MATHEMATICS OF COMPUTATION
Volume 67, Number 223, July 1998, Pages 1121–1147
S 0025-5718(98)00964-8

THE EFFICIENT COMPUTATION OF FOURIER TRANSFORMS ON THE SYMMETRIC GROUP

DAVID K. MASLEN

Abstract. This paper introduces new techniques for the efficient computation of Fourier transforms on symmetric groups and their homogeneous spaces. We replace the matrix multiplications in Clausen's algorithm with sums indexed by combinatorial objects that generalize Young tableaux, and write the result in a form similar to Horner's rule. The algorithm we obtain computes the Fourier transform of a function on S_n in no more than $\frac{3}{4}n(n-1)\,|S_n|$ multiplications and the same number of additions. Analysis of our algorithm leads to several combinatorial problems that generalize path counting. We prove corresponding results for inverse transforms and transforms on homogeneous spaces.

1. Introduction

The harmonic analysis of a complex function on a finite cyclic group is the expansion of that function in a basis of complex exponential functions. This is equivalent to the discrete Fourier transform of a finite data sequence, and may be computed efficiently using the fast Fourier transform algorithms of Cooley and Tukey [7] or their many variants (see e.g. [12]).

In the current paper we study the harmonic analysis of a function on the symmetric group. The analogues of the complex exponentials are the matrix entries of a complete set of irreducible complex matrix representations of S_n, called matrix coefficients, and the expansion of functions in this basis may be computed by a generalized Fourier transform on the symmetric group. We describe efficient algorithms for computing the harmonic analysis of a function on the symmetric group, or equivalently, its generalized Fourier transform. Thus our results may be considered a generalization of the fast Fourier transform to the symmetric group. We also present a related algorithm for the harmonic analysis of functions on homogeneous spaces.

Fourier transforms on finite groups have been studied by many authors. The books of Beth [1], Clausen and Baum [3], and the survey article [19] are general references for the computational aspects of these transforms. Rockmore [22] and Diaconis [9] contain discussions of the applications. For applications more specific to symmetric groups, see [8] and [11].

The computation of Fourier transforms on symmetric groups was first studied by Clausen [5] [6], and Diaconis and Rockmore [10], using approaches related to the one taken in the current paper; also see [4] for a detailed discussion of Clausen's

Received by the editor August 21, 1996 and, in revised form, April 23, 1997.
1991 *Mathematics Subject Classification.* Primary 20C30, 20C40; Secondary 65T20, 05E10.
An extended abstract summarizing these results appears in the FPSAC '97 proceedings volume.

algorithm and its implementation. Linton, Michler, and Olsson [16] use a different method that involves the decomposition of Fourier transforms taken at monomial representations.

The algorithms we develop in the current paper are refinements of Clausen's algorithm [5] for computing Fourier transforms on the symmetric group, and of algorithms developed to compute Fourier transforms on compact Lie groups [18]. To describe our main results, let f be a complex function on S_n and let ρ be an irreducible matrix representation of S_n given in Young's orthogonal form (see [15] for terminology). Then the Fourier transform of f at ρ is the matrix sum

$$\hat{f}(\rho) = \sum_{s \in S_n} f(s)\rho(s). \tag{1.1}$$

Computation of the transforms (1.1) at a complete set of irreducible representations in Young's orthogonal form gives us the harmonic analysis of f, because the scaled matrix entry $\frac{\dim \rho}{|S_n|}\left[\hat{f}(\rho)\right]_{ij}$ is the coefficient of the function $(s \mapsto [\rho(s)]_{ij})$ in the expansion of f in the basis of matrix coefficients. We prove the following theorem, which counts the maximum of the numbers of additions and multiplications required to compute a collection of Fourier transforms on S_n.

Theorem 1.1. *The Fourier transform of a complex function on the symmetric group S_n may be computed at a complete set of irreducible matrix representations in Young's orthogonal form in no more than $\frac{3n(n-1)}{4}|S_n|$ multiplications and the same number of additions.*

Note that since $|S_n| = n!$, the number of scalar operations counted in Theorem 1.1 is $O((\log |S_n|)^2 |S_n|)$. Although we have stated Theorem 1.1 for Young's orthogonal form, we actually prove a more general result that applies, e.g., to Young's seminormal form as well. Results on the complexity of the corresponding inverse Fourier transform follow immediately by considering the transpose of our algorithms.

Any complex function on a homogeneous space may also be considered to be a function on a group which is constant on cosets. In this way we may apply Fourier analysis on the group to functions on any homogeneous space. We prove the following theorem concerning the expansion of functions on homogeneous spaces.

Theorem 1.2. *The Fourier transform of a complex function on the homogeneous space S_n/S_{n-k} may be computed at a complete set of (class-1) irreducible matrix representations in Young's orthogonal form in no more than $\frac{3k(2n-k-1)}{4}|S_n/S_{n-k}|$ multiplications and the same number of additions.*

There are several novel features of our approach to the computation of Fourier transforms. One is the use of a kind of commutativity in the group algebra of the symmetric group that lets us replace an iterated group algebra product by a sequence of bilinear maps. This allows us to write an expression for the Fourier transform in a form similar to Horner's rule, and leads to an efficient algorithm.

Another interesting feature is the appearance of certain combinatorial objects that generalize Young tableaux. It is well known that Young tableaux may be associated with sequences of partitions, each obtained by adding a box to the Young diagram of the previous one. This corresponds to an upward walk in a partially ordered set called Young's lattice (see [24] and [25] for a discussion of combinatorial problems associated with these and other walks). In the current paper we encounter sequences of partitions that satisfy more general relations corresponding to the mapping of a multiply-connected graph into Young's lattice. In joint work with Dan Rockmore, such ideas have been generalized to apply to the computation of Fourier transforms on other finite groups [21].

The organization of the paper is as follows. Section 2 contains background from the theory of Fourier transforms on finite groups. Section 3 contains the proof of the main theorem modulo several lemmas that are proven in Section 4. In Section 5 we prove several combinatorial lemmas, and give an exact operation count for our algorithm. In Section 6 we turn our attention to homogeneous spaces, and finally, we conclude in Section 7.

Although we have tried to make the paper relatively self-contained, we do use a number of facts from representation theory that may be found in the books of Serre [23], James and Kerber [15], and Macdonald [17]. Background from the theory of computation of Fourier transforms may be found in the book of Clausen and Baum [3], and in the articles [20] and [19].

2. Fourier transforms on finite groups

The Fourier transform of a function on the symmetric group and the usual discrete Fourier transform of a finite data sequence are both special cases of Fourier transforms on finite groups. We refer the reader to Serre's book [23] for the relevant background from representation theory.

Definition 2.1 (Fourier transform)**.** Let G be a finite group and f be a complex-valued function on G.

1. Let ρ be a matrix representation of G. Then the Fourier transform of f at ρ, denoted $\hat{f}(\rho)$, is the matrix sum,
$$\hat{f}(\rho) = \sum_{s \in G} f(s)\rho(s). \tag{2.1}$$
2. Let $\mathcal{R}$ be a set of matrix representations of G. Then the Fourier transform of f on $\mathcal{R}$ is the direct sum,
$$\mathfrak{F}_{\mathcal{R}}(f) = \bigoplus_{\rho \in \mathcal{R}} \hat{f}(\rho) \in \bigoplus_{\rho \in \mathcal{R}} \mathrm{Mat}_{\dim \rho}(\mathbf{C}), \tag{2.2}$$
of Fourier transforms of f at the representations in $\mathcal{R}$.

Fast Fourier transforms, or FFTs, are algorithms for computing Fourier transforms efficiently.

Example 2.2. When $G = \mathbf{Z}/N\mathbf{Z}$ is a cyclic group, the irreducible representations are exactly the complex exponentials $\zeta_j(k) = e^{2\pi ijk/N}$ considered as 1×1 matrix-valued functions. The associated Fourier transform is the usual discrete Fourier transform, which may be computed by the fast Fourier transform algorithms of Cooley and Tukey [7] and others.

When defining the arithmetic complexity of computing a Fourier transform, we must allow for the possibility that the number of operations depends on the specific matrix representations used, and not just on their equivalence classes under change of bases. The reduced complexity is a related quantity, which is usually easier to work with.

Definition 2.3 (Complexity)**.** Let G be a finite group, and $\mathcal{R}$ be any set of matrix representations of G.

1. The complexity of the Fourier transform on the set $\mathcal{R}$, denoted $T_G(\mathcal{R})$, is the minimum number of arithmetic operations needed to compute the Fourier transform of f on $\mathcal{R}$ via a straight-line program for an arbitrary complex-valued function f defined on G.
2. The reduced complexity $t_G(\mathcal{R})$ is defined by
$$t_G(\mathcal{R}) = T_G(\mathcal{R})/\left|G\right|.$$

When there is no possibility of confusion, we will drop the '$\mathcal{R}$' in the notation for complexities and reduced complexities.

We will always define the number of arithmetic operations counted by Definition 2.3 to be the maximum of the number of complex multiplications and the number of complex additions, though for many of our algorithms these two numbers are the same. When the representations in $\mathcal{R}$ are unitary, all the multiplications occurring in our Fourier transform algorithms are by numbers of magnitude no greater than 1, so our results may be interpreted in terms of the 2-linear complexity of $\mathfrak{F}_{\mathcal{R}}$; see [3] Chapter 3. The recent book of P. Bürgisser, M. Clausen, and A. Shokrollahi [2] is a general reference for algebraic complexity theory that includes applications to Fourier transforms on groups.

A direct approach to computing a Fourier transform at a complete set of inequivalent irreducible matrix representations, using (2.1), gives the upper and lower bounds,
$$|G| - 1 \le T_G(\mathcal{R}) \le |G|^2.$$

2.1. **The group algebra.** Let G be a finite group. Then the group algebra $\mathbf{C}[G]$ is defined to be the space of all formal complex linear combinations of group elements, with the product defined by
$$\Big(\sum_{s\in G} f(s)s\Big)\cdot\Big(\sum_{t\in G} h(t)t\Big) = \sum_{s,t\in G} f(s)h(t)s\cdot t.$$
Elements of $\mathbf{C}[G]$ may be identified with functions on the group in the obvious way, and the algebra product corresponds to convolution of functions.

The most important case of Fourier transform arises when the set $\mathcal{R}$ is a complete set of inequivalent irreducible matrix representations of G. In this case the Fourier transform is an algebra isomorphism from the group algebra $\mathbf{C}[G]$, defined by functions on G, to a direct sum of matrix algebras,

$$\mathfrak{F}_{\mathcal{R}} : \mathbf{C}[G] \xrightarrow{\leftrightarrow} \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{\dim\rho}(\mathbf{C}) \tag{2.3}$$

Definition 2.4. Assume $\mathcal{R}$ is a complete set of inequivalent irreducible matrix representations of G. Then the inverse image of the natural basis of $\bigoplus_{\mathcal{R}} \mathrm{Mat}_{\dim\rho}(\mathbf{C})$ under the Fourier transform $\mathfrak{F}_{\mathcal{R}}$, is called the dual matrix coefficient basis for $\mathbf{C}[G]$ associated to $\mathcal{R}$.

Lemma 2.5 (cf. [5])**.** *The computation of the Fourier transform $\mathfrak{F}_{\mathcal{R}}f$ at a complete set of irreducible representations $\mathcal{R}$ is the same as computation of the sum*

$$\sum_{s\in G} f(s)s \tag{2.4}$$

in the group algebra, relative to the dual matrix coefficient basis associated to $\mathcal{R}$.

Proof. This holds by linear algebra, since by definition $\mathfrak{F}_{\mathcal{R}}$ is the change of basis map from functions on G represented by their function values to functions expressed in the dual matrix coefficient basis. □

For us, the group algebra is mainly a convenient notation for dealing with all irreducible representations of the group G at the same time. In particular, computation of a product $a \cdot b$ in the group algebra relative to the dual matrix coefficient basis is the same thing as computing the collection of matrix multiplications $\rho(a)\rho(b)$ for all ρ in $\mathcal{R}$. In Section 4 we shall identify the group algebra with its coordinate realization in the dual matrix coefficient basis. The problem we then face is to compute the sum (2.4) given the function values $f(s)$ and expressions for the group elements s in coordinates.

