DOUBLE COSET DECOMPOSITIONS AND COMPUTATIONAL HARMONIC ANALYSIS ON GROUPS

DAVID K. MASLEN AND DANIEL N. ROCKMORE

Abstract. In this paper we introduce new techniques for the efficient computation of a Fourier transform on a finite group. We use the decomposition of a group into double cosets and a graph theoretic indexing scheme to derive algorithms that generalize the Cooley-Tukey FFT to arbitrary finite groups. We apply our general results to special linear groups and low rank symmetric groups, and obtain new efficient algorithms for harmonic analysis on these classes of groups, as well as the two-sphere.

1. Introduction

The Fourier transform on a finite group is a generalization of the discrete Fourier transform of a finite data sequence. The analogues of the complex exponentials appearing in the discrete Fourier transform are the matrix elements of a complete set of irreducible complex matrix representations. The matrix elements form a basis for the space of complex valued functions on the group, and the Fourier transform describes the expansion of functions in this basis. From this point of view, the usual discrete Fourier transform is simply the Fourier transform on a cyclic group.

In the current paper we study algorithms for the efficient computation of Fourier transforms on arbitrary finite groups. On an abelian group the Fourier transform may be computed using the fast Fourier transform algorithms of Cooley and Tukey [17], and their many variants (see e.g. [26]). Our results may be considered a generalization of the Cooley-Tukey algorithm to nonabelian finite groups.

The computation of Fourier transforms on nonabelian finite groups was first studied by Willsky [53] in the context of filter design. Since then, various authors have treated solvable groups [8, 12, 49], supersolvable groups [4, 5], symmetric groups [13, 14, 38] and their wreath products [48], special linear groups [36], and also the general case of arbitrary finite groups [14, 22, 41]. The book [15] and the survey article [40] present overviews of this subject. For a discussion of applications see [18, 19, 20] and the survey article [47].

The algorithms we develop in this paper are examples of the "separation of variables" approach developed in [41, 42] and [40]. In addition to our general theorems on the computation of Fourier transforms, we have focused on the example of the special linear groups. Fourier transforms on this group are of interest for certain applications to error correcting codes [35]. Our main concrete result for special linear groups is as follows.

1991 *Mathematics Subject Classification.* 20C15; Secondary 65T10.

Key words and phrases. generalized Fourier transform, Bratteli diagram, Gel'fand-Tsetlin basis, Cooley-Tukey algorithm.

Rockmore is supported in part by ARPA as administered by the AFOSR under contract DOD F4960-93-1-0567 as well as NSF PFF Award DMS-9553134.

Theorem 1.1. *Let q be an odd prime power. Working in a system of Gel'fand-Tsetlin bases, for the irreducible representations of $SL_2(F_q)$, the Fourier transform of a complex valued function can be computed using no more than*

$$(\frac{1}{2}q + 3 + 16\log q) \cdot (q^3 - q)$$

scalar operations.

We prove several variations of this result, and also prove similar theorems for low order symmetric groups. In Section 4.3 we have included an application to the computation of expansions in spherical harmonics on grids with icosahedral symmetry.

A new feature of our approach to the computation of Fourier transforms is the indexing scheme that appears in the algorithms. According to this, the group, or time domain of our transform, is indexed using factorizations of group elements. The space of matrix elements is our frequency domain, and is indexed by paths in a graded directed graph called a Bratteli diagram. This provides an elegant generalization to other finite groups of the indexing scheme of the usual Cooley-Tukey algorithm.

The paper is organized as follows. In Section 2 we interpret the discrete Fourier transform and the Cooley-Tukey algorithm in terms of the representation theory of finite groups. In Section 3 we give a simplified treatment of the special linear group, in preparation for more general results. Section 4 is the core of the paper. We prove our most general result, and then specialize this to the case of symmetric groups. We briefly discuss Fourier transforms on the rotation group. In Section 5 we treat the special linear group in more detail, and prove the Theorem stated in this introduction. We conclude with Section 6 and two appendices which discuss the construction of the representations of $SL_2(F_q)$, and Gel'fand-Tsetlin bases.

2. Group theory and the Cooley-Tukey algorithm

The usual discrete Fourier transform of a finite data sequence is a special case of the Fourier transform on a finite group, which is defined in terms of the representations of the group. We do not intend to provide a self-contained introduction to the representation theory of finite groups, even though we do use a number of elementary facts from representation theory (see e.g., [50]). It is, however, useful to recall the following facts and definitions.

1. A (complex) matrix representation ρ of a finite group G is a matrix valued function on G that satisfies the group homomorphism property
$$\rho(gh) = \rho(g)\rho(h)$$
for all $g, h \in G$.
2. The size of the matrices $\rho(g)$ is called the dimension of ρ and is denoted d_ρ, and the space $\mathbf{C}^{d_\rho}$ on which the matrices act is called the representation space.
3. Each choice of basis for the representation space gives rise to an equivalent matrix representation, which is obtained by expressing the matrices $\rho(g)$ in the new basis.
4. The matrix entries $[\rho(g)]_{ij}$ of a matrix representation are complex valued functions on G, called matrix elements.

5. A set of matrix representations of G is called a complete set of inequivalent irreducible representations if and only if the collection of matrix elements of the representations, relative to an arbitrary choice of bases for the representations spaces, forms a basis for the space of complex functions on G.

With this terminology, we can now state the definition of a Fourier transform.

