Separation of Variables and the Computation of Fourier Transforms on Finite Groups, I

David K. Maslen,
Daniel N. Rockmore

Dept. of Math. and Comp. Sci.
Dartmouth College
Hanover, NH 03755

USA

Max-Planck-Institut für Mathematik
Gottfried-Claren-Str. 26
53225 Bonn

Germany

MPI / 95-10

Separation of Variables and the Computation of Fourier Transforms on Finite Groups, I*

David K. Maslen†
Max-Planck-Institut für Mathematik
53225 Bonn
Germany

Daniel N. Rockmore‡
Dept. of Math and Comp. Sci.
Dartmouth College
Hanover, NH 03755

November 15, 1994

Abstract

This paper introduces new techniques for the efficient computation of a Fourier transform on a finite group. We present a divide and conquer approach to the computation. The divide aspect uses factorizations of group elements to reduce the sum of products for the Fourier transform to simpler sums of matrix products and is the **separation of variables algorithm**. The conquer aspect is the final computation of matrix products which we perform efficiently using a special form of the matrices. This form arise from the use of subgroup-adapted representations and their structure when evaluated at elements which lie in the centralizers of subgroups in a subgroup chain. We present a detailed analysis of the matrix multiplications arising in the calculation and obtain easy-to-use upper bounds for the complexity of our algorithm in terms of representation theoretic data for the group of interest.

Our algorithm encompasses many of the currently known examples of fast Fourier transforms. We recover the best known fast transforms for some abelian groups, the symmetric groups and their wreath products, and the classical Weyl groups. Beyond this, we obtain greatly improved upper bounds for the general linear and unitary groups over a finite field, and for the classical Chevalley groups over a finite field.

This is part I of a two part paper. Part II will present a refinement of these techniques which results in further savings.

1 Introduction

Recently, increased attention has been paid to the problem of finding efficient algorithms for the computation of Fourier transforms on nonabelian groups. The abelian case has a long history, and since the publication of the Cooley-Tukey fast Fourier transform (FFT) [20] has been at the heart of digital signal processing (see for example [26, 3] and the many references contained therein). The nonabelian cases have also been motivated by applications. They have been found useful in new approaches to data analysis [22], VLSI design [10], the design of matched filters [36] and efficient group convolution algorithms [16, 44]. In the continuous setting, there are applications to computer vision, geophysics and climate modeling (cf. [25, 31]).

Apart from applications, these algorithms contribute to the understanding of the representation theoretic content of the fast Fourier transform. Although abelian groups have a unique Fourier transform, nonabelian groups have an

*A preliminary version of some of this work appears as an extended abstract, "Adapted Diameters and the Efficient Computation of Fourier Transforms of Finite Groups" in the *Proceedings of the 1995 ACM-SIAM Symposium on Discrete Algorithms.*

†Partially supported as a Shapiro Visitor while at Dartmouth.

‡This work supported in part by ARPA as administered by the AFOSR under contract DOD F4960-93-1-0567 as well as NSF DMS Award 9404275.

infinite number of Fourier transforms which correspond to different choices of bases for the irreducible representations of the group, G. The complexity of the group is defined as the least upper bound of the complexities of the algorithms computing Fourier transforms, over all choices of bases, and is bounded by $|G|^2$. This bound follows from a direct approach to the computation. It is conjectured that all finite groups have complexity $O(|G| \log^c |G|)$ (for some universal constant c), and this has already been proved for many different classes of nonabelian groups [17, 45, 46, 7].

We present a divide and conquer strategy for computing nonabelian Fourier transforms, which encompasses many known FFTs, and provides new fast algorithms in other cases. It has two main components. First, we use a set of factorizations of elements of G to write the matrix sum of products that defines the Fourier transform in terms of a sequence of sums of products which are easier to compute. We call this technique separation of variables and the corresponding algorithm is the separation of variables algorithm.

The second part of our strategy uses a subgroup chain for the group and the notion of a subgroup-adapted set of representations. When computing with a subgroup-adapted set of representations the matrix multiplications that occur in the separation of variables algorithm have a highly structured and sparse form and may therefore be computed efficiently. We provide a thorough analysis of the structure of these matrices and operation count of the corresponding matrix multiplications. The main tool used here is a form of Schur's Lemma which determines the structure of the representation matrix of a group element which commutes with a subgroup. The bulk of the new computational savings of this paper come from this use of commutativity. We believe this is a new contribution to the subject, although it does appear implicitly in the work of Clausen and Baum on the symmetric group [17] and that of Rockmore on wreath products [46].

Our techniques are quite general. We obtain upper bounds for the complexity of the Fourier transform of any group in terms of representation theoretic data. These bounds are expressed in terms of multiplicities of restrictions of irreducible representations from one subgroup to another. We thereby obtain a general procedure for bounding the complexity of the Fourier transform on a group which enables us to find explicit bounds even when the representation matrices are extremely complicated. Thus we derive both previously known and new results as part of a general theory, instead of using ad hoc techniques.

This paper does not use the full strength of the separation of variables approach, but despite this we recover the best known algorithms for many abelian groups, the symmetric groups, and their wreath products. Furthermore, we obtain new fast algorithms for matrix groups over finite fields. A more detailed analysis of the computation improves the results; that is the content of part II of this work, currently in preparation [39]. By dividing the work in this way we hope to present general results of interest without obscuring them with the technical machinery needed for more refined results.

We start the paper in section 2 with the definitions of Fourier transform, complexity, and adapted representation. Then, in section 3, we explain the previously known technique of reducing to subgroups. Section 4 forms the theoretical core of the paper; it contains the definition of the separation of variables algorithm, the analysis of matrix products, and the general complexity results that we use in our examples. Following this, section 5 develops results on the complexities of specific groups. We start it by deriving the Cooley-Tukey algorithm in the context of finite abelian groups, the results of Clausen and Baum [17] on the symmetric group, results on classical Weyl groups, and the results of Rockmore [46] on wreath products. We then give algorithms for the general linear and unitary groups over a finite field, and finish our examples with some results on classical Chevalley groups over finite fields. Finally we summarize the consequences of this work and indicate the contents of part II [39] of this paper.

Our bounds depend on some explicit knowledge of the restrictions to a subgroup and often involve the number of conjugacy classes in a group (i.e. the number of irreducible representations). For some of our results we need asymptotics for these quantities. To avoid interrupting the flow of the paper we have postponed this discussion to an appendix following the applications section (Section 5) of this paper.

Acknowledgement. Special thanks to Tom Hagedorn for explaining his interesting recent work on multiplicities for restricted representations. Thanks also to Herr Prof. Michael Clausen for some very helpful conversations.

2 Background

2.1 Nonabelian Fourier transforms

The familiar discrete Fourier transform (DFT) of a finite set of evenly spaced data and its efficient computation via the Cooley-Tukey fast Fourier transform [20] has a natural formulation in terms of the representation theory of

cyclic groups. This larger framework is necessary for posing the general problem of efficient computation for Fourier transforms on finite nonabelian groups. What follows is a brief review of the basic concepts and definitions necessary for the formulation of this problem. For a complete introduction to the representation theory for finite groups Serre's book [47] is a good reference.

Recall that a (complex) matrix representation of a finite group G is a map ρ from G into the group of $d \times d$ invertible matrices with complex entries, $GL_d(\mathbf{C})$, such that

$$\rho(st) = \rho(s)\rho(t)$$

for every $s, t \in G$. In this case d is called the degree or dimension of the representation ρ, and is denoted d_ρ and $V = \mathbf{C}^d$ is called the representation space of ρ.

Two representations ρ_1 and ρ_2 are said to be equivalent if they differ only by a change of basis, i.e. if there exists an invertible matrix A such that $\rho_1(s) = A^{-1}\rho_2(s)A$ for all $s \in G$. Notice that 1-dimensional matrix representations are uniquely determined by their equivalence class, whereas multidimensional representations have an infinite number of equivalent realizations.

A subspace $W \subset V = \mathbf{C}^d$ is said to be G-invariant if for all $s \in G$, $\rho(s)W \subset W$. The representation ρ is said to be irreducible if $V = \mathbf{C}^d$ has no G-invariant subspaces other than the trivial subspaces $\{0\}$ and V and reducible otherwise. Up to equivalence there are only a finite number of irreducible representations of any finite group — in fact there are as many as there are conjugacy classes in the group. Irreducible representations are the fundamental building blocks of all representations of a finite group. That is to say that any representation is equivalent to the direct sum of irreducible representations, where the direct sum of two representations is the matrix direct sum of the representations.

There are several equivalent definitions of the Fourier transform for a finite group [16, 10, 23]. The following is the most convenient for this paper.

Definition 1 (Fourier Transform) *Let G be a finite group and f be a complex-valued function on G.*

i. Let ρ be a matrix representation of G. Then the **Fourier transform of f at ρ,** *denoted $\hat{f}(\rho)$ is the matrix sum,*

$$\hat{f}(\rho) = \sum_{s \in G} f(s)\rho(s). \tag{1}$$

ii. Let $\mathcal{R}$ be a set of matrix representations of G. Then the **Fourier transform of f on $\mathcal{R}$** *is the set of Fourier transforms of f at the representations in $\mathcal{R}$.*

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

The most important case of a Fourier transform occurs when the set $\mathcal{R}$ is a complete set of inequivalent irreducible representations of G. In this situation we shall simply refer to such a calculation as the computation of a Fourier transform. A Fourier transform determines f through the Fourier inversion formula.

Theorem 2.1 (Fourier inversion formula) *(see e.g. [22], p. 13) Let G be a finite group, f a complex-valued function on G, and $\mathcal{R}$ a complete set of irreducible matrix representations of G. Then,*

$$f(s) = \frac{1}{|G|} \sum_{\rho \in \mathcal{R}} d_\rho \, trace\Big(\hat{f}(\rho)\rho(s^{-1})\Big) \tag{2}$$

where $d_\rho = dim(\rho)$.

Example: The "usual" discrete Fourier transform. The irreducible matrix representations of the cyclic group $\mathbf{Z}/n\mathbf{Z} = \{0, 1, \ldots, n-1\}$, are all one-dimensional. For each integer j with $0 \leq j \leq n-1$, define the representation ζ_j, by $\zeta_j(k) = \exp(\frac{2\pi ijk}{n})$ for $k \in \mathbf{Z}/n\mathbf{Z}$. The set of such representations is a complete set of inequivalent irreducible representations for $\mathbf{Z}/n\mathbf{Z}$ and the corresponding Fourier transform is usually known as the discrete Fourier transform. This computation is central to the subject of digital signal processing (cf. [43]).

The arithmetic complexity for computing a Fourier transform conceivably depends on the choice of basis for the irreducible representations. The notion of the **complexity** of a finite group provides a classification of finite groups according to the complexity of the most efficient algorithm to compute some such transform on the group.

Definition 2 (Complexity) *Let G be a finite group, and $\mathcal{R}$ any set of matrix representations of G. Let $T_G(\mathcal{R})$ 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. $T_G(\mathcal{R})$ is called the* **complexity of the Fourier transform for the set** $\mathcal{R}$. *Define the* **complexity of the group** G *to be*

$$C(G) = \min_{\mathcal{R}}\{T_G(\mathcal{R})\}$$

where $\mathcal{R}$ varies over all complete sets of inequivalent irreducible matrix representations of G.

The computational model used here is a common one in which an operation is defined as a single complex multiplication followed by a complex addition.

Elementary representation theory shows that the sum of the squares of the degrees of a complete set of irreducible representations of G is equal to $|G|$ (see e.g. [47], p. 18). Consequently direct computation of any Fourier transform gives the upper and lower bounds

$$|G| \leq C(G) \leq |G|^2$$

where the lower bound reflects the size of the input. As mentioned in Section 1, the techniques introduced in this paper show how structural properties of the group and a judicious choice of the set of representations $\mathcal{R}$, provide significantly better upper bounds for group complexity. When bounding $T_G(\mathcal{R})$ it is often easier to work with a related quantity, $t_G(\mathcal{R})$, called the **reduced complexity** and defined by

$$t_G(\mathcal{R}) = T_G(\mathcal{R})/|G| \tag{3}$$

This definition simplifies the statements and proofs of many following results.

Remark. Another common interpretation of the Fourier transform is as a change of basis for the group algebra $\mathbf{C}[G]$, from the basis of point masses on G to a basis of matrix coefficients coming from a complete set of inequivalent irreducible representations. When this point of view is adopted, the complexity of the Fourier transform can be measured as the c-linear complexity of the associated change of basis matrix [8]. The c-linear complexity of a group G is defined to be the minimum c-linear complexity of any such matrix for G. Assuming a choice of unitary representations (which is always possible) the results stated here can all be interpreted as statements about the 2-linear complexity of finite groups.

2.2 Adapted sets of representations

As remarked earlier, there are an infinite number of matrix representations equivalent to any given multidimensional matrix representation, all related by a change of basis. Even among equivalent representations the complexity of the associated Fourier transform might vary. For this reason and others, subgroup-adapted sets of representations have been found to be useful for efficiently computing Fourier transforms. Use of these representations permits the computation of a Fourier transform on a finite group, G, to be built up from the computation of several Fourier transforms on a chosen subgroup, H.

To briefly explain the idea, let H be a subgroup of a group G. An **H-adapted set of representations of G** has the property that when considered as representations of H via restriction, they may be constructed as matrix direct products of representations from a fixed complete set of inequivalent irreducible matrix representations of H. It is clear that for a function defined on H, the computation of the Fourier transforms at the (smaller) set of irreducible representations is is computationally equivalent to computing the Fourier transform at the set of restricted representations. As shown in [23] (which we explain in the next section), a Fourier transform on G always can be factored as a sum over a set of matrix multiplications against Fourier transforms at the restrictions of the representations to the subgroup H. By requiring that the restriction is H-adapted the computation of the Fourier transforms on H at the restricted representations is further reduced to several Fourier transforms on H.

Definition 3 (Subgroup-adapted representations) *Let G be a finite group and $\mathcal{R}$ be a set of matrix representations of G and let H be a subgroup of G. If ρ is a representation of G, let $\rho{\downarrow}H$ denote the representation of H obtained by restricting ρ to H. We say that $\mathcal{R}$ is* **H-adapted** *if there is a set $\mathcal{R}_H$ of inequivalent irreducible matrix representations of H such that the set of restricted representations*

$$(\mathcal{R}{\downarrow}H) = \{\rho{\downarrow}H \mid \rho \in \mathcal{R}\}$$

is a matrix direct sum of representations in $\mathcal{R}_H$.

Notice that if $\mathcal{R}$ is H-adapted, then the set $\mathcal{R}_H$ is uniquely determined by $\mathcal{R}$. When $H = G$, the property of being G-adapted allows us to reduce the computation of the Fourier transform of f on $\mathcal{R}$ to a Fourier transform on G at a set of inequivalent irreducible representations.

Lemma 2.2 *If $\mathcal{R}$ is a G-adapted set of matrix representations of G then $T_G(\mathcal{R}) = T_G(\mathcal{R}_G)$.*

Remark. The FFT algorithms presented in the following sections all assume the use of adapted sets of representations. The requirement of adaptability does not limit us, as any set of representations is equivalent to an adapted set of representations. To see this, it is easiest to work with the related concept of an **adapted basis** (also known as a Gelfand basis). A basis for a representation space is adapted to a subgroup if the matrix representation obtained by expressing the representation in coordinates for this basis is also adapted. Adaptedness for a set of bases is defined similarly. Adapted bases always exist and in fact, can always be constructed. Assuming that some complete set of irreducible matrix representations of G is known, then a change of basis can be computed so that the resulting set of representations are H-adapted for any fixed subgroup H.

To outline one such construction, we collect several previously known results. Babai and Rónyai [4] have shown that a complete set of irreducible representations of a finite group G can be constructed in polynomial time from the multiplication table of G. Further techniques from [4] or [5] provide efficient algorithms for decomposing representations into their irreducible constituents. By applying these results to the original set of representations restricted to the subgroup H, a complete set of irreducible representations for H is then found. A change of basis to insure that all representations of G are H-adapted is computed by the construction of certain projection operators. This last step is detailed in the fairly recent book of Fässler and Stiefel [27] which also provides a wealth of examples of uses of adapted bases in a variety of computational problems.