2.2. **Adapted representations.** In order to derive more efficient algorithms for computing Fourier transforms, we will need to place conditions on the set of matrix representations $\mathcal{R}$ used. We now define a property that allows us to relate the computation of a Fourier transform to a collection of Fourier transforms on a subgroup.

Definition 2.6 (Adapted representations)**.** Assume G is a finite group, and $\mathcal{R}$ is a set of matrix representations of G.

1. Assume K is a subgroup of G. Then $\mathcal{R}$ is K-adapted, if there is a set $\mathcal{R}_K$ of inequivalent irreducible matrix representations of K, such that for each $\rho \in \mathcal{R}$ the restricted representation $\rho\!\downarrow\! K$ is a matrix direct sum of representations in $\mathcal{R}_K$.
2. The set of representations $\mathcal{R}$ is adapted to the chain of subgroups,

$$G = K_n \geq K_{n-1} \geq \cdots \geq K_0 = 1, \tag{2.5}$$

provided that $\mathcal{R}$ is K_i-adapted for each subgroup in the chain.

Any restricted representation is always conjugate to a direct sum of irreducible representations by complete reducibility (cf. [23] Section 1.4). In Definition 2.6 we require the restricted representation to be equal to a matrix direct sum of irreducibles. Note that if $\mathcal{R}$ is K-adapted, then the set $\mathcal{R}_K$ is uniquely determined.

Systems of Gel′fand-Tsetlin bases are an equivalent concept to adapted sets of matrix representations. Let $\tilde{\mathcal{R}}$ be a set of finite dimensional representations of G. Then a collection of bases of the representation spaces of $\tilde{\mathcal{R}}$ (one basis for each representation) is called a system of Gel′fand-Tsetlin bases for $\tilde{\mathcal{R}}$ relative to the chain (2.5) if the set of matrix representations $\mathcal{R}$ obtained by writing the representations of $\tilde{\mathcal{R}}$ in coordinates relative to these bases is adapted to (2.5).

Systems of Gel′fand-Tsetlin bases were first defined in [13] for the calculation of the matrix coefficients of compact groups. The application to the efficient computation of Fourier transforms on finite groups was first noticed by Clausen [5], [6].

Example 2.7. If G is abelian, K is any subgroup of G, and $\mathcal{R}$ is any set of irreducible matrix representations of G, then $\mathcal{R}$ is K-adapted.

Example 2.8. Young's orthogonal form, and Young's seminormal form (see [15]) are both examples of complete sets of irreducible matrix representations for the symmetric group S_n, adapted to the chain of subgroups,

$$S_n > S_{n-1} > \cdots > S_1 = 1. \tag{2.6}$$

Since the restriction of representations from S_n to S_{n-1} is multiplicity free, the basis vectors of a system of Gel′fand-Tsetlin bases for the irreducible representations of S_n relative to (2.6) are determined up to scalar multiples. The corresponding sets of adapted representations are determined up to conjugation by diagonal matrices.

The dual matrix coefficient basis associated to a complete adapted set of inequivalent irreducible representations has particularly nice computational properties.

Definition 2.9. The dual matrix coefficient basis corresponding to a complete set of inequivalent irreducible representations adapted to the chain (2.5) is called a Gel′fand-Tsetlin basis for the group algebra $\mathbf{C}[G]$ relative to the chain (2.5).

We can now relate the computation of a Fourier transform at an adapted set of representations to a collection of Fourier transforms on a subgroup. This idea was first due to Beth, and was developed by Clausen [5], [6], and Diaconis and Rockmore [10]. Before giving a precise statement we must introduce some notation. Assume K is a subgroup of G, $\mathcal{R}$ is a K-adapted set of matrix representations of G, and Y is a subset of G. Then we let

(2.7)

$$m_G(\mathcal{R}, Y, K) = \frac{1}{|G|} \times \left\{ \begin{array}{l} \text{The minimum number of operations required to compute} \\ \sum_{y\in Y} y \cdot F_y \text{ in the Gel}'\text{fand-Tsetlin basis for } \mathbf{C}[G] \text{ associ-} \\ \text{ated to } \mathcal{R}, \text{ where each } F_y \text{ is an arbitrary element of } \mathbf{C}[K]. \end{array} \right.$$

Lemma 2.10 ([10] Proposition 1, [5], [6])**.** *Let K be a subgroup of G and let $\mathcal{R}$ be a complete K-adapted set of inequivalent irreducible matrix representations of G. Let $Y \subset G$ be a set of coset representatives for G/K. Then*

$$t_G(\mathcal{R}) \le t_K(\mathcal{R}_K) + m_G(\mathcal{R}, Y, K).$$

Proof. By Lemma 2.5, computation of $\mathfrak{F}_{\mathcal{R}} f$ is equivalent to computation of the following sum Σ in a Gel′fand-Tsetlin basis for the group algebra. We have

$$\begin{aligned} \Sigma \;=\; \sum_{s\in G} f(s)s \;&= \sum_{y\in Y}\sum_{k\in K} f(y\cdot k) y\cdot k \\ &= \sum_{y\in Y} yF_y, \end{aligned} \tag{2.8}$$

where for each $y \in Y$, $F_y = \sum_{k\in K} f_y(k)k \in \mathbf{C}[K]$, and $f_y(k) = f(y\cdot k)$.

We may therefore use the following procedure to compute the sum Σ. First compute the algebra elements $F_y \in \mathbf{C}[K]$ for all $y \in Y$, in the Gel′fand-Tsetlin basis of $\mathbf{C}[K]$ corresponding to $\mathcal{R}_K$, by means of Fourier transforms on K. This requires $|G/K|\, T_K(\mathcal{R}_K)$ scalar operations. The second step is to express the elements F_y in coordinates relative to the Gel′fand-Tsetlin basis of $\mathbf{C}[G]$. By Lemma 2.5, this is equivalent to finding the matrices $\hat{f}_y(\rho\downarrow K)$ for all $\rho \in \mathcal{R}$ given the matrices $\hat{f}_y(\tau)$ for all $\tau \in \mathcal{R}_K$. This does not require any arithmetic operations, since, by the adaptedness of $\mathcal{R}$, $\hat{f}_y(\rho\downarrow K)$ is a block diagonal matrix that may be built from the matrices $\hat{f}_y(\tau)$ by matrix direct sums. Finally we compute the sum Σ using

(2.8). By definition, this takes no more than $|G|\, m_G(\mathcal{R}, Y, K)$ scalar operations, as the elements F_y all lie in $\mathbf{C}[K]$. Thus we obtain

$$T_G(\mathcal{R}) \leq |G/K|\, T_K(\mathcal{R}_K) + |G|\, m_G(\mathcal{R}, Y, K).$$

Dividing by $|G|$ proves the lemma. □

2.3. **Harmonic analysis.** We now describe how to relate the harmonic analysis of a function to its Fourier transforms.

The dual matrix coefficient basis is not the same as the matrix coefficient basis referred to in the introduction. Instead, it is dual to the matrix coefficient basis under the bilinear form $(f, h) = \sum_{s \in G} f(s)h(s)$.

Assume $\mathcal{R}$ is a complete set of inequivalent irreducible matrix representations of G. Let $\{\rho_{ij}\}$ be the matrix coefficient basis and $\{\breve{\rho}_{ij}\}$ denote the dual basis, so $(\rho_{ij}, \breve{\rho}'_{i'j'}) = \delta_{\rho\rho'}\delta_{ii'}\delta_{jj'}$. Then by the Schur orthogonality relations, see [23] Section 2.2, $\breve{\rho}_{ij}(s) = \frac{\dim \rho}{|G|}\rho_{ji}(s^{-1})$. The coefficient of ρ_{ij} in the harmonic analysis of a function f is

$$(f, \breve{\rho}_{ij}) = \frac{\dim \rho}{|G|}\left[\hat{f}(\rho^\vee)\right]_{ij} = \frac{\dim \rho}{|G|}\left[\widehat{(f^\vee)}(\rho)\right]_{ji}, \tag{2.9}$$

where $f^\vee(s) = f(s^{-1})$, $\rho^\vee(s) = \rho(s^{-1})^T$, and $(\ )^T$ denotes transpose. Thus the harmonic analysis of f may be obtained by permuting the function values to get $f^\vee$, applying a Fourier transform on $\mathcal{R}$, reordering, and then rescaling the output by the factors $\frac{\dim \rho}{|G|}$.

The representation $\rho^\vee$ appearing in (2.9) is called the dual representation to ρ, and the set $\mathcal{R}^\vee = \{\rho^\vee : \rho \in \mathcal{R}\}$ is a complete set of irreducible representations that shares any adaptedness properties that $\mathcal{R}$ may have. By (2.9), the harmonic analysis of f may also be obtained by computing the Fourier transform of f on $\mathcal{R}^\vee$, and then scaling the output by $\frac{\dim \rho}{|G|}$. Clearly, any algorithms we develop for computing Fourier on finite groups may be applied to the computation of these transforms in at least two different ways.

When the representation matrices are all real and orthogonal, e.g., for Young's orthogonal form, then $\rho^\vee = \rho$ for each ρ in $\mathcal{R}$, and the harmonic analysis of f may be obtained directly from its Fourier transform on $\mathcal{R}$.

3. Fast transforms on the symmetric group

In this section we shall restate and prove Theorem 1.1 assuming the existence of certain bilinear maps with specific properties. We leave the construction of the bilinear maps to Section 4. In this way we hope to clarify the steps in the proof by giving the overall form of the proof first, and then filling in the technical details later.

Rewriting Theorem 1.1 in the language of adapted representations, gives us Theorem 3.1. For background on the representation theory of symmetric groups, we refer the reader to [15].

Theorem 3.1. *The Fourier transform of a complex function on the symmetric group S_n may be computed at a complete set of irreducible matrix representations of S_n adapted to the chain of subgroups*

$$S_n > S_{n-1} > \cdots > S_1 = 1 \tag{3.1}$$

in no more than $\frac{3n(n-1)}{4}|S_n|$ multiplications and the same number of additions.

Proof. We start by noting that if t_i is defined to be the transposition $(i-1\ i)$, then the group elements

$$t_2 \cdots t_n, \quad t_3 \cdots t_n, \quad \ldots, \quad t_n, \quad e,$$

form a complete set of coset representatives for S_n relative to S_{n-1}. Thus by Lemma 2.10, the problem of computing the Fourier transforms of a complex function at a set of adapted representations will be solved, if we can show how to compute sums of the form

$$\Sigma = \sum_{i=1}^{n} t_{i+1} \cdots t_n \cdot F_i \tag{3.2}$$

in a Gel′fand-Tsetlin basis for the group algebra relative to the chain (3.1), where the F_i are arbitrary elements of $\mathbf{C}[S_{n-1}]$.

We shall rearrange the sum (3.2) in a form similar to Horner's rule, and show that such a sum may be computed in no more than $\frac{3(n-1)}{2}|S_n|$ scalar operations, given the algebra elements $F_i \in \mathbf{C}[S_{n-1}]$ in the appropriate Gel′fand-Tsetlin basis. By Lemma 2.10, this relates the Fourier transform of a function on S_n to a collection of Fourier transforms on S_{n-1}, and allows us to prove the theorem inductively.

The key to rearranging the sum (3.2) is to permute the order in which the group algebra multiplications are performed. We claim that there is a sequence of bilinear maps $\underset{2}{*}, \ldots, \underset{n}{*}$, and spaces $V_1, \ldots, V_n$, $C_2, \ldots, C_n$, such that the following four properties hold.

Prop. 1. $V_1 = \mathbf{C}[S_{n-1}]$ and $V_n = \mathbf{C}[S_n]$. For $2 \le i \le n$,

$$C_i = \mathbf{C}[S_i] \cap \mathrm{Centralizer}(\mathbf{C}[S_{i-2}]),$$

and the map $\underset{i}{*} : V_{i-1} \times C_i \to V_i$ is bilinear.

Prop. 2. If $F \in V_1$ and $s_i \in C_i$ for $2 \le i \le n$, then

$$s_2 \cdot s_3 \cdots s_n \cdot F = \Big(\cdots \big(F \underset{2}{*} s_2\big) \underset{3}{*} s_3 \cdots\Big) \underset{n}{*} s_n.$$

Prop. 3. For each i with $2 \le i \le n$, the map

$$F \longmapsto \Big(\cdots \big(F \underset{2}{*} e\big) \underset{3}{*} e \cdots\Big) \underset{i}{*} e \in V_i$$

requires no arithmetic computation to apply.

