RECENT RESULTS IN GENERALIZED FAST FOURIER TRANSFORMS

ERIC J. MALM

1. Introduction

Spectral analysis methods play a crucial role in mathematics today and in the pure and applied sciences. For example, the study of Fourier series concerns itself with the decomposition of a periodic function f into a series of complex exponentials [11]:

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx} \tag{1.1}$$

The amplitudes c_n of these exponentials then constitute the spectrum of the function. Such decompositions have applications to linear partial differential equations, where the exponential functions act as a basis of eigenfunctions of the linear operator associated with the equation. The eigenvalue spectrum of the operator then determines how the solution to the equation depends on the initial and boundary conditions [11, 13]. This series decomposition can be extended to nonperiodic functions by allowing a continuous distribution of frequencies for the constituent complex exponentials, so that a function $f(x)$ decomposes as

$$f(x) = \int_{-\infty}^{\infty} g(\omega) e^{i\omega x}\, d\omega \tag{1.2}$$

The spectrum of amplitudes $g(\omega)$ is then the Fourier transform of the function $f(x)$. This terminology also illustrates that the Fourier transform itself is a map from one space of functions to another space of functions, possibly over a different domain.

Spectral methods, and the Fourier transform in particular, have significant applications in the physical sciences. In quantum mechanics, for instance, the eigenvalue spectrum of a Hermitian operator such as the Hamiltonian or the angular momentum operator determines the range of observable values for the physical quantity corresponding to that operator. The position-space wave function $\psi(x)$ of a particle describes the distribution of its possible positions states, and can be rewritten in terms of its momentum-space wave function $\hat{\psi}(p)$ in what is essentially a Fourier transform [24]:

$$\psi(x) = \frac{1}{\sqrt{2\pi\hbar}} \int_{-\infty}^{\infty} \hat{\psi}(p) e^{ipx/\hbar}\, dp.$$

The physical significance of these transforms arises from the natural duality between quantities such as position and momentum and energy and time, which also underlies the famous Heisenberg uncertainty relations $\Delta x \Delta p \geq \hbar/2$ and $\Delta E \Delta t \geq \hbar/2$.

From an engineering perspective, much of modern signal processing concerns determining not only the spectrum of a given input signal, but also how that spectrum will change when the signal passes through particular systems and how to design systems that amplify or isolate certain

Date: October 8, 2004.

portions of the spectrum. Of particular importance to signal processing is the Discrete Fourier Transform (DFT), which converts a function on N evenly spaced points to N amplitudes associated to certain frequencies. In keeping with the terminology above, these amplitudes are called the Fourier coefficients of the function. This transform allows discrete samples of a continuous signal to yield some information about the spectrum of that signal. Because computational methods operate primarily on such discrete data, the DFT has become ubiquitous in modern signal processing.

Naïve implementations of the DFT require $O(N)$ operations for the construction of each coefficient, resulting in an overall $O(N^2)$ algorithm for the DFT. This $O(N^2)$ complexity severely limits the application of the DFT to large data sets and motivates the search for more efficient implementations of the transform. Any such efficient implementation of the DFT is called a Fast Fourier Transform (FFT). Such FFTs trace back even to Gauss, who determined an efficient interpolation of a planetary orbit between n points from its interpolation on two sets of $n/2$ points. Modern FFTs are derived from the algorithm that Cooley and Tukey developed in the 1960s [5, 21], which computes a DFT on $N = pq$ points first by p transforms of length q and then by q transforms of length p, for a total of $pq(p+q)$ operations. Applied recursively in the case where $N = 2^n$, this algorithm yields a complexity of $O(N \log N)$ for the N-point DFT, a significant improvement over the naïve $O(N^2)$ implementation.

2. Group-Theoretical Fourier Transforms

As discussed above, given a continuous periodic function $f : \mathbb{R} \to \mathbb{C}$, we can determine some of its frequency spectrum by sampling the function at N points $x_0, x_1, \ldots, x_{n-1}$ over one of its periods and computing the DFT on those points. One of the key features of the DFT is that the amplitudes and relative phases of the coefficients it yields do not depend on the time period over which the function was sampled. In particular, this indicates that the DFT is invariant under cyclic shifts of the points $X = \{x_0, x_1, \ldots, x_{n-1}\}$. Such shifts correspond to the action of the group $\mathbb{Z}/n\mathbb{Z}$ on this set X. We can then replace the points of X with the elements of $\mathbb{Z}/n\mathbb{Z}$. Since $\mathbb{Z}/n\mathbb{Z}$ also shifts itself in the same manner, this replacement preserves that action of $\mathbb{Z}/n\mathbb{Z}$ on the set over which the function is defined. Finally, we can treat the function $f : \mathbb{Z}/n\mathbb{Z} \to \mathbb{C}$ as an element of the group $\mathbb{C}$-algebra $\mathbb{C}[\mathbb{Z}/n\mathbb{Z}]$, where the coefficient of $g \in \mathbb{Z}/n\mathbb{Z}$ is $f(g)$. Therefore, there exists a natural mapping of the function $f : X \to \mathbb{C}$ into the $\mathbb{C}[\mathbb{Z}/n\mathbb{Z}]$, and this mapping provides a representation-theoretic interpretation of the DFT and an avenue for its generalization to arbitrary groups.

Suppose G is a finite group, and let V be a $\mathbb{C}[G]$-module, so that V is a representation for G. Let $\hat{G}$ denote the set of h equivalence classes of irreducible representations of G, and let $\rho_1, \ldots, \rho_h$ be representatives of these equivalence classes. Then V decomposes as

$$V = \bigoplus_{i=1}^{h} V^i. \tag{2.1}$$

Furthermore, each V^i is isomorphic to a direct sum of n_i isomorphic copies of ρ_i, so that $V^i \cong n_i \rho_i$. These V^i are referred to as the isotypic subspaces of V and are unique for each $C[G]$-module V. In particular, $\mathbb{C}[G]$ is itself a $\mathbb{C}[G]$-module by left multiplication, and so it too decomposes into a direct sum of its irreducible representations. The isotypic components in this decomposition correspond to the minimal two-sided ideals of $\mathbb{C}[G]$, while the irreducible representations correspond to its minimal left ideals. Thus, each element $f \in \mathbb{C}[G]$ can be written as a unique sum of elements in the representations constituting this direct sum, and these elements provide a generalizations of the Fourier coefficients obtained by the usual DFT.