3 Coset decompositions and the Fourier transform

In previous work, adapted representations have already been used to speed the computation of Fourier transforms by factoring the computation through a subgroup [23]. The idea is to use the coset decomposition of elements in the group to relate a Fourier transform on G to Fourier transforms on a subgroup H. This may be thought of as the simplest example of the separation of variables technique (cf. Section 4).

To explain, let H be a subgroup of G and $Y \subset G$ be a set of coset representatives for G/H. Thus, G can be factored as the disjoint union of subsets $yH = \{yh \mid h \in H\}$ for all $y \in Y$. For any representation ρ of G we can use the relation $\rho(ab) = \rho(a)\rho(b)$ to produce a factorization of $\hat{f}(\rho)$ by

$$\begin{aligned} \hat{f}(\rho) &= \sum_{s \in G} f(s)\rho(s) \\ &= \sum_{y \in Y} \rho(y) \sum_{t \in H} f_y(t)\rho(t) \end{aligned} \tag{4}$$

where for each $y \in Y$, f_y is the function on H defined by $f_y(t) = f(yt)$ for all $t \in H$. Consequently, with the notation of (4) we can rewrite $\hat{f}(\rho)$ as a sum of Fourier transforms on H,

$$\hat{f}(\rho) = \sum_{y \in Y} \rho(y) \widehat{f_y}(\rho \downarrow H). \tag{5}$$

If we had computed the Fourier transform of f_y on $\mathcal{R} \downarrow H$ for a complete set of irreducible representations $\mathcal{R}$ of G and for all $y \in Y$, then the individual Fourier transforms $\widehat{f_y}(\rho \downarrow H)$ could be glued together by the "twiddle factors"[1] $\rho(y)$, to build each $\hat{f}(\rho)$ and thus the complete Fourier transform of f on $\mathcal{R}$.

In general, a restricted representation $\rho \downarrow H$ is reducible, even when ρ is irreducible, and is equivalent to the direct sum of a collection of irreducible representations of the subgroup H. The number of times any given equivalence class occurs in this decomposition is independent of the actual decomposition and is called its **multiplicity**. If $\rho \downarrow H$ is not only equivalent to, but also equal to a matrix direct sum of irreducibles, and all equivalent irreducible

[1] The terminology "twiddle factor" comes from the usual signal processing situation in which G is an abelian group. Then all irreducible representations are one-dimensional and the matrices $\rho(y)$ are simply roots of unity.

representations that occur in this sum are equal, then $\widehat{f}(\rho\downarrow H)$ can be constructed as a block diagonal matrix from the matrices of the appropriate Fourier transform of f_y on H. In the language of Section 2.2, this is precisely the condition that the set of representations $\mathcal{R}$, is H-adapted.

The discussion above directly yields an algorithm for computing the Fourier transform of any function f on G using any set of H-adapted representations of G:

(1) Choose a set of coset representatives Y for G/H, and for a fixed set $\mathcal{R}$ of H-adapted irreducible representations of G, and for each $y \in Y$ compute the Fourier transform of f_y on $\mathcal{R}_H$.

(2) For each $\rho \in \mathcal{R}$ build the restricted transforms $\widehat{f}_y(\rho\downarrow H)$. These will be block diagonal matrices with blocks given by the individual Fourier transforms of f_y at the representations of $\mathcal{R}_H$.

(3) Compute the products $\rho(y)\widehat{f}_y(\rho\downarrow H)$ and add them together.

To obtain an upper bound for the complexity of this basic algorithm it is useful to introduce some notation. Let $\mathcal{R}$ be a set of matrix representations of G and let Y be any subset of G. For each $\rho \in \mathcal{R}$ and $y \in Y$ let $F(y,\rho)$ be an arbitrary $d_\rho \times d_\rho$ matrix. Then we define

$$M_G(Y,\mathcal{R}) := \left\{ \begin{array}{l} \text{the minimum number of operations required} \\ \text{to compute the collection of sums,} \\ \quad \left\{ \sum_{y\in Y} \rho(y)F(y,\rho) | \rho \in \mathcal{R} \right\}. \end{array} \right. \tag{6}$$

Similarly, define a "reduced" version of (6) by

$$m_G(Y,\mathcal{R}) := \frac{M_G(Y,\mathcal{R})}{|G|}. \tag{7}$$

Theorem 3.1 *([23], Proposition 1) Let H be a subgroup of G and let $\mathcal{R}$ be a complete set of inequivalent irreducible H-adapted matrix representations of G. Let $Y \subset G$ be a set of coset representatives for G/H. Then with the notation of (6) and (7)*

$$T_G(\mathcal{R}) \leq \left|\frac{G}{H}\right| T_H(\mathcal{R}_H) + M_G(Y,\mathcal{R}) \tag{8}$$

or equivalently

$$t_G(\mathcal{R}) \leq t_H(\mathcal{R}_H) + m_G(Y,\mathcal{R}). \tag{9}$$

A better bound may be obtained using the block diagonality of $\widehat{f}_y(\rho\downarrow H)$. We take this into account in Sections 4.2 and 4.1.

The inequalities (8) and (9) can be viewed as recurrences which bound the complexity of a group in terms of the complexity of a subgroup. The recurrence may be iterated through a chain of subgroups for G. For example consider the chain of subgroups

$$G = K_n > K_{n-1} > \cdots > K_0. \tag{10}$$

We say that $\mathcal{R}$, a set of irreducible representations of G, is **adapted to the chain** (10) provided $\mathcal{R}$ is K_i-adapted for each subgroup K_i in the chain. Using the notation of Definition 3, this implies that each $\mathcal{R}_{K_i}$ is K_j-adapted for $j \leq i$. Theorem 3.1 now generalizes immediately.

Theorem 3.2 *Let G have the chain of subgroups* (10) *and for $i = 1,\ldots,n$, let Y_i be a set of coset representatives for K_i/K_{i-1}. If $\mathcal{R}$ is a set of matrix representations of G adapted to this chain, then*

$$t_G(\mathcal{R}) \leq t_{K_0}(\mathcal{R}_{K_0}) + \sum_{i=1}^{n} m_{K_i}(Y_i, \mathcal{R}_{K_i}). \tag{11}$$

When $G = H \times K$ is a direct product we get a special case of Theorem 3.1. The irreducible representations of G may all be obtained as tensor products of those of H and K, and the product basis constructed by the tensoring of a basis for the irreducible representations of H with those of K yields irreducible representations which are both H-adapted and K-adapted, up to a relabeling of the matrix rows and columns (cf. [11], Satz 5.8). If $\mathcal{R}'$, $\mathcal{R}''$ are sets of matrix representations of representations of H and K respectively then let $\mathcal{R}' \otimes \mathcal{R}''$ be the set of matrix tensor products of representations in $\mathcal{R}'$ with those in $\mathcal{R}''$.

Theorem 3.3 *(i) If $\mathcal{R}$, $\mathcal{R}'$ are sets of matrix representations of representations of H and K respectively, then*

$$t_{H\times K}(\mathcal{R}\otimes\mathcal{R}) \leq t_H(\mathcal{R}) + t_K(\mathcal{R}')$$

(ii) Let ρ be an irreducible K-adapted matrix representation of $H \times K$. Then there are irreducible matrix representations, ρ_H, ρ_K, of H and K respectively such that $\rho = \rho_H \otimes \rho_K$, as matrix representations, and hence ρ is also H-adapted.

(iii) Let $\mathcal{R}$ be a complete set of irreducible representations of $H \times K$. If $\mathcal{R}$ is both H-adapted and K-adapted then there are sets, $\mathcal{R}_H$, $\mathcal{R}_K$, of irreducible matrix representations of H and K respectively, such that $\mathcal{R} = \mathcal{R}_H \otimes \mathcal{R}_K$, as sets of matrix representations.

(iv) Let $\mathcal{R}$ be a set of irreducible matrix representations of a finite group G with center Z. Then $\mathcal{R}$ is Z-adapted. Therefore if $G = H \times K$ is a product of groups and H is abelian, then $\mathcal{R}$ is H-adapted.

Proof: (i) is a result of Beth [11]. (ii) and its corollaries, (iii) and (iv), are simple consequences of Schur's lemma (Lemma 4.2).

QED

Theorems 3.1 and 3.2 suggest that one approach to minimizing an upper bound of t_G, and hence T_G, is to attempt to efficiently evaluate sums of the form $\sum_{y\in Y} \rho(y)F(y)$, where the $F(y)$ are $d_\rho \times d_\rho$ matrices. Towards this end several possibilities are evident. The subgroup chain can be varied, as can the choice of coset representatives, so as to obtain matrices $\rho(y)$ with useful computational properties. Another idea is to attempt to use the properties of the matrix elements of $\rho(y)$ as special functions on the set Y. In this paper we explore the first approach.

Convention. Almost all of the results in remaining sections depend only on the adaptability of the representations and not the particular choice of adapted representation. For this reason explicit reference to a fixed $\mathcal{R}$ is often superfluous and we suppress this in much of the notation (e.g. we will write t_K for $t_K(\mathcal{R}_K)$ and $m_K(Y_i)$ for $m(Y, \mathcal{R}_K)$).

4 The main idea - Separation of variables

In this section we present the main new computational techniques for efficiently computing nonabelian Fourier transforms. We start by generalizing the approach of Section 3 to the separation of variables algorithm. This algorithm reduces the computation of a sum of products to other, potentially smaller, repeated sums of products. We then give a detailed analysis of the complexity of matrix multiplication when the matrices have a special structure related to a subgroup-adapted representation. These results on matrix multiplication produce the bulk of the new computational savings presented in this paper. The key idea here is that if representations are adapted to a subgroup, then any element in the centralizer of this subgroup is, by Schur's Lemma, guaranteed to have a sparse representation matrix. If coset representatives can be factored as products of such elements, then multiplication by the representation matrices of these coset representatives may be performed efficiently. When these elements are also contained in a proper subgroup of the group for which the representation remains adapted, the representation matrices are even sparser. Finally, we look at the effect of using a subgroup chain in this setting and present some general results on the complexities of our algorithms.

4.1 Sums of Products — the separations of variables idea

Let G be a finite group, Y a subset of G, ρ a matrix representation of G, and for each $y \in Y$, let $F(y)$ be a $d_\rho \times d_\rho$ matrix. In this section we focus on a method for computing sums of the form

$$\sum_{y\in Y} \rho(y)F(y). \tag{12}$$

This is a general setting which encompasses the algorithmic issues which we treat in this paper. For example, if we take $Y = G$ and $F(y) = f(y) \cdot I_{d_\rho}$[2], for some complex-valued function, f on G, then the sum (12) is $\hat{f}(\rho)$. If we let

[2]For any positive integer d, I_d will denote the $d \times d$ identity matrix.

Y be a set of coset representatives of a subgroup, $H < G$, and $F(y) = \widehat{f_y}(\rho \downarrow H)$, where $f_y(h) = f(y \cdot h)$ for $y \in Y$ and $h \in H$, then we are precisely in the setting of Theorem 3.1. Thus the results of this section may be applied both directly to the computation of Fourier transforms and indirectly in conjunction with the methods of Section 3.

We shall now define an algorithm for computing (12), which we call the **separation of variables algorithm**. Its definition depends on a choice of a set of words X, in elements of G, such that the associated set of group elements obtained by multiplying out the formal products is equal to Y. Let S denote the set of group elements which occur as symbols in the words of X, together with the identity, which we denote as e. For simplicity, assume that X has the same cardinality as Y. Thus, the words in X may be thought of as a choice of factorization for the elements of Y in terms of S. In practice S is usually choosen before X.

Let γ be the maximum length of any word in X. To avoid the need for special conventions to deal with the empty word, it is useful to assume that all words in X have length γ. This can be achieved by "padding" on the left with the identity if necessary. Call the resulting set of words X_0.

For each i define X_i to be the set of subwords of X_0 obtained by removing the rightmost i symbols from each word of X_0. Note that X_i is a set of words of length $\gamma - i$ in S. Let f be a complex-valued function on G and ρ a matrix representation of G of degree d. For w in X_0 define $F_0(w) = f(w)I_d$ where $f(w)$ is the value of f on the group element represented by w and I_d denotes the identity matrix of dimension $d = d_\rho$. The separation of variables algorithm proceeds in γ steps, computing for each i from 1 to γ, the recursively defined matrix-valued functions F_i on X_i,

$$F_i(w) = \sum_{s \in S, ws \in X_{i-1}} \rho(s) F_{i-1}(ws) \tag{13}$$

for any w in X_i. The algorithm completes by computing F_γ, which is, by the following lemma, the constant function whose value is the sum (12) with domain X_γ, consisting of only the empty word.

Lemma 4.1 *The separation of variables algorithm described above computes* $\sum_{y \in Y} \rho(y) F(y)$. *I.e., with all notation as above*

$$F_\gamma = \sum_{y \in Y} \rho(y) F(y).$$

Proof: We show by induction that for $0 \leq i \leq \gamma$,

$$\sum_{w \in X_i} \rho(w) F_i(w) = \sum_{y \in Y} \rho(y) F(y) \tag{14}$$

To start, note that (14) holds for $i = 0$ by the definition of X_0 and F_0. Now let $1 \leq i \leq \gamma$, and assume the induction hypothesis for $i - 1$. Then by (13)

$$\begin{aligned} \sum_{w \in X_i} \rho(w) F_i(w) &= \sum_{w \in X_i} \rho(w) \left[\sum_{w s_{i-1} \in X_{i-1}} \rho(s_{i-1}) F_{i-1}(w s_{i-1}) \right] \\ &= \sum_{w s_{i-1} \in X_{i-1}} \rho(w s_{i-1}) F_{i-1}(w s_{i-1}) \\ &= \sum_{w' \in X_{i-1}} \rho(w') F_{i-1}(w'). \end{aligned}$$

When $i = \gamma$ the only word in X_γ is the empty word, and $F_\gamma = \rho(e) F_\gamma$. This proves the lemma.

QED

The expression (13) shows the recursive nature of the separation of variables approach, as this sum may be rewritten in the same form as the original problem (12): by writing

$$F_i(s_\gamma \cdots s_i) = \sum_{s \in X_{i-1}(s_\gamma \cdots s_i)} \rho(s) F_{i-1}(s_\gamma \ldots s_i \cdot s) \tag{15}$$

where $X_{i-1}(s_\gamma \cdots s_i) = \{s \in S : s_\gamma \cdots s_i s \in X_{i-1}\}$, we reduce the original problem to γ subproblems of the same form. Hence we may apply the separation of variables algorithm to any of these subproblems, provided we first

choose a finer factorization of the elements $X_{i-1}(s_\gamma \cdots s_i)$. The separation of variables algorithm is the "divide" portion of a divide and conquer strategy for computing Fourier transforms; it reduces the computation of sums of products to the computation of other sums of products. Its construction only relies on having chosen factorizations for elements of the set Y. On the other hand, the "conquer" part of our strategy, which we treat in Section 4.2, uses subgroup chains and adapted bases.

It is easy to see how the separation of variables algorithm leads to the results of Section 3. Fix a subgroup $H < G$ and a set of coset representatives, Z, for G/H. Then let $Y = G$, and for any $y \in Y$, let $F(y) = f(y) \cdot I_{d_\rho}$. Let X be the set of all words, $z \cdot h$, of length two with $z \in Z$ and $h \in H$. Then for $z \in Z$ we have $X_0(z) = H$, $X_1 = Z$, and

$$F_1(z) = \sum_{h \in H} \rho(h) f(z \cdot h) = \widehat{f_z}(\rho \downarrow H). \tag{16}$$

When $i = 2$ we obtain

$$F_2 = \hat{f}(\rho) = \sum_{z \in Z} \rho(z) F_1(z) \tag{17}$$

and the separation of variables algorithm for computing $\widehat{f}(\rho)$ is exactly the algorithm considered in Section 3.

Separation of variables may be applied to the computation of both of the sums (16) and (17) by using factorizations of elements of H and of elements of Z respectively. The resulting composite algorithm is precisely the separation of variables algorithm for the set of words obtained by taking pairwise products of the padded words (i.e. the elements of X_0) used in both the algorithms for computing (16) and (17). This is a general property of the separation of variables technique; using it recursively is equivalent to using a single algorithm for a different set of words.