Prop. 4. Given $v_{i-1} \in V_{i-1}$, $s_i \in C_i$ and $v_i \in V_i$, we may compute $v_{i-1} \underset{i}{*} s_i + v_i$ in no more than $\frac{3(i-1)}{n}|S_n|$ multiplications and the same number of additions.

In order to simplify the presentation, we shall defer the construction of $\underset{i}{*}$ and the demonstration of Prop. 1–4 to Section 4, where they will follow from Lemmas 4.5–4.8 respectively. We have already chosen bases for the spaces V_1, V_n, and C_i (Gel′fand-Tsetlin bases); the spaces V_i, $1 < i < n$, will be constructed with a natural choice of basis, and it is with respect to these bases that the complexity statements Prop. 3 and Prop. 4 are to be interpreted.

Using Prop. 1 and Prop. 2, it is easy to rearrange (3.2) into a more manageable form,

$$
\begin{aligned}
\Sigma &= \sum_{i=1}^{n} e \cdots e \cdot t_{i+1} \cdots t_n \cdot F_i \\
&= \sum_{i=1}^{n} \Big(\cdots \big(\big(\cdots \big(F_i \underset{2}{*} e \big) \underset{3}{*} \cdots \big) \underset{i}{*} e \big) \underset{i+1}{*} t_{i+1} \cdots \Big) \underset{n}{*} t_n \\
&= \Bigg[\Bigg[\cdots \Bigg[\Big[F_1 \underset{2}{*} t_2 + F_2 \underset{2}{*} e \Big] \underset{3}{*} t_3 + \big(F_3 \underset{3}{*} e \big) \underset{3}{*} e \Bigg] \underset{4}{*} t_4 \\
&\qquad \cdots \qquad\qquad + \cdots \qquad\qquad \Bigg] \underset{n-1}{*} t_{n-1} \\
&\qquad\qquad + \big(\cdots \big(F_{n-1} \underset{2}{*} e \big) \underset{3}{*} \cdots \big) \underset{n-1}{*} e \Bigg] \underset{n}{*} t_n \\
&\qquad\qquad + \big(\cdots \big(F_n \underset{2}{*} e \big) \underset{3}{*} \cdots \big) \underset{n}{*} e.
\end{aligned} \tag{3.3}
$$

The algorithm for computing Σ given $F_1, \ldots F_n$ proceeds in the obvious way:

Stage 1. Let $G_1 = F_1$.
Stage i. Let $G_i = G_{i-1} \underset{i}{*} t_i + \big(\ldots \big(F_i \underset{2}{*} e \big) \underset{3}{*} \ldots \big) \underset{i}{*} e$, for $2 \le i \le n$.
Stage n. Let $\Sigma = G_n = G_{n-1} \underset{n}{*} t_n + F_n$.

A quick look at (3.3) verifies that $\Sigma = G_n$.

Assume that the F_i are given and the t_i have been precomputed relative to the Gel′fand-Tsetlin basis. Then Stage 1 requires no computation, and by Prop. 3 and Prop. 4 the computation of G_i from G_{i-1} and F_i at Stage i requires no more than $\frac{3(i-1)}{n} |S_n|$ scalar operations.

Adding the operation counts for all the stages shows that the computation of Σ given $F_1, \ldots, F_n$ takes no more than $\frac{3(n-1)}{2} |S_n|$ scalar operations. Thus by Lemma 2.10, the reduced complexities for the computation of Fourier transforms relative to Gel′fand-Tsetlin bases satisfy

$$
t_{S_n} \le t_{S_{n-1}} + \frac{3(n-1)}{2}. \tag{3.4}
$$

Applying (3.4) recursively shows that $t_{S_n} \le \frac{3n(n-1)}{4}$. Therefore the Fourier transform of a complex function on S_n may be computed in no more than $\frac{3n(n-1)}{4} |S_n|$ scalar operations. Prop. 4 easily implies that the number of multiplications required by our algorithm is the same as the number of additions. □

Remark 3.2. Clausen's algorithm [5] calculates the products $t_{i+1} \cdots t_n \cdot F_i$ occurring in (3.2) by matrix multiplication of the corresponding matrices in the order from right to left. By Lemma 4.1 equation (4.5), the matrices corresponding to to t_j are sparse so the product $t_j \cdot (t_{j+1} \cdots t_n \cdot F_i)$, $i < j$ may be computed efficiently given t_j and $t_{j+1} \cdots t_n \cdot F_i$ in a Gel′fand-Tsetlin basis.

Clausen's algorithm requires $\frac{(n+1)n(n-1)}{3} |S_n|$ scalar operations, so Theorem 3.1 represents an improvement of order a factor n.

Theorem 3.1 immediately gives us a method for computing inverse Fourier transforms as well. To see this, suppose that $\mathcal{R}$ is a complete set of inequivalent irreducible representations of the group G, and let $\mathfrak{D}$, $\mathfrak{I}$ be the maps

(3.5)
$$\mathfrak{D} : \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{(\dim\rho)}(\mathbf{C}) \longrightarrow \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{(\dim\rho)}(\mathbf{C}) : \bigoplus_{\rho\in\mathcal{R}} F(\rho) \longmapsto \bigoplus_{\rho\in\mathcal{R}} \frac{\dim\rho}{|G|} F(\rho),$$

(3.6)
$$\mathfrak{I} : \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{(\dim\rho)}(\mathbf{C}) \longrightarrow \bigoplus_{\varphi\in\mathcal{R}^\vee} \mathrm{Mat}_{(\dim\varphi)}(\mathbf{C}) : \bigoplus_{\rho\in\mathcal{R}} F(\rho) \longmapsto \bigoplus_{\varphi\in\mathcal{R}^\vee} F(\varphi^\vee),$$

where $\rho^\vee$ denotes the dual representation; see Section 2.3. Then $\mathfrak{F}^T_{\mathcal{R}^\vee}\mathfrak{I}\mathfrak{D}\mathfrak{F}_{\mathcal{R}} = I$, where $(\ )^T$ denotes transpose, and I is the identity transformation.

Theorem 3.3. *Assume $\mathcal{R}$ is a complete set of irreducible matrix representations of S_n adapted to the chain of subgroups* (3.1)*. For each $\rho \in \mathcal{R}$, let $F(\rho)$ be a complex $\dim\rho \times \dim\rho$ matrix. Then the inverse Fourier transform*

$$f(s) = \mathfrak{F}^{-1}_{\mathcal{R}}[\bigoplus_{\rho\in\mathcal{R}} F(\rho)](s) = \frac{1}{|S_n|}\sum_{\rho\in\mathcal{R}} (\dim\rho)\,\mathrm{Trace}(F(\rho)\rho(s^{-1})) \tag{3.7}$$

may be computed in no more than $\frac{3n(n-1)}{4}\,|S_n|$ scalar operations.

Proof. Equation (3.7) is simply the Fourier inversion formula; see [23] 6.2 Proposition 11. To compute the inverse transform $\mathfrak{F}^{-1}_{\mathcal{R}}$, first apply $\mathfrak{D}$, as defined by (3.5) with $G = S_n$, then apply $\mathfrak{I}$, and finally apply $\mathfrak{F}^T_{\mathcal{R}^\vee}$ using the transpose algorithm (see [3], Chapter 3) of the algorithm of Theorem 3.1 for computing the Fourier transform at the set of dual representations $\mathcal{R}^\vee$. The last step is possible because $\mathcal{R}^\vee$ is also adapted to the chain (3.1).

The map $\mathfrak{I}$ is a re-indexing map, and requires no arithmetic operations to apply. The Fourier transform algorithm of Theorem 3.1 has the same number of outputs as inputs, so by [3] Theorem 3.10, the transpose algorithm takes exactly the same number of scalar operations as the Fourier transform algorithm of Theorem 3.1. Application of $\mathfrak{D}$ requires at most an extra $|S_n|$ scalar operations, but the bound of Theorem 3.1 overestimates the complexity of the Fourier transform by at least this much (see the proofs of Lemma 4.8 and Lemma 5.3 in the following sections). □

Remark 3.4. If the representations in $\mathcal{R}$ are unitary, then $\mathfrak{F}^{-1}_{\mathcal{R}} = \mathfrak{F}^*_{\mathcal{R}}\mathfrak{D}$, where $(\ )^*$ denotes conjugate transpose. If the representations are orthogonal, e.g., Young's orthogonal form, then the conjugate transpose may be replaced by a transpose. For representations of the symmetric group, we may always find a diagonal transformation $\mathfrak{D}_{\mathcal{R}}$ such that $\mathfrak{F}^T_{\mathcal{R}}\mathfrak{D}_{\mathcal{R}}\mathfrak{F}_{\mathcal{R}} = I$ (see Example 2.8).

The transformation $\mathfrak{I}$ can be given a coordinate-free definition, but that requires a more sophisticated interpretation of the transposes.

4. Construction and properties of the bilinear maps

From now on, it is convenient for us to fix a complete set of irreducible matrix representations of S_n adapted to the chain of subgroups

$$S_n > S_{n-1} > \cdots > S_1 = 1. \tag{4.1}$$

The standard bases for the spaces of column vectors on which these representations act is then a system of Gel′fand-Tsetlin bases relative to (4.1). This also determines a Gel′fand-Tsetlin basis for the group algebra $\mathbf{C}[S_n]$. Unless explicitly stated otherwise, we shall always refer to this system of Gel′fand-Tsetlin bases, and this Gel′fand-Tsetlin basis for the group algebra.

To motivate our construction of the bilinear maps $\underset{i}{*}$ and the spaces V_i, we first investigate some explicit ways of writing a product of elements in the group algebra in coordinates. We start by noting that the irreducible representations of S_n are in one to one correspondence with partitions of n; see e.g., [17]. If α_n is a partition of n, then we denote the corresponding representation of S_n by Δ_{α_n}.

It is well known [15] that a system of Gel′fand-Tsetlin bases for representations of S_n relative to the chain of subgroups (4.1) may be indexed by a chain of partitions

$$\boldsymbol{\alpha} = \quad \alpha_n \longleftarrow \alpha_{n-1} \longleftarrow \cdots \longleftarrow \alpha_2 \longleftarrow \alpha_1 \longleftarrow (\alpha_0 = \phi) \tag{4.2}$$

where α_i is a partition of i, and $\alpha \longleftarrow \beta$ indicates that the partition β may be obtained from α by removing a single box, or equivalently that Δ_β occurs in the restriction of Δ_α to the symmetric group of one lower order. $\boldsymbol{\alpha}$ indexes the unique Gel′fand-Tsetlin basis vector for Δ_{α_n} which is contained in the isotypic subspace of type Δ_{α_i} under the action of S_i, for $1 \leq i \leq n$. Thus, a single chain of partitions determines an irreducible representation of S_n and a basis vector for that representation, whereas a pair of chains of partitions $\boldsymbol{\alpha}, \boldsymbol{\beta}$ with $\alpha_n = \beta_n$ determines an element of the Gel′fand-Tsetlin basis for the group algebra $\mathbf{C}[S_n]$.

The chain of partitions $\boldsymbol{\alpha}$ is equivalent to specifying a standard Young's tableau on a Young's diagram with n boxes (see [17]), so all our arguments involving chains of partitions could be rewritten in terms of Young's tableaux.

Convention 1. We shall identify the group algebra $\mathbf{C}[S_n]$ with its realization in coordinates relative to the Gel′fand-Tsetlin basis, indexed by pairs of chains of partitions. Thus, if G is an element of $\mathbf{C}[S_n]$ we shall denote its coordinates relative to the Gel′fand-Tsetlin basis by either $[G]_{\boldsymbol{\beta},\boldsymbol{\alpha}}$ or

$$G\begin{pmatrix}\beta_n\ \beta_{n-1}\ \dots\ \beta_1 \\ \quad\ \alpha_{n-1}\ \dots\ \alpha_1\end{pmatrix}, \tag{4.3}$$

where $\boldsymbol{\beta}$ is the chain of partitions indexing rows of Fourier transforms of G, and $\boldsymbol{\alpha}$ indexes columns. Note that we always have $\alpha_n = \beta_n$, which explains why α_n does not occur in (4.3).