Definition 2.1 (Fourier transform)**.** Let G be a finite group, and let f be a complex valued function on G. The Fourier transform of f is the collection of matrix sums

$$\hat{f}(\rho) = \sum_{g \in G} f(g)\rho(g) \tag{2.1}$$

where ρ runs over a complete set of inequivalent irreducible matrix representations of G.

Definition 2.1 depends on the choice of the set of matrix representations, but the Fourier transform matrices $\hat{f}(\rho)$ for the different choices are always related by a change of basis. We will always make clear the choice of basis.

The number of operations required to compute the Fourier transform also depends on the choice of bases for the representations. In whichever basis we are using, we will always denote the number of such operations by C_G, where we count the maximum of the number of additions and multiplications. In other words, C_G is the minimum number of operations needed to compute the Fourier transform in the given set of bases via a straight-line program, for an arbitrary complex valued function f.

Relative to any set of bases we have the upper and lower bounds

$$|G| - 1 \le C_G \le |G|^2,$$

where the upper bound is the number of operations required by a direct approach to the computation using Equation 2.1.

2.1. **The Cooley-Tukey algorithm.** Using Definition 2.1 we can interpret both the discrete Fourier transform and the Cooley-Tukey algorithm using representation theory.

Let N be a non-negative integer. Then the discrete Fourier transform of a finite data sequence $f_0, f_1, \ldots, f_{N-1}$ of length N is defined to be the sequence of complex numbers

$$F_k = \sum_{j=0}^{N-1} f_j e^{2\pi i jk/N}, \quad k = 0, 1, 2, \ldots, N-1. \tag{2.2}$$

To interpret this transform using group theory, we view the sequence f_j as defining a complex valued function f on the cyclic group $\mathbf{Z}/N\mathbf{Z}$. We shall write the elements of $\mathbf{Z}/N\mathbf{Z}$ multiplicatively, with generator x. In this notation, the elements of $\mathbf{Z}/N\mathbf{Z}$ are $1, x, x^2, \ldots, x^{N-1}$, and the function f associated to the sequence f_j is given by

$$f(x^j) = f_j, \quad j = 0, 1, \ldots, N-1.$$

The irreducible representations of $\mathbf{Z}/N\mathbf{Z}$ are one-dimensional and given by

$$\zeta_k(x^j) = e^{2\pi i jk/N}.$$

Substituting this into the group theoretic definition of a Fourier transform (2.1), shows that the group theoretic Fourier transform of the function f on $\mathbf{Z}/N\mathbf{Z}$ is given by

$$\hat{f}(\zeta_k) = \sum_{j=0}^{n-1} f(x^j) e^{2\pi i j k/N}. \tag{2.3}$$

This agrees with the usual definition (2.2) provided we take $F_k = \hat{f}(\zeta_k)$. The group element x^j corresponds to the time variable and the representation ζ_k corresponds to the frequency variable in the usual discrete Fourier transform.

The group theoretic approach to the Cooley-Tukey algorithm depends on the notion of restriction of representations. Recall that if H is a subgroup of the finite group G and ρ is a representation of G, then the restriction of ρ to H, denoted $\rho|H$, is the representation of H defined as the matrix valued function on H agreeing with ρ on H.

In the case of $\mathbf{Z}/N\mathbf{Z}$, the restriction of representations is easy to describe. Suppose $N = pq$ and let $H = p\mathbf{Z}/N\mathbf{Z} = \langle x^p \rangle$ (note that $H \cong \mathbf{Z}/q\mathbf{Z}$). Then

$$\zeta_k | H = \zeta_{k \bmod q}.$$

We may represent this dependence diagrammatically using a directed graded graph with three levels. At level 0 there is a single vertex 1, called the root vertex. The vertices at level one are labeled by the irreducible representations of H, and the vertices at level two are labeled by the irreducible representations of $\mathbf{Z}/N\mathbf{Z}$. We draw edges from the root vertex to all vertices at level one, and we draw an edge from a vertex at level one to a vertex at level two if and only if the representation at the tip restricts to the representation at the tail. The diagram we obtain is called the Bratteli diagram for the group-subgroup chain $\mathbf{Z}/N\mathbf{Z} > H > 1$. Figure 1 shows the situation for the chain $\mathbf{Z}/9\mathbf{Z} > 3\mathbf{Z}/9\mathbf{Z} > 1$.

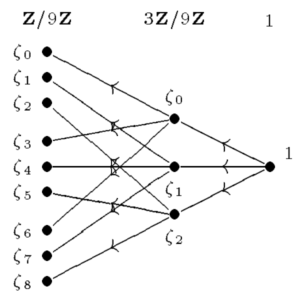

FIGURE 1. The Bratteli diagram for $\mathbf{Z}/9\mathbf{Z} > 3\mathbf{Z}/9\mathbf{Z} > 1$.

The Bratteli diagram provides us with a useful way of indexing the representations of $\mathbf{Z}/N\mathbf{Z}$ and H. By construction, there is a unique path from the root vertex to any other vertex in the diagram. Therefore, representations of H may be indexed by paths of length 1 starting at 1, and representations of $\mathbf{Z}/N\mathbf{Z}$ may be indexed by paths of length 2.

We are now ready to derive a simple case of the Cooley-Tukey algorithm.