These coefficients can be written explicitly by considering each irreducible representation ρ_i as a vector space of $\mathbb{C}$ of dimension equal to the degree d_i of the representation. Then each component of $f \in \mathbb{C}[G]$ corresponds to a vector in $\mathbb{C}^{d_i}$, each isotypic component of f corresponds to a matrix in $\mathbb{C}^{d_i \times d_i}$, and we can write our isomorphism $\mathbb{C}[G] = \bigoplus_{i=i}^{h} V^i \cong \bigoplus_{i=1}^{h} d_i \rho_i$ as

$$\mathbb{C}[G] \cong \bigoplus_{i=1}^{h} \mathbb{C}^{d_i \times d_i}. \tag{2.2}$$

This result restates Wedderburn's Theorem [5, 9]: every complex group algebra is isomorphic to a direct sum of matrix algebras. Because we have a choice of basis for each $\mathbb{C}^{d_i \times d_i}$, this isomorphism is not unique, and any choice of isomorphism D from $\mathbb{C}[G]$ into this direct sum is called a DFT for the group G. Combined with a choice of basis for $\mathbb{C}[G]$, we can reformulate the DFT as a $|G| \times |G|$ matrix from the coordinate representation of $f \in \mathbb{C}[G]$ to the coordinate representation of its transform $\hat{f} \in \bigoplus_{i=1}^{h} \mathbb{C}^{d_i \times d_i}$. This matrix form also indicates that the DFT is a linear transformation from the input function to the coefficients.

Given a particular DFT D for a group G and an element $f = \sum_{g \in G} f_g g \in \mathbb{C}[G]$, we can compute the Fourier coefficients in block j of the direct sum as

$$\hat{f} = \sum_{g \in G} f_g \rho_j(g), \tag{2.3}$$

where ρ_j is the irreducible representation corresponding to that block, expressed as a matrix in terms of the basis chosen for the DFT D.

2.1. **Abelian Groups.** We now illustrate how this concept of generalized Fourier transforms encompasses the familiar cases of spectral analysis discussed in Section 1, all of which then correspond to Fourier transforms on abelian groups. In the case of a finite abelian group G such as $\mathbb{Z}/N\mathbb{Z}$, each irreducible representation of G has degree 1, so that the isotypic subspaces of $\mathbb{C}[G]$ are all one-dimensional. Therefore, Wedderburn's Theorem states that

$$\mathbb{C}[G] \cong \bigoplus_{i=1}^{|G|} \mathbb{C}^{1 \times 1} \cong \mathbb{C}^{|G|}, \tag{2.4}$$

and the Fourier coefficients of $f \in \mathbb{C}[G]$ are simply the coordinates in $\mathbb{C}^{|G|}$ of the image of f under this isomorphism.

We can also reformulate the Cooley-Tukey FFT on $N = pq$ points in terms of a factorization of the group $\mathbb{Z}/N\mathbb{Z}$. In particular, we introduce the chain of subgroups $1 < \mathbb{Z}/p\mathbb{Z} < \mathbb{Z}/N\mathbb{Z}$. The analysis presented here closely follows that presented by Maslen and Rockmore [16]. The N irreducible representations of $\mathbb{Z}/N\mathbb{Z}$ are all one-dimensional and hence are equal to its irreducible characters. These characters are given by $\zeta_k(j) = \omega^{jk}$, where ω is a primitive Nth-root of unity. Thus, by the formula given above for the coefficients, the kth Fourier coefficient of an element $f = \sum_j f_j j \in \mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ is given by

$$X_k = \sum_{j=0}^{N-1} f_j \zeta_k(j) = \sum_{j=0}^{N-1} f_j \omega^{jk}. \tag{2.5}$$

We can reindex the sum into a double sum via the group factorization specified above. Each $j \in \mathbb{Z}/N\mathbb{Z}$ can be written as an element of a coset of the subgroup $q\mathbb{Z}/N\mathbb{Z} \cong \mathbb{Z}/p\mathbb{Z}$, so that $j = i_1 q + i_2$. Let A be a complete set of the representatives for the cosets of $q\mathbb{Z}/N\mathbb{Z}$ in $\mathbb{Z}/N\mathbb{Z}$. Then the sum above becomes

$$X_k = \sum_{a\in A}\sum_{b\in q\mathbb{Z}/N\mathbb{Z}} \zeta_k(a+b) f_{a+b} = \sum_{a\in A}\sum_{b\in q\mathbb{Z}/N\mathbb{Z}} \zeta_k(a)\zeta_k(b) f_{a+b} = \sum_{a\in A}\zeta_k(a)\sum_{b\in q\mathbb{Z}/N\mathbb{Z}} \zeta_k(b) f_{a+b}. \tag{2.6}$$

We note that, in the inner sum, ζ_k acts only on elements of the subgroup $q\mathbb{Z}/n\mathbb{Z}$, so we need only consider the restrictions $(\zeta_k \downarrow q\mathbb{Z}/N\mathbb{Z})$ of the characters ζ_k to this subgroup. Since

$$\zeta_k(i_1 q) = \omega^{i_1 q k} = (\omega^q)^{i_1 k},$$

and ω^q is a primitive pth-root of unity, these restrictions correspond exactly to the p irreducible characters χ_m of $\mathbb{Z}/p\mathbb{Z}$. Consequently, we need compute the inner sum only for $k \in \mathbb{Z}/p\mathbb{Z}$.

We now reformulate the calculation of the Fourier coefficients in two stages, at the expense of some storage space:

- We compute and store $f_1(a,k) = \sum_{b\in q\mathbb{Z}/N\mathbb{Z}} \zeta_k(b) f_{a+b}$ for all $a \in A$ and for all $k \in \mathbb{Z}/p\mathbb{Z}$. Each sum requires p operations, for a total of $p^2 q$ operations.
- We then compute $\sum_{a\in A} \zeta_k f_1(a,k)$ for all $k \in \mathbb{Z}/N\mathbb{Z}$ to obtain the Fourier coefficients. Each sum requires q operations, for a total of $Nq = pq^2$ operations.

The final operation count for these two stages is then $pq(p+q)$ operations, while the total count for the one-stage calculations given by Equation (2.5) is $N^2 = p^2q^2$. This algorithm therefore represents a significant improvement in complexity, and ultimately leads to the $O(N \log N)$ running time of the Cooley-Tukey FFT.

