^2$ bound to $O(|H| \log^4 |H|)$. Thus new classes of finite groups of close to minimal linear complexity are obtained. Applications to data analysis, signal processing and molecular spectroscopy are discussed.

1 Introduction

The fast Fourier transform (FFT) is arguably one of the most important classes of algorithms ever developed. It is fundamental to a tremendous range of computational applications in areas ranging from signal processing to fast computer arithmetic. The books [19, 35, 36] are a few good sources for many of the relevant techniques and applications of the FFT.

Group representation theory turns out to be a very useful framework for analyzing the FFT. A discrete Fourier transform (DFT) can be defined for any finite group G and requires $|G|^2$ operations if computed directly (cf. Section 2). In this setting, the now "classical" Cooley-Tukey FFT [12] and its variants are a family of algorithms which compute the DFT for any abelian group G in $O(|G| \log |G|)$ operations, where an operation is understood to be a complex multiplication followed by an addition.

Currently, in the context of various applications to data analysis and signal processing (cf. [9, 14] and the references therein), the problem of finding FFT's for nonabelian groups is being studied extensively. General techniques now exist for reducing the complexity of the

*This work supported in part by ARPA as administered by the AFOSR under contract DOD F4960-93-1-0567 as well as an NSF Math Sciences Postdoctoral Fellowship.

DFT for any finite group [15, 11, 7] with dramatic speed-ups possible for several classes of nonabelian groups. In particular, it is now known that many families of nonabelian groups also possess $O(|G|\log|G|)$ FFT's. This includes the symmetric groups [11], metabelian groups [10, 33] and more generally, supersolvable groups [5]. Baum and Clausen's recent book [9] is a good source for some of the recent work in this area. Other developments in the finite setting include new efficient algorithms for computing on quotients [16] as well as with sets of orthogonal polynomial [18, 30].

There has also been very recent work in the setting of continuous groups. Healy and Driscoll have discovered an FFT for band-limited functions on the 2-sphere [17, 22]. Maslen has been able to extend this to more general compact groups [29]. This work is already finding applications in areas such as geophysics and weather modelling.

In this paper FFT's are developed for finite groups obtained as wreath products by the symmetric group S_n (cf. Section 4). These algorithms use classical results from the representation theory of finite groups as well as some of the newer algorithmic techinques developed for FFT's on S_n [15, 11]. In particular, they depend on the construction of certain "adapted bases" for the irreducible representations of these groups (Section 3.2, Theorem 1).

The algorithms developed here have potential applications to data analysis, molecular spectroscopy and matched filter design. Wreath products are of interest in data analysis as the automorphism groups of nested designs. FFT's for wreath products will greatly facilitate the computation of the analysis of variance of data on a nested design for statistical models which assume symmetry invariant covariance [1, 21]. Similarly, efficient computation of spectral analysis for such data requires such FFT's [14]. Wreath products are also the symmetry groups of many classes of molecules - particularly so-called "non-rigid" molecules. Here the development of the adapted bases necessary for FFT's turns out to be useful for efficient computation for NMR spectroscopy [4]. Lastly, nonabelian FFT's have been proposed as potential approximants to Karhunen-Loeve filters [23, 24, 26]. By exhibiting a new class of nonabelian groups with close to optimal DFT algorithms, new possibilities for signal estimation and data compression are displayed.

The organization of this paper is as follows. Section 2 recalls the background and definitions relevant to construction of FFT's. The algorithms presented here depend on a use of the Mackey Subgroup Theorem which is reproved here as well (Section 2.3). Section 3 recalls the representation theory of wreath products and then constructs the suitably adapted bases for wreath products $G[S_n]$. These bases are building blocks for the main results of this paper, the FFT algorithms given in Section 4. Section 5 closes by explaining in some detail the applications of this work to data analysis as well as outlining the applications to molecular spectroscopy and signal processing.

Acknowledgement. Special thanks to Tom Sundquist for providing a patient and helpful sounding board during the evolution of these ideas.

2 Preliminaries

2.1 Wreath products

Let G be a group and $L < S_n$. Then L acts on G^n by translation, defined as

$$a^{\pi}(j) = a(\pi^{-1}(j))$$

for $a \in G^n$ and $\pi \in L$.

The **wreath product of** G **with** L, denoted as $G[L]$ is the set $G^n \times L$ with multiplication given by

$$[a; \pi] \cdot [b; \sigma] = [a \cdot b^{\pi}; \pi\sigma].$$

Let e denote the identity in G^n. The groups G^n and L are identified with the subgroups $\{[a; 1] \mid a \in G^n\}$ and $\{[e; \pi] \mid \pi \in L\}$ respectively. It is easy to check that G^n is a normal subgroup of $G[L]$ so that $G[L]$ is a semidirect product of G^n by L, denoted as $G^n \rtimes L$.

Let $S_{n-1} \times S_1 < S_n$ denote the subgroup of permutations which fix n. Then of particular importance will be the subgroup $G[L \cap S_{n-1} \times S_1] < G[L]$. Denote by G_i the copy of G in the i^{th} coordinate of $G^n < G[L]$, then $G[L \cap S_{n-1} \times S_1]$ may be identified with the subgroup $G[L \cap S_{n-1}] \times G_n < G[L]$. [1] Thus gives the partial tower of subgroups

$$G[L] > G[L \cap S_{n-1}] \times G > G[L \cap S_{n-1}].$$

For a thorough treatment of their combinatorial structure and generalizations of the wreath product construction see Wells' paper [37] as well as Kerber's monographs [27].

2.2 DFT's and FFT's

There are several equivalent definitions of the discrete Fourier transform for a finite group [9, 7, 15]. The following is the most convenient for this paper. Recall that the dual of G, denoted $\hat{G}$ is the collection of equivalence classes of irreducible complex matrix representations of G. Thus, a traversal of $\hat{G}$ is a complete set of inequivalent irreducible matrix representations.

Definition 1 *Let G be a finite group, and let $L(G)$ denote the $|G|$-dimensional complex vector space of functions defined on G. If ρ is a matrix representation G then the Fourier transform of f at ρ, denoted $\hat{f}(\rho)$ is the matrix sum,*

$$\hat{f}(\rho) = \sum_{s \in G} f(s)\rho(s).$$

Let $\mathcal{R}$ be a traversal of $\hat{G}$ (ie. $\mathcal{R}$ is a complete set of inequivalent complex irreducible matrix representations of G). The Discrete Fourier Transform (with respect to $\mathcal{R}$), denoted $DFT_{\mathcal{R}}(f)$ is the collection of individual Fourier transforms

$$DFT_{\mathcal{R}}(f) = \{\hat{f}(\rho) | \rho \in \mathcal{R}\}.$$

[1] Note that $G[L \cap S_{n-1}]$ is a subgroup of the wreath product $G[S_{n-1}]$, which is naturally identified with the subgroup of $G[S_n]$ under the identification of $S_{n-1} < S_n$ with the subgroup of elements fixing n. The context should make the various distinctions clear.

For a traversal $\mathcal{R}$ of $\hat{G}$ let $T(G,\mathcal{R})$ denote the minimum number of operations (defined to be a complex multiplication followed by an addition) needed to compute $DFT_{\mathcal{R}}(f)$ via a straight-line program for an arbitrary $f \in L(G)$. Similarly, let

$$T(G) = min_{\mathcal{R}}\{T(G,\mathcal{R})\}.$$

Remark 1. Another common interpretation of the DFT is as a change of basis for $L(G)$, from the basis of point masses on G, to a basis of matrix coefficients spanning $\hat{G}$. When this approach is adopted, the complexity of the DFT can be measured as the **c-linear complexity** of the DFT matrix [6]. The **c-linear complexity of a group** G, is defined to be the minimum c-linear complexity of any DFT matrix for G. The results stated here can all be translated into statements about the 2-linear complexity of wreath products. Thus Theorem 3 of this paper show that wreath products by symmetric groups have $O(|G|\log^4|G|)$ 2-linear complexity.

Fast Fourier transforms (FFT's) are algorithms for computing the DFT efficiently. One of the most common techniques for constructing such algorithms is through the use of **(subgroup)-adapted bases**.

Definition 2 *Let $H < G$, be a subgroup. Let $\mathcal{S} = \{\eta_1,\ldots,\eta_l\}$ and $\mathcal{R} = \{\rho_1,\ldots,\rho_h\}$ be sets of matrix representations for H and G respectively. Then $(G,\mathcal{R})$ is said to be $(H,\mathcal{S})$-* ***adapted****, or $\mathcal{R}$ is said to be $\mathcal{S}$-***adapted***, if for all $i = 1,\ldots,h$ and $y \in H$,*

$$\rho_i(y) = \eta_{i_1}(y) \oplus \cdots \oplus \eta_{i_m}(y)$$

*for some $\eta_{i_j} \in \mathcal{S}$. If the representations $\rho_i : G \longrightarrow GL(V_i)$ and bases e^i_j of V_i give rise to the $\mathcal{S}$-adapted representations, then the bases are also said to be $\mathcal{S}$-***adapted***.*

Remark 2. In particular, when $\mathcal{R}$ and $\mathcal{S}$ are traversals for $\hat{G}$ and $\hat{H}$ respectively, then adaptability is a key to the construction of FFT's for G. Choose coset representatives $\{s_1,\ldots,s_m\}$ for G/H. Let $f \in L(G)$. For each $i = 1,\ldots,m$ define functions $f_i \in L(H)$ by $f_i(h) = f(s_ih)$. Compute $DFT_{\mathcal{S}}(f_i)$ for each i. Thus, given $\rho \in \mathcal{R}$ as in (1), the Fourier transform $\hat{f}(\rho)$ can be computed as

$$\hat{f}(\rho) = \sum_{i=1}^{m} \rho(s_i) \begin{pmatrix} \hat{f}_i(\eta_1) & 0 & \cdots & 0 \\ 0 & \hat{f}_i(\eta_2) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \hat{f}_i(\eta_k) \end{pmatrix}. \tag{1}$$

So, with the notation defined in Definition 1, a recurrence is obtained,

$$T(G,\mathcal{R}) = \frac{|G|}{|H|}T(H,\mathcal{S}) + \sum_{\substack{i=1,\ldots,m \\ \rho\in\mathcal{R}}} M(\rho(s_i)) \tag{2}$$

where $M(\rho(s_i))$ is the number of operations needed to multiply $\rho(s_i)$ against the block diagonal matrix in (1).