The applications of Section 5 will always proceed by using coset representatives to obtain a coarse factorization of group elements and then refining this factorization by factoring the coset representatives themselves.

4.2 Products of pairs of matrices

The results introduced in Sections 3 and 4.1 have focused on rewriting the Fourier transform as a recursively structured summation of matrix products. This is the "divide" component of our divide and conquer strategy. In this section we consider conditions that will ensure that a matrix product involving $\rho(a)$ for a representation ρ and element a of G may be computed efficiently. This is the "conquer" portion of our divide and conquer strategy.

The main tool we use is a form of Schur's Lemma. This simple result pins down the structure of intertwining matrices for a given matrix representation.

Lemma 4.2 (Schur) *(see e.g. [47], p. 13) Let K be a subgroup of G and ρ a K-adapted representation of G such that $\rho = \eta_1 \oplus \cdots \oplus \eta_1 \oplus \cdots \oplus \eta_r \oplus \cdots \oplus \eta_r$ where $\eta_1, \ldots, \eta_r$ are inequivalent irreducible matrix representations of K, and η_i occurs with multiplicity m_i. Then the centralizer of the collection of matrices $\rho(K)$ is*

$$\left(\mathrm{Mat}_{m_1}(\mathbf{C}) \otimes I_{d_{\eta_1}}\right) \oplus \cdots \oplus \left(\mathrm{Mat}_{m_r}(\mathbf{C}) \otimes I_{d_{\eta_r}}\right) \tag{18}$$

where I_k denotes the $k \times k$ identity matrix, $\otimes$ the usual tensor product of matrices, and $\mathrm{Mat}_n(\mathbf{C})$ is the algebra of $n \times n$ matrices.

If $a \in G$ in the centralizer of a subgroup K, then its representation matrix, $\rho(a)$, is in the centralizer of $\rho(K)$. If ρ is a K-adapted representation, then $\rho(a)$ has the form 18 after some fixed permutation of rows and columns. We interpret this as saying the matrix $\rho(a)$ is sparse and as such can be multiplied efficiently against an arbitrary $d_\rho \times d_\rho$ matrix.

Corollary 4.3 *Let all notation be as in Lemma 4.2, and let a be a group element lying in the centralizer of K. Then for an arbitrary $d_\rho \times d_\rho$ matrix F, the product $\rho(a)F$ can be computed in at most $d_\rho \left(\sum_i d_{\eta_i} m_i^2\right)$ operations.*

Proof: The bound comes from considering the number of nonzero entries of the matrix $\rho(a)$. There are at most $\sum_i d_{\eta_i} m_i^2$ nonzero entries and each nonzero entry occurs at most d_ρ times — one for each column — in the expression for the matrix product $\rho(a)F$.

QED

When a is in a proper subgroup of G, Corollary 4.3 can be improved. To explain, let $H \geq K$ and let ρ and η be representations of H and K respectively. Define

$$\mathcal{M}(\rho, \eta) = \text{the multiplicity of } \eta \text{ in } \rho\downarrow K. \tag{19}$$

Also define

$$\mathcal{M}(H, K) = \max_{\rho,\eta} \mathcal{M}(\rho, \eta) \tag{20}$$

as ρ and η run over complete sets of irreducible representations of H and K respectively.

Corollary 4.4 *Let $H \geq K$ be subgroups of G, $\mathcal{R}$ a complete set of irreducible representations of G adapted to the chain $G \geq H \geq K$, and suppose that for each ρ in $\mathcal{R}$, $F(\rho)$ is a $d_\rho \times d_\rho$ matrix. Let a be in the centralizer of K in H. Then the set of matrix products $\{\rho(a) \cdot F(\rho) | \rho \in \mathcal{R}\}$ may be computed in at most $|G| \cdot \mathcal{M}(H, K)$ operations.*

Proof: For any ρ in $\mathcal{R}$, $\mathcal{M}(H, K)$ is an upper bound for the number of nonzero entries in any column of $\rho(a)$. Hence the number of operations needed to compute any entry of the matrix $\rho(a) \cdot F(\rho)$ is bounded by $\mathcal{M}(H, K)$. There are d_ρ^2 such entries so the computation of this matrix product takes $\mathcal{M}(H, K) d_\rho^2$ operations. Summing over all representations and using the relation $\sum_{\rho \in \mathcal{R}} d_\rho^2 = |G|$ gives the result.

QED

For most purposes the upper bounds of Corollaries 4.3 and 4.4 are all we require to get good bounds for group complexity. However, in some situations a more detailed analysis of the matrix multiplications is necessary. We shall now consider the multiplication of two matrices which are block diagonal according to some subgroup restrictions and also have the block scalar form (18), though possibly for different subgroups.

To state these results, let $G \geq H \geq K$, and let ρ be a representation adapted to this chain. We introduce the notation

$$\text{End}_K(\rho\downarrow H) = \text{span}_{\mathbf{C}}\,(\rho(H)) \cap \text{Centralizer}\,(\rho(K))$$

so $\text{End}_K(\rho\downarrow H)$ is the algebra matrices with block diagonal form according to $\rho\downarrow H$ that also have the form (18) up to a fixed permutation of rows and columns. In particular, if $a \in H$ is in the centralizer of K, then $\rho(a)$ is in $\text{End}_K(\rho\downarrow H)$.

Suppose $F_1 \in \text{End}_{K_1}(\rho\downarrow H_1)$ and $F_2 \in \text{End}_{K_2}(\rho\downarrow H_2)$, where the subgroup chains $H_1 \geq K_1$ and $H_2 \geq K_2$ both occur as subchains of some fixed subgroup chain of G for which ρ is adapted. We wish to examine the complexity of the matrix multiplication $F_1 \cdot F_2$. There are a number of special cases to consider corresponding to the different possible orderings of the subgroups H_1, K_1, H_2, K_2, in the subgroup chain. By exchanging F_1 and F_2 the number of cases under consideration is reduced from six to three. We shall consider one of these cases in detail and then indicate the adaptations needed to treat the other two.

Theorem 4.5 *Let $H_1 \geq H_2 \geq K_1 \geq K_2$ be a chain of subgroups of G, and let ρ be a representation of G adapted to this chain. Suppose that for $i = 1, 2$, $F_i \in \text{End}_{K_i}(\rho\downarrow H_i)$. Then the matrix multiplication $F_1 \cdot F_2$ can be computed in no more than*

$$\sum_{\rho_{H_1}, \rho_{H_2}, \rho_{K_1}, \rho_{K_2}} \mathcal{M}(\rho_{H_1}, \rho_{K_1}) \mathcal{M}(\rho_{H_2}, \rho_{K_2}) \mathcal{M}(\rho_{H_1}, \rho_{H_2}) \mathcal{M}(\rho_{H_2}, \rho_{K_1}) \mathcal{M}(\rho_{K_1}, \rho_{K_2}) \tag{21}$$

scalar operations, where for $L \in \{H_1, H_2, K_1, K_2\}$, the index ρ_L ranges over all irreducible representations of the subgroup L (up to equivalence) having nonzero multiplicity in $\rho\downarrow L$.

Proof: Both matrices F_1 and F_2 belong to $\text{End}_1(\rho\downarrow H_1)$ and are therefore block diagonal with blocks corresponding to the restriction of ρ to H_1. By considering the matrix multiplication one block at a time we may restrict ourselves to the case where $H_1 = G$ and ρ is an irreducible representation of G. Even with this reduction, the proof involves some tedious indexing of the rows and columns of an adapted representation, so to keep the length of our formulae to a reasonable size, we shall restrict ourselves to the case $K_2 = 1$; the general result is obtained by an analogous argument. So from now on we assume that $G_1 = G = H_1$, $G_2 = H_2$, $G_3 = K_1$ and $G_4 = K_2 = 1$.

In this situation it is useful to index the rows or columns of the chain-adapted representation ρ by a 5-tuple, $\Lambda = (\lambda_2, \rho_2, \lambda_3, \rho_3, \lambda_4)$, where for $i = 2$ or 3, ρ_i is an irreducible representation of G_i occurRing as a matrix direct

summand of $\rho_{i-1}\downarrow G_i$, where $\rho_1 = \rho$ and λ_i is a variable indexing the particular occurrences of ρ_i as a matrix direct summand of ρ_{i-1}. When $i = 4$, λ_4 simply indexes a basis of ρ_3 of dimension $\mathcal{M}(\rho_3, 1)$, where 1 is the trivial representation of the trivial group. Said differently, λ_2 indexes the blocks of $\rho\downarrow H_2$ which contain copies of ρ_2 and λ_3 indexes the blocks of $\rho_2\downarrow K_1$ which contain copies ρ_3. By Lemma 4.2 the entries of a matrix $F_1 \in \mathrm{End}_K(\rho\downarrow G)$ have the form

$$[F_1]_{\Lambda,\Lambda'} = f_1(\lambda_2, \lambda_2', \rho_2, \rho_2', \lambda_3, \lambda_3', \rho_3)\cdot\delta_{\rho_3,\rho_3'}\delta_{\lambda_4,\lambda_4'}$$

for some complex-valued function f_1, and the entries of a matrix, F_2, in $\mathrm{End}_1(\rho\downarrow H)$ have the form

$$[F_2]_{\Lambda,\Lambda'} = \delta_{\lambda_2,\lambda_2'}\cdot\delta_{\rho_2,\rho_2'}f_2(\rho_2, \lambda_3, \lambda_3', \rho_3, \rho_3', \lambda_4, \lambda_4')$$

for some complex-valued function f_2. Therefore, the expression for the matrix product entry $[F_1\cdot F_2]_{\Lambda,\Lambda'}$ is

$$\sum_{\lambda_3''} f_1(\lambda_2, \lambda_2', \rho_2, \rho_2', \lambda_3, \lambda_3'', \rho_3)\cdot f_2(\rho_2', \lambda_3'', \lambda_3', \rho_3, \rho_3', \lambda_4, \lambda_4'). \tag{22}$$

The variables appearing in the expression (22) range over values according to Diagram 1.

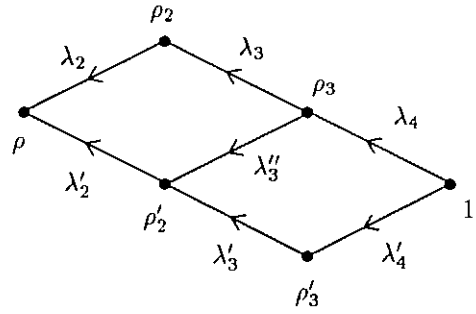

Diagram 1.

In Diagram 1 a directed edge $\beta\xleftarrow{\lambda}\alpha$ indicates that β is an irreducible representation which occurs as a matrix direct summand of the restriction of α, and that λ is a variable indexing the copy of β in this restriction. The number of operations required to compute the matrix product $F_1\cdot F_2$ is then bounded by the number of distinct ways of assigning values to $\lambda_2, \lambda_2', \rho_2, \rho_2', \lambda_3, \lambda_3', \lambda_3'', \rho_3, \rho_3', \lambda_4, \lambda_4'$ consistent with the conditions represented by Diagram 1.

To count this number, fix the two representations ρ_2' and ρ_3 and count the number of ways, if any, that the remaining variables may be assigned values in a manner consistent with the diagram. These variables may be collected into five sets corresponding to the five edges in Diagram 2; each set consists of the variables that label the path in Diagram 1 corresponding to the edge in Diagram 2.

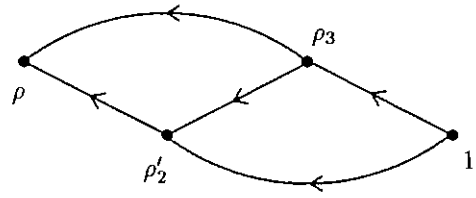

Diagram 2.

These sets of variables are $\{\lambda_2, \rho_2, \lambda_3\}$, $\{\lambda_2'\}$, $\{\lambda_3''\}$, $\{\lambda_4\}$, and $\{\lambda_3', \rho_3', \lambda_4'\}$. For a given choice of ρ_2' and ρ_3, the choices of values for variables in different sets are completely independent. Hence the choice of λ_2' is independent of the choice of λ_3''. Now consider the set of variables, $\{\lambda_2, \rho_2, \lambda_3\}$ which corresponds to the edge from ρ_3 to ρ in Diagram 2. Each different way of choosing values of these three variables, consistent with Diagram 1, corresponds to a choice of a copy of ρ_3 appearing as a matrix direct summand of the restriction of ρ to G_3, and hence there are $\mathcal{M}(\rho, \rho_3)$ possible choices. Similarly, the number of ways of choosing values for the variables in the set corresponding to any edge directed edge from α to β in Diagram 2 is $\mathcal{M}(\alpha, \beta)$. Therefore the total number of ways of assigning values to all variables in Diagram 1 is

$$\sum_{\rho_2', \rho_3}\prod_{\alpha\leftarrow\beta}\mathcal{M}(\alpha, \beta) \tag{23}$$

where the product in (23) is over all directed edges from β to α in Diagram 2. This is precisely

$$\sum_{\rho_2', \rho_3} \mathcal{M}(\rho, \rho_3)\mathcal{M}(\rho_2', 1)\mathcal{M}(\rho, \rho_2')\mathcal{M}(\rho_2', \rho_3)\mathcal{M}(\rho_3, 1)$$

Substituting $\rho_{H_1} = \rho$, $\rho_{H_2} = \rho_2'$, $\rho_{K_1} = \rho_3$ and $\rho_{K_2} = 1$ proves the theorem.

QED

The two other cases we need consider are

$$H_2 \geq H_1 \geq K_1 \geq K_2$$

and

$$H_1 \geq K_1 \geq H_2 \geq K_2.$$

The remaining three cases follow by symmetry. Extending the proof of Theorem 4.5 to these two other cases is strictly routine; the important difference is that other diagrams must be considered. For example, in the case $H_2 \geq H_1 \geq K_1 \geq K_2$, Diagram 2 must be replaced by

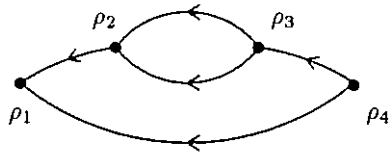

Diagram 3.

but the procedure for obtaining the complexity bound from the diagram is the same.

Theorem 4.6 *Let $H_1 \geq K_1$ and $H_2 \geq K_2$ be subgroups of G occuring in some chain of subgroups to which the representation ρ is adapted. Suppose that for $i = 1, 2$, F_i is a matrix in $\mathrm{End}_{K_i}(\rho{\downarrow}H_i)$.*

(i) When $H_2 \geq H_1 \geq K_1 \geq K_2$, the matrix multiplication $F_1 \cdot F_2$ can be computed in no more than

$$\sum_{\rho_{H_1}, \rho_{H_2}, \rho_{K_1}, \rho_{K_2}} \mathcal{M}(\rho_{H_1}, \rho_{K_1})\mathcal{M}(\rho_{H_2}, \rho_{K_2})\mathcal{M}(\rho_{H_2}, \rho_{H_1})\mathcal{M}(\rho_{H_1}, \rho_{K_1})\mathcal{M}(\rho_{K_1}, \rho_{K_2}) \tag{24}$$

scalar operations.

(ii) When $H_1 \geq K_1 \geq H_2 \geq K_2$, the matrix multiplication $F_1 \cdot F_2$ can be computed in no more than

$$\sum_{\rho_{G_1}, \rho_{G_2}, \rho_{G_3}, \rho_{G_4}} \mathcal{M}(\rho_{H_1}, \rho_{K_1})^2\mathcal{M}(\rho_{H_2}, \rho_{K_2})^2 \tag{25}$$

scalar operations.

For $L \in \{H_1, H_2, K_1, K_2\}$, the index ρ_L in the above sums ranges over all irreducible representations of the subgroup L (up to equivalence) having nonzero multiplicity in $\rho{\downarrow}L$.

Thus, Theorems 4.5 and 4.6 give exact operation counts for the appropriate matrix multiplications. It is useful to provide some notation for these counts.