Abelian groups other than the cyclic groups present similar DFTs and FFTs. We consider an example from Maslen and Rockmore [16], called the 2^n-factorial design, which corresponds to functions over $(\mathbb{Z}/2\mathbb{Z})^n$. Such a set of data might arise from the effects of n independent factors on the growth of a crop of plants, where each factor can assume either a high value or a low value. The irreducible representations of $(\mathbb{Z}/2\mathbb{Z})^n$ are, as above, all one-dimensional, and are given by $\chi_v(w) = (-1)^{\langle v,w\rangle}$, where $v, w \in (\mathbb{Z}/2\mathbb{Z})^n$ and $\langle v,w\rangle$ represents the usual inner product taken mod 2. We decompose the computation of the Fourier coefficients by Equation (2.3) into n stages according to the chain of subgroups

$$1 < \mathbb{Z}/2\mathbb{Z} < (\mathbb{Z}/2\mathbb{Z})^2 < \cdots < (\mathbb{Z}/2\mathbb{Z})^{n-1} < (\mathbb{Z}/2\mathbb{Z})^n$$

in a manner similar to the separation performed in the Cooley-Tukey FFT to yield a transform in $3 \cdot 2^n \log 2^n$ operations. The DFT under consideration here is equivalent the well-known Walsh-Hadamard transform, and the FFT algorithm we generate then represents a sparse factorization of the DFT matrix associated with the transform. Much work has already been done on FFTs for abelian groups. The key result, presented by Clausen and Baum [5], is that the combination of the methods of Cooley and Tukey and the so-called chirp-z and Rader transforms yields FFTs for all abelian groups G in fewer than $8|G|\log|G|$ operations.

We can even extend this group theoretic formulation to infinite groups, under certain restrictions which we detail below. Consider a function $f : \mathbb{R} \to \mathbb{C}$ that is 2π-periodic. We reformulate f as a function on the unit circle S^1, which is a compact group. The associated group algebra also has irreducible one-degree representations, although in this case there are an infinite number of them, indexed by $\mathbb{Z}$. If the elements S^1 are represented by the angle $\theta \in [0, 2\pi)$, then the nth irreducible representation is given by $\chi_n(\theta) = \frac{1}{2\pi} e^{-in\theta}$. Thus, the Fourier coefficients f_n are determined by the integral

$$f_n = \int_0^{2\pi} f(\theta)\chi_n(\theta)\, d\theta = \frac{1}{2\pi} \int_0^{2\pi} f(\theta) e^{-in\theta}\, d\theta, \tag{2.7}$$

In general, this integral exists because S^1 is compact and thus has finite volume under the measure $d\theta$. The inverse transform is as specified in Equation (1.1). We often require that the function f be band-limited, so that $f_n = 0$ for $|n| > B$ for some B. This restriction allows us to replace the infinite sum in Equation (1.1) with a finite sum from $-B$ to B. In this case, there exists a sampling method that allows the exact computation of the coefficients from $2B + 1$ samples of the function f [21]; the finite-case FFT then makes such computations efficient.

Finally, the Fourier transform over $\mathbb{R}$ detailed above in Equation (1.2) provides an example of a transform over a noncompact abelian group. As above, the Fourier coefficients are given by an integral, although this time over $\mathbb{R}$:

$$\hat{f}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} f(x) e^{-i\omega x}\, dx. \tag{2.8}$$

This result reflects that there are an infinite number of irreducible representations of $\mathbb{R}$, as in the S^1 case, although now they are indexed by $\mathbb{R}$ instead of by $\mathbb{Z}$. In order that these coefficients $\hat{f}(\omega)$ exist, we must place certain restrictions on f. These include that it be band-limited [21] or that it belong to a special class of function that ultimately decay faster than e^{-x^2} [13]. In the band-limited case, the Fourier coefficients have finite support, so as in the S^1 case, there exist sampling methods that then admit judicious use of the FFT to create efficient transforms.

2.2. **Nonabelian Groups.** We now explore the generalization of these Fourier transforms to the case of nonabelian groups G. Among the finite groups of particular interest are the symmetric groups S_n, the dihedral groups D_{2n}, solvable and supersolvable groups, and finite groups of Lie-type, including several types of matrix groups over finite fields [16, 21]. Such groups have applications that include the analysis of ranked data or the construction of error-correcting codes; see Section 5 for more details on the potential applications of spectral analysis on these groups.

Fourier transforms on finite nonabelian groups are even useful for understanding or manipulating the corresponding group algebras, as multiplication of elements in the group algebra corresponds to a convolution operation on the coefficients of the elements. In turn, this convolution becomes simple pointwise multiplication in the transform space, so the it can be computed through two DFTs, a multiplication, and an inverse DFT. Implemented naïvely, either approach requires $O(|G|^2)$ operations, but an FFT algorithm can reduce the complexity of the transform approach to at worst $O(|G|^{3/2})$ and in some cases $O(|G| \log^c |G|)$, depending on the efficiency of the FFT algorithm itself.

Such applications therefore motivate us to determine how efficient the FFTs for a given group can be. Such questions are usually stated in terms of the complexity $L_s(G)$ of a group G, which is defined to be the minimum of the complexities of all the possible DFT matrices associated with G [3, 5, 16]. A number of bounds on the complexity of nonabelian FFTs have already been established. Clausen [3] states that, for a finite group G, the complexity $L_s(G)$ of a generalized FFT on G is bounded above by

$$L_s(G) \leq \min_{\mathcal{C}} \{(s(\mathcal{C}) - l(\mathcal{C})) \cdot |G| + 7\sqrt{q(\mathcal{C})}|G|^{3/2}\},$$

where the minimum is taken over all possible chains $\mathcal{C}$ of subgroups $1 = G_0 < \cdots < G_n = G$ of G, where $l(\mathcal{C})$ is the length n of the chain, and where q and s are the maximum and sum of the indices

$[G_{i+1} : G_i]$ determined by the chain. While this is a significant improvement over the trivial bound of $2|G|^2$ operations, the existence of $O(|G| \log |G|)$ FFTs for abelian groups demonstrates that this is by no means a sharp bound. Also of interest are lower bounds on the complexity of an FFT, so that we can determine when we have an optimal algorithm. Clausen and Baum [5] state that $O(|G|)$ is the best lower bound that has been proved so far in computational models that allow arbitrarily large multiplications, although if limits are placed on those multiplications the lower complexity bound grows to $O(|G| \log |G|)$.