The recurrence (2), presented in [15] and [11], demonstrates the usefulness of adapted bases for constructing FFT's. The natural generalization of this approach is the notion of a **tower-adapted** basis (cf. [6], Chapter 7) in which traversals are determined for each subgroup within a tower of subgroups

$$G = G_0 > G_1 > \cdots > G_k = \{1\}$$

such that traversals are successively adapted. This allows for a recursive application of the recurrence (2), permitting the construction of very efficient FFT's. This technique has been the key to some of the FFT's developed for S_n [15, 11] as well as for supersolvable groups [5].

For computational purposes, strict subgroup-adaptability is not required. In particular, if a restricted representation differs from a direct sum only by a reordering of basis vectors, the complexity remains essentially unchanged. For this reason the following defintion is useful.

Definition 3 *Let all notation be as in Definition 2. If for each $\rho_i \in \mathcal{R}$, there exists a permutation matrix X_i such that for all $y \in H$,*

$$\rho_i(y) = X_i\,(\eta_{i_1}(y) \oplus \cdots \oplus \eta_{i_m}(y))\,X_i^{-1}$$

for some $\eta_{i_j} \in \mathcal{S}$ then $\mathcal{R}$ is said to be **essentially $\mathcal{S}$-adapted.**

The FFT's for wreath products which are constructed in this paper use in part both FFT's for the symmetric group, as well as FFT's for direct products. For the latter, Remark 2 indicates one method for obtaining efficient DFT's for direct products. For the results of Section 4 it is useful to recall this approach (cf. [7], Satz 5.8).

If A and B are finite groups with traversals $\mathcal{A}$ and $\mathcal{B}$ for $\widehat{A}$ and $\widehat{B}$ respectively, then a traversal for $A \times B$ is obtained using the set

$$\mathcal{A} \otimes \mathcal{B} = \{\rho \otimes \eta \mid \rho \in \mathcal{A},\ \ \eta \in \mathcal{B}\}. \tag{3}$$

Thus, if $\rho : A \longrightarrow GL(V)$ and $\eta : B \longrightarrow GL(W)$ with ordered bases $v_1, \ldots, v_m$ and $w_1, \ldots, w_p$ for V and W respectively, then $\rho \otimes \eta : A \times B \longrightarrow GL(V \otimes W)$ with ordered basis

$$\begin{gathered} v_1 \otimes w_1, v_1 \otimes w_2, \ldots, v_1 \otimes w_p, \\ \ldots, \\ v_m \otimes w_1, v_1 \otimes w_2, \ldots, v_m \otimes w_p. \end{gathered} \tag{4}$$

Proposition 1 *Let A, B be finite groups with traversals $\mathcal{A}, \mathcal{B}$ of $\widehat{A}$ and $\widehat{B}$ respectively.*

(i) $T(A \times B, \mathcal{A} \otimes \mathcal{B}) \leq |A| T(B, \mathcal{B}) + |B| T(A, \mathcal{A})$.

(ii) Let $A' < A$ and $B' < B$ with traversals $\mathcal{A}'$ and $\mathcal{B}'$ of $\widehat{A}'$ and $\widehat{B}'$ respectively, such that $\mathcal{A}$ is $\mathcal{A}'$-adapted and $\mathcal{B}$ is $\mathcal{B}'$-adapted. Then $\mathcal{A} \otimes \mathcal{B}$ is $\mathcal{A}' \otimes \mathcal{B}$-adapted and essentially $\mathcal{A} \otimes \mathcal{B}'$-adapted.

2.3 Induced representations and the Mackey Subgroup Theorem

If G is a group and $H < G$ is a subgroup then representations of G can be constructed from representations of H by **induction**. This technique is of particular relevance for this paper, as all irreducible representations of wreath products are obtained as induced representations (cf. Section 3). There are many equivalent definitions of induced representations. For the purposes of this paper, the following one is the most convenient.

Definition 4 *Let $\eta : H \longrightarrow GL(W)$ be a representation of G in a complex vector space W. Thus there is an action of the group algebra $\mathbb{C}H$ on W defined by*

$$\Big(\sum_{h\in H} a_h h\Big) w = \sum_{h\in H} (a_h \eta(h) w).$$

Define the new vector space

$$Ind_H^G(W) = \mathbb{C}G \otimes_{\mathbb{C}H} W. \tag{5}$$

Then G acts on $Ind_H^G(W)$ by translation, $s(g \otimes w) = sg \otimes w$. This action of G on $Ind_H^G(W)$ is **the representation of G induced by** *η and is denoted as $\eta \uparrow G$. Dually, the* **restriction** *of a representation $\rho : G \longrightarrow GL(V)$ defines a representation of a subgroup H. In this case, the representation is denoted as $\rho \downarrow H$ and the vector space $Res_H^G(V) = V$.*

One or both of the G-superscript or H-subscript may be omitted when the context makes them unnecessary.

Fix a set of right coset representatives $\{s_1 = 1, s_2, \ldots, s_m\}$ for G/H and basis $\{w_1, \ldots, w_l\}$ for W. Since the induced representation is defined as a tensor product over $\mathbb{C}H$, the elements

$$s_1 \otimes w_1, \ldots, s_1 \otimes w_l, \ldots, s_m \otimes w_1, \ldots, s_m \otimes w_l$$

define an ordered basis for $\mathrm{Ind}(W)$.

Definition 5 *Let all notation be as above. Then the ordered basis for Ind(W) given by*

$$s_1 \otimes w_1, \ldots, s_1 \otimes w_l, \ldots, s_m \otimes w_1, \ldots, s_m \otimes w_l$$

will be called **the standard basis** *for Ind(W) (with respect to s_i and w_j).*

To show that the representations constructed in Section 3 are suitably adapted, much use will be made of the Mackey Subgroup Theorem which relates induction and restriction. Statement of this result requires a little notation. If $d \in G$ and $H < G$ is a subgroup then let H^d denote the conjugate subgroup dHd^{-1}. If η is a representation of H in $GL(W)$, then let η^d denote the representation of H^d in $GL(W^d)$, where $W^d = W$ and the action of H^d on W^d is given by

$$\eta^d(x) = \eta(d^{-1}xd). \tag{6}$$

Theorem *(Mackey Subgroup Theorem, cf. [34], p. 58) Let $\eta : H \longrightarrow GL(W)$ be a representation of $H < G$ and let $K < G$ be another, possibly different, subgroup of G. Then*

$$(\eta \uparrow G) \downarrow K \sim \oplus_d (\eta^d \downarrow (H^d \cap K)) \uparrow K \tag{7}$$

where the direct sum is over a complete set of double coset representatives for $K\backslash G/H$ and $\sim$ denotes equivalence of representations.

Proof: The goal here is to decompose the action of K on the space $\mathrm{Ind}(W)$. Writing any element of $\mathrm{Ind}(W)$ as a sum of elements of the form $s_i \otimes w$ for coset representatives s_i of G/H, a simple proof relies upon a judicious choice of coset representatives. For this consider the partitioning of the right cosets G/H under the action of K by translation. Note that for $k \in K$, $kdH = dH$ if and only if $k \in H^d \cap K$. Consequently, if right coset representatives $\{t_i^{(d)} \mid i = 1, \ldots, n_d\}$ are chosen for $K/H^d \cap K$, then the coset representatives for G/H in the orbit of the coset dH are given by the set $\{t_i^{(d)} d \mid i = 1, \ldots, n_d\}$. Equivalently, these are the cosets of G/H whose disjoint union makes up the double coset KdH.

Consequently, a complete set of coset representatives for G/H is given by the set

$$\{t_i^{(d)} d \mid d \in K\backslash G/H; \quad i = 1, \ldots, n_d\}.$$

Thus, $\mathrm{Ind}(W)$ has a direct sum decomposition into K-invariant subspaces as

$$\mathrm{Ind}(W) = \oplus_{d \in K\backslash G/H} W_d$$

where

$$W_d = \mathrm{span}_{\mathbb{C}}\{t_i^{(d)} d \otimes w \mid i = 1, \ldots, n_d \quad w \in W\}. \tag{8}$$

All that remains is to determine the action of K on W_d. For this, consider the map $T : W_d \longrightarrow \mathrm{Ind}_{K\cap H^d}^{K}(\mathrm{Res}_{K\cap H^d}^{H^d}(W_d))$ determined by

$$t_i^{(d)} d \otimes w \mapsto t_i^{(d)} \otimes w.$$

Recall that the tensor product on the lefthand side of (8) is over $\mathbb{C}H$, while the righthand side is over $\mathbb{C}K \cap H^d$. Note further that the action on w on the righthand side is via translation after conjugation by d. It is then straightforward to check that T is a K-isomorphism. This completes the proof.

■

For computational purposes it is instructive to restate the Mackey Subgroup Theorem for a particular choice of basis. Suppose that $w_1, \ldots, w_l$ is an ordered basis for W. Furthermore, let $d_1(=1), \ldots, d_m$ be a fixed set of double coset representatives for $K\backslash G/H$ and let

$$t_1^{(j)} d_j H \coprod \cdots \coprod t_{n_j}^{(j)} d_j H = K d_j H.$$

Then consider the ordered basis for $\mathrm{Ind}(W)$ given by

$$\begin{gathered} t_1^{(1)} d_1 \otimes w_1, \ldots, t_1^{(1)} d_1 \otimes w_l, \ldots, t_{n_1}^{(1)} d_1 \otimes w_1, \ldots, t_{n_1}^{(1)} d_1 \otimes w_l, \ldots, \\ \cdots, \\ t_1^{(m)} d_m \otimes w_1, \ldots, t_1^{(m)} d_m \otimes w_l, \ldots t_{n_m}^{(m)} d_m \otimes w_1, \ldots, t_{n_m}^{(m)} d_m \otimes w_l. \end{gathered} \tag{9}$$

Then (in the notation of (8)) W_{d_j} is isomorphic to the $n_j \cdot l$-dimensional subspace spanned by $t_1^{(j)} d_j \otimes w_1, \ldots, t_{n_j}^{(j)} d_j \otimes w_l$. The ordered basis $w_1, \ldots, w_l$ determines an ordered basis for each W^{d_j} (cf. (6)) and therefore in the basis (9),

$$(\eta \uparrow G) \downarrow K = \left(\eta^{(d_1)} \downarrow (H^{d_1} \cap K)\right) \uparrow K \oplus \cdots \oplus \left(\eta^{(d_m)} \downarrow (H^{d_m} \cap K)\right) \uparrow K \tag{10}$$

where each induced representation on the righthand side of 10 is given with respect to the standard basis determined by $t_i^{(j)}$ and w_k (cf. Definition 5).