Theorem 2.2 (Cooley-Tukey, Gauss). *The Fourier transform of a complex valued function on $\mathbf{Z}/pq\mathbf{Z}$ may be computed using $(p+q)qp$ scalar operations.*

Proof. To compute the Fourier transform efficiently, we make substitutions which allow us to write the transform in a two-dimensional form. Let $Y = \{x^\alpha : 0 \leq \alpha \leq q-1\}$. Then Y is a set of coset representatives for $\mathbf{Z}/N\mathbf{Z}$ relative to H, and any element x^j of $\mathbf{Z}/N\mathbf{Z}$ may be written uniquely as a product

$$x^j = y \cdot h \tag{2.4}$$

where $y \in Y$, and $h \in H$. Thus we may use the substitution (2.4) to replace the time variable x^j by the pair (y, h) in the sum (2.3) for the Fourier transform.

We also make a substitution in the frequency domain, replacing the variable k, which indexes the representation, with a path r in the Bratteli diagram. We denote the first leg of this path by r_1, so that ζ_{r_1} denotes the restriction of the representation ζ_r to H.

The two substitutions imply the following formula for the representations

$$\zeta_r(y \cdot h) = \zeta_r(y)\zeta_{r_1}(h). \tag{2.5}$$

Substituting (2.5) into the equation for the Fourier transform gives us the desired two-dimensional equation for the Fourier transform

$$\hat{f}(\zeta_r) = \sum_{y \in Y} \zeta_r(y) \sum_{h \in H} f(y \cdot h)\zeta_{r_1}(h). \tag{2.6}$$

We now compute the Fourier transform in two stages.

Stage 1. Compute the sums

$$f^1(r_1, y) = \sum_{h \in H} f(y \cdot h)\zeta_{r_1}(h)$$

for each path r_1 of length one starting at the root in the Bratteli diagram, and for all $y \in Y$.

Stage 2. Compute the sums

$$\hat{f}(\zeta_r) = \sum_{y \in Y} f^1(r_1, y)\zeta_r(y)$$

for each path r of length 2 in the Bratteli diagram, where r_1 is the first leg of r.

Counting the number of operations required by this algorithm amounts to counting the indices that occur inside the sums at each stage. Thus, Stage 1 requires $|H|^2 |Y| = q^2 p$ operations while Stage 2 requires $|\mathbf{Z}/N\mathbf{Z}| |Y|$ operations. We have used the fact that any abelian group G has exactly $|G|$ representations. ☐

The group theoretic approach to the Cooley-Tukey algorithm is originally due to Beth [8]. We have presented this method because it generalizes to other finite groups, including nonabelian groups. However this approach does have some advantages even when considering the usual discrete Fourier transform. In particular we are free to choose any set of coset representatives for $(\mathbf{Z}/N\mathbf{Z})/H$ in place of the set Y. This leads to a uniform derivation of a large number of different fast Fourier transform algorithms.

Remark 2.3. The proof of Theorem 2.2 is easily extended to give a full proof of the Cooley-Tukey algorithm for cyclic groups of any order. If N has l prime factors, then we may find a chain of l subgroups $\mathbf{Z}/N\mathbf{Z} = H_l > H_{l-1} > \cdots > H_0 = 1$ and form a Bratteli diagram which describes the restrictions between successive subgroups in the chain (see e.g. Figure 2). We iterate the above idea and obtain a factorization of the group elements x^j in terms of coset representatives of H_i relative to H_{i-1}, for $i = 1, \ldots, l$. This leads to an l-stage algorithm for computing the Fourier transform, which is equivalent to the usual Cooley-Tukey algorithm (see [41]). This approach also works to give an efficient algorithm for any finite abelian group.

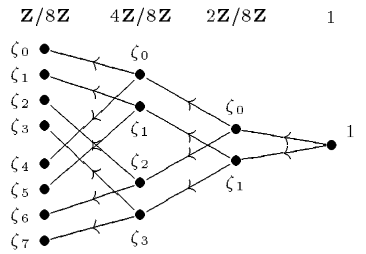

FIGURE 2. Bratteli diagram for $\mathbf{Z}/8\mathbf{Z} > 4\mathbf{Z}/8\mathbf{Z} > 2\mathbf{Z}/8\mathbf{Z} > 1$.

To summarize, we index the time domain (the group) using the factorization of group elements. We index the frequency domain by paths (concatenations of edges) in a Bratteli diagram that describes the restriction of representations. We index the computation at each stage in the Fourier transform algorithm using one less factor in the time domain, but paths of increasing length in the frequency domain, as compared with the previous stage.

Remark 2.4. The paper [1] gives a different representation theoretic interpretation of the Cooley-Tukey algorithm, relating it to the Weil-Brezin map.

2.2. **Algorithms for abelian groups.** The Cooley-Tukey algorithm by itself does not give an efficient algorithm for computing Fourier transforms on all abelian groups, or even on all cyclic groups. For cyclic groups of low order the methods due to Winograd [54] are more efficient, and for cyclic groups of prime order either the method of Rader [46] or the chirp-z transform due to Rabiner et. al. [45, 9] should be used. For a discussion of the average running times of methods based on these algorithms, see [21].

Baum, Clausen, and Tietz [7] have done a careful analysis of the number of operations required to compute a Fourier transform on a finite abelian group, and prove the following bound, which we will use repeatedly.

Theorem 2.5 (Baum, Clausen, Tietz). *The Fourier transform of a complex function on a finite abelian group G may be computed using no more than $8\,|G|\log|G|$ scalar operations. I.e.,*

$$C_G \le 8\,|G|\log|G|\,.$$

3. The special linear group

3.1. **A general approach.** We now study Fourier analysis on the special linear group. Let q be a prime power, and let F_q denote the finite field with q elements. Then the special linear group over F_q is the group of all 2×2 matrices with entries in F_q and determinant 1. I.e.,

$$SL_2(F_q) = \left\{ \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \end{pmatrix} \Big| x_{11}x_{22} - x_{12}x_{21} = 1, x_{ij} \in F_q \right\}.$$

This is a nonabelian group with $q^3 - q$ elements.