Better complexity bounds have been determined for specific families of groups. Clausen and Baum [5] cite a result that if G is a solvable group with a monomial DFT, then its complexity is less than $8.5|G| \log |G|$. Since all supersolvable groups meet this criterion, this result applies to them as well. Maslen [14] proves that the symmetric group S_n has an FFT involving $O(|S_n| \log^2 |S_n|)$ operations; further discussion of symmetric group FFTs occurs in Section 4.2. Maslen and Rockmore [16] also demonstrate that the complexity of $GL_n(F_q)$ is bounded above by $\frac{1}{2}2^{2q}q^{2n-2}|GL_n(F_q)|$.

As in the commutative case, these Fourier transforms on finite groups can be extended to a compact group G provided certain constraints apply [21]. In particular, the irreducible representation of G must be finite-dimensional, and square-integrable functions have a countable number of coefficients such that the Fourier decomposition converges. The Fourier coefficients are then computed as integrals over the group with respect to the Haar measure. Among the groups that meet these criteria are the classical compact Lie groups, such as $O(n)$, $SO(n)$, $U(n)$, $SU(n)$, and $Sp(n)$. Maslen [21] has made progress on bounds for transforms of band-limited functions on $U(n)$, $SU(n)$, and $Sp(n)$. Driscoll and Healy, furthermore, treat the 2-sphere S^2 as a homogeneous space of $SO(3)$ to construct an FFT that yields a spherical harmonic decomposition for a band-limited function on S^2.

In the noncompact case, Chirikjian [2] has made some progress with respect to Fourier transforms for the Euclidean motion group $SE(3) = SO(3) \rtimes R^3$, although a general theory of generalized FFTs on noncompact nonabelian groups has not yet been developed. Section 6 addresses current open questions in this and other aspects of FFT research.

3. Algorithmic Approaches to FFTs

3.1. **Decimation-in-Time Algorithms.** We now discuss different methods of constructing FFT algorithms. The majority of current FFT algorithms employ a decimation-in-time or separation of variables approach, in which the elements of the group G are factored according to a particular chain of subgroups $1 = G_0 < G_1 < \cdots < G_n = G$. As in the Cooley-Tukey case, the frequencies are then computed through a series of nested sums. The factorization that the subgroup chain affords reduces the total number of operations that must be performed to compute the sums at each stage.

This separation of variables is expressed graphically in the so-called Bratteli diagram for the chain of subgroups, which depicts how the irreducible representations of each G_i factor when restricted to G_{i-1}. We borrow the notation Ram uses in his explanation of these concepts [20]. In general, if V^ρ is an irreducible $\mathbb{C}[G_i]$-module, then $V^\rho \downarrow G_{i-1}$ will be a direct sum of irreducible $\mathbb{C}[G_{i-1}]$-modules, so that

$$V^\rho \downarrow G_{i-1} \cong \bigoplus_{\sigma \in \hat{G}_{i-1}} c^\rho_\sigma V^\sigma,$$

where $\hat{G}_{i-1}$ is the set of equivalence classes of irreducible representations of G_{i-1}, and c^ρ_σ is the multiplicity of the representation σ in ρ. In the Bratteli diagram, then, the irreducible representations of each subgroup in the chain are depicted as nodes arranged vertically, and c^ρ_σ arrows are drawn from each $\sigma \in \hat{G}_{i-1}$ to ρ. As an illustration, Figure 1 depicts a Bratteli diagram for S_4, which clearly shows, for example, that the restriction of the $(3,1)$ irreducible representation of S_4 to S_3 yields the direct sum of a (3) and a $(2,1)$ irreducible representation for S_3.

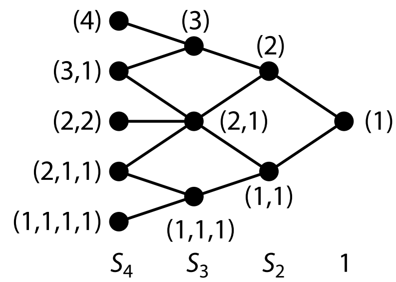

FIGURE 1. The Bratteli diagram for the subgroup chain $S_4 > S_3 > S_2 > 1$.

Consider a path drawn from $\sigma \in \hat{G}_i$ to $\rho \in \hat{G}_n$. This path then represents a G_i-invariant subspace of V^ρ with dimension d_σ. Consequently, the paths from $\emptyset$, the only irreducible representation of 1, to the representations of G represent one-dimensional subspaces and hence correspond to basis vectors for the irreducible representations of G.

In addition, the basis vectors that these paths in the Bratteli diagram represent then interact particularly well with the process of restricting representations to subgroups in the chain. Suppose that $\rho \in \hat{G}$ and V^ρ is an irreducible $\mathbb{C}[G_{n-1}]$ corresponding to ρ with basis $B^\rho = \{v_L\}$ determined by these paths in the diagram. Then restricting V^ρ to G_{n-1} yields a decomposition

$$V^\rho = V^{\sigma_1} \oplus V^{\sigma_2} \oplus \cdots \oplus V^{\sigma_k}$$

such that $\sigma_i \in \hat{G}_{n-1}$. Furthermore, B^ρ partitions into bases B^{σ_i} for these V^{σ_i} that correspond to the paths going from the $\emptyset$ representation at G_0 to σ_i. These bases B^{σ_i} then can be similarly partitioned upon the restriction of the v^{σ_i}s to G_{n-2}, and so on down the chain of subgroups. This basis B^ρ is called a seminormal basis, a Gel'fand-Tsetlin basis, or an adapted basis [16, 20]. This last name indicates that the basis is adapted to the chain of subgroups chosen above.

A seminormal basis for the representations of G provides a convenient form for the matrix representations of $\mathbb{C}[G]$ [20]. Each element in $\mathbb{C}[G]$ can be viewed as a $\mathbb{C}$-linear transformation on the irreducible $\mathbb{C}[G]$-modules V^ρ, and hence has a matrix representation for each choice of basis for these modules. In particular, if $\mathbb{C}^{d_\rho \times d_\rho}$ is the matrix algebra corresponding to $\rho \in \hat{G}$ with respect to the basis $\{v_k\}$, then its ijth coefficient indicates how v_j maps onto v_i. Because the seminormal basis partitions under restriction to $\mathbb{C}[G_i]$, it gives a matrix representation for $\mathbb{C}[G]$ that decomposes into a direct sum of matrix representations for $\mathbb{C}[G_i]$. Furthermore, we can index these matrix coefficients by pairs of basis vectors, or equivalently pairs of paths in the Bratteli diagram that have the same endpoints.