Thus, the Mackey Subgroup Theorem shows that if the coset representatives for G/H are chosen in a compatible fashion with those of $K/K \cap H^d$ for the set of double coset representatives $d \in K\backslash G/H$ (in the sense of the proof), then for each $k \in K$, $(\eta \uparrow G)(k)$ will be block diagonal, with blocks given by representations **equal** to the induced representations $(\eta^d \downarrow (H^d \cap K)) \uparrow K$.

Consider further the case in which $\eta^d \downarrow H^d \cap K$ is itself reducible. To simplify notation, assume the general situation in which $G > K > L$ and η is a matrix representation of K and $\eta \downarrow L = \tau_1 \oplus \cdots \oplus \tau_a$. The representations τ_i are not necessarily irreducible. Since the restriction is *equal* to the direct sum, this means that the basis for the representation η is given by an ordered set of vectors,

$$e_1^1, \ldots, e_{b_1}^1, \ldots, e_1^a, \ldots, e_{b_a}^a$$

where each subset $e_1^j, \ldots, e_{b_j}^j$ spans the L-invariant subspace for τ_j. It is an elementary fact that

$$(\eta \downarrow L) \uparrow G \sim \tau_1 \uparrow G \oplus \cdots \oplus \tau_a \uparrow G. \tag{11}$$

In particular, notice that if computed with respect to the standard basis for the induced representation (cf. Definition 5) this equivalence does not become an equality. In this basis the lefthand side of (11) is

$$\begin{gathered} s_1 \otimes e_1^1, \ldots, s_1 \otimes e_{b_1}^1, \ldots, s_1 \otimes e_{b_a}^a, \\ \cdots \\ s_m \otimes e_1^1, \ldots, s_m \otimes e_{b_1}^1, \ldots, s_m \otimes e_{b_a}^a \end{gathered}$$

while the righthand side of (11) is instead the ordered basis

$$\begin{gathered} s_1 \otimes e_1^1, \ldots, s_1 \otimes e_{b_1}^1, \\ \cdots, \\ s_m \otimes e_1^1, \ldots, s_m \otimes e_{b_1}^1, \\ \cdots \\ s_1 \otimes e_1^a, \ldots, s_1 \otimes e_{b_a}^a, \\ \cdots \\ s_m \otimes e_1^a, \ldots, s_m \otimes e_{b_a}^a. \end{gathered}$$

Thus, the change of basis required is a permutation matrix, which may be computed directly from the above equations. More explicitly, it is clear that the basis vectors occur in blocks indexed by the pairs (j,k) with $j = 1, \ldots, m$ and $k = 1, \ldots, a$. The (j,k) block then itself contains b_k objects. The pairs (j,k) are then in 1:1 ordered correspondence with the integers $1, \ldots, m \cdot a$ by $(j,k) \mapsto (j-1)a + k$. Under this correspondence, on blocks, the required permutation π is given by sending $(j-1)a + k \mapsto (k-1)m + j$ and then acting as the identity within blocks.

This discussion is summarized in the following Proposition.

Proposition 2 *Let $G > K > L$ and let η be a representation of K such that $\eta \downarrow L = \tau_1 \oplus \cdots \oplus \tau_a$. Then η is essentially $\{\tau_1, \ldots, \tau_a\}$-adapted with implied permutation matrix X_η constructed as above. Note that X_η depends only on the dimensions of the τ_i, η and $[G : L]$.*

3 Adapted representations for wreath products $G[S_n]$

The FFT's for wreath products developed here depend on the construction of irreducible representations of $G[S_n]$ adapted for the subgroup tower

$$G[S_n] > G \times G[S_{n-1}] > \cdots > G[S_{n-1}] > \cdots.$$

Assuming for now that the representations of $G[S_n]$ are all constructed as induced representations (cf. Section 3.1), then potentially, the Mackey Subgroup Theorem presents a possibility for construction of adapted representations, as equation (5) shows that restriction to $G \times G[S_{n-1}]$ is itself a direct sum of induced representations. In fact as will be shown, this is precisely the case (cf. Section 3.2).

While the results proved here are primarily for wreath products by the full symmetric group, under suitable hypotheses they can be applied to more general wreath products (cf. Section 3.3).

3.1 Representation theory of wreath products

In this section, the representation theory of wreath products $G[S_n]$ is summarized quickly. For a full account, Kerber's two volumes of lecture notes [27] are a complete and accessible source.

Let G be a finite group with $r_1, \ldots, r_h$ a traversal of $\hat{G}$ with $r_i : G \longrightarrow GL(V_i)$. Let α be any weak composition of n of exactly h parts, so that $\alpha = (\alpha_1, \ldots, \alpha_h)$ where $\alpha_i \geq 0$ and $\sum_i \alpha_i = n$. This is denoted as $\alpha \models n$ and α is said to have length h, denoted $l(\alpha) = h$. Define the **content** of α, denoted $c(\alpha)$ to be the subcomposition consisting of only the nonzero parts of α. Thus, if $1 \leq j_1 < \cdots < j_m \leq h$ are the indices of the nonzero parts of α, then $c(\alpha) = (\alpha_{j_1}, \ldots, \alpha_{j_m})$.

If $c(\alpha) = (\alpha_{j_1}, \ldots, \alpha_{j_m})$, then define

$$r^\alpha = r_{j_1}^{\otimes \alpha_{j_1}} \otimes \cdots \otimes r_{j_m}^{\otimes \alpha_{j_m}}.$$

Then r^α is an irreducible representation of G^n in $V^\alpha = V_{j_1}^{\otimes \alpha_{j_1}} \otimes \cdots \otimes V_{j_m}^{\otimes \alpha_{j_m}}$.

The irreducible representations of $G[S_n]$ are constructed by applying Clifford theory, which uses the action of $G[S_n]$ by conjugation on the representations of the normal subgroup $G^n \triangleleft G[S_n]$ (cf. [13]). Under this action the representations of G^n are partitioned into orbits for which a complete set of representatives are given by the representations r^α as α runs over all weak compositions of n with h parts. The stabilizer of r^α in S_n is the Young subgroup

$$S_\alpha = S_{\alpha_{j_1}} \times \cdots \times S_{\alpha_{j_m}}$$

of permutations which permute the first α_{j_1} indices among themselves, the next α_{j_2} among themselves and so on.

Notice that S_α operates naturally on V^α by permuting components of the multitensor within tensor products of equal vector spaces,

$$\pi(v_1 \otimes \cdots \otimes v_n) = v_{\pi^{-1}(1)} \otimes \cdots \otimes v_{\pi^{-1}(n)}.$$

This action allows for the representation r^α to be extended in a well defined manner to $G[S_\alpha]$, which is still denoted by r^α, by

$$r^\alpha([f;\pi]) = r^\alpha(f)\left(v_{\pi^{-1}(1)} \otimes \cdots \otimes v_{\pi^{-1}(n)}\right). \tag{12}$$

In order to complete the description of the irreducible representations of $G[S_n]$ it is necessary to recall the representation theory of the symmetric group. The irreducible representations of S_n are indexed by partitions of n. Partitions are distinct from compositions in that all parts are strictly positive and are listed in weakly decreasing order. If μ is a partition of n, denoted $\mu \vdash n$, then let ρ_μ denote the associated irreducible representation of S_n. James' monograph [25] gives a very readable account of the theory.

Keeping the above notation, let $\lambda = (\lambda^{(1)}, \ldots, \lambda^{(m)})$ denote a vector of partitions with $\lambda^{(i)} \vdash \alpha_{j_i}$. Then

$$\rho_\lambda = \rho_{\lambda^{(1)}} \otimes \cdots \otimes \rho_{\lambda^{(m)}},$$

will denote this irreducible representation of S_α.

The irreducible representation ρ_λ extends trivially to an irreducible representation of $G[S_\alpha]$, (still denoted as ρ_λ) by

$$\rho_\lambda([f;\pi]) = \rho_\lambda(\pi). \tag{13}$$

The representations $r^\alpha \otimes \rho_\lambda$ defined by

$$r^\alpha \otimes \rho_\lambda([f;\pi]) = r^\alpha([f;\pi]) \otimes \rho_\lambda(\pi)$$

are irreducible representations of $G[S_\alpha]$.

Theorem *(cf. [27], Vol I, Ch. 2.) The collection of induced representations*

$$\left\{(r^\alpha \otimes \rho_\lambda) \uparrow G[S_n] \;\middle|\; \alpha \models n, l(\alpha) = h, c(\alpha) = (\alpha_{j_1}, \ldots, \alpha_{j_m}), \lambda = (\lambda^{(1)}, \ldots, \lambda^{(m)}), \lambda^{(i)} \vdash \alpha_{j_i}\right\}$$

are a complete set of irreducible representations of $G[S_n]$.

Remark 3. Babai and Rónyai have shown that a traversal for $\hat{G}$ can always be constructed in deterministic polynomial time [3]. In the special case of S_n, highly efficient algorithms exist to construct particular sets of representations, *Young's orthogonal form* or *seminormal form*, which are used in the algorithms described in the coming sections. For example, closed form expressions exist for representations at the pairwise adjacent transpositions and from these all necessary representations of S_n may be obtained [25]. Thus, using the standard basis for the induced representation (cf. Definition 5) the necessary input for the wreath product FFT's can be constructed efficiently.