Definition 4.1 *Let $H_1 \geq K_1$ and $H_2 \geq K_2$ be a chain of subgroups of G and let ρ be a representation of G. Define $C(\rho; H_1, K_1; H_2, K_2)$ to be*

1. *the sum (21) when $H_1 \geq H_2 \geq K_1 \geq K_2$,*
2. *the sum (24) when $H_2 \geq H_1 \geq K_1 \geq K_2$, and*
3. *the sum (25) when $H_1 \geq K_1 \geq H_2 \geq K_2$.*

We extend this definition to include the three other possible arrangements of H_1, H_2, K_1, K_2 in the subgroup chain, by the symmetry condition

$$C(\rho; H_1, K_1; H_2, K_2) = C(\rho; H_2, K_2; H_1, K_1)$$

It is clear that $C(\rho; H_1, K_1; H_2, K_2)$ is an upper bound for the complexity of the matrix multiplication of a matrix in $\mathrm{End}_{K_1}(\rho{\downarrow}H_1)$ with a matrix in $\mathrm{End}_{K_2}(\rho{\downarrow}H_2)$, whatever the order of the subgroups appear the subgroup chain. The next theorem gives another useful bound.

Theorem 4.7 *Let $H_1 \geq K_1$ and $H_2 \geq K_2$ be subgroups of G occuring as subchains of some chain of subgroups to which the representation ρ is adapted. Let $G_1 \geq G_2 \geq G_3 \geq G_4$ be the rearrangement of the H_i and K_i into a single chain. Then*

$$C(\rho; H_1, K_1; H_2, K_2) \leq \mathcal{M}(G_2, G_3) \cdot \sum_{\rho_{G_1}, \rho_{G_4}} \mathcal{M}(\rho_{G_1}, \rho_{G_4})^2$$

where ρ_{G_1} ranges over inequivalent irreducible representations of G_1 having nonzero multiplicity in ρ, and ρ_{G_4} ranges over inequvalent irreducible representations of G_4 having nonzero multiplicity in ρ_{G_1}.

Proof: For simplicity we only consider the case when $H_1 \geq H_2 \geq K_1 \geq K_2$; all the other cases use a similar line of proof. First note that if ρ is a representation of G, ρ_K is a representation of a subgroup of G, and H is a subgroup of G containing K, then

$$\mathcal{M}(\rho, \rho_K) = \sum_{\rho_H} \mathcal{M}(\rho, \rho_H)\mathcal{M}(\rho_H, \rho_K)$$

By Theorem 4.6 we may bound $C(\rho; H_1, K_1; H_2, K_2)$ as follows

$$\begin{aligned}
C(\rho; H_1, K_1; H_2, K_2) &= \sum_{\rho_{H_1}, \rho_{H_2}, \rho_{K_1}, \rho_{K_2}} \mathcal{M}(\rho_{H_1}, \rho_{K_1})\mathcal{M}(\rho_{H_2}, \rho_{K_2})\mathcal{M}(\rho_{H_1}, \rho_{H_2})\mathcal{M}(\rho_{H_2}, \rho_{K_1})\mathcal{M}(\rho_{K_1}, \rho_{K_2}) \\
&\leq \mathcal{M}(H_2, K_1) \sum_{\rho_{H_1}, \rho_{K_2}} \left(\sum_{\rho_{H_2}} \mathcal{M}(\rho_{H_1}, \rho_{H_2})\mathcal{M}(\rho_{H_2}, \rho_{K_2}) \right) \left(\sum_{rho_{K_1}} \mathcal{M}(\rho_{H_1}, \rho_{K_1})\mathcal{M}(\rho_{K_1}, \rho_{K_2}) \right) \\
&= \mathcal{M}(H_2, K_1) \sum_{\rho_{H_1}, \rho_{K_2}} \mathcal{M}(\rho_{H_1}, \rho_{K_2})^2
\end{aligned}$$

QED

The diagrammatic techniques introduced in the proof of Theorem 4.5 may be generalized and used to prove even better complexity results for finite groups than those given in this paper. This approach to Fourier transforms on finite groups is explained in the sequel [39]. In particular, an appropriate setting for discussing multiplication of block scalar matrices is a tower of multi-matrix algebras (cf. [34]).

4.3 Complexity of the algorithm

We now combine the ideas of Sections 4.1 and 4.2 to obtain some general upper bounds for the complexity of a Fourier transform. Assume all notation is as in Section 4.1, so that for a fixed subset $Y \subset G$, X is a set of words from a subset $S \subset G$, whose products equal Y, X_0 is obtained by padding the words of X with copies of the identity element on the left until they all have the same length γ, and X_k is obtained from X_0 by deleting k symbols from the right of each word. Furthermore, let X^i denote the set of words obtained from words of X by deleting the $\gamma - i$ leftmost symbols.

Let $K_n \geq \cdots \geq K_0$ be a chain of subgroups of G, and assume that ρ is adapted to to this chain. Given any $g \in G$, define the indices $c^+(g)$ and $c^-(g)$ by

$$\begin{aligned}
K_{c^+(g)} &= \text{the smallest subgroup in the chain containing } g \text{ and,} \\
K_{c^-(g)} &= \text{the largest subgroup of } K_{c^+(g)} \text{ in the chain which commutes with } g.
\end{aligned} \tag{26}$$

So $\rho(g) \in \mathrm{End}_{K_{c^+(g)}}(\rho{\downarrow}K_{c^-(g)})$. Let b_0^+ and b_0^- be such that $F(y) \in \mathrm{End}_{K_{b_0^+}}(\rho{\downarrow}K_{b_0^-})$ for each $y \in Y$. Then for any i between 0 and γ we let

$$\begin{aligned} b_i^+ &= \max\{b_0^+, c^+(g) : g \in X^i\} \\ b_i^- &= \min\{b_0^-, c^-(g) : g \in X^i\} \end{aligned}$$

By Definition 4.1 the number of operations needed to perform the matrix product $\rho(s_0) \cdot F_{i-1}(s_n \cdots s_1)$ appearing in the algorithm (13) is no more than $C(\rho; K_{c^+(s_0)}, K_{c^-(s_0)}; K_{b_{i-1}^+}, K_{b_{i-1}^-})$ operations.

Theorem 4.8 *Let ρ be a matrix representation of G which is adapted to a chain of subgroups, $K_n \geq \cdots \geq K_0$. Let $Y \subset G$ and for each $y \in Y$ let $F(y) \in \mathrm{End}_{K_{b_0^+}}(\rho{\downarrow}K_{b_0^-})$. Then the sum (12) may be calculated in less than*

$$\sum_{k=0}^{\gamma-1} \sum_{\substack{s_n \cdots s_0 \in X_k \\ s_0 \neq e}} C(\rho; K_{c^+(s_0)}, K_{c^-(s_0)}; K_{b_k^+}, K_{b_k^-}) \tag{27}$$

scalar operations.

Proof: By the definition of C (Definition 4.1) and Theorem 4.5, the sum (27) is an upper bound for the number of scalar operations needed for all the matrix multiplications occurring in the separation of variables algorithm. We now have to include the matrix additions as well. The proof of Theorem 4.5 shows that for each nonzero entry of the matrix products, the number of scalar multiplications used in the computation of that entry is one more than the number of scalar additions. When we include the scalar additions used to compute the matrix additions occurring in our algorithm we see that the total number of additions used is still no greater than the total number of multiplications. Hence the complexity of the algorithm is the same as the total number of multiplications and is bounded by the sum (27).

QED

We now give several simpler bounds that are direct corollaries of Theorem 4.8 and the results of Section 4.2. For this, it is useful to introduce the **multiplicity function** $\mathcal{M}$, defined on G. For a fixed chain of subgroups $G \geq K_n \geq \cdots \geq K_0$ define

$$\mathcal{M}(g) = \begin{cases} \mathcal{M}(K_{c^+(g)}, K_{c^-(g)}) & \text{if } g \neq 1, \\ 0 & \text{if } g = 1. \end{cases} \tag{28}$$

For any subset $S \subset G$ define

$$\mathcal{M}(S) = \max_{s \in S} \mathcal{M}(s). \tag{29}$$

Corollary 4.9 *Let $\mathcal{R}$ be a complete set of inequivalent irreducible matrix representations for G, adapted to the subgroup chain $K_n \geq \cdots \geq K_0$. Let H be a subgroup of G, Y a complete set of coset representatives for G/H, and X a set of factorizations of elements of Y in terms of elements from a subset $S \subset G$. Let γ be the maximum length of any word in X. Then*

$$m(\mathcal{R}, Y, H) \leq \sum_{k=0}^{\gamma-1} \sum_{s_n \cdots s_0 \in \tilde{X}_k} \mathcal{M}(s_0) \tag{30}$$

$$\leq \mathcal{M}(S) \left[\sum_{k=0}^{\gamma-1} \left| \tilde{X}_k \right| \right] \tag{31}$$

where $\tilde{X}_k$ is obtained from X by deleting k elements from the right of each word and then deleting all occurrences of the identity as symbols in these words.

Proof: This is an immediate consequence of Theorem 4.8, Theorem 4.7 and the definition of $\mathcal{M}$.

QED

Remarks. 1. Applications. Corollary 4.9 is the primary result for the applications of Section 5. It has the virtue of simplicity, but when $\mathcal{R}$ is H-adapted, it does not use the block diagonal form of the Fourier transforms of the restricted representations on H. To take this into account, Theorem 4.8 must be used directly.

2. General results. Corollary 4.9 might be useful in the search for general results on the complexity of Fourier transforms on any finite group, possibly improving on the general bounds of Clausen [19], or those of Diaconis and Rockmore [23]. As a first step in this direction, let $l_{H,S}(y)$ be the minimum non-negative integer, l, such that y is in the same coset of G/H as some product of l elements of S. The generating function $\mathcal{P}_{G/H,s}(t) = \sum_{y \in Y} t^{l_{H,s}(y)}$ is independent of the choice of coset representatives and is sometimes called the **Poincaré polynomial** of G/H with respect to S. Note that $\mathcal{P}'_{G/H,S}(1) = \sum_{y \in Y} l_{H,S}(y)$, is the sum of the lengths of minimal coset representatives for G/H.

Corollary 4.10 *Let $\mathcal{R}$ be a complete set of inequivalent irreducible matrix representations for G, adapted to the subgroup chain $K_n \geq \cdots \geq K_0$. Let H be a subgroup of G, and Y a set of minimal coset representatives for G/H, relative to the subset, S of G. Then*

$$\begin{aligned} m(\mathcal{R}, Y, H) &\leq \mathcal{M}(S).\mathcal{P}'_{G/H,S}(1) \\ &\leq \mathcal{M}(S).\gamma.\left|\frac{G}{H}\right| \end{aligned}$$

where γ is the maximum length of any element of Y in S.

Notice that this is a general upper bound, depending only on a set of generators for a finite group, and a subgroup chain.

3. Adapted diameters. In order to use Corollary 4.10 in conjunction with Theorem 3.2, we must assume that for each i, a set of coset representatives for K_i/K_{i-1} can be expressed in terms of $S \cap K_i$. In this case we say that S is a **generating set for the chain of subgroups** (32).

$$G = K_m \geq \cdots \geq K_0. \tag{32}$$

When the subgroup chain contains both the whole group, G, and the trivial subgroup, 1, a generating set for the chain is called a **strong generating set** for G with respect to the chain of subgroups (32). Strong generating sets arise naturally in the context of many algorithmic issues in computational group theory [48]. In particular, fast algorithms for their construction for stabilizer subgroup chains in permutation groups are a cornerstone for many important techniques [2].

Using the bounds of Corollary 4.10 in Theorem 3.2, we obtain an upper bound on the complexity of G in terms of the quotient sizes $|K_i/K_{i-1}|$, multiplicity data $\mathcal{M}(S)$ and combinatorial data in the form of the maximum lengths needed to construct the coset representatives at each level. This last aspect is nicely encapsulated in the notion of the **adapted diameter** of a group with respect to a generating set for a given chain of subgroups (cf. [40] for details).

4. Choosing the generating set or subgroup chain. The complexity bounds of Theorem 4.8, Corollary 4.9, and Theorem 3.1 only depend on the choice of subgroup chain and on the choice of factorization for group elements, e.g. they do not depend on the choice of a particular adapted basis. We now discuss some ideas which guide these choices with the aim of minimizing the complexity (27) of the separation of variables algorithm.

For a fixed factorization, refining the subgroup chain always decreases the bound (27). This is because the complexity for the matrix product, $C(\rho; H_1, L_1; H_2, L_2)$, is decreased if we increase L_1 or L_2 or if we decrease H_1 or H_2. Refining the subgroup chain therefore decreases $C(\rho; K_{c^+(s_0)}, K_{c^-(s_0)}; K_{b^+_{i-1}}, K_{b^-_{i-1}})$. Of course, this is also changing the original problem, as we must assume our representations are adapted to the new subgroup chain, so that Theorem 4.9 applies; this is an additional hypothesis.

It is conceivable that for a given group, a natural chain of subgroups may be given; in this case we are faced with the problem finding a factorization of group elements that makes the separation of variables algorithm efficient. If we plan to apply separation of variables recursively through the chain,

$$G = K_n \geq K_{n-1} \geq \cdots \geq K_0 = 1, \tag{33}$$

then the factorization we use must be a refinement of a factorization using coset representatives, and the set of generators, S, for the factorization is necessarily a strong generating set (cf. Remark 3).

We now construct a strong generating set with minimal $\mathcal{M}(S)$. For any subgroups $H \geq L$ in the subgroup chain, this set will also minimize the quantity

$$\max_{s \in S} C(\rho; K_{c^+(s)}, K_{c^-(s)}; H, L) \tag{34}$$

over all strong generating sets for the chain (33). We start by defining $S(0) = K_0 = 1$, which clearly solves this problem for the trivial group. Then we define $S(i) = S \cap K_i$ inductively by

$$S(i) = S(i-1) \cup (K_i \cap \text{Centralizer}(K_j))$$

where j is chosen to be maximal with respect to the property that $S(i)$ generates K_i. By induction, $S(i)$ is a strong generating set for the chain $K_i \geq \cdots \geq K_0$ and $S = S(n)$ minimizes both $\mathcal{M}(S)$ and (34) amongst strong generating sets. Note that we do not need to calculate any restriction multiplicities to find this generating set.

Minimizing $\mathcal{M}(S)$ places a restriction on the generators which increases the lengths of factorizations. In practice it seems that the advantage of smaller multiplicities outweighs the disadvantage of long factorizations.

The converse problem is to construct a subgroup chain from a generating set so the complexity of the separation of variables algorithm is small. Suppose now, that we are given a minimal generating set, S. Then an arbitrary ordering of elements of S as $s_1, \ldots, s_n$, defines a subgroup chain via $K_i = \langle s_1, \ldots s_i \rangle$. It is clear that $c^+(s_i) = i$ for this subgroup chain. If we draw a graph with vertices corresponding to elements of S and edges between elements that do not commute then $c^-(s_i)$ can be read straight from the graph as the largest j such that s_i is not connected to any of $s_1, \ldots s_j$ by an edge. Ordering S corresponds to labeling the vertices of this graph with numbers from 1 to n. Finding an ordering of S such that the numbers $c^+(s_i) - c^-(s_i)$ are minimized is related to the problem of drawing the graph in a form which is "close" to a chain.

5 Applications

The results of Section 4 may be immediately applied to derive useful upper bounds for the complexities of many families of finite groups. We first show how our general machinery reobtains the best known FFT's for some abelian groups, the symmetric groups and their wreath products. We then move on to derive new results for some of the families of classical groups over finite fields as well as their various generalizations.

Our usual approach is via Corollary 4.9. Thus in each situation we require a chain of subgroups with accompanying sequence of coset representatives. For families of groups which nest naturally (e.g. symmetric groups, general linear groups) the subgroup chains contain the nesting and we get a recursive description of the algorithm. To take full advantage of Corollary 4.9 the coset representatives should admit a factorization in terms of a generating set such that the value of $\mathcal{M}$ on the generators is small.

5.1 Finite abelian groups

Applications in digital signal processing and data analysis motivated the need for a fast cyclic discrete Fourier transform (cf. the example of Section 2.1) and more generally a fast Fourier transform on any abelian group [26, 43]. Application of Corollary 4.9 immediately gives us some well-known results bounding the complexity of the Fourier transform on any finite abelian group.