Decimation-in-time algorithms rely upon writing elements of the group G as elements of the double cosets of G_{n-1}, where the coset representatives are drawn from a fixed set A [16]. These representatives are subsequently represented as elements of the double cosets of G_{n-2} and so on until an entire factorization of the group with respect to this chain is reached. Then, just as in the algebraic approach to the Cooley-Tukey algorithm presented above, the computation of the Fourier coefficients can be approached in stages relating to the chosen chain of subgroups. Finally, the seminormal basis ensures that the representations in the sum will restrict to direct sums of representations of subgroups, reducing the number of terms in each sum.

3.2. **Decimation-in-Frequency Algorithms.** Decimation-in-frequency algorithms present an approach to these FFTs that is essentially dual to the decimation-in-time approach. In particular, the same concepts of Bratteli diagrams and seminormal bases are used, but in these algorithms the frequency space is systematically decomposed according to the irreducible representations of the chosen chain of subgroups.

We illustrate this approach with an algebraic description of the Gentleman-Sande FFT [15, 21]. In order to do so, we first discuss the notion of a separating set for a representation V of a group G. Recall that V decomposes into isotypic subspaces V^{ρ_i}, as in Equation (2.1). Consider a set of simultaneously diagonalizable linear transformations $\{T_1, \dots, T_k\}$ on V such that the eigenspaces of the T_js are direct sums of the isotypic components of V. Then applying each T_j to each V^{ρ_i}s yields a list $c_i = (\lambda_{i1}, \dots, \lambda_{ik})$ of the eigenvalues of each T_j on V^{ρ_i}. If $c_i = c_j$ implies that $V^{\rho_i} = V^{\rho_j}$, we say that the T_is form a separating set for V. Essentially, then, the T_is suffice to distinguish among the isotypic components of V.

Consider now the case of a group algebra $\mathbb{C}[G]$ acting on itself as a left-$\mathbb{C}[G]$ module. Then $\mathbb{C}[G]$ is also a representation of G, as discussed in Section 2, and the elements of $\mathbb{C}[G]$ act as linear transformations of $\mathbb{C}[G]$. Thus, we can represent a separating set for $\mathbb{C}[G]$ as a collection of elements of $\mathbb{C}G$. One such separating set is the collection of centrally primitive idempotents $\{e_1, \dots, e_h\}$ that correspond to the two-sided ideals of $\mathbb{C}G$, as e_i has eigenvalue 1 on the ith isotypic component and 0 elsewhere. These idempotents form a basis for the space of conjugacy class sums in $\mathbb{C}G$, so any linear combination of class sums is also a diagonalizable linear transform on $\mathbb{C}[G]$, and any set of them can be diagonalized simultaneously [23]. This result also recommends the set of class sums as a separating set for $\mathbb{C}[G]$ [15].

We now address the DFT from a separating set perspective. We borrow the algebraic formulation of the conventional DFT as an isomorphism of $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ from Section 2.1. Since each of the representations of $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ is one-dimensional, so are the isotypic subspaces of $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$. Hence, they correspond to the Fourier coefficients we seek to recover. The representations are given by $\zeta_j(i) = \omega^{ij}$, where ω is a primitive Nth root of unity. Furthermore, the conjugacy class sum $T_1 = \bar{1}$ separates these isotypic components, since $\zeta_j(\bar{1}) = \omega^j$ and each of these is distinct for distinct j. Thus, each isotypic component V_j corresponds to the eigenspace of T_1 with eigenvalue ω^j. To isolate the Fourier coefficients of $f \in \mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$, we then compute the projections of f onto these spaces. Doing so by the projection formula

$$f_i = \frac{d_i}{|G|} \sum_{g \in G} \chi_i(g)^* \rho(g) \tag{3.1}$$

given in [15] or [23] requires $O(N)$ operations for each of the N coefficients, however, which yields an $O(N^2)$ algorithm.

The Gentleman-Sande FFT, and decimation-in-frequency algorithms in general, take advantage of a chain of subgroups of G to compute the Fourier coefficients in a series of projections, just as the decimation-in-time algorithms construct the coefficients in a series of sums. The idea in the Gentleman-Sande FFT is to first consider the effect of the class sum $T_q = \omega^q$, which is a separating set for $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ as a $\mathbb{C}[q\mathbb{Z}/N\mathbb{Z}] \cong \mathbb{C}[\mathbb{Z}/p\mathbb{Z}]$-module. Thus, we first project $f \in \mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ onto the eigenspaces $W_0, \dots, W_{q-1}$ of T_q, each of which then consists of a direct sum of q isotypic subspaces of $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$:

$$W_k = V_k \oplus V_{k+p} \oplus \cdots \oplus V_{k+(q-1)p}. \tag{3.2}$$

Each projection then takes only $O(pq)$ operations to compute, for a total of $O(p^2q)$ operations. Then the projections via T_1 onto the $N = pq$ isotypics of $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$ as a $\mathbb{C}[\mathbb{Z}/N\mathbb{Z}]$-module take only $O(q)$ operations per coefficient, for a total of $O(pq^2)$ operations [15]. The overall complexity is $O(pq(p+q))$, the same as for the Cooley-Tukey FFT.

Decimation-in-frequency algorithms for a finite group G follow a pattern similar to that of the Gentleman-Sande FFT. If G is nonabelian, however, some of the isotypic subspaces of $\mathbb{C}[G]$ will have dimension greater than one and will no longer correspond directly to the Fourier coefficients of $\mathbb{C}[G]$. Recalling that pairs of paths in the Bratteli diagram correspond to the seminormal basis vectors of $\mathbb{C}[G]$, any set of linear transformations of $\mathbb{C}[G]$ that differentiates between these paths suffices to determine the coefficients.