3.2 Adapted bases for wreath products by S_n

The Mackey Subgroup Theorem shows the way to producing $G \times G[S_{n-1}]$-adapted bases for the irreducible representations of $G[S_n]$. It turns out that constructing these bases depends primarily on two things.

(a) *A traversal for S_n which is S_{n-1}-adapted.*
(b) *A "good" choice of double coset representatives for $G[S_{n-1}] \times G \backslash G[S_n]/G[S_\alpha]$.*

To satisfy (a) it is well known that the irreducible representations of S_n, in Young's orthogonal (resp. seminormal) form are S_{n-1}-adapted for Young's orthogonal (resp. seminormal) form for S_{n-1}.

For (b), let the $c(\alpha) = (\alpha_{j_1}, \ldots, \alpha_{j_m})$. For $i = 1, \ldots, m$ define $\partial_i(\alpha) \models n-1$ by

$$\partial_i(\alpha)_k = \begin{cases} \alpha_k & \text{if } k \neq j_i \\ \alpha_k - 1 & \text{if } k = j_i \end{cases} \tag{14}$$

and extend this definition to S_α by defining

$$\partial_i S_\alpha = S_{\partial_i \alpha}. \tag{15}$$

Proposition 3 *Let $\alpha \models n$ with the j_i as above.*

(i) Double coset representatives for $G \times G[S_{n-1}] \backslash G[S_n]/G[S_\alpha]$ are given by the elements

$$\{(1), (n\ n{-}1\ \cdots\ n{-}\alpha_{j_m}), \ldots, (n\ n-1 \cdots n{-}\alpha_{j_m}{-}\alpha_{j_{m-1}}{-}\cdots{-}\alpha_{j_2})\} \tag{16}$$

viewed in the natural way as elements of $G[S_n]$.

(ii) Furthermore, for $i = 1, \ldots, m$,

$$\partial_i S_\alpha = S_{n-1} \cap S_\alpha^{d_i}$$

where $d_m = (1)$ and for $1 \leq i < m$, $d_i = (n\ n-1 \cdots n{-}\alpha_{j_m}{-}\cdots{-}\alpha_{j_{i+1}})$.

Proof: To prove (i), first note that it is enough to find double coset representatives for $S_{n-1}\backslash S_n/S_\alpha$. Any such set of representatives, viewed as elements of $G[S_n]$ will comprise a set of representatives for $G[S_{n-1}] \times G\backslash G[S_n]/G[S_\alpha]$.

For this, it is useful (and equivalent) to instead consider the decomposition of the cosets $S_{n-1}\backslash S_n$ under the action of S_α by

$$S_{n-1}x \ \stackrel{y}{\mapsto}\ S_{n-1}xy$$

for $y \in S_\alpha$. The set

$$\{(1), (n\ n-1), \ldots (n\ n-1\ \ldots 1)\} \tag{17}$$

is a set of coset representatives for $S_{n-1}\backslash S_n$. Consequently, the coset representatives comprising $S_{n-1}(n\ \ldots\ i)S_\alpha$ are given by $(n \ldots i)y$ where y are coset representatives for

$$S_\alpha/S_\alpha \cap (n \ldots i)^{-1} S_{n-1}(n \ldots i) = S_\alpha/\partial_l S_\alpha$$

where $i \in \{\alpha_{j_1} + \cdots + \alpha_{j_{l-1}} + 1, \ldots, \alpha_{j_1} + \cdots + \alpha_{j_l}\}$. Under this correspondence it is easy to check that the coset representatives (17) partition into orbits with representatives given by (16).

Proof of (ii) requires a simple calculation.

∎

Remark 4. Proposition 3 indicates a recursive method for generating the cosets of S_n/S_α which are compatible with the double coset decomposition $S_{n-1}\backslash S_n/S_\alpha$.

For

$$\{d_1, \ldots, d_m\} = S_{n-1}\backslash S_n/S_\alpha,$$

defined by

$$d_k = \begin{cases} (1) & \text{for } k = m \\ (n \ldots n - \alpha_{j_m} - \cdots - \alpha_{j_{k+1}}) & \text{for } 1 \leq k < m, \end{cases} \tag{18}$$

the coset representatives of S_n/S_α whose cosets comprise the double coset $S_{n-1}d_iS_\alpha$ are computed as the collection of products yd_i where the elements y run over coset representatives of $S_{n-1}/(S_{n-1} \bigcap d_iS_\alpha d_i^{-1})$, which by Proposition 3 is the set of coset representatives for $S_{n-1}/\partial_i S_\alpha$, thereby "reducing" the problem to one for S_{n-1}. To illustrate this construction with an example, conside the problem of finding compatible coset reps for $S_4/S_2 \times S_2$.

Following the proof of Proposition 3, representatives for $S_3\backslash S_4/S_2 \times S_2$ are given by

$$(1) \quad \text{and} \quad (432).$$

Consequently, coset representatives for $S_4/S_2 \times S_2$ are given by the sets

$$\{y(1) \mid y \in S_3/S_2 \times S_1\}$$
$$\text{and}$$
$$\{y(432) \mid y \in S_3/S_1 \times S_2\}$$

since $S_2 \times S_1 = S_3 \cap (S_2 \times S_2)$ and $S_1 \times S_2 = S_3 \cap \left[(432)\,(S_2 \times S_2)\,(432)^{-1}\right]$. Proceeding recursively, coset representatives for $S_3/S_2 \times S_1$ and $S_3/S_1 \times S_2$ are chosen by determining the representatives for $S_2\backslash S_3/S_2 \times S_1$ and $S_2\backslash S_3/S_1 \times S_2$ respectively. Again following Proposition 3,

$$\{(1), (321)\} = S_2\backslash S_3/S_2 \times S_1$$

and

$$\{(1), (32)\} = S_2\backslash S_3/S_1 \times S_2$$

respectively.

The coset representatives $S_2/S_1 \times S_1 = S_2 = \{(1), (12)\}$ and $S_2/S_2 = \{(1)\}$, are completely determined. Thus,

$$\{(1)(1), (1)(321), (12)(321)\} = S_3/S_2 \times S_1$$

and

$$\{(1)(1), (12)(1), (1)(32)\} = S_3/S_1 \times S_2.$$

So that finally,

$$\begin{aligned} S_4/S_2 \times S_2 &= \{(1)(1)(1), (1)(321)(1), (12)(321)(1), (1)(1)(432), (12)(1)(432), (1)(32)(432)\} \\ &= \{(1), (321), (12)(321), (432), (12)(432), (32)(432)\}. \end{aligned}$$

This sort of decomposition is neatly encoded as a tree (Figure 1). As is clear, paths from the root to any leaf correspond to a unique coset representative.

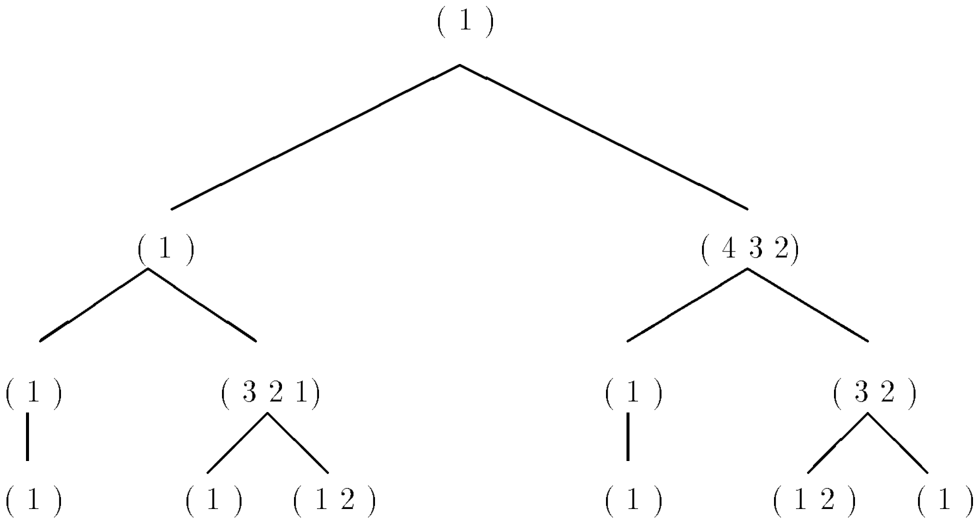

Figure 1: Tree representation of coset construction for $S_4/S_2 \times S_2$.

To construct the appropriate subgroup adapted bases the following conventions will be followed.

Convention 1: *Coset representatives for $G[S_n]/G[S_\alpha]$ are chosen to be in $S_n < G[S_n]$ and compatible with the double coset representatives for $G[S_{n-1}] \times G \backslash G[S_n]/G[S_\alpha]$ as indicated by Proposition 3.*

Convention 2: *All representations of S_α are given as tensor products of matrices in Young's orthogonal form.*

Convention 3: *Induced representations are computed in the standard basis (cf. Definition 4) with respect to the coset representatives described in Convention 1 and the bases for the appropriate tensor products (cf. (4)).*

With these conventions, the following theorem is easily proved.