Convention 2. An element F of $\mathbf{C}[S_{n-1}]$ can be written in coordinates relative to the restricted Gel′fand-Tsetlin basis for $\mathbf{C}[S_{n-1}]$. When we do this, we shall denote the coordinates by

$$F\begin{pmatrix}\beta_{n-1}\ \beta_{n-2}\ \dots\ \beta_1 \\ \quad\ \alpha_{n-2}\ \dots\ \alpha_1\end{pmatrix}.$$

Alternatively, F may be considered as an element of $\mathbf{C}[S_n]$, and expanded in the Gel′fand-Tsetlin basis for that algebra. Fortunately, these two notations are easily reconciled by Lemma 4.1, which follows. In particular, moving from one realization to another is simply a re-indexing process and does not require any arithmetic computation.

Recall that we defined the spaces C_i, $2 \le i \le n$, to be the centralizer algebras

$$C_i = \mathbf{C}[S_i] \cap \mathrm{Centralizer}(\mathbf{C}[S_{i-2}]).$$

Elements of the spaces C_i and $\mathbf{C}[S_{n-1}]$ have a very special form when written in the Gel′fand-Tsetlin basis for $\mathbf{C}[S_n]$.

Lemma 4.1. *Assume that $2 \le i \le n$, that $s_i \in C_i$ and $F \in \mathbf{C}[S_{n-1}]$. Then, relative to a Gel′fand-Tsetlin basis for the group algebra $\mathbf{C}[S_n]$, the elements s_i and F have the forms*

$$[s_i]_{\boldsymbol{\beta},\boldsymbol{\alpha}} = \delta_{\alpha_{n-1},\beta_{n-1}} \cdots \delta_{\alpha_i,\beta_i} \cdot P^i_{s_i}\begin{pmatrix} \beta_i & \beta_{i-1} \\ \alpha_{i-1} & \alpha_{i-2} \end{pmatrix} \cdot \delta_{\alpha_{i-2},\beta_{i-2}} \cdots \delta_{\alpha_1,\beta_1}, \tag{4.4}$$

$$[F]_{\boldsymbol{\beta},\boldsymbol{\alpha}} = \delta_{\alpha_{n-1},\beta_{n-1}} \cdot F\begin{pmatrix} \beta_{n-1} & \beta_{n-2} \ldots & \beta_1 \\ & \alpha_{n-2} \ldots & \alpha_1 \end{pmatrix}, \tag{4.5}$$

where P is a complex function of the variables indicated.

Proof. These are standard facts about Gel′fand-Tsetlin bases; see e.g., [14] Proposition 2.3.12 for a proof in different notation. Equation (4.5) follows immediately from the definition of adaptedness to S_n and S_{n-1}, since it describes the correct block diagonal matrices. Iterating (4.5) shows that an element H of $\mathbf{C}[S_i]$, $i \le n-1$, has the form

$$[H]_{\boldsymbol{\beta},\boldsymbol{\alpha}} = \delta_{\alpha_{n-1},\beta_{n-1}} \cdots \delta_{\alpha_i,\beta_i} \cdot H\begin{pmatrix} \beta_i & \beta_{i-1} \ldots & \beta_1 \\ & \alpha_{i-1} \ldots & \alpha_1 \end{pmatrix}$$

in the Gel′fand-Tsetlin basis. The general form of an element of $\mathbf{C}[S_n]$ which commutes with $\mathbf{C}[S_i]$ is easily found by solving the equations $[AH - HA]_{\boldsymbol{\beta},\boldsymbol{\alpha}} = 0$ as H runs over the basis for $\mathbf{C}[S_i]$. $\square$

Remark 4.2. Lemma 4.1 shows us that C_i is isomorphic to the space of complex functions of the partition-valued variables $\beta_i, \beta_{i-1}, \alpha_{i-1}, \alpha_{i-2}$, where these variables are constrained to satisfy the relation

$$\begin{tikzcd} \beta_i & \beta_{i-1} \arrow[l, tail] \\ \alpha_{i-1} \arrow[u, tail] & \alpha_{i-2} \arrow[l, tail] \arrow[u, tail] \end{tikzcd} \tag{4.6}$$

This isomorphism may be given as $s_i \mapsto P^i_{s_i}$, which requires no computation relative to a Gel′fand-Tsetlin basis for the subgroup chain (4.1).

Example 4.3. A particularly relevant case of Lemma 4.1 equation (4.4) is when the complete adapted set of irreducible matrix representations is Young's orthogonal form, and $s_i = t_i = (i-1\ i)$. In that case there is an explicit formula for $P^i_{t_i}$, first determined by A. Young.

For any two boxes b_1 and b_2 in a Young diagram, we define the axial distance from b_1 to b_2 to be $d(b_1, b_2)$, where

$$d(b_1, b_2) = \mathrm{row}(b_1) - \mathrm{row}(b_2) + \mathrm{column}(b_1) - \mathrm{column}(b_2).$$

Thus, $d(b_1, b_2)$ is positive if b_1 lies to the right and upwards from b_2, and negative if b_1 lies to the left and downwards from b_2.

Now suppose that $\beta_i, \beta_{i-1}, \alpha_{i-1}, \alpha_{i-2}$ are partitions which satisfy (4.6). Then the skew diagrams of $\beta_i - \beta_{i-1}$ and $\beta_{i-1} - \alpha_{i-2}$ each consist of a single box, and the axial distance $d(\beta_i - \beta_{i-1}, \beta_{i-1} - \alpha_{i-2})$ is simply the signed length of the hook in β_i starting at one box and ending one box before the other. The formula for $P^i_{t_i}$ may now be stated as

$$(4.7)\qquad P^i_{t_i}\begin{pmatrix}\beta_i & \beta_{i-1}\\ \alpha_{i-1} & \alpha_{i-2}\end{pmatrix} = \begin{cases} d(\beta_i-\beta_{i-1},\beta_{i-1}-\alpha_{i-2})^{-1} & \text{if } \alpha_{i-1}=\beta_{i-1},\\ \sqrt{1-d(\beta_i-\beta_{i-1},\beta_{i-1}-\alpha_{i-2})^{-2}} & \text{if } \alpha_{i-1}\neq\beta_{i-1}.\end{cases}$$

For a proof of this formula, in slightly different notation, see [15], Chapter 3. The constraints (4.6) imply that $P^i_{t_i}$ given by (4.7) is symmetric in α_{i-1} and β_{i-1}.

Now we may give an expression for the product $s_2\cdots s_n\cdot F$ in the Gel′fand-Tsetlin basis.

Lemma 4.4. *Assume $F\in\mathbf{C}[S_{n-1}]$ and $s_i\in C_i$ for $2\le i\le n$. Then, relative to the Gel′fand-Tsetlin basis for the group algebra $\mathbf{C}[S_n]$, the element $s_2\cdots s_n\cdot F$ may be expressed as*

$$(4.8)\qquad [s_2\cdots s_n\cdot F]_{\boldsymbol{\beta},\boldsymbol{\gamma}} = \sum_{\alpha_{n-2},\dots,\alpha_1} F\begin{pmatrix}\gamma_{n-1} & \alpha_{n-2}\dots\alpha_1\\ & \gamma_{n-2}\dots\gamma_1\end{pmatrix}\cdot\prod_{i=2}^{n}P^i_{s_i}\begin{pmatrix}\beta_i & \beta_{i-1}\\ \alpha_{i-1} & \alpha_{i-2}\end{pmatrix}$$

where the partitions $\alpha_j,\beta_j,\gamma_j$ satisfy the relations (4.9), *and $\alpha_{n-1}=\gamma_{n-1}$.*

(4.9)
$$\begin{tikzcd}
\beta_n & \beta_{n-1} \arrow[l] & \beta_{n-2} \arrow[l] & \cdots \arrow[l] & \beta_3 \arrow[l] & \beta_2 \arrow[l] & \beta_1 \arrow[l] \\
\gamma_{n-1} \arrow[u] & \alpha_{n-2} \arrow[l] \arrow[u] & \alpha_{n-3} \arrow[l] \arrow[u] & \cdots \arrow[l] & \alpha_2 \arrow[l] \arrow[u] & \alpha_1 \arrow[l] \arrow[u] & \alpha_0 \arrow[l] \arrow[u] \arrow[dl] \\
 & \gamma_{n-2} \arrow[ul] & & \cdots \arrow[ll] & & \gamma_1 \arrow[ll] &
\end{tikzcd}$$

Proof. This follows by multiplying the algebra elements $s_2,\dots,s_n,F$ in coordinates, using the expressions (4.4) and (4.5). □

4.1. **Definition of the spaces and maps.** For $1\le i\le n$ we define V_i to be the space of complex functions of the form

$$(4.10)\qquad G\begin{pmatrix}\beta_i & \dots & \beta_1\\ \alpha_{n-2} & \dots & \alpha_{i-1}\\ \gamma_{n-1} & \dots & \gamma_1\end{pmatrix}$$

where $\alpha_{i-1},\dots,\alpha_{n-2}$, $\beta_1,\dots,\beta_i$, and $\gamma_1,\dots,\gamma_{n-1}$, are partitions satisfying the restriction relations (4.11).

(4.11)
$$\begin{tikzcd}
 & & & \beta_i & \beta_{i-1} \arrow[l] & \cdots \arrow[l] & \beta_1 \arrow[l] \\
\gamma_{n-1} & \alpha_{n-2} \arrow[l] & \cdots \arrow[l] & \alpha_{i-1} \arrow[l] \arrow[u] & & & \phi \arrow[u] \arrow[dl] \\
 & \gamma_{n-2} \arrow[ul] & & \cdots \arrow[ll] & & \gamma_1 \arrow[ll] &
\end{tikzcd}$$

When $i=1$ or $i=n$, a collection of partitions satisfying (4.11) is equivalent to specifying a pair of standard Young's tableaux of the same shape, and the spaces we get are $\mathbf{C}[S_{n-1}]$ and $\mathbf{C}[S_n]$ respectively, using Convention 1. (In the case of V_1 note that the variable β_1 can only assume one possible value.) This justifies the definitions $V_1=\mathbf{C}[S_{n-1}]$ and $V_n=\mathbf{C}[S_n]$.

Notice that the spaces V_i, $2 \le i \le n-1$, come equipped with a natural choice of basis given by indicator functions which are each 1 at exactly one point (choice of sequences of partitions) and zero elsewhere. When $i = 1$ or $i = n$, these are exactly the Gel′fand-Tsetlin bases.

The bilinear maps $\underset{i}{*}$ are now easy to define. Assume that $2 \le i \le n$, that $G_{i-1} \in V_{i-1}$, and that $s_i \in C_i$. Then we define $G_{i-1} \underset{i}{*} s_i \in V_i$ by

(4.12)
$$\left[G_{i-1} \underset{i}{*} s_i\right] \begin{pmatrix} \beta_i & \dots & \beta_1 \\ \alpha_{n-2} & \dots & \alpha_{i-1} \\ \gamma_{n-1} & \dots & \gamma_1 \end{pmatrix} = \sum_{\alpha_{i-2}} G_{i-1} \begin{pmatrix} \beta_{i-1} & \dots & \beta_1 \\ \alpha_{n-2} & \dots & \alpha_{i-2} \\ \gamma_{n-1} & \dots & \gamma_1 \end{pmatrix} \cdot P^i_{s_i} \begin{pmatrix} \beta_i & \beta_{i-1} \\ \alpha_{i-1} & \alpha_{i-2} \end{pmatrix}$$

where α_{i-2} satisfies (4.13).

(4.13)
$$\begin{tikzcd} \beta_i & \beta_{i-1} \arrow[l] \\ \alpha_{i-1} \arrow[u] & \alpha_{i-2} \arrow[l] \arrow[u] \end{tikzcd}$$

Notice that in going from G_{i-1} to $G_{i-1} \underset{i}{*} s_i$ we remove a dependence on α_{i-2} and add a dependence on β_i.

4.2. **Properties of the bilinear maps.** We now prove a sequence of lemmas corresponding to the properties Prop. 1–4, required by the proof of Theorem 3.1.