3.3. **Convolution Algorithms.** Other algorithms have been used in the case of the Cooley-Tukey FFT to increase the efficiency of transforms on $\mathbb{Z}/p\mathbb{Z}$ for large prime p. These groups have no nontrivial subgroups, so they are susceptible to neither the decimation-in-time nor the decimation-in-frequency approaches described above. One useful algorithm in this setting is the Rader transform [5, 16], which relates the DFT on p points to a convolution on $(\mathbb{Z}/p\mathbb{Z})^\times$, which is a cyclic group of order $p-1$. If $p-1$ contains a number of small prime factors, these convolutions themselves are efficient by the usual Cooley-Tukey methods and in turn provide an FFT for these p points. Similarly, the chirp-z transform uses a different change of variables to relate the DFT to a convolution on a larger cyclic group. If this cyclic group has order equal to a power of 2, this convolution can again be performed efficiently by Cooley-Tukey.

4. The Symmetric Group

We now address the representations of the symmetric groups, S_n, and the FFTs these representations have afforded to date. In order to determine such FFTs, we must first

- characterize the irreducible representations of S_n,
- identify a chain of subgroups of S_n such that the irreducible representations of S_n restrict to these subgroups conveniently, and
- determine a basis for $\mathbb{C}S_n$ that is adapted to this chain of subgroups.

Much of this has been done classically, although in recent years reformulations and generalizations of these classical approaches have emerged.

4.1. **Representations of S_n.** Much of the classical work on the representations of the symmetric group was performed by Alfred Young in the late 1920s [22]. James and Kerber [12] modernizes Young's approach significantly and remains a canonical reference on this classical characterization of these representations, and the following material draws heavily from their analysis. At the center of Young's formulation are partitions, Young diagrams, and Young tableaux. A proper partition α of n, written $a \vdash n$, is any nonincreasing sequence $\{\alpha_i\}$ of integers that sum to n; for example, $(4, 2, 1, 1, 0, \dots)$ is a proper partition of 8. Then the Young diagram corresponding to $\alpha \vdash n$ is a left-aligned arrangement of n boxes such that the ith row contains α_i boxes. Finally, a Young tableaux is formed from a Young diagram by placing the numbers 1 through n in the boxes of the diagram. If the numbers increase left to right across rows and down columns, the tableaux is standard.

Each tableaux λ determines a Young subgroup S_λ in the following way. Let S_{λ_i} be the subgroup of S_n that fixes everything except the $|\lambda_i|$ elements in the ith row of the tableaux; this subgroup is isomorphic to $S_{|\lambda_i|}$. Then the Young subgroup S_λ is the direct product of all these S_{λ_i}s, which itself is isomorphic to a direct product of all the nontrivial S_{λ_i}s:

$$S_\lambda = \prod_{i=1}^{\infty} S_{\lambda_i} \cong S_{\lambda_1} \times \cdots \times S_{\lambda_k}.$$

Two Young subgroups with the same diagram (or shape) are then isomorphic.

The proper partitions of n determine the irreducible representations of S_n as follows. As a side note, any field is a splitting field for S_n [12, 19] so we need only consider $\mathbb{Q}[S_n]$ instead of $\mathbb{C}[S_n]$. Let α be a partition of n, and let α' be the complementary partition, such that the ith row of the Young diagram associated to α' contains the number of boxes in the ith column of α. We can construct such a partition graphically by taking the transpose of the Young diagram associated with α. Let IS_α denote the trivial representation of the Young subgroup S_α, and let $AS_{\alpha'}$ denote the alternating representation of $S_{\alpha'}$. Then

$$i(IS_\alpha \uparrow S_n, AS_{\alpha'} \uparrow S_n) = 1,$$

where i is the intertwining number, which measures the dimensionality of the space of $\mathbb{Q}[S_n]$-homomorphisms from $IS_\alpha \uparrow S_n$ to $AS_{\alpha'} \uparrow S_n$ [1]. Since this quantity is one, these two representations share a single copy of an irreducible representation of S_n, which we represent by $[\alpha]$. These representations in fact determine all such irreducibles up to isomorphism. Furthermore, the dimensionality of $[\alpha]$ as a $\mathbb{Q}$-vector space is equal to the number f_α of standard tableaux in the shape α; this is easily shown from the path-algebraic formulation given by Ram [20] or by Okounkov and Vershik [19], which we discuss below.

Finally, it is possible to construct matrix representations of these permutations that are adapted to the natural chain

$$1 < S_2 < S_3 < \cdots < S_{n-1} < S_n \tag{4.1}$$

of subgroups of S_n [12], although the details of the process are complicated and not especially pertinent here. The basis affording this matrix representation is called the Young seminormal basis, and the matrix representation itself is then called the Young seminormal form. Typically, the matrix representations are constructed only for the transpositions $(i-1\ \ i)$, $1 < i \leq n$, as these suffice to generate S_n. This chain of subgroups and associated seminormal basis is precisely what is needed for the construction of FFTs on S_n.

More recently, alternate constructions of these representations have emerged. For example, Jucys and Murphy [17, 18] independently identified elements of $\mathbb{Q}[S_n]$ (now called Jucys-Murphy elements in their honor [19, 20]) that generate these matrix representations in a far simpler fashion. These elements arise as the differences of the class sums of the transpositions in S_{i+1} and S_i. Thus, the first Jucys-Murphy element is $M_1 = (1\,2)$, the second $M_2 = (1\,3) + (2\,3)$, and so on. These elements have other desirable properties: for example, they generate a maximal commutative subalgebra of $\mathbb{Q}[S_n]$, and they are simultaneously diagonalizable in the Young seminormal basis.

In addition, Ram [20] and Okounkov and Vershik [19] present path-algebraic formulations of the seminormal bases for S_n. Their formulations rely on the property that S_n has a multiplicity-free character graph. In particular, if V^ρ is an irreducible representation of S_i, then its restriction to S_{i-1} yields a decomposition into irreducible representations of S_{i-1} that is multiplicity free, so that for each $\sigma \in \hat{S}_{i-1}$, there exists at most one isomorphic copy of V^σ in V^ρ. Thus, the Bratteli diagram for the chain of subgroups in Equation (4.1) contains only single edges between each level of nodes, so for each $\rho \in \hat{S}_n$, each path from $\emptyset$ to ρ is uniquely determined by the list of representations it passes through in the graph. As discussed above, these paths correspond to basis vectors in the seminormal basis for the irreducible representations of S_n, so these basis vectors are also determined by such lists.