As in the derivation of the Cooley-Tukey algorithm, we start by factorizing group elements in terms of elements coming from subgroups. Define subgroups $N, T < SL_2(F_q)$ by

$$N = \left\{ \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} \middle| x \in F_q \right\}, \quad T = \left\{ \begin{pmatrix} \alpha & 0 \\ 0 & \alpha^{-1} \end{pmatrix} \middle| \alpha \in F_q^\times \right\}.$$

These are abelian subgroups, isomorphic to the additive group of F_q and the multiplicative group of nonzero elements $F_q^\times$, respectively. Thus, T is cyclic of order $q-1$. The matrix w defined by

$$w = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

also plays an important rôle in the construction of Fourier transform algorithms.

The next lemma is elementary.

Lemma 3.1 (Bruhat decomposition). *Let g be any element of $SL_2(F_q)$. Then g has a unique factorization into exactly one of the following forms.*

$$g = \begin{cases} t \cdot n & t \in T, n \in N, \quad or \\ n_2 \cdot w \cdot t \cdot n_1 & n_1, n_2 \in N, t \in T \end{cases}$$

The factorization of Lemma 3.1 is unique. We shall sometimes find it more useful to use the redundant factorization $SL_2(F_q) = N \cdot A \cdot N$, where

$$A = T \cup wT. \tag{3.1}$$

Fourier analysis on the special linear group is more complicated than Fourier analysis on cyclic groups, mainly because the group is nonabelian. A list of the irreducible representations is given in Table 2 in Appendix A (see also Figures 3 and 4 in this section). There is always at least one irreducible representation of dimension greater than one. The Fourier transform of a function will therefore depend on the bases chosen for the irreducible representations. We shall always use Gel'fand-Tsetlin bases, which are defined for $SL_2(F_q)$ as follows.

Definition 3.2. A system of Gel'fand-Tsetlin bases for the irreducible representations of $SL_2(F_q)$, relative to N, is a collection of bases for these representations, one for each representation, such that the line spanned by each basis vector is N-invariant.[1]

Remark 3.3. The condition that the line spanned by v is N-invariant is equivalent to the existence of a one-dimensional representation $\chi_v : N \longrightarrow \mathbf{C}$ such that $\rho(n)v = \chi_v(n)$ for all $n \in N$. Thus, the matrices $\rho(n)$ for all $n \in N$ are diagonal with diagonal entries $\chi_v(n)$ for each of the Gel'fand-Tsetlin basis vectors v.

The vectors in a system of Gel'fand-Tsetlin bases may be indexed by paths in a directed graded graph in a way similar to the indexing used in the Cooley-Tukey algorithm. To define the graph, draw a vertex for each of the irreducible representations of $SL_2(F_q)$, N and 1; the vertex labeled by the trivial representation of 1 is at level 0, those for the representations of N are at level 1 and finally those for the representations of $SL_2(F_q)$ are at level 2. For each Gel'fand-Tsetlin basis vector v, draw a directed edge starting at the representation χ_v of N (in the notation of

[1] This is a special case of Definition 4.4 of the next section.

the remark) and ending at the representation of $SL_2(F_q)$ to which v belongs. This describes the restriction of representations from $SL_2(F_q)$ to N. Finally, draw edges from the trivial representation of the trivial group to each representation of N. The resulting graph is the Bratteli diagram of the chain of groups, $SL_2(F_q) > N > 1$.

The overall structure of the Bratteli diagram depends on whether q is even or odd. The case of q even is simpler, so we present this first in Figure 3. In this, and subsequent figures, we have used a vertical line through a single vertex to indicate that a vertex ranges over a set of irreducible representations of a particular type. We have also indicated here that there are two families of irreducible representations of $SL_2(F_q)$, the principal series and the discrete series. Thus the vertex labeled 1_N in Figure 3 should be interpreted as having 2 edges connected to each of $\frac{q-2}{2}$ principal series representations of dimension $q+1$ and one edge to each of the remaining two principal series representations. Similarly, there is exactly one edge between any of the individual $\frac{q}{2}$ discrete series representation of $SL_2(F_q)$ and any of the $q-1$ nontrivial irreducible representations of N.

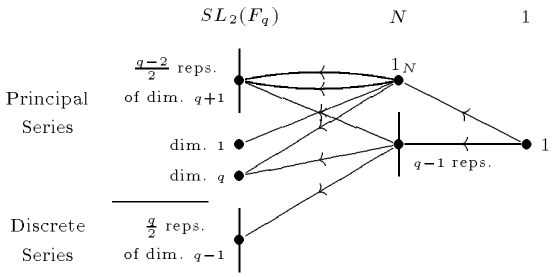

FIGURE 3. The Bratteli diagram for $SL_2(F_q) > N > 1$, q even. A vertex on a vertical line ranges over a set of representations.

The Bratteli diagram for $SL_2(F_q) > N > 1$ for odd q is shown in Figure 4. The particular examples of $q = 2$ and $q = 3$ are shown in Figures 5, and 7. respectively in Section 3.2. Note that for $q > 3$, the Bratteli diagram has multiple edges between some vertices.