Theorem 1 *Let $\mathcal{R}$ denote a traversal for G and all notation be as above. For each $\alpha \models n$, $l(\alpha) = h$, $c(\alpha) = (\alpha_{j_1}, \ldots, \alpha_{j_m})$ and $\lambda = (\lambda^{(1)}, \ldots, \lambda^{(m)})$ with $\lambda^{(i)} \vdash \alpha_{j_i}$ let the irreducible matrix representations $r^\alpha \otimes \rho_\lambda \uparrow G[S_n]$ be constructed according to Conventions (1)-(3). Denote this traversal of $\widehat{G[S_n]}$ as $\mathcal{Q}_n$. Then $(G[S_n], \mathcal{Q}_n)$ is essentially $(G[S_n] \times G, \mathcal{Q}_{n-1} \otimes \mathcal{R})$-adapted.*

Proof: Up to a reindexing of the representations of G, it is enough to consider only the case in which $c(\alpha) = (\alpha_1, \ldots, \alpha_m)$. Then let

$$\begin{aligned}
\eta &= r^\alpha \otimes \rho_\lambda \\
&= r_1^{\alpha_1} \otimes \cdots r_k^{\alpha_k} \otimes \cdots r_m^{\alpha_m} \otimes \rho_{\lambda(1)} \otimes \cdots \rho_{\lambda(m)}.
\end{aligned}$$

By the choice of coset representatives for each of these induced representations, the proof simply amounts to showing that for each double coset representative d of $G[S_{n-1}] \times G \backslash G[S_n]/G[S_\alpha]$ (chosen according to Proposition 3), the matrix representations $\eta^d \downarrow G[S_\alpha]^d \cap G[S_{n-1}] \times G$ are suitably adapted. That is, it is enough to show that these restrictions decompose into a direct sum of the appropriate form which is then induced up to $G[S_{n-1}] \times G$.

With d_k defined as in (18)

$$\begin{aligned} G[S_\alpha]^{d_k} \cap (G[S_{n-1}] \times G) &= G[\partial_k S_\alpha] \times G \\ &= G[S_{\alpha_1} \times \cdots \times S_{\alpha_k - 1} \times \cdots \times S_{\alpha_m}] \times G \end{aligned} .$$

It is now straightforward to check that

$$\begin{aligned} &\eta^{d_k} \downarrow G[S_\alpha]^{d_k} \cap (G[S_{n-1}] \times G) = \\ &\qquad r_1^{\alpha_1} \otimes \cdots r_k^{\alpha_k - 1} \otimes \cdots r_m^{\alpha_m} \otimes \rho_{\lambda(1)} \otimes \cdots (\rho_{\lambda(j)} \downarrow S_{\alpha_k - 1}) \otimes \cdots \rho_{\lambda(m)} \otimes r_k. \end{aligned}$$

Since $\rho_{\lambda(k)}$ is $S_{\alpha_k - 1}$ adapted, η^{d_k} is essentially adapted for $(G[S_\alpha]^{d_k} \cap G[S_{n-1}] \times G, \mathcal{Q}_{\partial_k \alpha} \otimes \mathcal{R})$ where $\mathcal{Q}_{\partial_k \alpha}$ denotes the irreducibles for $G[S_\alpha]^{d_k} \cap G[S_{n-1}]$ given by $r^{\partial_k \alpha} \otimes \rho_\lambda$. Thus, applying the Mackey Subgroup Theorem and Proposition 1, Theorem 1 is proved.

■

Corollary 1 *Let the notation be as above and let $H < G$ be a subgroup. If the r_i are chosen so that they are essentially H-adapted for some traversal of $\hat{H}$, then the irreducible representations $r^\alpha \otimes \rho_\lambda \uparrow G[S_n]$ will be essentially $G[S_{n-1}] \times H$-adapted as well.*

Proof: Applying Theorem 1 and Proposition 2 yields this result.

3.3 Extension to general wreath products

The irreducible representations of an arbitrary wreath products $G[L]$ for $L < S_n$ are constructed almost exactly as those for $G[S_n]$. Keeping the notation of the previous sections, then for $\alpha \models n$, define $L_\alpha = L \cap S_\alpha$. If ρ is any irreducible representation of L_α, then as in (12) and (13) the representations r^α and ρ extend to representations of $G[L_\alpha]$. Clifford theory implies that the induced representation

$$r^\alpha \otimes \rho \uparrow_{G[L_\alpha]}^{G[L]}$$

is irreducible.

Theorem 2 *([27], Vol. I, Theorem 5.20) Let all notation be as above. Then the set of irreducible representations*

$$\{r^\alpha \otimes \rho \uparrow G[L] \mid \alpha \models n, \rho \in \widehat{L_\alpha}\}$$

is a complete set of inequivalent irreducible representations of $G[L]$.

Thus, construction of the irreducible representations for $G[L]$ is exactly analogous to the particular case of $G[S_n]$. Consequently, it is natural to attempt to mimic the construction of Theorem 1 to form $G[L \cap S_{n-1}] \times G$-adapted bases. It is clear that all that is necessary is that for each double coset representative d of $G[L \cap S_{n-1}] \times G \backslash G[L]/G[L_\alpha]$, the representations

$$(r^\alpha \otimes \rho)^d \downarrow G[L_\alpha] \cap G[L \cap S_{n-1}] \times G$$

is adapted. While such bases always exist, their construction may be difficult. Success in their construction would then pave the way to FFT's as obtained in Section 4. Note that in the case in which L is abelian, then this is the situation of an abelian extension and the methods of [33] can be used to obtain adapted bases as well as FFT's.

4 FFT's for wreath products

The complexity results presented here depend mainly on careful use of the structure of the matrices of the induced representations in their standard bases (cf. Definition 5) as well as various properties of Young's orthogonal form.

If η is a representation of $H < G$, of dimension d_η, then in the standard basis, for every $s \in G$, $(\eta \uparrow G)(s)$ is block monomial with $d_\eta \times d_\eta$ blocks. Consequently, multiplication of $(\eta \uparrow G)(s) \cdot M$ for any $\frac{|G|}{|H|}d_\eta \times \frac{|G|}{|H|}d_\eta$ complex matrix M requires at most $\left(\frac{|G|}{|H|}\right)^2$ multiplications of $d_\eta \times d_\eta$ matrices. If these blocks have additional structure, then more savings can be obtained.

4.1 General results

Using the adapted bases constructed in Section 3.2, FFT's for $G[S_n]$ can now be obtained.

Theorem 3 *Recall the notation of equation (2) and Definitions 1 and 2. Let G be a finite group with h inequivalent irreducible representations. Fix a traversal $\mathcal{R}$ for $\hat{G}$. Let $\mathcal{Q}_n$ denote the traversal for $\widehat{G[S_n]}$ constructed in according to Theorem 1. Let $\mathcal{R} \otimes \mathcal{Q}_{n-1}$ denote the traversal for $(\widehat{G[S_{n-1}]} \times G)$ given tensoring the matrices of $\mathcal{R}$ and $\mathcal{Q}_{n-1}$. Let $T(n) = T(G[S_n], \mathcal{Q}_n)$ Then the the recurrence*

$$T(n) \leq \Big[n\, T(G, \mathcal{R}) \cdot \mid G[S_{n-1}] \mid\Big] + \Big[n\, T(n-1) \cdot \mid G \mid\Big] + \Big[n^3 2^h \cdot \mid G[S_n] \mid\Big] \tag{19}$$

holds.

Proof: Let $r_i \in \mathcal{R}$ be of dimension d_i. Then for any $s \in G[S_n]$, $r^\alpha \otimes \rho_\lambda \uparrow G[S_n](s)$ is a $\binom{n}{\alpha_1 \ldots \alpha_h} \times \binom{n}{\alpha_1 \ldots \alpha_h}$ block monomial matrix with $d_\alpha d_\lambda \times d_\alpha d_\lambda$ blocks where $d_\alpha = d_1^{\alpha_1} \cdots d_h^{\alpha_h}$ and $d_\lambda = d_{\lambda(i)} \cdots d_{\lambda(h)}$

Fix coset representatives $s_i = (n \ldots n-i+1)$ for $G[S_n]/G[S_{n-1}] \times G$. Following Remark 2, since the bases being used are (up to conjugation by a permutation matrix) $G[S_{n-1}] \times G$-adapted, the first step in computing $DFT_{\mathcal{Q}}(f)$ is to compute for each $i = 1, \ldots, n$, $DFT_{\mathcal{R}\otimes\mathcal{Q}_{n-1}}(f_i)$ for $f_i \in L(G[S_{n-1}] \times G)$ defined by $f_i(h) = f(s_i h)$. By applying Proposition 1 to $G[S_{n-1}] \times G$, the first two terms of (19) are obtained.

Having computed each of $DFT(f_i)$, all that is left is to compute the summations

$$\sum_i r^\alpha \otimes \rho_\lambda \uparrow G[S_n](s_i) \cdot \widehat{f}_i\Big((r^\alpha \otimes \rho_\lambda \uparrow G[S_n]) \downarrow G[S_{n-1}] \times G\Big). \tag{20}$$

To compute (20) efficiently, two observations are crucial. First, note that since the coset representatives for $G[S_n]/G[S_\alpha]$ are chosen to be in $S_n < G[S_n]$, each nonzero block of $r^\alpha \otimes \rho_\lambda \uparrow G[S_n](s_i)$ will be of the form $r^\alpha \otimes \rho_\lambda(\sigma)$ for some $\sigma \in S_\alpha$. Consequently, these blocks will evaluate to matrices of of the form $I_{d_\alpha} \otimes \rho_\lambda(\sigma)$ where I_{d_α} denotes the $d_\alpha \times d_\alpha$ identity matrix.

Second, ρ_λ is a tensor product of matrices in Young's orthogonal form. Now recall that for any partition $\mu \vdash m$, the matrices $\rho_\mu(j j+1)$ for $1 \leq j \leq m-1$ are sparse with at most

$2d_\mu$ nonzero entries (eg. [25]). Thus, any $\rho_\mu(\pi)$ for any $\pi \in S_m$ can be written as a product of at most m^2 such sparse matrices. Consequently,

$$\rho_\lambda(\sigma) = \rho_{\lambda^{(1)}}(\sigma_1) \otimes \cdots \otimes \rho_{\lambda^{(h)}}(\sigma_h) \tag{21}$$

will be a product of at most n^2 sparse matrices, each with at most $2^h d_\lambda$ nonzero entries. So that finally, each matrix $r^\alpha \otimes \rho_\lambda(\sigma)$ will be a product of at most n^2 sparse matrices each with at most $2^h d_\lambda d_\alpha$ nonzero entries.