Ram [20] then constructs his path algebra from pairs of paths with common endpoints in $\hat{S}_n$, which as discussed above correspond to the seminormal basis vectors of the matrix representation of $\mathbb{C}[S_n]$. In particular, the path algebra has a basis $\{E_{LM}\}$, where L and M have common endpoints, with multiplication defined by

$$E_{ST}E_{PQ} = \delta_{TP}E_{SQ}. \tag{4.2}$$

Thus, the basis vector E_{ST} in the algebra acts as a linear operator that maps the vector v_T corresponding to T onto the vector v_S corresponding to path S. This relation then determines an isomorphism between the algebra and $\mathbb{C}[S_n]$ such that the E_{ST}s map to the elements of the seminormal basis for the matrix representation of $\mathbb{C}[S_n]$. This analysis rests only on the property that the group have a multiplicity-free character graph, in which case we call the group an MC-group. If the chain

$$1 = G_0 < G_1 < \cdots < G_{n-1} < G_n = G$$

provides G with such a character graph, then $\mathbb{C}[G]$ also has a seminormal basis that can be indexed in this fashion. Such MC-groups include the Weyl groups WB_n, WF_4, WE_6, and WE_7. The same holds for any semisimple algebra H exhibiting a similar chain of subalgebras that generates a multiplicity-free character graph. The Iwahori-Hecke algebras constitute one notable family of such algebras that is closely related to these Weyl groups.

To construct these representations for $\mathbb{C}[G]$ explicitly, Ram supposes the existence of central elements $z_{k,j}$ for each subalgebra $\mathbb{C}[G_k]$ arising from the chain of subgroups of G. Since each $z_{k,j}$ is central in $\mathbb{C}[G_k]$, its action on an irreducible representation $\rho^{(k)}$ of G_k will be to scale it by some value $c_{k,j}(\rho^{(k)})$. These central elements then assign a set of weights $\{c_{k,j}(\rho^{(k)})\}$ to each node in each path in the Bratteli diagram, and hence to each vector in the seminormal basis. If this list of weights is distinct for distinct paths, then these lists of weights uniquely determine each seminormal basis vector.

Returning to the symmetric group, Ram identifies the class sum of transpositions

$$z_k = \sum_{1 \le i < j \le k} (i\, j) \tag{4.3}$$

as a central element in the algebra $\mathbb{C}S_k$, and shows that the eigenvalues of these elements at the irreducible representations of the S_is suffice to distinguish paths in the diagram and hence seminormal basis vectors. Furthermore, it suffices to consider the differences $m_k = z_k - z_{k-1}$, which are precisely the Jucys-Murphy elements. In this way, the values of the Jucys-Murphy elements on the irreducible representations of the S_is yield enough information to construct the Young seminormal basis for $\mathbb{C}S_n$. In addition, analogous elements for other groups or other algebras provide similar seminormal representations, which Ram details in subsequent sections of his paper [20].

Finally, this path algebra affords a generalization of the Young tableaux used in the construction of the symmetric group to other MC-groups. Ram [20] explicitly constructs such a system of tableaux for WB_n, while Clausen [4] explores this notion for supersolvable groups as well. In essence, moving from node σ at level i to node ρ at level $i+1$ of the Bratteli diagram adds a box to the tableau of σ according to some rule. For the symmetric group, this rule is that the resulting tableau must represent a proper partition of $i+1$, while for the Weyl group WB_n, a pair of partitions is constructed, each of which must remain proper. Furthermore, the paths from $\emptyset$ to σ in level i correspond to standard tableaux for the shape determined by σ, constructed by placing k in the box added during the move from level $k-1$ to level k. This correspondence also elegantly explains why the degree of each irreducible representation of S_n equals the number of standard tableaux in the associated shape.

4.2. **FFTs on S_n.** To date, significant work has been done on FFTs on the symmetric group from a decimation-in-time perspective. Clausen and Baum produced the first promising results in 1989 with a proof that the complexity of S_n is bounded above by $\frac{1}{2}(n^3+n^2)n!$ operations [3] and in 1993 with a fairly explicit implementation of both a DFT and an inverse DFT for S_n, each requiring that number of operations [6]. Their results arise from a sparse factorization of the DFT matrix based on Young's seminormal form at each S_i in the chain specified in Equation (4.1) [6, 16].

Maslen's 1998 paper [14] improves upon this bound with a decimation-in-time algorithm yielding a DFT on S_n in fewer than $\frac{3}{4}n(n-1)n!$ operations. His method relies on a separation of variables at the scalar level, rather than the matrix separation that Clausen and Baum employ. The commutativity of these scalars allows more sophisticated rearrangement of the sums involved in constructing the Fourier coefficients. This rearrangement entails a more complex indexing scheme based on the paths in the Bratteli diagram rather than only on the subgroups in the chain.

5. Applications

Applications for generalized Fourier transforms abound in engineering, mathematics, and the physical and social sciences. Fourier transforms on the symmetric group have natural applications to the spectral analysis of ranked data, for example. Each voter effectively creates a permutation in S_n by ranking their n candidates, so that the final tallies of votes yield a function on S_n which can be analyzed using the generalized Fourier transforms described above. Diaconis [7] identifies the decomposition of $\mathbb{C}S_n$ into its isotypic components as the key to understanding the effects of candidates on ranking preferences. Such transforms have been applied to partially ranked data as well [21].

The symmetric group is not the only finite group on which Fourier analysis presents applications. The group $SL_n(F_p)$ of two-by-two matrices with determinant one over the finite field F_p has applications in coding theory, particularly with respect to low-density parity check codes, and in graph theory [21]. Maslen, Orrison, and Rockmore [15] discuss applications of generalized Fourier analysis to the study of distance-transitive graphs; while their analysis includes examples that relate primarily to the symmetric group, other groups could also be used in this context. In addition, transforms on other finite groups may yield lossy data compression algorithms with better performance than such standards as JPEG, which is based on the Discrete Cosine Transform [5]. Finally, such transforms have applications to quantum mechanics and quantum computing. In particular, Shor's quantum factoring algorithm relies on transforms on the cyclic group $(\mathbb{Z}/n\mathbb{Z})^{\times}$, and it is conjectured that generalized FFTs may provide an efficient quantum algorithm for the graph isomorphism problem [21].