Lemma 4.5 (Prop. 1). *The map* $\underset{i}{*} : V_{i-1} \times C_i \to V_i$ *is bilinear.*

Proof. This follows from the bilinearity of (4.12), and the linearity of the coordinatizing map P^i. $\square$

Lemma 4.6 (Prop. 2). *Assume* $F \in V_1$, *and* $s_i \in C_i$ *for* $2 \le i \le n$. *Then*

$$s_2 \cdot s_3 \cdots s_n \cdot F = \left(\cdots \left(F \underset{2}{*} s_2\right) \underset{3}{*} s_3 \cdots \right) \underset{n}{*} s_n. \tag{4.14}$$

Proof. Rearranging (4.8) in Lemma 4.4 shows that

(4.15) $\quad [s_2 \cdots s_n \cdot F]_{\boldsymbol{\beta},\boldsymbol{\gamma}}$
$$= \sum_{\alpha_{n-2}} \left[\cdots \left[\sum_{\alpha_2} \left[\sum_{\alpha_1} \left[F \begin{pmatrix} \gamma_{n-1} & \alpha_{n-2} & \dots & \alpha_1 \\ & \gamma_{n-2} & \dots & \gamma_1 \end{pmatrix} \cdot P^2_{s_2} \begin{pmatrix} \beta_2 & \beta_1 \\ \alpha_1 & \phi \end{pmatrix} \right] \cdot P^3_{s_3} \begin{pmatrix} \beta_3 & \beta_2 \\ \alpha_2 & \alpha_1 \end{pmatrix} \right] \cdot \right.\right.$$
$$\left.\left. \cdot P^4_{s_4} \begin{pmatrix} \beta_4 & \beta_3 \\ \alpha_3 & \alpha_2 \end{pmatrix} \right] \cdots \right] \cdot P^n_{s_n} \begin{pmatrix} \beta_n & \beta_{n-1} \\ \alpha_{n-1} & \alpha_{n-2} \end{pmatrix}.$$

The right hand side of (4.15) is exactly the composition of bilinear maps

$$\left(\dots \left(F \underset{2}{*} s_2\right) \underset{3}{*} \dots \right) \underset{n}{*} s_n.$$

The summation over α_{i-2} corresponds to the application of $\underset{i}{*}$. We have not written the summation over α_0 explicitly, because the only partition on 0 boxes is ϕ. Similarly, one could omit the sum on α_1, as that is trivial too. $\square$

Lemma 4.7 (Prop. 3)**.** *Assume $2 \le i \le n$ and $F \in \mathbf{C}[S_{n-1}]$. Then we have the following expression for $\big(\dots\big(F \underset{2}{*} e\big) \underset{3}{*} \dots\big) \underset{i}{*} e$ in coordinates.*

$$
\tag{4.16}
\begin{aligned}
\Big[\big(\dots\big(F \underset{2}{*} e\big) \underset{3}{*} \dots\big) \underset{i}{*} e\Big]
&\begin{pmatrix} \beta_i & \dots & \beta_1 \\ \alpha_{n-2} & \dots & \alpha_{i-1} \\ \gamma_{n-1} & \dots & \gamma_1 \end{pmatrix} \\
&= \delta_{\alpha_{i-1}\beta_{i-1}} F \begin{pmatrix} \gamma_{n-1} \; \alpha_{n-2} \dots \alpha_{i-1} \; \beta_{i-2} \dots \beta_1 \\ \gamma_{n-2} \dots \; \dots \; \dots \; \dots \gamma_1 \end{pmatrix}.
\end{aligned}
$$

This requires no arithmetic computation; it is simply a re-indexing operation.

Proof. Equation (4.16) follows by using the definition (4.12) repeatedly, and noting that P_e^i has a particularly simple form,

$$
P_e^i \begin{pmatrix} \beta_i & \beta_{i-1} \\ \alpha_{i-1} & \alpha_{i-2} \end{pmatrix} = \delta_{\alpha_{i-1}\beta_{i-1}}. \qquad \square
$$

Before proving Prop. 4 we introduce notation which lets us give an exact count of the number of operations we use to apply the bilinear maps $\underset{i}{*}$. We prove the exact count in this section, but defer the proof of the bound $\frac{3(i-1)}{n}\,|S_n|$ to the next section.

Equation (4.12), which defines the bilinear maps, has a combinatorial indexing scheme that generalizes Young's tableaux. The left hand side of that formula involves sequences of partitions $\gamma_1, \dots, \gamma_{n-1}$, $\beta_1, \dots, \beta_i$, $\alpha_{i-2}, \dots, \alpha_{n-2}$ (with α_j a partition of j etc.), which satisfy the relations

$$
\tag{4.17}
\begin{tikzcd}
 & & & \beta_i & \beta_{i-1} \arrow[l] & \cdots \arrow[l] & \beta_1 \arrow[l] \\
\gamma_{n-1} & \alpha_{n-2} \arrow[l] & \cdots \arrow[l] & \alpha_{i-1} \arrow[l] \arrow[u] & \alpha_{i-2} \arrow[l] \arrow[u] & & \phi \arrow[u] \arrow[dl] \\
 & \gamma_{n-2} \arrow[ul] & & \cdots \arrow[ll] & & \gamma_1 \arrow[ll] &
\end{tikzcd}
$$

Let $\mathcal{F}_i^n$ denote the number of such sequences.

The number of arithmetic operations taken by our algorithm may be expressed in terms of $\mathcal{F}_i^n$, and the combinatorial lemmas proven in Section 5 allow us to further express this count in terms of $\mathcal{F}_i^i$, which we bound. Lemma 4.8, summarizes the end result.

Lemma 4.8 (Prop. 4)**.** *Assume that $2 \le i \le n$, and that $v_{i-1} \in V_{i-1}$, $s_i \in C_i$ and $v_i \in V_i$, are given. Then we may compute $v_{i-1} \underset{i}{*} s_i + v_i$ in no more than $\mathcal{F}_i^n \le \frac{3(i-1)}{n}\,|S_n|$ multiplications and the same number of additions.*

Proof. Let $\mathcal{G}_i^n$ denote the number of sequences of partitions $\gamma_1, \dots, \gamma_{n-1}$, $\beta_1, \dots, \beta_i$, $\alpha_{i-1}, \dots, \alpha_{n-2}$ which satisfy the relations (4.11). Clearly $\mathcal{G}_i^n = \dim V_i$.

Calculating $v_{i-1} \underset{i}{*} s_i$ using (4.12) directly takes $\mathcal{F}_i^n$ scalar multiplications and $\mathcal{F}_i^n - \mathcal{G}_i^n$ scalar additions. Adding v_i to the result requires an additional $\mathcal{G}_i^n$ additions. Therefore the computation of $v_{i-1} \underset{i}{*} s_i + v_i$ takes a total of $\mathcal{F}_i^n$ multiplications and $\mathcal{F}_i^n$ additions. The bound $\mathcal{F}_i^n \le \frac{3(i-1)}{n}\,|S_n|$ is proven in Lemma 5.3. $\square$

We have now verified all four properties of $\underset{i}{*}$. This completes the proof of Theorem 3.1, except for the combinatorial Lemma 5.3.

5. Combinatorial lemmas

We now turn to the combinatorial lemmas needed to complete the proof of Theorem 3.1. First we introduce some notation which is useful for counting chains of partitions.

Assume that $i \geq j$, that α is a partition of i, and that β is a partition of j. Then let $\mathcal{M}(\alpha, \beta)$ denote the number of sequences of partitions $\alpha_j, \dots, \alpha_i$ such that

$$(\alpha = \alpha_i) \longleftarrow \alpha_{i-1} \longleftarrow \cdots \longleftarrow \alpha_{j+1} \longleftarrow (\alpha_j = \beta) \ .$$

The function $\mathcal{M}$ has a number of other equivalent definitions.

$$\begin{aligned}\mathcal{M}(\alpha, \beta) &= \text{multiplicity of } \Delta_\beta \text{ in the restriction of } \alpha \text{ to } S_j \\ &= \text{number of standard tableaux on the skew diagram } \alpha - \beta \\ &= \text{number of ways of removing boxes from } \alpha \text{ to get } \beta.\end{aligned}$$

These numbers are a special case of the Kostka numbers [17] and are usually denoted $K_{\alpha-\beta,(1^{|\alpha|-|\beta|})}$, although [15] writes $k_{\alpha/\beta,(1^{|\alpha|-|\beta|})}$. We have chosen our notation to emphasize the properties of this function which come from its interpretation as restriction multiplicities (cf., [20]). In this paper we will only use the formal properties of $\mathcal{M}$ and a few special values. In particular, it is easily shown ([14] Corollary 2.3.2) that if α, β are partitions of i and j respectively, and $j \leq k \leq i$, then

$$\mathcal{M}(\alpha, \beta) = \sum_{\alpha_k} \mathcal{M}(\alpha, \alpha_k)\mathcal{M}(\alpha_k, \beta), \tag{5.1}$$

where α_k ranges over all partitions of k.

We shall also use the notation $d_\alpha = \mathcal{M}(\alpha, \phi)$. Thus d_α is the dimension of the representation Δ_α, and may be calculated using the famous hook-length formula of Frame, Robinson, and Thrall (see [17] or [15]).

Recall that $\mathcal{F}_i^n$ denotes the number of sequences of partitions $\gamma_1, \dots, \gamma_{n-1}$, $\beta_1, \dots, \beta_i$, $\alpha_{i-2}, \dots, \alpha_{n-2}$ (with α_j a partition of j etc.), which satisfy the relations (4.17).

Lemma 5.1.

$$\begin{aligned}\mathcal{F}_i^n = \sum_{\substack{\beta_i, \beta_{i-1} \\ \gamma_{n-1}, \alpha_{i-1}, \alpha_{i-2}}} &\mathcal{M}(\gamma_{n-1}, \alpha_{i-1})\mathcal{M}(\alpha_{i-1}, \alpha_{i-2})\mathcal{M}(\beta_i, \alpha_{i-1}) \\ &\cdot \mathcal{M}(\beta_i, \beta_{i-1})\mathcal{M}(\beta_{i-1}, \alpha_{i-2}) d_{\beta_{i-1}} d_{\gamma_{n-1}},\end{aligned} \tag{5.2}$$

where $\alpha_j, \beta_j, \gamma_j$ range over partitions of j.

Proof. We count the sequences satisfying (4.17) as follows. First choose the partitions $\alpha_{i-2}, \alpha_{i-1}, \beta_{i-1}, \beta_i, \gamma_{n-1}$ subject only to the restrictions that α_{i-2} is a partition of $i-2$, etc. Then the number of ways of choosing the chain of partitions from α_{i-1} to γ_{n-1} is $\mathcal{M}(\gamma_{n-1}, \alpha_{i-1})$. Similarly, the number of ways of choosing the chain of partitions $\phi, \gamma_1, \dots$ from ϕ to γ_{n-1} is $d_{\gamma_{n-1}}$, and the number of ways of choosing the chain of partitions from ϕ to β_{i-1} is $d_{\beta_{i-1}}$. Furthermore, these choices are independent given the choices of γ_{n-1}, α_{i-1} and β_{i-1}. Finally we note that the choice of $\alpha_{i-2}, \alpha_{i-1}, \beta_{i-1}, \beta_i$ is only consistent with (4.17) when the product

$$\mathcal{M}(\alpha_{i-1}, \alpha_{i-2})\mathcal{M}(\beta_i, \alpha_{i-1})\mathcal{M}(\beta_i, \beta_{i-1})\mathcal{M}(\beta_{i-1}, \alpha_{i-2})$$

is nonzero. This product is always either 0 or 1, so the number of sequences satisfying (4.17) may be found by summing the product

$$\mathcal{M}(\gamma_{n-1}, \alpha_{i-1})d_{\gamma_{n-1}}d_{\beta_{n-1}}\mathcal{M}(\alpha_{i-1}, \alpha_{i-2})\mathcal{M}(\beta_i, \alpha_{i-1})\mathcal{M}(\beta_i, \beta_{i-1})\mathcal{M}(\beta_{i-1}, \alpha_{i-2})$$

over all choices of $\alpha_{i-2}, \alpha_{i-1}, \beta_{i-1}, \beta_i, \gamma_{n-1}$. □

Suppose β is a partition. Then let $\operatorname{jmp}(\beta)$ denote the number of jumps in the Young diagram of β. For example, if $\beta = (4, 3, 3, 1, 1)$, then $\operatorname{jmp}(\beta) = 3$.