Thus, any individual block multiplication in (20) can be performed in at most $n^2 2^h (d_\lambda d_\alpha)^2$ operations. Doing this for each block gives an upper bound of at most $\binom{n}{\alpha_1 \ldots \alpha_h}^2 n^2 2^h (d_\lambda d_\alpha)^2$ operations for each product in (20). Summing over all allowable α and λ gives

$$\begin{aligned} n \sum_{\substack{\alpha \models n \\ l(\alpha)=h}} \sum_{\lambda^{(i)} \vdash \alpha_i} n^2 2^h \binom{n}{\alpha_1 \ldots \alpha_h}^2 (d_\lambda d_\alpha)^2 &= 2^h n^3 (n!) \sum_{\substack{\alpha \models n \\ l(\alpha)=h}} \binom{n}{\alpha_1 \ldots \alpha_h} d_\alpha^2 \\ &= 2^h n^3 (n!)(d_{\alpha_1}^2 + \cdots + d_{\alpha_h}^2)^n \\ &= 2^h n^3 |G[S_n]|. \end{aligned}$$

■

Iterating Theorem 2 yields the following corollary.

Corollary 2 *With the notation of Theorem 2,*

$$T(n) \leq |G[S_n]| \cdot \left[\frac{T(G)}{|G|} \cdot \frac{n(n+1)}{2} + 2^h \frac{n^2(n+1)^2}{4} + 1 \right]$$

In particular, since $T(G) \leq |G|^2$,

$$T(n) \leq |G[S_n]| \cdot \left[|G| \cdot \frac{n(n+1)}{2} + 2^h \frac{n^2(n+1)^2}{4} + 1 \right].$$

4.2 FFT's for wreath products of abelian groups and symmetric groups

The cases of wreath products by abelian or symmetric groups are of great interest for their uses in applications to chemistry and statistics. As much is known regarding the linear complexity of these groups, Theorem 2 simplifies.

The Cooley-Tukey algorithm [12] and its variants show that $T(A) = O(|A| \log |A|)$ for any abelian group A. Clausen's work [11] shows that $T(S_m) = O(|S_m| \log^3 |S_m|)$. So applying Corollary 2 to the more general setting $T(G) = O(|G| \log^c |G|)$ and simplifying gives the following result.

Theorem 4 *Let G be such that $T(G) = O(|G| \log^c |G|)$. Let $T(n)$ denote the number of operations needed to compute the DFT for an arbitrary function $f \in L(G[S_n])$ using the bases constructed by Theorem 1.*

$$T(n) \leq O(|\, G[S_n] \,| \cdot \log^b |\, G[S_n] \,|) \tag{22}$$

where $b = \max\{4, c\}$. In particular, for $G = S_m$ or G abelian, (11) holds for $b = 4$.

It is clear that these techinques can be applied to iterated wreath products, for which analogous results would be obtained. These iterated wreath products are of great interest as the automorphism groups of nested designs and spherically homogeneous rooted trees (rooted trees with constant branching at each level).

Corollary 3 *Let G be the automorphism group of a spherically homogeneous tree (cf. Section 5.1). Then*

$$T(G) \leq O(|G| \log^4 |G|).$$

4.3 Fast Fourier inversion and convolution

Given the $DFT_{\mathcal{R}}(f)$, the function $f \in L(G)$ may be recovered via the **Fourier inversion formula**,

$$f(s) = \frac{1}{|G|} \sum_{\rho \in \mathcal{R}} \operatorname{trace}\Big(\hat{f}(\rho)\rho(s^{-1})\Big). \tag{23}$$

Efficient computation of the inverse Fourier transform is the recovery of the values of f from the initial data of $DFT_{\mathcal{R}}(f)$ and the $\mathcal{R}$.

Following Remark 1, the FFT may be viewed as an efficient change of basis algorithm, and in particular a straight-line program for the task. (It is not too hard to see that the FFT explained in Theorem 2 can be written as a matrix factorization, if the FFT chosen for G is described as such.) As the matrix which computes the inverse transform is effectively the "transpose" of the DFT, (cf. [6], Theorem 1), the complexity of a fast inversion performed by transposing the FFT is of the same order ([8], [6], Theorem 3). This is recorded in the following Theorem.

Theorem 5 *The bounds stated in Theorems 2 and 3, also give upper bounds on the complexity of computing the inverse Fourier transform in the various cases.*

A different constructive approach to the computation of Fourier inversion may be found in [32]. Application of the methods explained there in concert with the sparse factorization of the representations at the coset representatives will yield a similar complexity result.

Fast forward and inverse transforms imply the existence of a fast convolution algorithm as well [11, 15]. That is, recall that for functions $f, f' \in L(G)$, the **convolution** of f and f', denoted $f \star f'$ is defined as

$$f \star f'(s) = \sum_{t \in G} f(st^{-1}) f'(t). \tag{24}$$

It is easy to check that the Fourier transform turns convolution into (matrix) multiplication,

$$\widehat{f \star f'}(\rho) = \hat{f}(\rho) \cdot \hat{f}'(\rho). \tag{25}$$

Consequently, to compute $DFT(f \star f')$, first compute $DFT(f)$ and $DFT(f')$. Compute the set of matrix products $\hat{f}(\rho) \cdot \hat{f}'(\rho)$ for all irreducible matrix representations ρ. Finally perform the fast inversion on the set of products to recover the convolved function. Thus, for convolution of functions in $L(G[S_n])$, the following complexity estimate is obtained.

Theorem 6 *Convolution of arbitrary functions in $L(G[S_n])$ can be computed in at most*

$$O(\mid G[S_n] \mid \cdot \left[|G| \cdot \frac{n(n+1)}{2} + 2^h \frac{n^2(n+1)^2}{4} + 1 \right] + \sum_{\substack{\alpha_1+\cdots+\alpha_h=n \\ \lambda^{(i)} \vdash \alpha_i}} \left(\binom{n}{\alpha_1 \ldots \alpha_h} d_\alpha d_\lambda \right)^3)$$

operations.

5 Applications

As mentioned, the techniques described here have potential application to data analysis, molecular spectroscopy as well as various aspects of signal processing. These possibilities derive from the use of the subgroup adapted bases constructed in Section 3. The following section sketches these applications and gives pointers to the relevant literature for a more thorough treatment.

5.1 Computational methods for nested designs

The applications to data analysis come from the necessity of finding efficient algorithms to compute projections for data on homogeneous spaces for iterated wreath products, for which **nested designs** comprise a large class of examples.

From a graph theoretic point of view, nested designs may be described as **spherically homogeneous trees**, which are rooted trees with constant branching at a fixed distance from the root. If the tree has height n and is such that all nodes at distance k from the root having m_k descendants, then the tree is said to **have type m** $= (m_1, \ldots, m_n)$ and is denoted as $X_{\mathbf{m}}$.

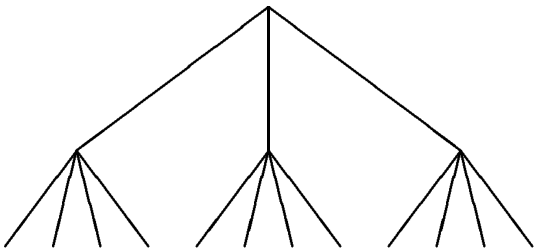

Figure 2. The spherically homogeneous tree of type (3,4), $X_{(3,4)}$.

The automorphism group of $X_{\mathbf{m}}$, denote as $G_{\mathbf{m}}$ is an iterated wreath product. More precisely, if $G_{(m_2,\ldots,m_n)}$ is the automorphism group of $X_{(m_2,\ldots,m_n)}$, then

$$G_{(m_1,m_2,\ldots,m_n)} = G_{(m_2,\ldots,m_n)}[S_{m_1}].$$

Statisticians are familiar with the objects $X_{\mathbf{m}}$ as nested designs. To give an explicit example, consider the tree $X_{(k,m)}$ with automorphism group $S_m[S_k]$. Each of the k nodes at distance 1 from the root correspond to a block containing m plots (in natural correspondence with the children of these nodes). The relevant statistical automorphism group allows permutations by blocks followed by any permutation within a block (cf. [1, 31]), so is precisely $S_m[S_k]$. Repeated nesting in turn is equivalent to further branching of the tree with corresponding wreath product as symmetry group. For a fuller description of the statistical motivation for these designs see [31].

Let $L(X_{\mathbf{m}})$ denote the vector space of complex-valued functions on the leaves of $X_{\mathbf{m}}$. Thus, for $\mathbf{m} = (m_1, \ldots, m_n)$, $L(X_{\mathbf{m}})$ has dimension $m_1 \cdots m_n$. Data on $X_{\mathbf{m}}$ is naturally viewed as an element of $L(X_{\mathbf{m}})$. Various subspaces of $L(X_{\mathbf{m}})$ have natural data analytic interpretations. Analysis and estimation for data on $X_{\mathbf{m}}$ often takes the form of computing the projections of the data vector onto various orthogonal subspaces of $L(X_{\mathbf{m}})$. If the symmetry group $G_{\mathbf{m}}$ is taken into account in some fashion, then additionally the subspace decomposition is required to be $G_{\mathbf{m}}$-invariant. To cite but a few examples, this is the case in the spectral analysis approach (cf. [14], esp. Chapters 7 and 8 and [20]), ANOVA under the assumption of symmetry-invariant covariance (cf. [21]) as well as ANOVA under the assumption of randomized treatment assignment with symmetry (cf. [1]).

The orthogonal decomposition of $L(X_{\mathbf{m}})$ under the assumption of $G_{\mathbf{m}}$-invariance takes the form of the isotypic decomposition with possible further decomposition into irreducibles.

In [16] the problem of computing projections onto homogeneous spaces of the form $L(G/H)$ for finite groups $G > H$ was studied. Standard character theory constructs the projections by summing over the group - an unfeasible approach for all but the smallest groups. In [16] it was shown that projections may be obtained much more efficiently if adapted bases are used.

In particular, suppose that $\rho_1, \ldots, \rho_k$ are all irreducible representations of G that occur in $L(G/H)$ and $\eta_1 \ldots, \eta_l$ are all those that occur in the restrictions $\rho_1 \downarrow H, \ldots, \rho_k \downarrow H$. Then if bases are chosen so that the set $\{\rho_1, \ldots, \rho_k\}$ is essentially $\{\eta_1, \ldots, \eta_l\}$-adapted, the projections can be computed much efficiently.