Theorem 5.1 *Let A be a finite abelian group whose order has the prime factorization $|A| = p_1^{r_1} \ldots p_m^{r_m}$. Then for any complete set of irreducible representations $\mathcal{R}$ of A,*

$$C_A \leq T_A(\mathcal{R}) \leq |A| \sum_{i=1}^{m} r_i p_i.$$

Proof: Since A is abelian, the irreducible representations of A are all one-dimensional. Thus, the unique complete set of irreducible representations is adapted with respect to any chain of subgroups of G. Let $S = A$ be the generating set for A. As all representations of A are one-dimensional, $\mathcal{M}(S) = 1$ with respect to any chain of subgroups. Let

$A = K_n > \cdots > K_0 = \{1\}$ be any chain of subgroups of A. For a fixed i let Y_i be any complete set of coset representatives for K_i/K_{i-1} and let $X = Y_i$ be the set of trivial factorizations of elements of Y_i (*i.e.* each element in Y_i is represented by the one element word consisting of itself). Clearly, $X_1 = \emptyset$ so that $m(\mathcal{R}_{K_i}, Y_i, K_{i-1}) \leq |Y_i|$, by Corollary 4.9. Applying Theorem 3.2 then yields

$$t_A \leq \sum_{i=1}^{n} \frac{|K_i|}{|K_{i-1}|}. \tag{35}$$

The right-hand side of (35) is a sum of divisors of $|A|$ whose product is equal to $|A|$. Such a sum is minimized precisely when each term $|K_i| / |K_{i-1}|$ is prime. This type of chain can always be found in an abelian group and any chain of subgroups of A may be refined to such a chain. Hence the theorem is proved.

QED

This is essentially the derivation of the well-known Cooley-Tukey FFT [20]. Note that when $|A| = 2^n$ we find that $\mathcal{C}(A) \leq n \cdot 2^n = |A| \log_2 |A|$. For primes greater than 2 other techniques have been discovered for further optimizing the discrete Fourier transform (see e.g. [26]). For any abelian group A, $\mathcal{C}(A) \leq 8\,|A| \log_2 |A|$ (cf. [9]).

5.2 FFTs for S_n and other Weyl groups

Applications in data analysis as well as the analysis of certain random walks related to card shuffling (cf. [22]) have motivated recent work related to FFT's for the symmetric group. For a survey of some approaches to these algorithms see [17]. In this section we show how the most efficient known algorithm due to Clausen (cf. [17]) can be rederived by our general approach and then show how our techniques extend directly to the other Weyl groups.

For the symmetric group we use the natural chain of subgroups

$$S_n > S_{n-1} > \cdots > S_1 = \{1\} \tag{36}$$

where S_k is identified with the subgroup of S_n of elements fixing each of the points $k+1, \ldots, n$. This chain has a natural generalization in the other Weyl groups.

Theorem 5.2 *Let S_n denote the symmetric group on n elements. If $\mathcal{R}$ is any complete set of irreducible representations of S_n adapted to the chain of subgroups (36). Then*

$$\mathcal{C}(S_n) \leq T_{S_n}(\mathcal{R}) \leq \frac{(n+1)n(n-1)}{3} \cdot n!. \tag{37}$$

Proof: Take as generating set the pairwise-adjacent transpositions, $S = \{t_2, \ldots, t_n\}$, where t_j denotes the transposition $(j-1, j)$. Note that

$$\begin{cases} t_j \in S_j \text{ and} \\ t_j \text{ commutes with } S_k \text{ for } k < j-1. \end{cases}$$

Thus, in the notation of Section 4.1

$$\begin{aligned} K_{c^+(t_j)} &= S_j \quad \text{and} \\ K_{c^-(t_j)} &= S_{j-2}. \end{aligned}$$

Furthermore, it is easily derived from the combinatorics of Young tableaux and "Young's rule" (cf. [32], p. 51) that the maximum multiplicity occurring in the restriction of any irreducible representation from S_k to S_{k-2} is 2, i.e. $\mathcal{M}(S_k, S_{k-2}) = 2$, so that $\mathcal{M}(t_j) = 2$. Lastly, note that coset representatives for S_n/S_{n-1} of minimal length in the generating S are given by the elements

$$\begin{aligned} Y &= \{1, t_n, t_{n-1}t_n, \ldots, t_2 \cdots t_n\} \\ &= \{1, (n\ n-1), (n-2\ n-1\ n), \ldots, (1\ \cdots\ n)\} \end{aligned}$$

If we let X be the corresponding set of words, then in the notation of Section 4.1

$$X_k = \{t_{n-k}, \ldots, t_2 \cdots t_{n-k}\}$$

and the longest product in X has length $\gamma = n-1$. Therefore

$$\sum_{k=0}^{\gamma-1} |X_k| = \frac{n(n-1)}{2}.$$

Plugging this data into Corollary 4.9 gives the recurrence $t_{S_n} \leq t_{S_{n-1}} + n(n-1)$ which is easily iterated to finish the proof.

QED

Remark. The bound of Theorem 5.2 is on the order of $n!(\log_2 n!)^3$. In this case the representations given by Young's orthogonal form or Young's seminormal form (cf. [32], p. 114) are adapted for the chain of subgroups (36) for S_n. The resulting algorithm is the best known for computing a Fourier transform on S_n [17].

The above discussion for S_n generalizes naturally to all Weyl groups. The pairwise adjacent transpositions correspond to simple reflections, and the chain of subgroups (36) is the corresponding chain of parabolic subgroups. From this point of view the coset representatives we use are quite natural: they are the minimal coset representatives for the chain of parabolic subgroups, and their factorization comes from the Bruhat order by taking subwords of the unique minimal coset representative of maximal length. In this language (The book [35] is a good reference for the basic material) the generalization of Theorem 5.2 to the Weyl groups B_n and D_n is straightforward.

We shall consider the chains of parabolic subgroups

$$\begin{aligned} B_n &> B_{n-1} > \cdots, \\ D_n &> D_{n-1} > \cdots, \end{aligned} \tag{38}$$

and the generating sets consisting of the simple reflections. The minimal coset representatives with respect to (38) are well-known as are explicit expressions for the corresponding Poincare series. There are explicit formulae for the multiplicities of the restrictions of the classical Weyl groups to any parabolic subgroup in terms of the Littlewood-Richardson coefficients .

The results we obtain for the groups B_n and D_n are superseded by the results on wreath products in the next section (cf. Theorem 5.6). However, the techniques used here illustrate the combinatorial methods used in our construction of FFTs on Chevalley groups (cf. Section 5.6).

Theorem 5.3 *Consider the Weyl groups B_n and D_n. If $\mathcal{R}_B$ and $\mathcal{R}_D$ are complete sets of irreducible representations of B_n and D_n respectively, each adapted to the appropriate chain of subgroups (38). Then*

(i)
$$\mathcal{C}(B_n) \leq T_{B_n}(\mathcal{R}_B) \leq \frac{(n+1)n(4n-1)}{3} \cdot |B_n|$$

and

(ii)
$$\mathcal{R}(D_n) \leq T_{D_n}(\mathcal{R}_D) \leq \frac{4(n+1)n(n-1)}{3} \cdot |D_n|.$$

Before we prove Theorem 5.3 we state some lemmas which provide the data needed to apply Corollary 4.9 to this situation.

Lemma 5.4 *(i) The maximum multiplicity occurring in a restriction of any irreducible representation of S_n to S_{n-1}, B_n to B_{n-1}, or D_n to D_{n-1} is 2, i.e. $\mathcal{M}(S_n, S_{n-1}), \mathcal{M}(B_n, B_{n-1}), \mathcal{M}(D_n, D_{n-1}) \leq 2$.*

(ii) The maximum dimension of a representation of $D_3 \cong S_4$ is 3.

Proof:

(i) It is well-known that for the restriction of an irreducible representation of S_n to S_{n-1} is multiplicity-free (see e.g. [32]) as is that of B_n to B_{n-1} (see e.g. [55]). The result for D_n follows easily from that of B_n, and the fact that D_n is of index 2 in B_n.

(ii) This follows from the hook formula, see ([32], p. 77).

QED

The minimal coset representatives and the sums of their lengths may be found using the following lemma.

Lemma 5.5 *(cf. [35]) Let $\mathbf{W}$ be a Weyl group with S its set of simple reflections. For any subset $J \subset S$ let $\mathbf{W}_J$ denote the corresponding parabolic subgroup. Let $\mathcal{P}_{\mathbf{W}/\mathbf{W}_J,S}(t)$ denotes the Poincare polynomial of $\mathbf{W}/\mathbf{W}_J$ in the variable t. Then the sum of the lengths of the minimal coset representatives of $\mathbf{W}/\mathbf{W}_J$ is given by*

$$\mathcal{P}'_{\mathbf{W}/\mathbf{W}_J,S}(1) = \frac{1}{2}\left|\frac{\mathbf{W}}{\mathbf{W}_J}\right|[N_S - N_J]$$

where P' denotes the derivative with respect to t and where N_S and N_J are the numbers of reflections in $\mathbf{W}$, $\mathbf{W}_J$ and hence the lengths of the longest elements in $\mathbf{W}$ and $\mathbf{W}_J$ respectively. In addition the minimal coset representatives for $\mathbf{W}/\mathbf{W}_J$ and their minimal factorizations all occur as subwords of a minimal factorization for $w_S w_J$, where w_S is the longest element in $\mathbf{W}$ and w_J is the longest word in $\mathbf{W}_J$.

In Table 1 we summarize the data required to bound the complexities for the Weyl groups.

$\mathbf{W}$	$\mathbf{W}_J$	$\mathcal{M}(S)$	$\lvert\mathbf{W}\rvert$	N_S	$\mathcal{P}'_{\mathbf{W}/\mathbf{W}_J,S}(1)$
S_n	S_{n-1}	2	$n!$	$\frac{1}{2}n(n-1)$	$\frac{1}{2}n(n-1)$
B_n	B_{n-1}	2	$2^n n!$	n^2	$n(2n-1)$
D_n	D_{n-1}	3	$2^{n-1}n!$	$n(n-1)$	$2n(n-1)$

Table 1: Combinatorial data for the Weyl groups.

It is now straightforward to use to obtain recursive bounds for the reduced complexities of these chains of groups.

Proof: [of Theorem 5.3] From the data in Table 1, Corollary 4.10, and Theorem 3.1, we obtain the recurrences $t_{B_n} \leq t_{B_{n-1}} + 2n(2n-1)$ and $t_{D_n} \leq t_{D_{n-1}} + 6n(n-1)$. Iterating the recurrence for t_{B_n} gives the result for that series of groups, but for t_{D_n} we need a more careful count.

Let $s_1, \ldots, s_n$ denote the simple reflections for D_n, in the order shown in diagram 4. Then $\mathcal{M}(s_i) = 2$ for $i \geq 4$, $\mathcal{M}(s_3) = 3$ and $\mathcal{M}(s_i) = 1$ for $i = 1$ or $i = 2$. The maximal minimal coset representative for D_n/D_{n-1} is $s_n \cdots s_3 s_2 s_1 s_3 \cdots s_n$ and the minimal coset representatives have minimal factorizations given as follows.

$$1,\ s_n,\ s_{n-1}s_n,\ \ldots,\ s_3\cdots s_n,\ s_2 s_3 \cdots s_n,\ s_1 s_3 \cdots s_n,\ s_1 s_2 s_3 \cdots s_n,\ s_3 s_1 s_2 s_3 \cdots s_n,\ \ldots,\ s_n \cdots s_3 s_2 s_1 s_3 \cdots s_n \quad (39)$$

The number of times s_3 occurs in these words is exactly equal to the number of times s_1 and s_2 occur in total, so the average value of $\mathcal{M}$ over all occurrences of symbols in the set of minimal factorizations is 2. The sum of the lengths of the minimal coset representatives of D_n/D_{n-1} is $2n(n-1)$. Therefore if we let X be equal to the set of words (39), then we have

$$\sum_{k=0}^{2n-2} \sum_{a_n \cdots a_0 \in X_k} \mathcal{M}(a_0) = 4n(n-1).$$

Applying Corollary 4.9 and Theorem 3.1 gives us $t_{D_n} \leq t_{D_{n-1}} + 4n(n-1)$. Solving this recurrence completes the proof.

QED

We have already given the minimal coset representatives for S_n/S_{n-1} and D_n/D_{n-1}. For D_n/D_{n-1} they are

$$1,\ s_n,\ s_{n-1}s_n,\ \ldots,\ s_1 \cdots s_n,\ s_2 s_1 \cdots s_n,\ \ldots,\ s_n \cdots s_1 \cdots s_n$$

where $s_1, \ldots s_n$ are the simple reflections of B_n labelled according to diagram 4.

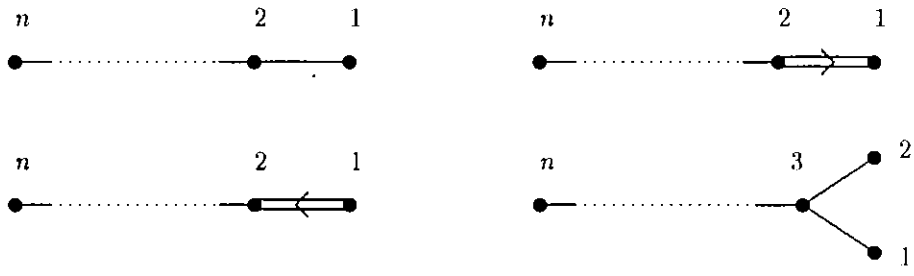

Diagram 4: Labelling the simple roots.

5.3 Wreath products of the symmetric group

For wreath products of the form $G[S_n]$, a decomposition similar to that used for Weyl groups is used. Wreath products are of interest in data analysis as the symmetry groups of nested designs [42] and in structural chemistry as the automorphism groups of non-rigid molecules [54]. They are often studied as the automorphism groups of graphs obtained by "composition" (cf. [30]).

Abstractly, $G[S_n]$ has the structure of a semidirect product $G^n \times S_n$ in the following way. Elements of this group may be described by pairs $(f;\sigma)$ where $f : \{1,\ldots,n\} \longrightarrow G$, and S_n acts on G^n by

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

for $\pi \in S_n$ and $f \in G^n$. Multiplication is defined by

$$(f;\pi)\cdot(g;\sigma) = (f\cdot g^\pi;\pi\sigma)$$

where $f\cdot g^\pi(\mathrm{j}) = f(j)g^\pi(j)$. In this notation it is clear that both S_n and G^n are naturally identified with subgroups of $G[S_n]$ and that under such an identification G^n is a normal subgroup and the group so defined is a semidirect product of these subgroups. It is not too difficult to see that such a construction makes sense for any permutation group $H < S_n$. A thorough but accessible treatment of wreath products may be found in [37].

A slight modification of the techniques used in Section 6.2 for the symmetric group yields comparable results for their wreath products. In this case we will use the chain of subgroups

$$G[S_n] > G\times G[S_{n-1}] > G[S_{n-1}] > \cdots \tag{40}$$

where $G[S_{n-1}] < G[S_n]$ denotes the subgroup of elements $(f;\sigma)$ for which σ lies in S_{n-1} and $f(n)$ is the identity element of G.

Theorem 5.6 *Let $G[S_n]$ denote the wreath product of S_n by the finite group G and let d_G denote the maximum dimension of an irreducible representation of G. Let $\mathcal{R}$ is any complete set of irreducible representations of $G[S_n]$ adapted to the chain of subgroups (40). Then,*

$$C(G[S_n]) \le T_{G[S_n]}(\mathcal{R}) \le |G[S_n]|\left[\frac{(n+1)n(n-1)}{3}(d_G)^2 + nt_G\right].$$

Proof: Note that coset representatives for $G[S_n]/G\times G[S_{n-1}]$ can be chosen to be the same as for S_n/S_{n-1} so that these coset representatives can be written as words in S, the set of pairwise-adjacent transpositions in S_n. The transposition t_j lies in $G[S_j]$ and commutes with $G[S_{j-2}]$. So if we use the chain of subgroups

$$G[S_j] > G\times G[S_{j-1}] > G[S_{j-1}] > G\times G[S_{j-2}] > G[S_{j-2}]$$

and the fact that the restriction of representations from $G[S_j]$ to $G\times G[S_{j-1}]$ is multiplicity-free (see e.g. [37]) be find that $\mathcal{M}(S)$ is $2d_G^2$, for d_G the maximum dimension of an irreducible representation of G. Using the minimal coset representatives for S_n/S_{n-1} as coset representatives for $G[S_n]/G\times G[S_{n-1}]$ we obtain the relation

$$\begin{aligned} t_{G[S_n]} &\le t_{G\times G[S_{n-1}]} + n(n-1)(d_G)^2 \\ &\le t_{G[S_{n-1}]} + t_G + n(n-1)(d_G)^2. \end{aligned}$$

Applying this inequality recursively proves the theorem.