Lemma 5.2.
1. $\mathcal{F}_i^n = \frac{(n-1)!}{(i-1)!}\mathcal{F}_i^i$.
2. $\mathcal{F}_i^i = (i-1)\cdot(i-1)! + \sum_{\beta_{i-1}} \operatorname{jmp}(\beta_{i-1})^2 d_{\beta_{i-1}}^2$, *where* β_{i-1} *ranges over partitions of* $i-1$.

Proof. 1. follows immediately from (5.2) by Frobenius reciprocity, since for any α_{i-1} we have

$$\sum_{\gamma_{n-1}} \mathcal{M}(\gamma_{n-1}, \alpha_{i-1})d_{\gamma_{n-1}} = \dim \operatorname{Ind}_{S_{i-1}}^{S_{n-1}} \Delta_{\alpha_{i-1}} = |S_{n-1}/S_{i-1}| \cdot d_{\alpha_{i-1}}.$$

For 2. we start with the sum (5.2) in the case $n = i$, and split it into two parts, distinguishing the cases where $\alpha_{i-1} \neq \beta_{i-1}$ and $\alpha_{i-1} = \beta_{i-1}$. If α_{i-1} and β_{i-1} are distinct partitions of $i-1$ which are both obtained from β_i by removing a box, then they jointly determine β_i (and α_{i-2}), since the boxes removed from β_i to get to these two partitions are distinct. Thus the contribution to $\mathcal{F}_i^i$ from terms with $\alpha_{i-1} \neq \beta_{i-1}$ may be written as

$$\sum_{\substack{\alpha_{i-2} \\ \alpha_{i-1}\neq\beta_{i-1}}} \mathcal{M}(\beta_{i-1}, \alpha_{i-2})\mathcal{M}(\alpha_{i-1}, \alpha_{i-2})d_{\alpha_{i-1}}d_{\beta_{i-1}} \tag{5.3}$$
$$= \sum_{\substack{\alpha_{i-2} \\ \alpha_{i-1},\beta_{i-1}}} \mathcal{M}(\beta_{i-1}, \alpha_{i-2})\mathcal{M}(\alpha_{i-1}, \alpha_{i-2})d_{\alpha_{i-1}}d_{\beta_{i-1}} - \sum_{\beta_{i-1},\alpha_{i-2}} \mathcal{M}(\beta_{i-1}, \alpha_{i-2})^2 d_{\beta_{i-1}}^2.$$

Using Frobenius reciprocity and (5.1), the first term of (5.3) may be evaluated as

$$\begin{aligned} &\sum_{\alpha_{i-1},\alpha_{i-2}} \mathcal{M}(\alpha_{i-1},\alpha_{i-2})d_{\alpha_{i-1}} \sum_{\beta_{i-1}} \mathcal{M}(\beta_{i-1}, \alpha_{i-2})d_{\beta_{i-1}} \\ &\qquad = \sum_{\alpha_{i-1},\alpha_{i-2}} \mathcal{M}(\alpha_{i-1}, \alpha_{i-2})d_{\alpha_{i-1}} \cdot \dim \operatorname{Ind}_{S_{i-2}}^{S_{i-1}} \Delta_{\alpha_{i-2}} \\ &\qquad = \sum_{\alpha_{i-1},\alpha_{i-2}} \mathcal{M}(\alpha_{i-1}, \alpha_{i-2})d_{\alpha_{i-1}} \left|S_{i-1}/S_{i-2}\right| d_{\alpha_{i-2}} \\ &\qquad = (i-1)\sum_{\alpha_{i-1}} d_{\alpha_{i-1}}^2 = (i-1)\left|S_{i-1}\right|. \end{aligned} \tag{5.4}$$

The second term of (5.3), including the minus sign, is $-\sum_{\beta_{i-1}} \operatorname{jmp}(\beta_{i-1})d_{\beta_{i-1}}^2$.

On the other hand, if $\alpha_{i-1} = \beta_{i-1}$, then the only conditions on β_i and α_{i-2} are that they may be obtained from β_{i-1} by adding or removing a box, respectively. In this case, given β_{i-1}, there are $\operatorname{jmp}(\beta_{i-1}) + 1$ ways of choosing β_i, and $\operatorname{jmp}(\beta_{i-1})$

ways of choosing α_{i-2}. Thus the contribution to $\mathcal{F}_i^i$ from terms with $\alpha_{i-1} = \beta_{i-1}$ is

$$\sum_{\beta_{i-1}} \operatorname{jmp}(\beta_{i-1})(\operatorname{jmp}(\beta_{i-1})+1)d^2_{\beta_{i-1}}. \qquad \square$$

Lemma 5.3.

$$\mathcal{F}_i^n \leq \frac{3(i-1)}{n} |S_n| .$$

Proof. In light of Lemma 5.2, it suffices to show that for any partition β_i of i, we have $\operatorname{jmp}(\beta_i)^2 \leq 2i$. Let $a = \operatorname{jmp}(\beta_i)$. By deleting rows and columns from the Young diagram of β_i, we may obtain a new partition with fewer boxes, but the same number of jumps, and the Young diagram of this new partition can be made to have a staircase form, i.e., the new partition is exactly $(a, a-1, \dots, 1)$. For an example, see Figure 1. The number of boxes in the staircase $(a, a-1, \dots, 1)$ is $\frac{1}{2}a(a+1)$, which shows that $a(a+1) \leq 2i$. $\square$

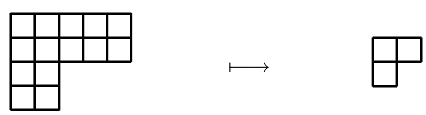

FIGURE 1. Removing rows and columns to obtain a staircase.

Remark 5.4. The same techniques used to prove Lemma 5.2 part 1 also show that

$$\dim V_i = \frac{i}{n} |S_n| .$$

The analogous problem of finding an explicit formula for $\mathcal{F}_i^n$ in closed form, if one exists, appears to be much more difficult.

5.1. **Exact operation counts.** Lemma 4.8 allows us to give an exact expression for the complexity of our Fourier transform algorithm on S_n, which we may evaluate using the combinatorial lemmas.

Theorem 5.5. *The Fourier transform of a complex function on the symmetric group S_n may be computed at a complete set of irreducible matrix representations of S_n adapted to the chain of subgroups*

$$S_n > S_{n-1} > \cdots > S_1 = 1$$

in no more than

$$\left(\sum_{k=2}^{n} \frac{1}{k} \sum_{i=2}^{k} \frac{1}{(i-1)!} \mathcal{F}_i^i \right) \cdot |S_n|$$

multiplications and the same number of additions.

Proof. By Lemma 4.5 and the proof of Theorem 3.1, we know that the number of multiplications (or additions) required by our algorithm is

$$\sum_{k=2}^{n} \frac{n!}{k!} \sum_{i=2}^{k} \mathcal{F}_i^k = \sum_{k=2}^{n} \sum_{i=2}^{k} \frac{n!(k-1)!}{k!(i-1)!} \mathcal{F}_i^i.$$

We have used Lemma 5.2 to simplify the result. $\square$

Lemma 5.2 allows us to calculate $\mathcal{F}_i^i$, and hence the exact complexity of our algorithm for computing Fourier transforms on S_i, for small values of i. We have done this for $1 \le i \le 50$, which certainly includes all cases where the algorithm might ever be implemented. In Table 1 we display these values for $1 \le i \le 12$, where T_{S_i} denotes the number of additions (or the number of multiplications) taken by the algorithm for computing Fourier transforms on S_i.

TABLE 1. Exact sequence and operation counts.

i	$\sum_{\beta_i} \mathrm{jmp}(\beta_i)^2 d_{\beta_i}^2$	$\mathcal{F}_i^i$	T_{S_i}	$\frac{1}{\lvert S_i\rvert} T_{S_i}$
1	1	0	0	0.0
2	2	2	2	1.0
3	18	6	16	2.7
4	78	36	130	5.4
5	474	174	1088	9.1
6	4004	1074	9792	13.6
7	32404	8324	96452	19.1
8	290558	67684	1034656	25.7
9	2924922	613118	12029342	33.1
10	33884848	6190842	150941204	41.6
11	416578024	70172848	2037003932	51.0
12	5485499312	855662824	29442867576	61.5

It is interesting to note that for $i \le 50$, the reduced complexity $T_{S_i}/|S_i|$ is bounded above by $\frac{1}{2}i(i-1)$, and their ratio lies close to 1 for i in this range.

5.2. **Remarks concerning the combinatorial lemmas.** Lemmas 5.1–5.3 have some simple generalizations, which become important when we extend the algorithm for Fourier transforms on the symmetric group to other finite groups and semisimple algebras. The main observation is that Young's lattice may be replaced by other Bratteli diagrams.

Let $\mathbf{N}$ denote the nonnegative integers. A **Bratteli diagram** (see [14] and [26]) is a connected $\mathbf{N}$-graded multigraph such that

(i) Each level (vertices with the same grading) has only finitely many vertices, and finitely many edges connected to it.
(ii) Edges only connect adjacent levels, and if two adjacent levels are non-empty, then the bipartite graph, consisting of those two levels and the edges connecting them, is connected.
(iii) The zeroth level contains a unique vertex, denoted ϕ.

Given any Bratteli diagram and two vertices α, β in the diagram, we let $\mathcal{M}(\alpha, \beta)$ denote the number of upward paths in the diagram from β to α. As before, we let $d_\alpha = \mathcal{M}(\alpha, \phi)$.

The definition of $\mathcal{F}_i^n$ is easy to generalize to any Bratteli diagram which has at least $n+1$ levels: given such a Bratteli diagram, we define $\mathcal{F}_i^n$ to be the number of grading-preserving maps from a graded graph of the form (4.17) into the Bratteli diagram. Each such map not only sends $\alpha_j, \beta_j, \gamma_j$ into vertices of level j, but also sends each edge into an edge of the Bratteli diagram.

With these definitions of $\mathcal{M}$ and $\mathcal{F}_i^n$, the statement of Lemma 5.1 holds with the only change being that $\alpha_j, \beta_j, \gamma_j$ now range over vertices at level j in the Bratteli diagram. To generalize Lemma 5.2, we need to place some extra conditions on our Bratteli diagrams.

Any Bratteli diagram is uniquely associated to a chain of semisimple algebras, called path algebras (see [14] 2.3.11). The i^{th} path algebra A_i has dimension $\dim A_i = \sum_{\beta_i} d_{\beta_i}^2$, where β_i ranges over vertices at level i, and A_i contains A_{i-1} as a subalgebra. A Bratteli diagram is **locally free** if A_i is free over A_{i-1} for $i \geq 1$.

A Bratteli diagram is **multiplicity free** if $\mathcal{M}$ is either 0 or 1 for any two adjacent vertices.

Lemma 5.6. 1. *Assume $n \geq i \geq 1$. Then for any locally free Bratteli diagram, $\dim A_i$ is an integer multiple of $\dim A_{i-1}$, and*

$$\mathcal{F}_i^n = \frac{\dim A_{n-1}}{\dim A_{i-1}} \mathcal{F}_i^i.$$

2. *For any locally free, multiplicity free Bratteli diagram,*

$$\mathcal{F}_i^i = \frac{(\dim A_{i-1})^2}{\dim A_{i-2}} + \sum_{\beta_{i-1}} (c^+(\beta_{i-1}) - 1) c^-(\beta_{i-1}) d_{\beta_{i-1}}^2$$

where β_{i-1} ranges over vertices at level $i-1$, and $c^+(\beta_{i-1}), c^-(\beta_{i-1})$ are the number of edges from β_{i-1} to levels i and $i-2$ respectively.

Proof. If A_i is free over A_{i-1}, then $\dim A_i / \dim A_{i-1}$ is the size of a basis for A_i as an A_{i-1}-module. For the rest of the lemma, start with equation (5.2), which holds for any Bratteli diagram, and follow the proof of Lemma 5.2 with partitions replaced by vertices and the operation of adding a box replaced by an upward step in the Bratteli diagram. Frobenius reciprocity still holds, and the locally free property implies that

$$\sum_{\alpha_k} \mathcal{M}(\alpha_k, \beta_j) d_{\alpha_k} = \frac{\dim A_k}{\dim A_j} d_{\beta_j}$$

for all vertices β_j at level j, where $j \leq k$ and α_k ranges over vertices at level k. See [21]. □

Example 5.7. Any differential poset [24] [25] is a locally free, multiplicity free Bratteli diagram. The Bratteli diagram of a tower of group algebras is locally free.