Let r be a path in the Bratteli diagram of $SL_2(F_q) > N > 1$, starting at 1 and ending at a representation of $SL_2(F_q)$. We denote by χ_r the irreducible representation of N which the path r passes through after one step.

Lemma 3.4. *Let ρ be an irreducible representation of $SL_2(F_q)$, and let $n_1, n_2 \in N$, $a \in A$. Then relative to a Gel'fand-Tsetlin basis for ρ, the matrix entries of $\rho(n_1 a n_2)$ have the form*

$$[\rho(n_2 \cdot a \cdot n_1)]_{r,s} = [\rho(a)]_{rs}\chi_r(n_2)\chi_s(n_1) \tag{3.2}$$

where r, s denote maximal paths in the Bratteli diagram of $SL_2(F_q) > N > 1$.

Proof. As remarked, the definition of Gel'fand-Tsetlin bases for $SL_2(F_q) > N > 1$ states that the matrices $\rho(n)$ are diagonal in these bases for all $n \in N$. Therefore $[\rho(n)]_{rs} = \delta_{rs}\chi_r(n)$, where δ_{rs} is 1 if $r = s$, and 0 otherwise. The result now follows from the matrix equation $\rho(n_2 \cdot a \cdot n_1) = \rho(n_2)\rho(a)\rho(n_1)$. $\square$

A formula for $[\rho(a)]_{rs}$ is given in Appendix A.

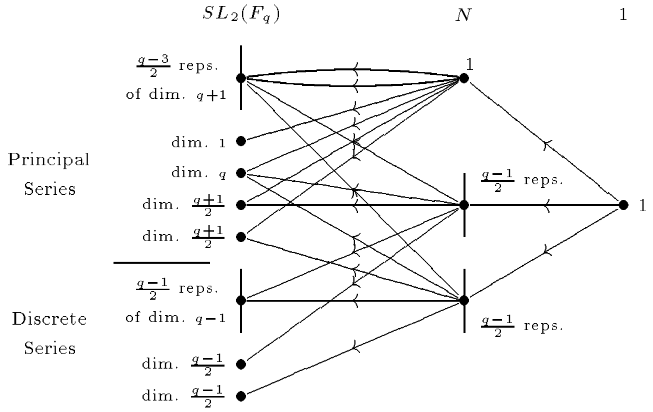

FIGURE 4. The Bratteli diagram for $SL_2(F_q) > N > 1$, q odd. A vertex on a vertical line ranges over a set of representations.

We now give two simple bounds on the number of operations required to compute Fourier transforms on $SL_2(F_q)$. We improve these results in Section 5.

Theorem 3.5 ([40] Corollary 3.30)**.** *Working in a system of Gel'fand-Tsetlin bases, relative to N, for the irreducible representations for $SL_2(F_q)$, the Fourier transform of a complex valued function can be computed in no more than $(4q-2)\cdot(q^3-q)$ operations.*

Theorem 3.6. *Working in a system of Gel'fand-Tsetlin bases, relative to N, for the irreducible representations for $SL_2(F_q)$, the Fourier transform of a complex valued function can be computed in no more than $(2(q-1)+24\log[q])\cdot(q^3-q)$ operations.*

Proof of Theorems 3.5 and 3.6. The two different results come from variations on the same basic algorithm. Let f be a complex valued function on $SL_2(F_q)$. Define a function f^0 on the Cartesian product $N\times A\times N$, by

$$f^0(n_2,a,n_1)=\begin{cases} f(n_2\cdot a\cdot n_1) & \text{if } a\in w\cdot T, \text{ or } n_2=1,\\ 0 & \text{otherwise.}\end{cases} \tag{3.3}$$

By Lemma 3.4 the entries of Fourier transform matrix $\hat{f}(\rho)$ have the factorization

$$[\hat{f}(\rho)]_{rs}=\sum_{a\in A}\sum_{n_2\in N}\sum_{n_1\in N}[\rho(a)]_{rs}\chi_r(n_2)\chi_s(n_1)f^0(n_2,a,n_1), \tag{3.4}$$

where r,s denote paths in the Bratteli diagram, starting at 1 and ending at ρ.

We now evaluate (3.4), by summing over n_1,n_2, and a in three stages.

Stage 1. Compute

$$f^1(n_2,a,\chi)=\sum_{n_1\in N}f^0(n_2,a,n_1)\chi(n_1) \tag{3.5}$$

for all $a\in A$, all $n_2\in N$, and for all (one-dimensional) irreducible representations χ of N.

Stage 2. Compute

$$f^2(a,\chi,\chi') = \sum_{n_2\in N} f^1(n_2,a,\chi)\chi'(n_2) \tag{3.6}$$

for all $a \in A$, and all irreducible representations χ, χ' of N.

Stage 3. Compute

$$[\hat{f}(\rho)]_{r,s} = \sum_{a\in A} f^2(a,\chi_s,\chi_r)\,[\rho(a)]_{r,s}\,. \tag{3.7}$$

for all irreducible representations ρ of $SL_2(F_q)$, and paths r, s in the Bratteli diagram, starting at 1 and ending at ρ.

A rough count shows that the algorithm just described requires no more than $6q^3(q-1)$ multiplications, with the same bound on the number of additions. In particular, the number of possible pairs of paths ending at the same irreducible representation is $|SL_2(F_q)|$, so Stage 3 of the algorithm requires $(2q-2)q(q^2-1)$ operations.