QED

Remark. Given a subgroup chain for G it is possible to construct a chain of subgroups of $G[S_n]$ refining the chain (40). Bases adapted to the subgroup chain (40) have been constructed and the above discussion recovers the best known algorithm for wreath products of the form $G[S_n]$ (cf. [46]).

5.4 A new FFT for the general linear group over a finite field

Let $GL_n(q)$ denote the group of invertible $n \times n$ matrices with entries in the field of q elements where q is a prime power. For data analysis, these groups and their generalizations are of interest as the automorphism groups of the many designs based on finite geometries and codes (see e.g. [1]). Throughout this section all matrix groups are assumed to be over $\mathbf{F}_q$, the finite field of q elements. Thus, $GL_n \equiv GL_n(q)$, etc.

To apply the results of Section 5 to these groups, we will consider the chain of subgroups

$$GL_n > P_n > GL_{n-1} \times GL_1 > GL_{n-1} > \cdots > GL_1 \tag{41}$$

where P_n is the subgroup of GL_n of all block matrices of the form

$$\left(\begin{array}{c|c} * & * \\ \hline 0 \ \ldots \ 0 & * \end{array}\right) \tag{42}$$

and $GL_k \times GL_1$ is identified with the subgroup of block diagonal matrices of the form $\mathrm{Diag}(A, b_k, I_{n-k+1})$ with A in GL_k and b_k in GL_1 and I_r denoting the $r \times r$ identity matrix. [3]

Theorem 5.7 *Let $\mathcal{R}$ be any complete set of irreducible representations of $GL_n(q)$ adapted to the chain of subgroups (41). There is a positive constant, K, independent of n and q, such that for any $n \geq 2$, $q \geq 2$,*

$$C(GL_n(q)) \leq T_{GL_n(q)}(\mathcal{R}) \leq \frac{1}{2} 2^{2n} q^{2n-2}(1 + Kq^{-3}) \cdot |GL_n(q)| \,. \tag{43}$$

We postpone the proof of Theorem 5.7 in order to first collect the preliminary results necessary for applying Corollary 4.9. As before, we seek generators for the successive sets of coset representatives for which the values of $\mathcal{M}$ are low.

Let $E_{i,j}$ be the matrix that is zero everywhere except for a 1 in the i,j entry. For any x in $\mathbf{F}_q$ define $X_{i,j}(x) = I + xE_{i,j}$ when $i \neq j$, and let $X^*_{i,i}(x) = I + (x-1)E_{i,i}$. Also let t_i denote the transposition matrix $E_{i-1,i} + E_{i,i-1}$. These elements generate GL_n [33] and will serve as our generating set.

Factorizations of coset representatives of GL_n/P_n are easily derived from the Bruhat decomposition for GL_n (see [33, 15]) Those for $P_n/(GL_n \times GL_1)$ may be derived using some simple matrix algebra.

Lemma 5.8 *(i)* $GL_n = \coprod_{k=1}^{n} (X_{k,k+1}t_{k+1}) \cdots (X_{n-1,n}t_n).P_n$.

(ii) $P_n = X_{n-1,n}(X_{n-2,n-1}t_{n-1}) \cdots (X_{1,2}t_2).t_3 \cdots t_{n-1}.X_{n-1,n}(1).(GL_{n-1} \times GL_1)$.

We now need to calculate the the value of $\mathcal{M}$ on the elements $X_{i,i-1}$. $X_{i-1,i}$, and t_i. As a first step note that **all these elements are in GL_i and commute with GL_{i-2}**. Hence we must bound $\mathcal{M}(GL_n, GL_{n-2})$. Furthermore, $X^*_{i,i}$ lies in the centre of $GL_{n-1} \times GL_1$.

Lemma 5.9 *(i) The maximum multiplicity occuring in the restriction of any representation GL_n to GL_{n-1} no more than 2^{n-1}.*

(ii) There is a constant, $K > 0$, such that for any $n \geq 1$ and $q \geq 2$, the number of conjugacy classes of $GL_n(q)$ is less than $q^n + Kq^{n-3}$.

(iii) The maximum multiplicity occuring in the restriction of any representation of GL_n to GL_{n-2} is less than $2^{2n-3}(q^{n-1} + Kq^{n-4})$.

Proof: (i) follows straight from the paper of Thoma [50]. (ii) follows more or less directly from the asymptotics of Stong [49]. For the sake of completeness we prove this here, but postpone the proof to the appendix which follows this section. (iii) then follows from (i) and (ii) by noting that $\mathcal{M}(GL_n, GL_{n-2})$ is bounded by the product of the number of representations of GL_{n-1} with $\mathcal{M}(GL_n, GL_{n-1})$ and $\mathcal{M}(GL_{n-1}, GL_{n-2})$.

QED

[3] In general, it will be useful to adopt the standard notation that if $B_1, \ldots, B_r$ are square matrices of dimensions $d_1, \ldots, d_r$, then let $\mathrm{Diag}(B_1, \ldots, B_r) = (B_1 \oplus \cdots \oplus B_r)$ denote the block diagonal matrix with i^{th} block equal to B_i.

The following corollary is an immediate consequence of Lemma 5.9.

Corollary 5.10 *(i)* $\mathcal{M}X^*_{i,i} = 1$.

(ii) Let $g \in GL_i$ *and commute with* GL_{i-2}. *Then* $\mathcal{M}(g) \leq 2^{2i-3}(q^{i-1} + q^{i-4})$.

We are now ready to prove Theorem 5.7.

Proof: [of Theorem 5.7.] Applying Corollary 4.9 to the factorization of the first part of Lemma 5.8, gives us that

$$\begin{aligned} t_{GL_n} &\leq t_{P_n} + \sum_{k=2}^{n} 2^{2k-3} q^{-1}(q^k + Kq^{k-3}) \cdot \left(\sum_{l=1}^{k-1} q^l \right) \\ &\leq \frac{4}{15} 2^{2n} q^{2n-2}(1 + Kq^{-3}). \end{aligned}$$

Applying the same result to the second factorization of Lemma 5.8 gives

$$t_{P_n} \leq t_{GL_{n-1} \times GL_1} + q^{n-1}\mathcal{M}(GL_n, GL_{n-2}) + \sum_{k=2}^{n-1} q^k \mathcal{M}(GL_k, GL_{k-2}) + \sum_{k=3}^{n} q \cdot \mathcal{M}(GL_k, GL_{k-2}).$$

By Theorem 3.3

$$t_{GL_{n-1} \times GL_1} \leq t_{GL_{n-1}} + t_{GL_1},$$

so for $n \geq 2$ we obtain

$$t_{P_n} \leq t_{GL_{n-1}} + t_{GL_1} + \frac{1}{5} 2^{2n} q^{2n-2}(1 + Kq^{-3}).$$

Now we use these inequalities recursively. In the case of GL_1 we use the naive bound of $q - 1$ for t_{GL_1}. A careful look at the inequalities above shows that we have dropped several negative terms along the way, and that these terms dominate all the t_{GL_1} terms that appear. Thus we may ignore the t_{GL_1} terms that appear during the recursion and at the bottom of the recursion. Summing all the other terms that appear gives the final result

$$t_{GL_n} \leq \frac{1}{2} 2^{2n} q^{2n-2}(1 + Kq^{-3}).$$

QED

Remarks. 1. The constant. There is nothing particularly special about the exponent -3 appearing in the factor $(1 + Kq^{-3})$. We have shown, using a computer algebra package, that this can be replaced by a factor of $(1 + Kq^{-k})$ for $k \leq 200$ and we conjecture that in fact we may take $K = 0$. This conjecture has been verified by computer for $2 \leq n \leq 200$.

2. Further improvements. By improving the bound for t_{GL_2} we can improve on Theorem 5.7. Application of the results of [38] show that $t_{GL_2} \leq 200q \log q$. In fact, a generalization of our methods, applied to the appropriate subgroup chain of GL_2 shows that t_{GL_2} may be bounded by $5q - 3$; for details see [39].

3. Variations of the algorithm. There is of course nothing canonical about either the generators chosen here for GL_n or the subgroup chain. It seems highly likely that better choices for either are possible. Always, commutativity will need to be exploited and here it may be necessary to effectively compute the centralizers of various subsets of elements. Towards this end, recent advances in computational group theory for matrix groups [6] may prove useful.

4. Other work. The problem of finding an efficient algorithm for computing a Fourier transform for $GL_n(q)$ was first considered in [41]. There an algorithm is proposed which uses "models" (direct sums of induced one-dimensional representations which contain each irreducible of the group exactly once) to compute a Fourier transform for GL_n. In so doing the algorithm proceeds in two parts: (1) Computing the Fourier transform at reducible representations which are given by monomial matrices and then (2) applying projection operators to these reducible matrices in order to obtain collection of unique irreducible Fourier transforms. Some simple asymptotics for the bounds they obtain yield an estimate for the complexity of their algorithm to be

$$O(|GL_n(q)| q^{\frac{n^2-2n}{4}}).$$

5. Direct approach. It is also necessary to compare our algorithm with the algorithm which uses the subgroup chain but does not factor the coset representatives and thus performs direct matrix multiplication of the twiddle factors. Straightforward analysis then shows that such an algorithm yields an upper bound which depends on the maximum degree of an irreducible representation of GL_n, which is of the order of $q^{\frac{1}{2}(n^2-n)}$. This direct algorithm gives an upper bound of

$$O(nq^{\frac{1}{2}(n^2-3n)}|GL_n(q)|).$$

5.5 The unitary group over a finite field

Let $U_n(q^2)$ denote the group of unitary $n \times n$ matrices with entries in the field of q^2 elements, relative to the field automorphism of order 2, where q is some prime power. We shall often abbreviate this to U_n. To simplify our calculations we shall always assume that q is odd. We consider the chain of subgroups

$$U_n > U_{n-1} > \cdots > U_1 \tag{44}$$

where U_k is identified with the subgroup $\mathrm{Diag}(U_k, I_{n-k})$ of U_n.

Theorem 5.11 *Let $\mathcal{R}$ be any complete set of irreducible representations of $U_n(q^2)$ adapted to the chain of subgroups (44). There is a positive constant, K, such that for any $n \geq 2$, $q \geq 2$,*

$$C(U_n(q^2)) \leq T_{U_n(q^2)}(\mathcal{R}) \leq |U_n(q)| \frac{12}{7} q^{3n-3}(1 + 2q^{-1} + 4q^{-2} + Kq^{-3}). \tag{45}$$

We shall first prove the following weaker but simpler result:

Claim: With all notation as in Theorem 5.11,

$$t_{U_n} \leq \frac{8}{7} q^{3n-2}(1 + 2q^{-1} + 4q^{-2} + Kq^{-3}). \tag{46}$$

To prove the Claim we proceed as in the case of GL_n and obtain a factorization of any element of U_n as a product of matrices which are either diagonal or have a single 2×2 block with ones on the diagonal elsewhere. The multiplicity results we will need are given in the following lemma.

Lemma 5.12 *(i) The maximum multiplicity occuring in the restriction $U_n \downarrow U_{n-1}$ is 1.*

(ii) There is a constant, $K > 0$, such that for any $n \geq 1$ and $q \geq 2$, the number of conjugacy classes of $U_n(q^2)$ is less than $q^n + 2q^{n-1} + 4q^{n-2} + Kq^{n-3}$.

(iii) The maximum multiplicity occuring in the restriction $U_n \downarrow U_{n-2}$ is less than $q^n + 2q^{n-1} + 4q^{n-2} + Kq^{n-3}$.

Proof: (i) is a result of Hagedorn [29]. (ii) is proved in the appendix which follows this section. (iii) is a direct consequence of (i) and (ii).

QED

So as to not unduly interrupt the flow of the section the necessary factorization of coset representatives of U_n/U_{n-1} is obtained using some simple geometry in the appendix which follows this section. To state the result succinctly, we let $u_i(x_1, x_2)$ be the block diagonal matrix with 1's on the diagonal except for a 2×2 block of the form

$$\begin{pmatrix} -x_2^q & x_1 \\ x_1^q & x_2 \end{pmatrix}$$

This matrix is in $U_n(q^2)$ provided that $x_1^{1+q} + x_2^{1+q} = 1$.

Lemma 5.13 *Let N be the group homomorphism on $F^\times$ given by $N(\alpha) = \alpha^{1+q}$ and let R be a complete set of coset representatives for $F^\times/\ker N$. Then every coset of U_n/U_{n-1} has at least one coset representative of the form $\varepsilon \cdot a_2 \cdots a_n$, where ε is an element of F satisfying $\varepsilon^{1+q} = 1$ and for $2 \leq i \leq n-1$, the matrix a_i has one of the following forms*

(A) $a_i = u_i(r,x)$ *for some* $r \in R$, $x \in F$ *such that* $r^{1+q} + x^{1+q} = 1$.

(B) $a_i = t_{i+1}u_i(r,x)$ *for some* $r \in R$, $x \in F$ *such that* $r^{1+q} + x^{1+q} = 1$.

(C) $a_i = t_i t_{i+1} u_i(r, r\delta)$, *where* r *is the unique element of* R *with* $r^{1+q} = \frac{1}{2}$, *and* $\delta \in F$ *satisfies* $\delta^{1+q} = 1$.

The factor a_n *has the form (A).*

We can now prove the claim.

Proof: [of the Claim.] Applying Corollary 4.9 to the factorization of Lemma 5.13, shows that

$$\begin{aligned} t_{U_n} &\leq t_{U_{n-1}} + q^{n-1}(q^n - (-1)^n)\left[1 + g_{n-1}(q^2) + 3\sum_{k=2}^{n-1} g_{k-1}(q^2)\right] \\ &\leq t_{U_{n-1}} + 4q^{3n-2}(1 + 2q^{-1} + 4q^{-2} + K_1 q^{-3}) \end{aligned}$$

for some constant K_1, where $g_k(q^2)$ denotes the number of conjugacy classes of $U_n(q^2)$. Using this inequality recursively and noting that $t_{U_1} \leq q + 1$ gives the result (46).

QED

With these preliminary results in hand we now easily prove Theorem 5.11.

Proof: [of Theorem 5.11.] The improvement on the Claim comes from looking at the matrix multiplications in the separation of variables algorithm more carefully. Suppose we are computing the Fourier transform at the adapted irreducible representation, ρ. At some point in the algorithm we will calculate matrix products of the form $\rho(a_n) \cdot \hat{h}(\rho{\downarrow}U_{n-1})$, where $a_n \in U_n$ commutes with U_{n-2} and $\hat{h}(\rho{\downarrow}U_{n-1})$ is in $(\mathrm{End}\, V_\rho)^1_{U_{n-1}}$. To obtain the complexity result (46) we used the bound of $\mathcal{M}(a_n)d_\rho^2$ for the complexity of such a matrix multiplication—a bound which comes without assuming any special form of the matrix $\hat{h}(\rho{\downarrow}U_{n-1})$. However, we could get a better result by using part Theorem 4.7 to bound the complexity of that matrix multiplication:

$$C(\rho; U_n, U_{n-2}; U_{n-1}, 1) \leq d_\rho^2$$

Using this new complexity gives us

$$\begin{aligned} t_{U_n} &\leq t_{U_{n-1}} + q^{n-1}(q^n - (-1)^n)\left[1 + 1 + 3\sum_{k=2}^{n-1} g_{k-1}(q^2)\right] \\ &\leq t_{U_{n-1}} + 6q^{3n-3}(1 + 2q^{-1} + 4q^{-2} + K_2 q^{-3}) \end{aligned}$$

Using this bound recursively proves the theorem.

QED

5.6 Chevalley groups

The techniques used to compute a Fourier transform in GL_n may be extended in a relatively straightforward manner to Chevalley groups and other finite groups of Lie type. We refer the reader to the book of Carter [14] for definitions. We limit the current discussion to the classical Chevalley groups although the techniques generalize in a natural way to other finite groups of Lie type.