Several other combinatorial results for Young's lattice also extend to locally free Bratteli diagrams: In particular, the theorems of Stanley ([24] Theorem 3.7 and [25] Theorem 2.7) which count the number of paths in a differential poset, which start and end at ϕ, hold in this more general setting. See [21] for a more detailed treatment of these and similar results.

6. Homogeneous spaces

The harmonic analysis of a function on a homogeneous space is an important special case of harmonic analysis on groups. If K is a subgroup of the finite group G, then the associated spherical functions on the space G/K are defined to be the right K-invariant matrix coefficients on G viewed as functions on G/K. The harmonic analysis of a function on G/K is the expansion of that function in a basis of associated spherical functions, and may be computed by means of a Fourier transform on the homogeneous space. We direct the reader to [20] for background on the computation of Fourier transforms on homogeneous spaces.

Definition 6.1 (Fourier transform)**.** Let G be a finite group with subgroup K, and let f be a complex-valued function on G/K. Then the Fourier transform of f at a K-adapted matrix representation of G, or a K-adapted set of matrix representations of G, is defined to be the Fourier transform of the right K-invariant function $\tilde{f}$ on G defined by

$$\tilde{f}(g) = \frac{1}{|K|} f(gK).$$

We shall denote the Fourier transform at a representation ρ, or a set of representations $\mathcal{R}$, by $\hat{f}(\rho)^K$ and $\mathfrak{F}^K_{\mathcal{R}} f$ respectively.

The factor $\frac{1}{|K|}$ appearing in the definition of Fourier transform on homogeneous spaces ensures that the Fourier transform on the trivial homogeneous space K/K is trivial, and not multiplication by $|K|$. This will not affect our complexity results, but it does make the theory a bit tidier.

It is important to note that the only matrix entries of $\mathfrak{F}^K_{\mathcal{R}} f$ which may be nonzero are those entries in columns corresponding to K-invariant basis vectors. Moreover, the Fourier transform relative to a complete K-adapted set of inequivalent irreducible representations of G is an isomorphism from the space of functions on G/K to the space obtained by ignoring those columns which do not correspond to K-invariant vectors

A representation of G is said to be of class-1 with respect to K if it contains a nontrivial K-invariant vector. If desired, we could restrict ourselves to class-1 representations when discussing Fourier transforms on homogeneous spaces.

Remark 6.2. Let $\mathcal{R}$ be a complete K-adapted set of inequivalent irreducible representations of G, let $\rho \in \mathcal{R}$ be class-1 with respect to K, and let f be a complex function on G/K. Then the coefficient of the associated spherical function ρ_{ij_0} in the harmonic analysis of f is

$$\frac{\dim \rho}{|G/K|} \left[\hat{f}(\rho^\vee)^K\right]_{ij_0}$$

where $\rho^\vee$ denotes the dual of ρ, and j_0 indexes the right K-invariant columns of ρ. Clearly the harmonic analysis of f may be found by computing the Fourier transform $\mathfrak{F}^K_{\mathcal{R}^\vee} f$ relative to the set of dual representations, and then scaling the output by the factors $\frac{\dim \rho}{|G/K|}$.

Of course, if the group is S_n and $\mathcal{R}$ is Young's orthogonal form, then taking the dual has no effect.

The complexity and reduced complexity of the Fourier transform on a homogeneous space were defined in [20] by analogy with the group case.

Definition 6.3 (Complexity)**.** Let G be a finite group with subgroup K, and let $\mathcal{R}$ be any K-adapted set of matrix representations of G.

1. Let $T_{G/K}(\mathcal{R})$ denote the minimum number of operations needed to compute the Fourier transform of f on $\mathcal{R}$ via a straight-line program for an arbitrary complex-valued function f defined on G/K.
2. Let $t_{G/K}(\mathcal{R}) = T_{G/K}(\mathcal{R}) / |G/K|$.

$T_G(\mathcal{R})$ is called the complexity of the Fourier transform on G/K for the set $\mathcal{R}$, and $t_{G/K}(\mathcal{R})$ is called the reduced complexity.

The complexity always satisfies the inequalities

$$|G/K| - 1 \leq T_{G/K}(\mathcal{R}) \leq |G/K|^2.$$

When there is no possibility of confusion, we will drop the '$\mathcal{R}$' in the notation for complexities and reduced complexities.

In order to compute Fourier transforms on homogeneous spaces efficiently it suffices to see how the algorithms we have already developed for groups simplify when applied to a right invariant function. In [20] it was shown that for a large class of algorithms the bounds on the group reduced complexity t_G also apply to the homogeneous space reduced complexity $t_{G/K}$, so complexity results for homogeneous spaces could be obtained with essentially no extra work.

This is true in the current case as well. For instance, we shall show that if $\mathcal{R}$ is an adapted set of representations of S_n, the homogeneous space reduced complexities satisfy

$$t_{S_n/S_{n-k}}(\mathcal{R}) \leq t_{S_{n-1}/S_{n-k}} + \frac{3(n-1)}{2}. \tag{6.1}$$

Notice that this has the same form as equation (3.4) of Section 3. Applying (6.1) recursively and noting that $t_{S_{n-k}/S_{n-k}} = 0$ will give us Theorem 1.2 of the introduction.

We now restate and prove Theorem 1.2 using the terminology of adapted representations.

Theorem 6.4. *The Fourier transform of a complex function on the homogeneous space S_n/S_{n-k} may be computed at a complete set of (class-1) irreducible matrix representations of S_n adapted to the chain of subgroups*

$$S_n > S_{n-1} > \cdots > S_1 = 1 \tag{6.2}$$

in no more than $\frac{3k(2n-k-1)}{4}|S_n/S_{n-k}|$ scalar operations.

Proof. The result follows by chasing through the algorithm for computing Fourier transforms on S_n to see how it simplifies when applied to a right S_{n-k}-invariant function on S_n. We will simply indicate how to change the proofs already given in the group case to the current situation.

First we note that if f is a right S_{n-k}-invariant function on S_n, then the corresponding element of $\mathbf{C}[S_n]$ is invariant under multiplication by elements of $\mathbf{C}[S_{n-k}]$ on the right. In particular, if $k \geq 1$ then this also holds for the elements $F_y \in \mathbf{C}[S_{n-1}]$ that occur when the proof of Lemma 2.10 is applied to the subgroup S_{n-1} of S_n. Therefore, we must bound the number of operations required to compute any sum of the form

$$\Sigma = \sum_{i=1}^{n} t_{i+1} \cdots t_n \cdot F_i \tag{6.3}$$

in a Gel′fand-Tsetlin basis for the group algebra relative to the chain of subgroups (6.2), where the F_i are arbitrary right S_{n-k}-invariant elements of $\mathbf{C}[S_{n-1}]$.

In Sections 3 and 4 we showed that relative to a Gel′fand-Tsetlin basis for (6.2), the sum (6.3) has the following expression in coordinates,

$$[\Sigma]_{\boldsymbol{\beta},\boldsymbol{\gamma}} = \sum_{i=1}^{n} \sum_{\alpha_{n-2},\ldots,\alpha_1} F_i \begin{pmatrix} \gamma_{n-1} & \alpha_{n-2} & \ldots & \alpha_1 \\ & \gamma_{n-2} & \ldots & \gamma_1 \end{pmatrix} \cdot \prod_{j=2}^{i} \delta_{\alpha_{j-1}\beta_{j-1}} \cdot \prod_{j=i+1}^{n} P_{t_j}^{j} \begin{pmatrix} \beta_j & \beta_{j-1} \\ \alpha_{j-1} & \alpha_{j-2} \end{pmatrix} \tag{6.4}$$

where the partitions α_j, β_j, γ_j satisfy the relations (4.9), and $\alpha_{n-1} = \gamma_{n-1}$. This follows from Lemma 4.4, the proof of Lemma 4.7, and (6.3). Then, by equation (3.3) and the proof of Lemma 4.8, we were able to show that this sum could be computed in $\sum_{i=2}^{n} \mathcal{F}_i^n$ scalar operations, where $\mathcal{F}_i^n$ is the number of sequences of partitions satisfying the relations described by (4.17).

Now suppose that each F_i is invariant under right multiplication by elements of S_{n-k}. Then the coordinate of F

$$[F]_{\boldsymbol{\alpha},\boldsymbol{\gamma}} = F \begin{pmatrix} \gamma_{n-1} & \alpha_{n-2} & \ldots & \alpha_1 \\ & \gamma_{n-2} & \ldots & \gamma_1 \end{pmatrix}$$

is only nonzero when γ_{n-k} is the partition $(n-k)$ with a single row, i.e., the corresponding representation $\Delta_{\gamma_{n-k}}$ is the trivial representation of S_{n-k}. Therefore we only need to compute (6.4) in those cases where $\gamma_{n-k} = (n-k)$, and the number of operations required to do this is $\sum_{i=1}^{n} \tilde{\mathcal{F}}_i^n$, where $\tilde{\mathcal{F}}_i^n$ is the number of sequences of partitions which have $\gamma_{n-k} = (n-k)$ and satisfy the relations (4.17) as well.

Following through the arguments of Lemma 5.1 in the case where $\gamma_{n-k} = (n-k)$, it is easy to see that an expression for $\tilde{\mathcal{F}}_i^n$ may be obtained from the expression (5.2) for $\mathcal{F}_i^n$ by replacing the factor $d_{\gamma_{n-1}}$ by $\mathcal{M}(\gamma_{n-1},(n-k))$. Splitting the resulting sum in two according to the cases $\alpha_{i-1} \neq \beta_{i-1}$ and $\alpha_{i-1} = \beta_{i-1}$, and using (5.1) leads to the expression (6.5) for $\tilde{\mathcal{F}}_i^n$, in the notation of Section 5.

$$\begin{aligned} \tilde{\mathcal{F}}_i^n = &\sum_{\substack{\gamma_{n-1},\beta_{i-1} \\ \alpha_{i-2}}} \mathcal{M}(\gamma_{n-1},\alpha_{i-2})\mathcal{M}(\beta_{i-1},\alpha_{i-2}) d_{\beta_{i-1}} \mathcal{M}(\gamma_{n-1},(n-k)) \\ &+ \sum_{\gamma_{n-1},\beta_{i-1}} \mathcal{M}(\gamma_{n-1},\beta_{i-1})\,\mathrm{jmp}(\beta_{i-1})^2 d_{\beta_{i-1}} \mathcal{M}(\gamma_{n-1},(n-k)) \end{aligned} \tag{6.5}$$

By Frobenius reciprocity and (5.1), the first term of (6.5) may be evaluated as

$$\begin{aligned} &\sum_{\gamma_{n-1},\alpha_{i-2}} \mathcal{M}(\gamma_{n-1},\alpha_{i-2})\mathcal{M}(\gamma_{n-1},(n-k)) \left[\sum_{\beta_{i-1}} \mathcal{M}(\beta_{i-1},\alpha_{i-2}) d_{\beta_{i-1}} \right] \\ &\qquad = \sum_{\gamma_{n-1},\alpha_{i-2}} \mathcal{M}(\gamma_{n-1},\alpha_{i-2})\mathcal{M}(\gamma_{n-1},(n-k)) \left[\dim(\mathrm{Ind}_{S_{i-2}}^{S_{i-1}} \Delta_{\alpha_{i-2}}) \right] \\ &\qquad = |S_{i-1}/S_{i-2}| \sum_{\gamma_{n-1},\alpha_{i-2}} \mathcal{M}(\gamma_{n-1},(n-k))\mathcal{M}(\gamma_{n-1},\alpha_{i-2}) d_{\alpha_{i-2}} \\ &\qquad = (i-1) \sum_{\gamma_{n-1}} \mathcal{M}(\gamma_{n-1},(n-k)) d_{\gamma_{n-1}} \\ &\qquad = (i-1)\,|S_{n-1}/S_{n-k}| = \frac{i-1}{n} \cdot |S_n/S_{n-k}|\,. \end{aligned} \tag{6.6}$$

Notice that the last step follows from Frobenius reciprocity applied to the trivial representation of S_{n-k} induced up to S_{n-1}.