In order to prove the Theorems we need to count the number of operations more carefully. We give two variations on the same basic approach. The first variation is to account for the zeroes occurring in the definition of f^0. At Stage 1, we need only compute the sum (3.5) when $a \in w\cdot T$, or $n_2 = 1$. Thus Stage 1 requires $q^2(q^2-1)$ operations. At Stage 2 we can simplify the sum (3.6) by observing that $f^1(n_2,a,\chi)$ is zero whenever both $a \notin wT$ and $n_2 \neq 1$. Thus Stage 2 also requires $q^2(q^2-1)$ operations. We make no simplification at Stage 3, which still requires $(2q-2)q(q^2-1)$ operations. This proves Theorem 3.5

The second variation of the algorithm is to use standard fast Fourier transforms to improve the calculation at Stages 1 and 2. Both (3.5) and (3.6) have the form of Fourier transforms on the group N. Stage 1 may be computed using $q(q-1)$ Fourier transforms on the abelian group N, one Fourier transform for each (n_2,a) with $a \in w\cdot T$ or $n_2 = 1$. Stage 2 may be computed using $q(2q-2)$ Fourier transforms on N, corresponding to all possible values of a and χ. We keep Stage 3 the same.

Finally, since the group N is abelian, application of Lemma 2.5 shows that for any q a Fourier transform on N may be computed in $8q\log q$ operations. Using this bound with the second variation of the algorithm proves Theorem 3.6.

□

3.2. **Two examples: $SL_2(F_2)$, $SL_2(F_3)$.** The simplest instance of a special linear group is $SL_2(F_2)$. For this example, we let

$$n = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

The subgroup N has two elements and T is the trivial subgroup. The group $SL_2(F_2)$ is isomorphic to the symmetric group S_3, and under this isomorphism, N is isomorphic to to S_2. This implies that the Bratteli diagram of $SL_2(F_2) > N > 1$ is simply Young's lattice for partitions with $1, 2$, or 3 boxes (see eg. [33]). This is shown in Figure 5. The elements of $SL_2(F_2)$ are exactly $1, n, w, wn, nw, nwn$.

For $SL_2(F_2)$ our algorithms are not very efficient, but they do illustrate the general case. Consider the data-flow diagram for the computation of Fourier transforms on $SL_2(F_2)$ shown in Figure 6. Starting with the inputs at the left hand side of the diagram, the data flows along the edges from left to right. At each edge

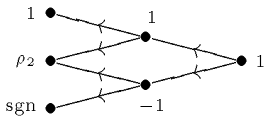

FIGURE 5. The Bratteli diagram for $SL_2(F_2) > N > 1$.

the data is multiplied by a precomputed factor (a matrix entry) which is shown on the diagram; if no factor is shown the data is left unchanged. At each vertex the incoming data is summed to form the next collection of inputs. In Figure 6 we show the precomputed factors for the matrix representations of $SL_2(F_2)$ defined in Appendix A. These are exactly the matrices of representations in Young's seminormal form for S_3 (cf. [33]).

We also indicate those edges and vertices which are omitted in the first variation of our algorithm on account of zero input values. These parts of the diagram are shown as dotted lines and hollow vertices. Therefore, the solid part of the diagram is just the data-flow diagram for the algorithm of Theorem 3.5. Including the dotted or hollow parts of the Figure 6 as well gives the data flow diagram for the algorithm of Theorem 3.6. Note, in particular, that $f^0(n,1,\cdot) = f^1(n,1,\cdot) = 0$.

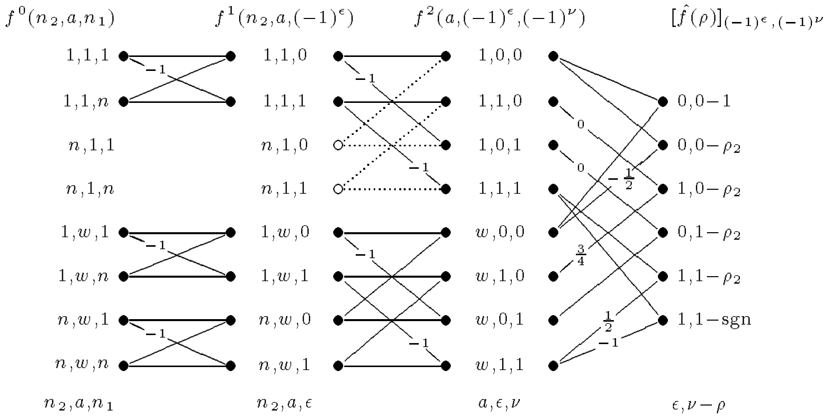

FIGURE 6. Data-flow diagram for $SL_2(F_2)$ transform.

It is easy to distinguish the Fourier transforms on F_2 as the "butterfly" operations in the first two stages in the diagram. Observe that at Stage 2, some of the butterfly operations include both solid and dotted edges, which shows that we must include the dotted edges if we wished to compute the second stage using Fourier transforms on F_2. For larger values of q this second stage would be performed more efficiently using fast Fourier transforms, but for $SL_2(F_2)$ there is no advantage to be gained.

For $q = 2$, the analysis of the algorithms of Theorem 3.5 and Theorem 3.6 give upper bounds of 28 and 26 operations respectively as compared to 36 operations for a direct approach to the computation. These bounds assume the number of operations for an F_2 transform is 2. ($8q \log q$ is clearly an overestimate here.) However, it is clear from Figure 6 that these are both overstatements of the number of operations used by our algorithms. Thus, both Theorems 3.5 and 3.6 miss the fact that several of the precomputed factors at the Stage 3 are zero, and also that certain other precomputed factors are 1. If we take this into account, our approach requires 12 additions and 10 multiplications, i.e., 12 operations.