Fourier transforms on nonabelian compact groups also have significant applications. The spherical harmonics are orthogonal functions on the unit sphere S^2 that yield a series decomposition for functions on S^2 analogous to that provided by the Fourier transform on S^1. Such decompositions have applications in physics, where they play a key role in describing the distributions of electrons in atomic orbitals [10, 24]. In addition, any frequency analysis of spherically distributed data rests on these spherical harmonic functions. Such analysis arises in global circulation modeling, control theory, and computer vision models, for example [16]. As mentioned above, Driscoll and Healy [8] present an efficient algorithm for the computation of spherical harmonics for band-limited functions on S^2 through the analysis of FFTs on the group $SO(3)$, which acts transitively on S^2.

Applications of Fourier transforms on noncompact groups may even exist. Chirikjian and Kyatkin [2] present their analysis of transforms on the Euclidean motion group $SE(3)$ in order to describe the configuration space for certain robotic arms. Such transforms may also apply to the configuration space of proteins as they fold into their appropriate forms and hence may provide a convenient means of describing these folded states [21].

6. Open Questions

We conclude with a number of open questions and directions for future development in the fields of generalized FFTs and noncommutative harmonic analysis. Many of these derive from papers by Maslen and Rockmore [16, 21].

- Although certain classes of groups present $O(|G| \log |G|)$ or $O(|G| \log^2 |G|)$ FFTs algorithms, there exists no universal $O(|G| \log^c |G|)$ complexity bound on generalized FFTs on finite groups. One approach to this problem may involve the determination of FFTs for all the groups in the classification of finite simple groups. In particular, this goal requires better FFTs for finite groups of Lie type and for matrix groups.
- It remains to be seen if decimation-in-frequency algorithms can be generalized to match the level of progress that has been made with decimation-in-time algorithms. These decimation-in-frequency formulations are particularly appealing because their theory more closely reflects the module-theoretic underpinnings of group representation theory.
- Similarly, no noncommutative analogues of the important Rader and chirp-z transforms are currently known. If they exist, such analogues may relate transforms between groups that have no group-subgroup relationship.
- FFTs for groups seem to rest mainly on the semisimplicity of the group algebra $\mathbb{C}[G]$. Because of this result, it seems likely that there exist FFTs for band-limited functions on all semisimple Lie groups.
- Much of the theory of generalized Fourier transforms on noncompact groups such as $SE(n)$ is in its initial stages. Such transforms would require the development of suitable sampling algorithms for these groups as a first step.
- The recursive nature of many of the known FFTs algorithms suggests that there exist effective parallel implementations of these algorithms. Decimation-in-frequency algorithms in particular would seem to admit parallel implementations because of the explicit separation of frequency space they entail.

References

1. William A. Adkins and Steven H. Weintraub, *Algebra: An approach via module theory*, Graduate Texts in Mathematics, no. 136, Springer, New York, 1992.
2. Gregory S. Chirikjian and Alexander B. Kyatkin, *Engineering applications of noncommutative harmonic analysis: With emphasis on rotation and motion groups*, CRC Press, Boca Raton, FL, 2001.
3. Michael Clausen, *Fast generalized fourier transforms*, Theoretical Computer Science **67** (1989), 55–63.
4. ______, *Elements of a general algebraic theory of standard tableaux*, Algebraic Combinatorics and Applications (Berlin) (A. Betten, A. Kohnert, R. Laue, and A. Wassermann, eds.), Springer-Verlag, 2001, pp. 67–78.
5. Michael Clausen and Ulrich Baum, *Fast fourier transforms*, BI-Wissenschaftsverlag, Mannheim, Germany, 1993.

6. ———, *Fast fourier transforms for symmetric groups: Theory and implementation*, Mathematics of Computation **61** (1993), no. 204, 833–847.
7. Persi Diaconis, *A generalization of spectral analysis with application to ranked data*, The Annals of Statistics **17** (1989), no. 3, 949–979.
8. James R. Driscoll and Dennis M. Healy, Jr., *Computing fourier transforms and convolutions on the 2-sphere*, Advances in Applied Mathematics **15** (1994), 202–250.
9. David Dummit and Richard Foote, *Abstract algebra*, John Wiley and Sons, New York, 1999.
10. Robert Eisberg and Robert Resnick, *Quantum physics of atoms, molecules, solids, nuclei, and particles*, John Wiley and Sons, New York, 1985.
11. Kenneth I. Gross, *On the evolution of noncommutative harmonic analysis*, Am. Math. Monthly **85** (1978), 525–548.
12. Gordon James and Adalbert Kerber, *The representation theory of the symmetric group*, Addison-Wesley, Reading, MA, 1981.
13. J. David Logan, *Applied partial differential equations*, Undergraduate Texts in Mathematics, Springer-Verlag, New York, 1998.
14. David K. Maslen, *The efficient computation of fourier transforms on the symmetric group*, Mathematics of Computation **67** (1998), no. 223, 1121–1147.
15. David K. Maslen, Michael E. Orrison, and Daniel N. Rockmore, *Computing isotypic projections with the Lanczos iteration*, SIAM J. Matrix Anal. Appl. **25** (2004), no. 3, 784–803.
16. David K. Maslen and Daniel N. Rockmore, *Generalized FFTs - a survey of some recent results*, DIMACS Series in Discrete Mathematics and Computer Science **28** (1997), 183–237.
17. G. E. Murphy, *A new construction of young's seminormal representation of the symmetric groups*, Journal of Algebra **69** (1981), 287–297.
18. ———, *The idempotents of the symmetric group and nakayama's conjecture*, Journal of Algebra **81** (1983), 58–265.
19. Andrei Okounkov and Anatoly Vershik, *A new approach to represention theory of symmetric groups*, Selecta Mathematica **2** (1996), no. 4, 581–605.
20. Arun Ram, *Seminormal representations of Weyl groups and Iwahori-Hecke algebras*, Proc. London Math. Soc. **73** (1997), 99–133.
21. Daniel N. Rockmore, *Recent progress and applications in group FFTs*, NATO Advanced Study Institute on Computational Noncommutative Algebra and Applications, July 2003.
22. Bruce E. Sagan, *The symmetric group: Representations, combinatorial algorithms, and symmetric functions*, Wadsworth & Brooks/Cole, Pacific Grove, CA, 1991.
23. Jean-Pierre Serre, *Linear representations of finite groups*, Graduate Texts in Mathematics, no. 42, Springer, New York, 1977.
24. John S. Townsend, *A modern approach to quantum mechanics*, University Science Books, Sausolito, CA, 2000.