The second term of (6.5) is bounded by

$$|S_{n-1}/S_{n-k}| \times \max_{\beta_{i-1}} \mathrm{jmp}(\beta_{i-1})^2,$$

where $\mathrm{jmp}(\beta_{i-1})$ is the number of jumps in the partition β_{i-1}. By the arguments of Lemma 5.3 we already know that the max in this expression is bounded by $2(i-1)$, so the second term of (6.5) is bounded by $\frac{2(i-1)}{n}|S_n/S_{n-k}|$.

Adding the bounds for the two terms of (6.5) shows us that

$$\tilde{\mathcal{F}}_i^n \leq \frac{3(i-1)}{n}\left|S_n/S_{n-k}\right|,$$

and hence that the reduced complexities for the computation of Fourier transforms on homogeneous spaces satisfy

$$t_{S_n/S_{n-k}} \leq t_{S_{n-1}/S_{n-k}} + \frac{1}{|S_n/S_{n-k}|}\sum_{i=2}^{n} \tilde{\mathcal{F}}_i^n \leq t_{S_{n-1}/S_{n-k}} + \frac{3(n-1)}{2}. \tag{6.7}$$

Applying (6.7) recursively with $t_{S_{n-k}/S_{n-k}} = 0$ shows that $t_{S_n/S_{n-k}} \leq \frac{3k(2n-k-1)}{4}$, and hence that the Fourier transform of a complex function on S_n/S_{n-k} may be computed in no more than $\frac{3k(2n-k-1)}{4}|S_n/S_{n-k}|$ multiplications and the same number of additions. □

Remark 6.5. Theorem 6.4 is an improvement on the result of Maslen and Rockmore ([20], Theorem 6.5), which was obtained by applying Clausen's algorithm [5] to a right invariant function on the symmetric group. They showed that the Fourier transform of a complex function on S_n/S_{n-k} could be computed at an adapted set of representations in no more than $k\left(n^2 - kn + \frac{1}{3}(k^2-1)\right)|S_n/S_{n-k}|$ scalar operations.

As in the case of transforms on groups, Theorem 6.4 gives us a method for computing inverse transforms, with no extra work. Suppose that $\mathcal{R}$ is a complete K-adapted set of inequivalent irreducible representations of the finite group G, and let $\mathfrak{D}_K$ be the map

$$\begin{aligned} \mathfrak{D}_K : \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{(\dim\rho)}(\mathbf{C}) &\longrightarrow \bigoplus_{\rho\in\mathcal{R}} \mathrm{Mat}_{(\dim\rho)}(\mathbf{C}) \\ \bigoplus_{\rho\in\mathcal{R}} F(\rho) &\longmapsto \bigoplus_{\rho\in\mathcal{R}} \frac{\dim\rho}{|G/K|} F(\rho). \end{aligned} \tag{6.8}$$

Then $(\mathfrak{F}^K_{\mathcal{R}^\vee})^T \mathfrak{I}\mathfrak{D}_K \mathfrak{F}^K_{\mathcal{R}} = I$, where $\mathfrak{I}$ is the re-indexing map defined by (3.6). Therefore, inverse Fourier transform algorithms may be obtained from Fourier transform algorithms by taking the transpose algorithm and scaling the input.

When $G = S_n$ and the representations are in Young's orthogonal form, we have $(\mathfrak{F}^K_{\mathcal{R}})^T \mathfrak{D}_K \mathfrak{F}^K_{\mathcal{R}} = I$, where K is any subgroup in the chain (6.2).

The preceding discussion gives us Theorem 6.6, which we state without further proof.

Theorem 6.6. *Assume $\mathcal{R}$ is a complete set of* (*class-1*) *irreducible matrix representations of S_n adapted to the chain of subgroups* (6.2). *For each $\rho \in \mathcal{R}$, let*

$F(\rho)$ be a complex $\dim\rho \times \dim\rho$ matrix with zeroes in those columns which are not S_{n-k}-invariant columns of ρ. Then the inverse Fourier transform

$$f(sS_{n-k}) = \left[(\mathfrak{F}_{\mathcal{R}}^{S_{n-k}})^{-1} [\bigoplus_{\rho\in\mathcal{R}} F(\rho)] \right] (sS_{n-k}) = \frac{1}{|S_n/S_{n-k}|} \sum_{\rho\in\mathcal{R}} (\dim\rho) \operatorname{Trace}(F(\rho)\rho(s^{-1})) \tag{6.9}$$

may be computed in no more than $\frac{3k(2n-k-1)}{4}\,|S_n/S_{n-k}|$ scalar operations.

Remark 6.7. Diaconis and Rockmore [11] discuss the computation of isotypic projections of functions on homogeneous space. They suggest a direct method equivalent to the composition of a Fourier transform followed by a truncation followed by an inverse Fourier transform, and which takes $|G/K|^2$ scalar operations.

The current techniques can be applied to efficiently compute the isotypic projections of functions on the space S_n/S_{n-k}. First compute the Fourier transform of the function on S_n/S_{n-k} with respect to Young's orthogonal form, by the method of Theorem 6.4. Next truncate those parts of the transform which do not correspond to representations of the chosen type ρ. Finally, compute an inverse Fourier transform by multiplying by $\frac{\dim\rho}{|S_n/S_{n-k}|}$, and then applying the transpose of the algorithm of Theorem 6.4, again with respect to Young's orthogonal form. Note that for some applications the final inverse transform may not be necessary.

7. Conclusion

Although the results presented in this paper are specific to the symmetric group, the techniques used to obtain them are much more general. The use of Gel′fand-Tsetlin bases, the choice of factorizations for group elements or coset representatives, and the rearrangement of sums similar to Horner's rule are all well known tools for computing Fourier transforms on finite groups [3] and compact Lie groups [18]. Together they form the basis for the general 'separation of variables' method for constructing Fourer transform algorithms [20] [21].

The construction of the bilinear maps in Section 4 may also be generalized to any finite group (or semisimple algebra). Given a system of Gel′fand-Tsetlin bases, a collection of products of group elements, and a permutation, there is a well defined sequence of bilinear maps that allow the products to be rearranged (cf., Prop. 2) according to the chosen permutation. The spaces on which the bilinear maps are defined are associated to diagrams generalizing (4.9) (4.11) (4.17), and formulae for the number of operations needed to apply these maps can be read off the diagrams, in terms of restriction multiplicities. In joint work with Dan Rockmore [21], such ideas have been systematically developed, and applied to the computation of Fourier transforms on a variety of groups and algebras.

The methods used in Section 4 also raise some combinatorial questions. Walks generalizing (4.11) have already been studied on Young's lattice and other posets [24] [25], but the appearance of multiply-connected configurations, e.g., (4.9) and (4.17), appears to be a new phenomenon; also see [18] and [20]. More generally, one may consider mappings from any graded diagram into a Bratteli diagram. Such objects appear in the construction of Fourier transform algorithms on other finite groups [21], and the complexity of the algorithms may again be obtained by counting the mappings. There are always expressions for the numbers of these objects, generalizing equation (5.2) for $\mathcal{F}_i^n$, but it is not clear when these expressions may be evaluated in closed form. We do not even know the answer for (5.2) itself.

Finally, we should note that, even for the symmetric groups, the problem of computing Fourier transforms is far from completely solved. In particular, some applications [8] [11] require the transform to be computed at representations which are adapted relative to other chains of parabolic subgroups. Although Clausen's algorithm and the algorithms in this paper may both be adapted to these new situations, the results are less convincing.

We have not yet implemented the algorithms in this paper. Because of their close relationship with Clausen's algorithm, we expect these algorithms to be stable and efficient, in practice as well as in theory.

Acknowledgements

I would like to thank Persi Diaconis, Dan Rockmore, and Micheal Clausen for their encouragement, comments, and advice. I would also like to thank Institut des Hautes Études Scientifiques and Universiteit Utrecht, which supported me during the writing of this paper.

References

[1] T. Beth, *Verfahren der schnellen Fourier–Transformation*, Teubner Studienbücher, Stuttgart, 1984. MR **86g:**65002

[2] P. Bürgisser, M. Clausen, A. Shokrollahi, *Algebraic Complexity Theory*, Springer-Verlag, Berlin, 1996. CMP 97:10

[3] M. Clausen and U. Baum, *Fast Fourier transforms*, Wissenschaftsverlag, Mannheim, 1993. MR **96i:**68001

[4] ______, *Fast Fourier transforms for symmetric groups, theory and implementation*, Math. Comp. **61**(204) (1993), 833–847. MR **94a:**20028

[5] M. Clausen, *Fast generalized Fourier transforms*, Theoret. Comput. Sci. **67** (1989), 55–63. MR **91f:**68081

[6] ______, *Beiträge zum Entwurf schneller Spektraltransformationen*, Habilitationsschrift, Fakultät für Informatik der Universität Karlsruhe (TH), 1988.

[7] J. W. Cooley and J. W. Tukey, *An algorithm for machine calculation of complex Fourier series*, Math. Comp. **19** (1965), 297–301. MR **31:**2843

[8] P. Diaconis, *A generalization of spectral analysis with applications to ranked data*, Ann. Stat. **17** (1989), 949–979. MR **91a:**60025

[9] ______, *Group representations in probability and statistics*, IMS, Hayward, CA, 1988.

[10] P. Diaconis and D. Rockmore, *Efficient computation of the Fourier transform on finite groups*, J. Amer. Math. Soc. **3**(2) (1990), 297–332. MR **92g:**20024

[11] ______, *Efficient computation of isotypic projections for the symmetric group*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. **11**, L. Finkelstein and W. Kantor (eds.), 1993, 87–104. MR **94g:**20022

[12] D. Elliott and K. Rao, *Fast transforms: algorithms, analyses, and applications*, Academic, New York, 1982. MR **85e:**94001

[13] I. Gel′fand and M. Tsetlin, *Finite dimensional representations of the group of unimodular matrices*, Dokl. Akad. Nauk SSSR **71** (1950), 825–828 (Russian). MR **12:**9j

[14] F. Goodman, P. de la Harpe, and V. Jones, *Coxeter graphs and towers of algebras*, Springer-Verlag, New York, 1989. MR **91c:**46082

[15] G. James and A. Kerber, *The representation theory of the symmetric group*, Encyclopedia of Mathematics and its Applications, vol. 16, Addison-Wesley, Reading MA, 1981. MR **83k:**20003

[16] S. Linton, G. Michler, and J. Olsson, *Fourier transforms with respect to monomial representations*, Math. Ann. **297** (1993), 253–268. MR **94i:**20015

[17] I. Macdonald, *Symmetric functions and Hall polynomials*, Oxford University Press, 1979. MR **84g:**05003

[18] D. Maslen, *Efficient computation of Fourier transforms on compact groups*, J. Fourier Anal. Appl. (to appear).

[19] D. Maslen and D. Rockmore *Generalized FFTs - a survey of some recent results*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Groups and Computation, II, L. Finkelstein and W. Kantor (eds.), (1996), 183–237. CMP 97:11

[20] ———, *Separation of variables and the efficient computation of Fourier transforms on finite groups, I*, J. Amer. Math. Soc. **10** (1) (1997). MR **97i:**20019

[21] ———, *Separation of variables and the efficient computation of Fourier transforms on finite groups, II*, (preprint).

[22] D. Rockmore, *Applications of generalized FFTs*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Groups and Computation, II, L. Finkelstein and W. Kantor (eds.), (1996). CMP 97:11

[23] J. Serre, *Linear representations of finite groups*, Springer-Verlag, New York, 1977. MR **56:**8675

[24] R. Stanley, *Differential Posets*, J. Amer. Math. Soc. **1**(4) (1988), 919–961. MR **89h:**06005

[25] R. Stanley, *Variations on differential posets*, in Invariant Theory and Tableaux (ed. D. Stanton), IMA Vol. Math. Appl. **19**, Springer, New York, 1990, 145–165. MR **91h:**06004

[26] A. Vershik and S. Kerov, *Locally semisimple algebras. Combinatorial theory and the K_0 functor*, Itogi Nauki i Tekhniki, Sovremennye Problemy Matematiki, Noveishie Dostizheniya, Vol. 26 (1985), 3–56. MR **88h:**22009

Institut des Haute Études Scientifiques, Le Bois-Marie, 35 Route de Chartres, 91440, Bures-sur-Yvette, France

Current address: Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098SJ, Amsterdam, The Netherlands

E-mail address: `maslen@cwi.nl`