The next easiest case is the Fourier transform on $SL_2(F_3)$. This group is still solvable, and so does not fully represent the general case. Nevertheless, with 24 elements in the group, the Bratteli and data-flow diagrams are as large as we can comfortably represent on a single page. The Bratteli diagram of $SL_2(F_3)$ is shown in Figure 7, and the data-flow diagram is shown in Figure 8. We omit the multiplication factors in the data-flow diagram to save space. As before, the different algorithms are described by reading Figure 8 with or without the dotted or hollow parts.

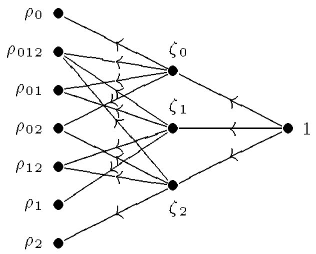

FIGURE 7. The Bratteli diagram for $SL_2(F_3) > N > 1$.

Careful examination of the precomputed factors implied in Figure 8 shows that the Fourier transform on $SL_2(F_3)$ can actually be performed in 132 additions and 122 multiplications, i.e., 132 operations. This sort of careful counting is not important asymptotically.

3.3. **Related work.** In the context of giving an algorithm for computing the spectra of various Cayley graphs on $SL_2(F_q)$, the first efficient algorithm for computing Fourier transforms on special linear groups was developed by Lafferty and Rockmore [36]. Their approach requires $O(q^4 \log q)$ operations.

The adjacency matrix can be viewed as the Fourier transform at the regular representation of the characteristic function of the generating set. Consequently, the spectrum of the Cayley graph is the union of the spectra of the individual Fourier transforms. This reduces the complexity of computing the spectrum from $|SL_2(F_q)|^2$ operations to $O(q^4)$ operations. In the case in which the generating set is small relative to $|SL_2(F_q)|$ a fast transform does not affect the overall complexity of the eigenvalue computation. However, for generating sets of larger than $O(q)$, the algorithms presented here allow the entire computation to require no more than $O(q^4)$ operations. So, in general, the eigenvalues of any Cayley graph on $SL_2(F_q)$ can be computed in $O(q^4)$ operations.

In further work, Lafferty and Rockmore show how an efficient Fourier transform on $SL_2(F_q)$ may be used to give a fast encoding (subquadratic) for a new class of error correcting codes built on Cayley graphs of $SL_2(F_q)$ [35]. The algorithms of this paper give the fastest such encoding known to date.

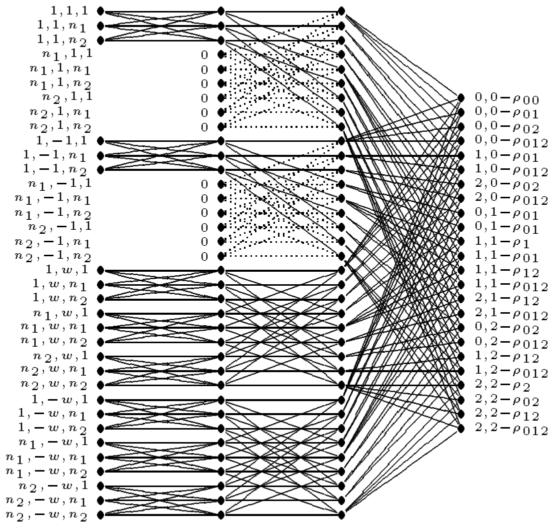

FIGURE 8. The data-flow diagram for $SL_2(F_3)$. Precomputed factors are omitted from the edges.

Additionally, fast Fourier transforms on $SL_2(F_q)$ can be used for efficient computation of spherical harmonic expansions for functions on the two-sphere sampled on grids with icosahderal symmetry. We discuss this in Section 4.3.

4. Fast transforms and double cosets

The arguments of Section 3 generalize to any finite group, provided we generalize our definition of Gel'fand-Tsetlin basis, and replace the Bruhat decomposition of Lemma 3.1 by a factorization of the group elements using double cosets.

Definition 4.1 (Double cosets)**.** Let G be a finite group, and let H and K be subgroups of G.

1. For any group element g in G, the double coset containing g, relative to H and K, is defined to be the set of products

$$HgK = \{hgk | h \in K \text{ and } k \in K\}. \tag{4.1}$$

2. Any element of HgK is called a double coset representative for HgK.
3. The set of all double cosets of G is called the double coset space of G, and is denoted $H\backslash G/K$.
4. A subset of G containing exactly one element from each double coset, relative to H and K, is called a complete set of double coset representatives for $H\backslash G/K$.

Lemma 4.2. *Let G be a finite group, let H and K be subgroups of G, and let A be a complete set of double coset representatives for $H\backslash G/K$.*

1. *Any two double cosets of G relative to H and K are either disjoint or equal. The double cosets of G relative to H and K partition G into disjoint subsets.*

2. *Any group element $g \in G$ may be written in the form*

$$g = h \cdot a \cdot k$$

for some $h \in H$, $a \in A$, and $k \in K$.