As usual, let $A_n(q)$, $B_n(q)$, $C_n(q)$, $D_n(q)$ denote the simply connected forms of the Chevalley groups over a finite field with q elements. Any Chevalley group, G, has a subgroup chain analogous to (41), where P_{n-1} is replaced by a maximal parabolic subgroup $GL_{n-1} \times GL_1$ by its reductive part, and GL_{n-1} by the semisimple part of the parabolic subgroup. More specifically, we shall label the simple roots of a rank n group from 1 to n in the order shown in Diagram 4. Then P_k will denote the parabolic subgroup corresponding to the set of simple roots labeled from 1 to k with reductive part L_k and semisimple part G_k (not to be confused with the exceptional group G_2). For any Chevalley group G the chain of subgroups we shall use in the construction of a fast Fourier transform on G, will always be

$$G_n \geq P_{n-1} \geq L_{n-1} \geq G_{n-1} \geq \cdots \geq G_1 \tag{47}$$

Theorem 5.14 *There exist positive constants K_n, such that for any $n \geq 2$, $q \geq 2$,*

(i) $T_{A_n(q)} \leq K_n q^{2n+1}.|A_n(q)|$

(ii) $T_{B_n(q)} \leq K_n q^{5n-3}.|B_n(q)|$

(iii) $T_{C_n(q)} \leq K_n q^{5n-3}.|C_n(q)|$

(iv) $T_{D_n(q)} \leq K_n q^{5n-6}.|D_n(q)|$, *for* $n \geq 4$, *and* $T_{D_3(q)} \leq K_3 q^{10}.|D_3(q)|$, *for* $n = 4$

where the complexities are taken with respect to a complete set of representations adapted to the chain (47).

We shall give the proof of Theorem 5.14 after we have collected some lemmas on multiplicities and factoring elements in these groups.

We refer the reader to [14] for all the relevant notation. For any root α in the root system of G_n, we let X_α denote the corresponding root subgroup. We also let s_α denote the corresponding involution in the Weyl group, and let n_α be an element of N mapping onto s_α where N comes from the BN-pair for G_n. We shall denote the simple roots $\alpha_1, \ldots, \alpha_n$ according to Diagram 4. With the exception of the root α_3 of D_3, we know that X_{α_i} and n_{α_i} lie in G_i and commute with G_{i-2}. Consequently, the construction of an FFT using a factorization in terms of the X_{α_i} or n_{α_i}, will require that we understand the maximum multiplicity $\mathcal{M}(G_i, G_{i-2})$.

Lemma 5.15 *Let $G \downarrow K$ be one of the restrictions, $A_n(q) \downarrow A_{n-2}(q)$, $B_n(q) \downarrow B_{n-2}(q)$, $C_n(q) \downarrow C_{n-2}(q)$ or $D_n(q) \downarrow D_{n-2}(q)$. Then for fixed n the maximum multiplicity $\mathcal{M}(G, K)$ is bounded by a function of q of the form $O(q^{\sigma(G,K)})$, where*

$$\sigma(G, K) = \frac{1}{2}\left[\dim G - \operatorname{rank} G - \dim K - \operatorname{rank} K\right].$$

Proof: This is proved in the appendix following this section using an argument due to Tom Hagedorn. See also [29].

QED

The other piece of information we need concerns the factorization of coset representatives in terms of the elements X_{α_i} and n_{α_i}.

Lemma 5.16 *Let G be a simply connected Chevalley group with Weyl group $\mathbf{W}$ and let J be any subset of the set of simple roots of $\mathbf{W}$. Let $\mathbf{W}_J$ denote the parabolic subgroup corresonding to J and $\mathbf{W}^J$ the the set of minimal coset representatives for $\mathbf{W}/\mathbf{W}_J$. We let N, N_J denote the number of positive roots of $\mathbf{W}$, $\mathbf{W}_J$ respectively. Also let P_J denote the parabolic subgroup of G corresponding to J, let L_J and U_J be its reductive and maximal normal unipotent parts, let $Z(L_J)$ be the center of L_J and G_J be the semisimple part of L_J. Then*

(i)

$$G = \left[\coprod_{\substack{w \in \mathbf{W}^J \\ w = s_{\beta_1} \cdots s_{\beta_k}}} (X_{\beta_1} n_{\beta_1}) \ldots (X_{\beta_k} n_{\beta_k}) \right] \cdot P_J$$

where the $w = s_{\beta_1} \cdots s_{\beta_k}$ is a reduced expression for w in terms of simple reflections.

(ii) $P_J = U_J \cdot L_J$ *and* $|U_J| = q^{N - N_J}$.

(iii) If G is not of type G_2, then there is a sequence, $\beta_1, \ldots, \beta_m$ of simple roots such that $U_J \subseteq \prod_i X_{\beta_i}$ over any field of odd characteristic.

(iv) $L_J = Z_{L_J}.G_J$ *and* $|L_J/G_J| = (q-1)^{\operatorname{rank} G - |J|}$.

Proof: (i), (ii) and (iv) follow from the first two chapters in Carter's book [14]. (iii) follows frmo the Steinberg commutator relations in the form given in ([15],Theorem 12.1.1).

QED

Proof: [of Theorem 5.14.] We let N_k denote the number of positive roots of G_k. First we assume that $n \geq 2$ in the cases where G_n is of type A, B, or C, and $n \geq 4$ in the case that G has type D. From the lemma it is clear that $|G_n/P_n|$ is a polynomial of degree $N_n - N_{n-1}$ in q and that any coset of G_n/P_n has a coset representative of length no more than $N_n - N_{n-1}$ in the generators $(X_\alpha n_\alpha)$. Therefore

$$\begin{aligned} t_{G_n} &\leq (N_n - N_{n-1}) \left|\frac{G_n}{P_n}\right| \mathcal{M}(G_n, G_{n-2}) + t_{P_n} \\ &\leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}) + t_{P_n}. \end{aligned}$$

Now let U_n denote the maximal normal unipotent subgroup of P_n and let γ_n be such that U_n is contained in a product of no more than γ_n simple root subgroups (independent of q), then

$$\begin{aligned} t_{P_n} &\leq \gamma_n |U_n| \mathcal{M}(G_n, G_{n-2}) + t_{L_n} \\ &\leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}) + t_{L_n} \\ &\leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}) + t_{G_{n-1}} \end{aligned}$$

and therefore $t_{G_n} \leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}) + t_{G_{n-1}}$. A quick glance at Table 2 verifies that for all the series of groups, $N_n - N_{n-1} + \sigma(G_n, G_{n-2})$ is an increasing function of n, and hence that

$$t_{G_n} \leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}) + t_{G_1}$$

for the series A, B, C. For these three series, $G_1 = A_1(q)$ and hence t_{G_1} is bounded by $O(q^3)$ using a naive method of calculating a Fourier transform. For the D series of groups and $n \geq 4$ we have

$$t_{D_n} \leq O(q^{5n-6}) + t_{D_3}$$

and t_{D_3} may be bounded by $O(q^{10})$ using similar techniques. Hence we see that when $n \geq 2$, in the A, B, or C case, or $n \geq 4$ in the D case, we have

$$t_{G_n} \leq O(q^{N_n - N_{n-1} + \sigma(G_n, G_{n-2})}).$$

QED

G_n	N_n	$\sigma(G_n, G_{n-2})$	$N_n - N_{n-1} + \sigma(G_n, G_{n-2})$
$A_n(q)$	$\frac{1}{2}n(n+1)$	$n+1$	$2n+1$
$B_n(q)$	n^2	$3n-2$	$5n-3$
$C_n(q)$	n^2	$3n-2$	$5n-3$
D_n	n^2-n	$3n-4$	$5n-6$

Table 2: Combinatorial data for Chevalley groups.

Appendix A: Proofs of the technical lemmas

Now we shall indicate the proofs of some lemmas used in the explicit calculations of Section 5. These concern estimates of the number of the number of conjugacy classes of the general linear and unitary groups, the derivation of the factorization for coset representatives of U_n/U_{n-1}, and bounds for the multiplicity of restrictions between Chevalley groups.

A.1 Conjugacy classes

The generating functions for the number of conjugacy classes of $GL_n(q)$, the number of canonical forms of $n \times n$ matrices over $\mathbf{F}_q$ and the number of conjugacy classes of $U_n(q^2)$ are closely related. Define

$$F_\alpha(q,t) = \prod_{n=1}^{\infty} \frac{1+\alpha t^n}{1-qt^k}$$

and let $f_n(q;\alpha)$ be the coefficient of t^n in the expansion of $F_\alpha(q,t)$ considered as a power series in t. Then by results of [49] and [53] $f_n(q;-1)$ is the number of conjugacy classes of $GL_(q)$, $f_n(q;0)$ is the number of canonical forms of $n \times n$ matrices over $\mathbf{F}_q$, and $f_n(q;1)$ is the number of conjugacy classes of $U_n(q^2)$; The first result we need to bound $f_n(q;\alpha)$ is an asymptotic result due to Stong.

Lemma 5.17 (Stong)

$$f_n(q;-1) = q^n + \frac{1}{2}\left[\frac{1}{1-q^{\frac{1}{2}}} + (-1)^n \frac{1}{1+q^{\frac{1}{2}}}\right] q^{\frac{n}{2}} + O(q^{\frac{n}{3}})$$

as n tends to infinity, for fixed q.

Proof: $F_{-1}(q,t)$ is a meromorphic function of t in $|t| < 1$ with isolated poles at the k-th roots of q^{-1} for $k \geq 1$. The asymptotics come from considering the behavior at the poles q^{-1}, $q^{-\frac{1}{2}}$, and $q^{-\frac{1}{3}}$. See [49] for details.

QED

Corollary 5.18 *Define*

$$B_\alpha(t) = \prod_{k=1}^{\infty} \frac{1+\alpha t^k}{1-t^k}$$

Then $B_\alpha(t)$ is an analytic function of t in $|t| < 1$, provided that $\alpha \geq -1$. We have $F_\alpha(q,t) = B_\alpha(t)F_{-1}(q,t)$, and hence

$$f_n(q;\alpha) = B_\alpha(q^{-1})q^n + \frac{1}{2}\left[\frac{B_\alpha(q^{-\frac{1}{2}})}{1-q^{\frac{1}{2}}} + (-1)^n \frac{B_\alpha(-q^{-\frac{1}{2}})}{1+q^{\frac{1}{2}}}\right] q^{\frac{n}{2}} + O(q^{\frac{n}{3}})$$

as n tends to infinity, for fixed q and fixed $\alpha \geq -1$.

Proof: The residues of $F_\alpha(q,t)$ at $q^{-\frac{1}{k}}$ differ from those of $F_{-1}(q,t)$ by a factor of $B_\alpha(q^{-\frac{1}{k}})$.

QED

To obtain more useful bounds for $f_n(q;\alpha)$ we now consider some explicit formulae. Let $P(n,k)$ denote the number of partitions of n into k parts. For any nonnegative integer, m, define

$$h_m(q;\alpha) = \begin{cases} q^m + \alpha q^{m-1} & \text{when } m \geq 1 \text{ and} \\ 1 & \text{when } m = 0. \end{cases} \tag{48}$$

Next, for any partition, $\mu = 1^{m_1}2^{m_2}3^{m_3}\ldots$, define

$$h_\mu(q;\alpha) = \prod_{i \geq 1} h_{m_i}(q;\alpha).$$

Then is is clear that

$$f_n(q;\alpha) = \sum_{\mu \vdash n} h_\mu(q;\alpha).$$

If we now define

$$f_{n,k}(q;\alpha) = \sum_{\mu \in P(n,k)} h_\mu(q;\alpha)$$

then $f_n(q;\alpha) = \sum_{k=1}^{n} f_{n,k}(q;\alpha)$, and it is easy to see that $f_{n,k}$ satisfies the recurrence relation

$$f_{n,k} = q f_{n-1,k-1} + f_{n-k,k} + \alpha f_{n-k,k-1}. \tag{49}$$

This in turn shows that for $k > \frac{n}{2}$ and $p = n - k + 1$ we have

$$f_{n,k}(q;\alpha) = q^{n-2p+1} f_{2p-1,p}(q;\alpha)$$

so the high order coefficients of the polynomial $f_n(q;\alpha)$ are independent of n.

Theorem 5.19 *There is a constant K_α independent of n and q such that*

$$f_n(q;\alpha) \leq q^n + (\alpha+1)q^{n-1} + 2\cdot(\alpha+1)q^{n-2} + (\alpha+1)(\alpha+3)q^{n-3} + K_\alpha q^{n-4}$$

for any $n \geq 1$ and $q \geq 2$.

Proof: First we prove the case where $\alpha \geq 0$. In this case $f_n(q;\alpha)$ is a polynomial in q of degree n with positive coefficients. Using the recurrence relation (49) we see that for fixed n and α

$$f_n(q;\alpha) \leq q^n + 2.(\alpha+1)q^{n-1} + (\alpha+1)(\alpha+3)q^{n-3} + p_{n-4}(q;\alpha))$$

where $p_{n-4}(q;\alpha)$ is a polynomial of degree $n-4$ in q. By Lemma 5.17 we know that for each α and q there is a constant $K_\alpha(q)$ such that

$$f_n(q;\alpha) \leq B_\alpha(q^{-1})q^n + K_\alpha(q)q^{\frac{n}{2}}.$$

But $q^{4-n}\cdot p_{n-4}(q;\alpha)$ is a decreasing function of q, so

$$\begin{aligned} p_{n-4} &\leq q^{n-4}.\frac{B_\alpha(2^{-1})2^n + K_\alpha(2)2^{n/2}}{2^{n-4}} \\ &\leq K_\alpha \cdot q^{n-4} \end{aligned}$$

where $K_\alpha = 16\Big(B(2^{-1}) + K_\alpha(2)\Big)$.

Now consider the case when α is less than zero. In this case it's clear that $f_{n,k}(q;\alpha) \leq f_{n,k}(q;0)$. Hence

$$\begin{aligned} f_{n,k}(q;\alpha) &\leq \left(\sum_{k=n-3}^{n} f_{n,k}(q;\alpha)\right) + p_{n-4}(q;0) \\ &\leq \left(\sum_{k=n-3}^{n} f_{n,k}(q;\alpha)\right) + K_0.q^{n-4} \end{aligned}$$

QED

The same techniques can be used to find any finite number of the highest degree coefficients of $f_n(q;\alpha)$. In the case $\alpha = -1$ we conjecture that $f_n(q;-1)$ has coefficient 0 in all degrees strictly greater than $\frac{n}{2}$ but strictly less than n. This has been verified for $n \leq 400$ and is equivalent to the following statement.

Conjecture. In the ring of formal power series in q^{-1}, $\sum_{p=1}^{\infty} q^{-2p} f_{2p-1,p}(q;-1) = q^{-1}$.

A.2 Coset representatives for U_n/U_{n-1}

The group $U_{n-1}(q^2)$ acts transitively on the unitary unit n-sphere, consisting of all column vectors $(x_1,\ldots,x_n)^T$ with entries in $\mathbf{F}_{q^2}$ such that $\sum_{k=1}^{n} x_k^{1+q} = 1$. The stabilizer of the point $(0,\ldots,0,1)^T$ is U_{n-1}. To obtain a factorization of coset representatives according to Lemma 5.13, it suffices to show how to use the inverses of elements of the forms (A), (B), or (C), referred to in that lemma to rotate an arbitrary vector in the unitary sphere onto $(0,\ldots,0,1)^T$. We assume we are working in odd characteristic.

As in the statement of Lemma 5.13, we let N denote the group homomorphism $N(\alpha) = \alpha^{1+q}$. N is an epimorphism onto the group of nonzero elements of the subfield of q elements. We let R be a complete set of coset representatives of $\mathbf{F}_{q^2}^{\times}/\ker N$.

Now consider an arbitrary element, $x = (x_1,\ldots x_n)^T$ of the unitary unit sphere. If the vector (x_1,x_2) has nonzero unitary norm, then choose an element, $s \in R$ such that $N(s) = x_1^{1+q} + x_2^{1+q}$. Hence $(x_1/s, x_2/s)$ is a unit vector and so by the transitivity of U_2 on the unitary 2-sphere, it is clear that we can choose $y \in \mathbf{F}_{q^2}$ and $r \in R$ such that $u_2(r,y)^{-1}$ maps (x_1,x_2) onto a multiple of $(0,1)$.