Theorem *(cf. [16], Theorem 2.5) Let $L(G/H)$ give a multiplicity-free representation of G. Then assuming that irreducibles which occur in $L(G/H)$ are H-adapted, then the complete set of isotypic projections for $f \in L(G/H)$, can be computed using at most $|G/H|^2$ additions and $|H\backslash G/H|^2 \cdot |G/H|$ additional operations.*

To apply this result in the context of $X_{\mathbf{m}}$, note first that $G_{\mathbf{m}}$ acts transitively on the leaves of $X_{\mathbf{m}}$. If $G^0_{\mathbf{m}}$ denotes the stabilizer of a given leaf, then it is clear that if $\mathbf{m}' = (m_1, \ldots, m_n)$ and $\mathbf{m} = (m_0, m_1, \ldots, m_n)$, then the stabilizer subgroups have a recursive structure given by

$$G^0_{(m_0, m_1, \ldots, m_n)} = G^0_{\mathbf{m}'} \cong S_{m_0 - 1}[G_{\mathbf{m}}] \times G^0_{\mathbf{m}}.$$

The first factor reflects the arbitrary automorphisms of the subtree not containing the distinguished leaf, while the second factor indicates that automorphisms of the subtree containing that leaf, must stabilize the leaf.

Theorem 7 *Let $\mathcal{R}_{\mathbf{m}}$ denote the traversal of $\widehat{G_{\mathbf{m}}}$ for $\mathbf{m} = (m_0, m_1, \ldots, m_n)$ formed using Conventions (1)-(3). Let $\mathcal{S}^0_{\mathbf{m}}$ denote the traversal for $G^0_{\mathbf{m}}$ constructed according to Conventions (1)-(3). Let $\mathbf{m}' = (m_1, \ldots, m_n)$. Then $(G_{\mathbf{m}}, \mathcal{R}_{\mathbf{m}})$ is essentially $(G^0_{\mathbf{m}}, \mathcal{R}_{(\mathbf{m}')} \otimes \mathcal{S}^0_{\mathbf{m}'})$-adapted.*

Proof: This is proved by induction on n. For $n = 0$ this is precisely the claim that Young's orthogonal form for S_{m_0} is $S_{m_0} - 1$-adapted.

Let $n > 0$. By Theorem 1, $\mathcal{R}_{\mathbf{m}}$ is essentially $S_{m_0-1}[G_{\mathbf{m}'}] \times G_{\mathbf{m}'}$-adapted. By induction, $\mathcal{R}_{\mathbf{m}}$ is essentially $G^0_{\mathbf{m}'}$-adapted. Applying Corollary 1 completes the proof.

■

Theorem 8 *Let* $\mathbf{m} = (m_0, \ldots, m_{n+1})$. *The representation of* $G_{\mathbf{m}}$ *on* $L(X_{\mathbf{m}})$ *is multiplicity-free.*

Proof: As above, let $\mathbf{m}' = (m_1, \ldots, m_n)$. Then recall that

$$G_{\mathbf{m}} = G_{\mathbf{m}'}[S_{m_0}].$$

and the stabilizer of a leaf of the tree $X_{\mathbf{m}}$, denoted $G^0_{\mathbf{m}}$ has the recursive structure

$$G^0_{\mathbf{m}} \cong G_{\mathbf{m}'}[S_{m_0-1}] \times G^0_{\mathbf{m}'}.$$

Thus,

$$X_{\mathbf{m}} \cong G_{\mathbf{m}}/(G_{\mathbf{m}'}[S_{m_0-1}] \times G^0_{\mathbf{m}'}).$$

The proof that $L(X_{\mathbf{m}})$ affords a multiplicty-free representation of $G_{\mathbf{m}}$ proceeds by induction on n. If $n = 0$, then this is simply the well known statement that the representation of S_{m_0} on S_{m_0}/S_{m_0-1} is multiplicity-free (see e.g. [25]).

Let $n > 0$. By Frobenious reciprocity, a representation of $G_{\mathbf{m}}$ occurs in $L(X_{\mathbf{m}})$ if and only if its restriction to $G^0_{\mathbf{m}}$ contains the trivial representation. For a given representation τ of $G_{\mathbf{m}}$, this will be denoted as

$$\tau \downarrow G^0_{\mathbf{m}} \ni 1.$$

In order to prove that the representation in $L(X_{\mathbf{m}})$ is multiplicity-free it is enough to prove that the sum of the degrees of the irreducible representations which occur in $L(X_{\mathbf{m}})$ is equal to $\dim_{\mathbb{C}}(L(X_{\mathbf{m}})) = m_0 m_1 \cdots m_n$.

For each r an irreducible representation of $G_{\mathbf{m}'}$ consider then the irreducible representation of $G_{\mathbf{m}'}[S_{m_0-1}] \times G_{\mathbf{m}'}$ defined by

$$\eta_r = 1^{m_0-1}_{G_{\mathbf{m}'}} \otimes r \otimes 1_{S_{m_0-1}}.$$

Then (cf. Section 2.2)

$$\eta_r \uparrow G_{\mathbf{m}}$$

is irreducible. Since $(\eta_r \uparrow G_{\mathbf{m}}) \downarrow G_{\mathbf{m}'}[S_{m_0-1}] \times G_{\mathbf{m}'}$ contains η_r, then $\eta_r \uparrow G_{\mathbf{m}} \downarrow G_{\mathbf{m}'}[S_{m_0-1}] \times G^0_{\mathbf{m}'} \ni 1$ if and only if $r \downarrow G^0_{\mathbf{m}'} \ni 1$.

Consider then the set

$$\{\eta_r \uparrow G_{\mathbf{m}} \mid r \downarrow G^0_{\mathbf{m}'} \ni 1\}.$$

Notice that this set is multiplicity-free and contains only representations contained in $L(X_{\mathbf{m}})$. Furthermore, if $r \downarrow G^0_{\mathbf{m}'} \ni 1$, then r occurs in the multiplicity-free (by induction) representation of $L(X_{\mathbf{m}'})$.

Thus, to complete the proof of the Theorem, all that is left is to show that the sum of dimensions is equal to $m_0 m_1 \cdots m_n$. If d_r denotes the dimension of the irreducible representation r, then

$$\dim(\eta_r \uparrow G_{\mathbf{m}}) = m_0 d_r.$$

Consequently,

$$\begin{aligned}\sum_{r\downarrow G^0_{\mathbf{m}'}\ni 1} \dim(\eta_r \uparrow G_{\mathbf{m}}) &= \textstyle\sum_{r\downarrow G^0_{\mathbf{m}'}\ni 1} m_0 d_r \\ &= m_0 \textstyle\sum_{r\downarrow G^0_{\mathbf{m}'}\ni 1} d_r \\ &= m_0(m_1\cdots m_n)\end{aligned}$$

where the last line follows by induction.

■

Using the work of [16], Theorem 7 implies the following result.

Corollary 4 *For $f \in L(X_{\mathbf{m}})$, the isotypic decomposition can be computed in $O(X^2_{\mathbf{m}})$ operations.*

5.2 Further directions

Potential applications for the algorithms developed in this paper may also be found in molecular spectroscopy as well as in signal processing. Detailed study of these possibilities is in progress. The general ideas are sketched below.

5.2.1 Molecular spectroscopy

Historically, group theory and representation theory have helped to simplify the analysis of many physical systems (cf. [28]). The general idea is that symmetry-invariant properties are interesting and lead to physically intersting results. The symmetry of interest may be nuclear (e.g. spin-symmetry) time-symmetry, or geometric symmetry, just to name a few examples.

For molecular spectroscopy, the notion of symmetry is encapsulated in the requirement that a group G of unitary matrices commute with a particular Hamiltonian $\mathcal{H}$, the choice of which depends on the assumed physical dynamics. In many cases of interest it is enough to restrict attention to finite groups and thus, the Hamiltonian becomes an interwining operator for a particular matrix representation of a finite group G.

Assume that $\mathcal{H}$ acts on a vector space V which supports a representation space for G. Representation theory can be used to determine the objects of interest, the eigenvalues and eigenvectors of $\mathcal{H}$. On the one hand, any eigenspace of $\mathcal{H}$ will be a G-invariant subspace of V. As a partial converse to this, understanding the representation of G on V can also simplify the computation of the eigenvalues and eigenvectors of $\mathcal{H}$. For this, consider the isotypic decomposition of V with respect to G,

$$V = V_0 \oplus V_1 \oplus \cdots \oplus V_n.$$

As the distinct V_i have no representations in common, $\mathcal{H}$ must preserve the different V_i. Thus, in some basis, $\mathcal{H}$ becomes block diagonal and the eigenvectors must themselves be within isotypics. In particular, note that if a given V_i is itself irreducible, then Schur's Lemma implies that every element of V_i will be an eigenvector for $\mathcal{H}$.

In particular, direct products of wreath products of the form $\mathbb{Z}/m\mathbb{Z}[S_n]$ and their iterates have been identified as the symmetry groups of "non-rigid" molecules [38] and their induced permutation action on the atoms in these molecules is shown to commute with the NMR Hamiltonian [4]. In the most elementary setting, this permutation action is effectively the action of $\mathbb{Z}/m\mathbb{Z}[S_n]$ on the leaves of the spherically homogeneous tree $X_{m,n}$ (cf. Figure 2). Analysis of this action is analogous to that of Section 5.1. Consequent block diagonalization of any operator commuting with this action follows quickly, allowing then for possible simplification of the determination of normal modes for the Hamiltonian. A more detailed analysis of this sort of application is in progress.

5.2.2 Group filters

It has been suggested that applications of generalized FFT research in signal processing may come from the construction of "nonabelian group filters". Intriguing recent work of Holmes [23], [24], Trachtenberg, Karpovsky [26] and others have proposed that certain DFT's for nonabelian (noncommutative) groups may be employed as potential approximants of the Karhunen-Loeve transform for particular classes of stochastic signals. To formulate the problem, suppose a given problem of signal estimation involves N-dimensional data. Let G denote any group of order N, with a fixed ordering of group elements $G = \{s_1, \ldots, s_N\}$ and transversal of $\hat{G}$, $\mathcal{R} = \{\rho^{(1)}, \ldots, \rho^{(h)}\}$, with $\rho^{(i)}$ of dimension d_i. Since $\sum_i d_i^2 = |G|$, consider the $|G| \times |G|$ matrix $\mathcal{F}_{G,\mathcal{R}}$ with entries given by the the matrix coefficients $\rho_{jk}^{(i)}(s_l)$. The matrix $\mathcal{F}_{G,\mathcal{R}}$ has columns indexed by group elements and rows by the multi-index i, j, k. Thus, the first d_1^2 rows will contain the matrix coefficients $\rho_{jk}^{(1)}(s_l)$, the second d_2^2 the matrix coefficients $\rho_{jk}^{(2)}(s_l)$ and so on. It is straightforward to see that for any $h \in L^2(G)$, the product

$$\mathcal{F}_{G,\mathcal{R}} \cdot \begin{pmatrix} h(s_1) \\ \vdots \\ h(s_N) \end{pmatrix} \tag{26}$$

computes $DFT_{\mathcal{R}}(h)$, with consecutive blocks of d_i^2 entries of the product (5) consisting of the entries in $\hat{f}(\rho^{(i)})$. The Fourier inversion formula (23) shows that $\mathcal{F}_{G,\mathcal{R}}$ is invertible.

A **group filter** (from G) is the application to a data vector f of an operator of the form

$$\mathcal{F}_{G,\mathcal{R}}^{-1} \cdot D \cdot \mathcal{F}_{G,\mathcal{R}}. \tag{27}$$

In (6), D is an $N \times N$ block diagonal matrix with i^{th} diagonal block an arbitrary $d_i \times d_i$ matrix. This ordering assumes that $\mathcal{F}_{G,\mathcal{R}}$ has the form described above. The filter is said to be nonabelian if G is nonabelian. The application of this sequence of operators has the following interpretation: view the data or signal f as a function on G, then application of the group filter (6) is equivalent to computing the convolution of f with the function h which has $\hat{h}(\rho^{(i)}) = D_i$, the i^{th} diagonal block of D (cf. Section 4.3).

If G is abelian then each $\rho^{(i)}$ is unique and one-dimensional. In this case the Fourier transform is computed quickly and stably by the usual abelian FFT. However, if G is nonabelian then there are an infinite number of choices for $\mathcal{R}$. The choice is usually dictated by computational considerations. However, as discussed in Sections 1 and 4, it is now known that for a given N, there may be many nonabelian groups whose DFT's enjoy the same computational advantages as the abelian FFT. This opens the door for many new potential avenues of investigation for matched filters using nonabelian group transforms. In particular, the results of Section 4 show wreath products are a possible sources of efficient group filters. Given the computational advantages of wreath product transforms, it is of interest to determine classes of stochastic signals for which the wreath product FFT's act as nearly optimal filters.

5.2.3 Data compression

Holmes also points out [23] that nonabelian FFT's provide interesting possibilities for data compression. To explain his ideas, again assume N-dimensional data is given by a signal x of constant (although not necessarily stationary) variance and let Σ_x denote the associated covariance matrix. Then for any $N \times N$ matrix U consider

$$\Sigma_x(U) = U^* \Sigma_x U.$$

When $U = \mathcal{F}_{G,\mathcal{R}}$ (recall the notation of Section 3.2.3) write $\Sigma_x(G, \mathcal{R})$ for $\Sigma_x(\mathcal{F}_{G,\mathcal{R}})$ for any group G of order N and transversal $\mathcal{R}$ of $\hat{G}$.

Compression is measured by the geometric mean of the diagonal of $\Sigma_x(U)$. This is known to be minimal for U a discrete Karhunen-Loeve matrix for Σ_x. Instead taking $U = \mathcal{F}_{G,\mathcal{R}}$, one arrives at perhaps suboptimal compression schemes, but ones which may enjoy certain computational advantages over Karhunen-Loeve. Consequently, given the complexity results of Section 4, it is now of interest to understand the compression characteristics of wreath product FFT's.

References

[1] R. Bailey. Strata for randomized experiments, *J. Roy. Statist. Soc. Ser. B*, **55**, 27-78 (1991).

[2] R. Bailey, C. E. Praeger, C. A. Rowley, and T. Speed. Generalized wreath products of permutation groups, *Proc. London Math. Soc.* **47**(3), 69-82 (1983).

[3] L. Babai and L. Rónyai. Computing irreducible representations of finite groups. *Math. Comp.* **55**, 705-722 (1990).

[4] K. Balasubramanian. Graph theoretical characterization of NMR group nuclear spin species and the construction of symmetry adapted NMR spin functions. *J. Chem. Phys.* **73** No. 7, 3321-3337 (1980).

[5] U. Baum. Existence and efficient construction of fast Fourier transforms for supersolvable groups. *Computational Complexity*, **1**,235-256 (1991).

[6] U. Baum and M. Clausen. Some lower and upper complexity bounds for generalized Fourier transforms and their inverses. *SIAM J. Comput.* **20** No. 3, 451-459 (1991).

[7] T. Beth. *Verfaren der schnellen Fourier-Transformation*, Teubner Studienbücher, Stuttgart (1984).

[8] N. Bshouty, M. Kaminski, and D. Kirkpatrick. Addition requirements for matrix and transposed matrix products. *J. of Algorithms* **9**, 354-364 (1988).

[9] M. Clausen and U. Baum. *Fast Fourier Transforms*, Wissenschaftsverlag, Manheim (1993).

[10] M. Clausen. Fast Fourier transforms for metabelian groups. *SIAM J. Comput.* **18**, 584-593 (1989).

[11] M. Clausen. Fast generalized Fourier transforms. *Theor. Comp. Sci.* **67**, 55-63 (1989).

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

[13] C. Curtis and I. Reiner. *Representation Theory of Finite Groups and Associative Algebras*, Wiley Classics Library, J. Wiley and Sons, NY (1988).

[14] P. Diaconis. *Group Representations in Probability and Statistics*, IMS, Hayward, CA (1988).

[15] P. Diaconis and D. Rockmore. Efficient computation of the Fourier transform on finite groups. *Journal of the American Mathematical Society* **3** No. 2, 297-332 (1990).

[16] P. Diaconis and D. Rockmore. Efficient computation of isotypic projections for the symmetric group. *DIMACS Series in Disc. Math., Vol. 11* (eds. L. Finkelstein and W. Kantor), 87-104 (1993).

[17] J. R. Driscoll and D. Healy. Computing Fourier transforms and convolutions on the 2-sphere. *Adv. in Appl. Math.*, To appear.

[18] J. R. Driscoll, D. Healy and D. Rockmore. Fast spherical transforms for distance transitive graphs. *Technical Report, Department of Mathematics and Computer Science, Dartmouth College* (1989).

[19] D. F. Elliott and K. R. Rao. *Fast Transforms: Algorithms, Analyses, and Applications.* Academic, New York (1982).

[20] P. Fortini. *Representations of Groups and the Analysis of Variance.* Ph.D. Thesis, Dept. of Statistics, Harvard University.

[21] E. J. Hannan. Group representations and applied probability. *J. Appl. Prob.* **2** (1965), 1-68.

[22] D. Healy, S. Moore, and D. Rockmore. Efficiency and stability issues in the numerical computation of Fourier transforms and convolutions on the 2-sphere. *Technical Report, Department of Mathematics and Computer Science, Dartmouth College* (1993).

[23] R. Holmes, Mathematical foundations of signal processing, II. *Technical Report 781, MIT, Lincoln Laboratory* (1987).

[24] R. Holmes, Signal processing on finite groups, *Technical Report 873, MIT, Lincoln Laboratory* (1990).

[25] G. D. James. *The Representation Theory of the Symmetric Groups.* Lecture Notes in Mathematics., Vol. 682, Springer-Verlag, Berlin (1978).

[26] M. Karpovsky and E. Trachtenberg, Filtering in a communication channel by Fourier transforms over finite groups, in *Spectral Techniques and Fault Detection*, M. Karpovsky (ed.) Academic Press, NY, pp. 179-212 (1985).

[27] A. Kerber. *Representations of Permutations Groups I, II* Lecture Notes in Mathematics, Vols. 240 and 495, Springer-Verlag, Berlin (1971) and (1975).

[28] P. Kramer and M. D. Cin (eds.), *Groups, Systems and Many-Body Physics.*, Vieweg, and Sohn Verlagsgesellschaft mbH, Braunschweig, (1980).

[29] D. Maslen. *Fast Transforms and Sampling for Compact Groups*, Ph.D. Thesis, Department of Mathematics, Harvard University, (1993).

[30] S. Moore, D. Healy, and D. Rockmore. Symmetry stabilization for polynomial evaluation and interpolation. *Lin. Alg. Appl.* **192**, 249-299 (1993).

[31] J. A. Nelder. The analysis of randomized experiments with orthogonal block structure, I and II. *Proc. of the Royal Acad., Series A* **283**, 147-162 and 163-178 (1965).

[32] D. Rockmore. Efficient computation of Fourier inversion for finite groups. *Jour. of the ACM* January (1994).

[33] D. Rockmore. Fast Fourier analysis for abelian group extensions. *Adv. in Appl. Math.* **11**, 164-204 (1990).

[34] J. P. Serre. *Linear Representations of Finite Groups.* Springer-Verlag, New York (1977).

[35] R. Tolimieri, M. An, and C. Lu. *Algorithms for Discrete Fourier Transform and Convolution.* Springer-Verlag, New York (1989).

[36] C. Van Loan. *Computational Framework for the Fast Fourier Transform*, SIAM, Philadelphia (1992).

[37] C. Wells. Some applications of the wreath product construction. *Amer. Math. Monthly* **83**, 317-338 (1976).

[38] C. M. Woodman, The symmetry groups of non-rigid molecules as semi-direct products, *Mol. Phys.* (6) **19** (1970), 753-780.