In the case where $x_1^{1+q} + x_2^{1+q} = 0$ it is possible that either $x_1^{1+q} + x_3^{1+q}$ or $x_2^{1+q} + x_3^{1+q}$ is nonzero. In the first case multiplying x by t_3 brings the vector into a form where the vector of the first two components has a nonzero

norm, and in the second case multiplication by $t_3 \cdot t_2$ achieves this. If neither of these three cases hold, then the vector (x_1, x_2, x_3) must be zero (this requires that the characteristic is not 2). Therefore it is always possible to map x onto a vector with zero first component, using the inverse of a matrix of form (A), (B) or (C), provided n is greater than 2.

Now we may apply the same method to map x onto a vector with the first two entries zero, the first three, and so on. Finally we ontain a vector with only the last two entries nonzero. Clearly we can use the inverse of an element of form (A) to map this vector onto a vector with only the last entry nonzero. As the vector so obtained is a unit vector, it must have the form $(0, \ldots 0, \varepsilon)$ for some ε with $\varepsilon^{1+q} = 1$.

A.3 Multiplicities of restrictions in Chevalley groups

Now we shall prove Lemma 5.15 on the multiplicities of restrictions in Chevalley groups. The proof we use was suggested by Tom Hagedorn and follows the line of argument of his thesis [29]. We shall limit ourselves to a brief sketch of this argument. As usual we shall always assume the characteristic is odd.

First we note that if $G \geq H \geq K$, then $\sigma(G,K) = \sigma(G,H) + \sigma(H,K) + \operatorname{rank} H$ where σ is as in Lemma 5.15. But

$$\mathcal{M}(G,K) \leq \mathcal{M}(G,H)\mathcal{M}(H,K)\left|\hat{H}\right|$$

where $\hat{H}$ denotes the set of equivalence classes of irreducible representations of H. The bounds we need follow from the same result for the restriction of irreducible representations from $A_n(q)$ to $A_{n-1}(q)$, $B_n(q)$ to $B_{n-1}(q)$, $C_n(q)$ to $C_{n-1}(q)$ or $D_n(q)$ to $D_{n-1}(q)$.

The problem of bounding multiplicities can also be reduced, as follows, to bounding the pairing of a Deligne-Lusztig character of G, restricted to H, with a Deligne-Lusztig character of H: let us say that a linear combination is bounded if the number of terms may be bounded independently of q and the coefficients may also be bounded independently of q. Then for any irreducible character, χ, of G (or of H) there is a bounded linear combination of Deligne-Lusztig characters which is the character of a representation containing χ.

We shall now let G and H denote connected reductive algebraic groups of classical type over an algebraically closed field of odd characteristic, and we let F be a Frobenious map. Suppose T, T' are F-stable maximal tori of G and H respectively and θ, θ' are irreducible characters of T^F and $(T')^F$ respectively. As usual, $R_{T,\theta}$ denotes the Deligne-Lusztig character associated to T and θ (cf. [21] for the complete definitions). Then the pairing of $R_{T,\theta}\downarrow H^F$ with $R_{T',\theta'}$ has the form

$$\left\langle R_{T,\theta}\downarrow H^F, R_{T',\theta'}\right\rangle = \sum_s \sum_u \sum_{w,w'} a(s,w,w',u) \frac{Q_{T_w}^{C_G^0(s)}(u) Q_{T'_{w'}}^{C_H^0(s)}(u)}{\left|C_{C_H^0(s)^F}(u)\right|} \tag{50}$$

where s varies over H^F conjugacy classes of elements in $(T')^F$, u varies over unipotent conjugacy classes of the connected centralizer $C_H^0(s)^F$; T_w, $T'_{w'}$ are F-stable maximal tori in $C_G^0(s)$, $C_H^0(s)$ respectively, and the Q's are Green polynomials. For a given s, u, T_w, $T'_{w'}$, the term

$$\frac{Q_{T_w}^{C_G^0(s)}(u) Q_{T'_{w'}}^{C_H^0(s)}(u)}{\left|C_{C_H^0(s)^F}(u)\right|} \tag{51}$$

is a function of q that can be bounded by $O(q^{\sigma(G,H)-\alpha(s,u)})$ where

$$\alpha(s,u) = \frac{1}{2}\left[\dim G - \dim H - \left(\dim C_{C_G^0(s)}(u) - \dim C_{C_H^0(s)}(u)\right)\right]$$

and for any given s the inner summations in (50) are a bounded linear combination of terms of the form (51); $a(s,w,w',u)$ may also be bounded independently of s. Note that there are only finitely many (a number bounded independently of q) different forms that the term (51) can have given G and H.

Therefore, to obtain a bound for the pairing (50) it suffices to bound $\alpha(s,u)$ and then determine how many s this bound applies to.

To bound $\dim C_{C_G^0(s)}(u) - \dim C_{C_H^0(s)}(u)$ we can reduce to the case where G and H come from one of the series of classical groups: $SL(n)$, $SO(2n+1)$, $Sp(2n)$, or $SO(2n)$. In this case the connected centralizer, $C_H^0(s)$ is determined up to isomorphism by the characteristic polynomial of s considered as an element of H; up to isogeny it is simply a product of groups corresponding to different eigenvalues of s. The characteristic polynomial of s considered as an element of G may be obtained from its characteristic polynomial as an element of H by multiplying by either 1 or 2 factors of $(t-1)$; 1 factor in the case of restricting from A_n to A_{n-1} and 2 in the cases of restricting from B_n to B_{n-1}, C_n to C_{n-1} and D_n to D_{n-1}. Hence the centralizer, $C_G^0(s)$, only differs from $C_H^0(s)$ in the factor that corresponds to the eigenvalue 1 of s. Having obtained the form of the centralizers, the formulas in [14] p. 398 (see also the article of Springer and Steinberg in [12]), may be used to compute the dimensions of centralizers of unipotent elements, in order to bound $\dim C_{C_G^0(s)}(u) - \dim C_{C_H^0(s)}(u)$ in terms of the multiplicity, m, of 1 as an eigenvalue of s. We call this bound β_m.

Hence we can bound $\alpha(s,u)$ from below by a function, α_m, of m and the number of s in $(T')^F$ with a given m can be bounded by $O(q^{\gamma_m})$ for some easily determined function γ_m. To prove the theorem we need only verify that $\alpha_m - \gamma_m \geq 0$ for all possible values of m. We present this verification in the form of a table.

Restriction	β_m	α_m	γ_m	$\alpha_m - \gamma_m$	Maximum m
$A_n \downarrow A_{n-1}$	$2m+1$	$n-m$	$\max\{n-m, 0\}$	1 or 0	n
$B_n \downarrow B_{n-1}$	$2m+1$	$2n-m$	$n-1-\frac{m-1}{2}$	$n-\frac{m}{2}-\frac{1}{2}$	$2n-1$
$C_n \downarrow C_{n-1}$	$2m+3$	$2n-m-2$	$n-1-\frac{m}{2}$	$n-\frac{m}{2}-1$	$2n-2$
$D_n \downarrow D_{n-1}$	$2m+1$	$2n-m-2$	$n-1-\frac{m}{2}$	$n-\frac{m}{2}-1$	$2n-2$

Table 3: Verification of Lemma 5.15.

For the proof to make sense for $D_2 \downarrow D_1$ we have to replace D_1 by a 1 dimensional torus. We have now proved the lemma.

6 Further improvements and directions

Theorem 4.8 and Corollary 4.9 are particularly easy to use but are by no means the best results possible. We no briefly describe some of the improvements we have obtained, which will appear in the second part of this paper [39].

6.1 Variations on the main results

In many cases, further savings can be realized in the the Fourier transform is treated as a collection of scalar equations rather than as a matrix equation. The separation of variables idea still applies to the scalar setting, but now a recursive sum of products of numbers, as opposed to matrices, is obtained. These products may be computed in any order. Consequently, the scalar separation of variables algorithm possesses a flexibility which is not present in the matrix separation of variables algorithm: the ability to choose the order in which the factors are summed over. Roughly speaking, this flexibility allows us to sum over factors with a low value of $\mathcal{M}$ first, successively building the complete computation. In practice the first summations we perform occur the most times in the separation of variables algorithm (in the matrix case, this amounts to saying that the sets X_k get smaller as k increases), so by ensuring these sums are done quicker, we make the whole algorithm more efficient.

The sums that occur in the scalar separation of variables algorithm are generalizations of the sum (22), and the factors that appear are indexed by collections of representations which satisfy relations generalizing the relations represented by Diagram 1. The diagrammatic methods used in the proof of Theorem 4.5 generalize to this situation, so complexity bounds for the new algorithms may be obtained explicitly. A useful combinatorial tool here is to treat the indices as injections from the diagrams describing the relations into the Bratteli diagram for the subgroup chain. The explicit expressions for the complexity of the algorithm has a form similar to, but generalizing, the expressions in Theorems 4.5, 4.6, and 4.8.

We use the techniques just described to refine the results we have already obtained in section 5. For example, we get a better bound for the complexity of the Fourier transform on $GL_n(q)$ using the same bases as in Section 5.

Theorem 6.1 *For any n, there is a positive constant, K_n, such that*

$$T_{GL_n(q)} \leq K_n q^n |GL_n(q)|$$

for any q greater than or equal to 2.

Similar improvements hold for the unitary groups and Chevalley groups. We also prove a general theorem, bounding t_G in terms of the complexities of two subgroups and the number of double cosets. This result works particularly nicely when the subgroups are abelian, and in that case it yields new results for $SL_2(q)$ and for the symmetric groups.

6.2 Homogeneous spaces

For many statistical applications, data on homogeneous spaces is of primary interest, rather than data on the full group. In brief, a homogeneous space for a finite group is simply a set on which the group acts transitively as permutations. A common example is the action of the finite affine group on point-line pairs and more generally, the action of an automorphism group of a design on its block-point pairs. In this case generalizations of the "usual" analysis of variance for data on such sets require the computation of projections of the data vector onto group-invariant subspaces.

The scalar separation of variables algorithm generalizes easily to the context of homogeneous spaces. This is in contrast to the techniques of section 4, which do not improve on a naive algorithm (such as a directly computed matrix-vector product). The idea in the improved algorithms is to write the associated spherical functions of the homogeneous space as a sum of products, with a small number of terms in the sum. The separation of variables algorithm then amounts to calculating the inner product of a function and an associated spherical function by summing over one factor in the product at a time. This provides speed-ups of the most efficient algorithms currently known (cf. [24] and references therein). This material will also appear in [39].

References

[1] E. F. Assmus Jr. and J. D. Key, *Designs and Their Codes*, Cambridge Univ. Press, Cambridge, 1992.

[2] L. Babai, E Luks, and A. Seress, Fast management of permutation groups, *Proc. 28th IEEE FOCS* (1988), pp. 272-282.

[3] L. Auslander and R. Tolmieri, Is computing with the fast Fourier transform pure or applied mathematics?, *Bulletin of the A. M. S.*, **1** (New Series), (1979), 847-897.

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

[5] L. Babai, K. Freidl and M. Stricker, Decomposition of *-closed algebras in polynomial time, *Proc. of 18th ACM ISSAC* (1993), pp. 86-94.

[6] R. Beals. and L. Babai, Las Vegas algorithms for matrix groups, *Proc. 34th IEEE FOCS* (1993), pp. 427-436.

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

[8] U. Baum and M. Clausen, Some lower and upper complexity bounds for generalized Fourier transforms and their inverses, *SIAM J. Comput.* **20**(3) (1991), 451-459.

[9] U. Baum, M. Clausen and B. Tietz, Improved upper complexity bounds for the discrete Fourier transform, *AAECC* **2** (1991) 35-43.

[10] T. Beth, On the computational complexity of the general discrete Fourier transform, *Theoretical Computer Science* **51** (1987), 331-339.

[11] T. Beth, *Verfahren der schnellen Fourier-Transformation*, Teubener Studienbücher, Stuttgart, 1984.

[12] A. Borel, R. Carter, C. Curtis, N. Iwahori, T. Springer and R. Steinberg, *Seminar on Algebraic Groups and Related Finite Groups*, Lecture Notes in Mathematics Volume 131, Springer-Verlag, NY, 1970.

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

[14] R. Carter, *Finite Groups of Lie Type: Conjugacy Classes and Characters.* Wiley-Interscience, NY, 1985.

[15] R. Carter, *Simple Groups of Lie Type.* Wiley-Interscience, NY, 1989.

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

[17] M. Clausen and U. Baum, Fast Fourier transforms for symmetric groups, theory and implementation, *Math. Comp.* **61**(204) (1993), 833-847.

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

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

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

[21] P. Deligne and G. Lusztig, Representations of reductive groups over finite fields, *Ann. Math.* **103** (1976), 103-161.

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

[23] P. Diaconis and D. Rockmore, Efficient computation of the Fourier transform on finite groups, *J. of the A.M.S.* **3**(2) (1990), 297-332.

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

[25] J. R. Driscoll and D. Healy. Computing Fourier transforms and convolutions on the 2-sphere. (Extended abstract) *Proc.* 34^{th} *IEEE FOCS*, (1989) pp. 344-349; (*Adv. in Appl. Math.*, **15** (1994), 202-250.

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

[27] A. Fässler and E. Stiefel, *Group Theoretical Methods and Their Applications.* Birkhäuser, Boston. MA, 1992.

[28] L. Finkelstein and W. Kantor, *Groups and Computation*, DIMACS Series in Disc. Math., Vol. 11, AMS, 1993.

[29] T. Hagedorn. *Multiplicities in Restricted Representations.* Ph.D. Thesis, Department of Mathematics, Harvard University, 1994.

[30] F. Harary, *Graph Theory*, Addison-Wesley, Reading, MA, 1972.

[31] D. Healy, D. Maslen, S. Moore, and D. Rockmore, Applications of a fast convolution algorithm on the 2-sphere, *Technical Report, Department of Mathematics and Computer Science, Dartmouth College*, 1994.

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

[33] G. D. James, *Representations of General Linear Groups*, LMS Vol. 94, Cambridge Univ. Press, Cambridge, 1984.

[34] F. Goodman, P. de la Harpe and V. Jones, *Coxeter Graphs and Towers of Algebras*, Spinger-Verlag, NY, 1989.

[35] H. Hiller, *The Geometry of Coxeter Groups*, Research Notes in Mathematics Vol 54, Pitman, 1982.

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

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

[38] J. Lafferty and D. Rockmore, Fast Fourier analysis for SL_2 over a finite field and related numerical experiments, *J. Exp. Math.* **1** (1992) 115-139.

[39] D. Maslen and D. Rockmore, Separation of variables and the efficient computation of Fourier transforms on finite groups, II. In preparation.

[40] D. Maslen and D. Rockmore, Adapted diameters and the efficient computation of Fourier transforms of finite groups, *Proceedings of the 1995 ACM-SIAM Symposium on Discrete Algorithms.* To appear.

[41] S. Linton, G. Michler, and J. Olsson, Fourier transforms with respect to monomial representations, *Math. Ann.* **297** (1993), 253-268.

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

[43] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing.* Prentice Hall, NJ, 1989.

[44] D. Rockmore, Efficient computation of Fourier inversion for finite groups, *J. of the A.C.M.* **41**(1) (1994), 31-66.

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

[46] D. Rockmore, Fast Fourier transforms for wreath products, *Technical Report PMA-TR94-176, Dept. of Mathematics, Dartmouth College*, 1994.

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

[48] C. C. Sims, Computational methods in the study of permutation groups, *in: Computational Problems in Abstract Algebra*, J. Leech, ed., Pergamon Press 1970, pp. 169-183.

[49] R. Stong, Some asymptotic results on finite vector spaces, *Adv. in Appl. Math.* 9 (1988), 167-199.

[50] E. Thoma, Die Einschrankung der Charaktere von $GL(n,q)$ auf $GL(n-1,q)$, *Math. Zeit.*, **119** (1971) 321-338.

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

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

[53] G. E. Wall, On the conjugacy classes in the unitary, symplectic and orthogonal groups, *J. Austral. Math. Soc.* (1963), 1-62.

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

[55] A. Zelevinsky. *Representations of Finite Classical Groups* Lecture Notes in Mathematics Vol 869, Springer-Verlag, NY, 1981.