Proof. It suffices to show that if $g_1, g_2 \in G$, and $g_2 \in Hg_1K$, then $Hg_2K = Hg_1K$. Suppose $g_2 \in Hg_1K$. Then $Hg_2K \subseteq H \cdot Hg_1K \cdot K = Hg_1K$. The result now follows by symmetry. ☐

We have already observed, in the case of $SL_2(F_q)$, that the restriction of an irreducible representation to a subgroup is not always irreducible. In general, if ρ is a matrix representation of the finite group G, and K is a subgroup of G, then $\rho|K$ is equivalent to a matrix direct sum of irreducible representations of K (see e.g. [50] Chapter 1). If ν is an irreducible representation of K, then the number of representations equivalent to ν occurring in a matrix direct sum of irreducible representations equivalent to $\rho|K$ is called the multiplicity of ν in $\rho|K$. These multiplicities are conveniently represented by a Bratteli diagram.

Definition 4.3 (Bratteli diagram)**.** Assume $G = K_n \geq \cdots \geq K_0 = 1$ is a chain of subgroups of the finite group G. Then the Bratteli diagram for this chain is the directed graded graph defined as follows.

1. For $i = 0, \ldots, n$, the vertices at level i are labeled by the inequivalent irreducible representations of K_i. At level 0 there is a unique vertex, 1.
2. A vertex β at level $i-1$, is connected to a vetex α at level i by $\mathcal{M}(\alpha, \beta)$ edges (all from β to α) where

$$\mathcal{M}(\alpha, \beta) = \text{ the multiplicity for } \beta \text{ in } \alpha|K_{i-1}.$$

This generalizes the definitions of Bratteli diagram given in Sections 2 and 3 for cyclic groups and special linear groups.

As we have seen in Sections 2 and 3, Bratteli diagrams are important for organizing the computation of a Fourier transform because the paths in a Bratteli diagram index a system of Gel'fand-Tsetlin bases for the group G. We shall describe a path by giving a list of its edges. Thus if r is a path in a Bratteli diagram, from level l to level m, we will denote it by $r_{l+1}r_{l+2}\cdots r_m$, where each r_i is an edge from level $i-1$ to level i. Also, let ρ^{r_i} denote the irreducible representation of K_i at the tip of the directed edge r_i.

Definition 4.4 (Gel'fand-Tsetlin bases)**.** Assume $G = K_n \geq \cdots \geq K_0 = 1$ is a chain of subgroups of the finite group G. Then a system of Gel'fand-Tsetlin bases for the irreducible representations of G, relative to the chain of subgroups, is a collection of bases for these representations indexed by paths from level 0 to level n in the Bratteli diagram for the chain of subgroups, such that:

1. The basis vectors for the irreducible representation ρ of G are indexed by paths from the vertex 1 at level 0 to the vertex ρ at level n.
2. If $1 \leq i \leq n$ and $k_i \in K_i$, then the entries of $\rho(k_i)$ in the basis for ρ have the following form.

$$[\rho(k_i)]_{r,s} = \delta_{r_{i+1}\cdots r_n, s_{i+1}\cdots s_n} \left[\rho^{r_i}(k_i)\right]_{r_1\cdots r_i, s_1 \cdots s_i}, \tag{4.2}$$

where δ_{xy} is 1 or 0 according to whether x is or is not equal to y.

3. If $\rho^{r_i} = \nu$, then the matrices $[\rho^{r_i}(k_i)]_{r_1\cdots r_i, s_1\cdots s_i}$ (where we allow $r_1 \cdots r_i$ and $s_1 \cdots s_i$ vary over all paths such that $\rho^{r_i} = \rho^{s_i} = \nu$ and k_i varies over K_i), form an irreducible representation of K_i equivalent to ν.

If K is any subgroup of G, then a system of Gel'fand-Tsetlin bases relative to K is defined to be a system of Gel'fand-Tsetlin bases relative to the chain $G \geq K \geq 1$.

Gel'fand-Tsetlin bases are also sometimes called "subgroup-adapted" bases.

Remark 4.5. Conditions 2. and 3. of Definition 4.4 imply that in a Gel'fand-Tsetlin basis, if ρ is an irreducible representation of G, then the restricted representation $\rho|K_i$ is not only equivalent to, but is equal to a direct sum of irreducible representations of K_i. Furthermore, any two equivalent summands will in fact be equal, and moreover, if ρ' is another irreducible representation of G, then equivalent summands in ρ and ρ' will be equal.

By condition 3. of the Definition 4.4, a system of Gel'fand-Tsetlin bases determines a complete set of matrix representations for any subgroup K_i in the chain. Any set of bases for the irreducible representations of K_i that gives this set of matrix representations is said to be associated to the system of Gel'fand-Tsetlin bases of G.

Lemma 4.6. *Assume $G = K_n \geq \cdots \geq K_0 = 1$ is a chain of subgroups of the finite group G.*

1. *Then there exists a system of Gel'fand-Tsetlin bases for the irreducible representations of G, relative to the chain of subgroups.*
2. *If the Bratteli diagram for the chain of subgroups contains no multiple edges, then the system of Gel'fand-Tsetlin bases is unique up to multiplication of each basis vector by a nonzero scalar.*

Proof. See Appendix B. □

We shall mainly be concerned with the case of subgroup chains of the form $G > K > 1$. However, the double coset approach to computing Fourier transforms requires that we use bases that are Gel'fand-Tsetlin bases relative to two subgroups H and K simultaneously. It is not always possible to find such a set of bases, but there are two important cases where we are guaranteed that such bases exist.

Lemma 4.7. *Assume G is a finite group and H, K are subgroups of G. Assume that either*

1. *$H \geq K$, or $K \geq H$, or*
2. *all elements of H commute with all elements of K.*

Then there exists a system of Gel'fand bases for the irreducible representations of G relative to H, that is simultaneously a system of Gel'fand-Tsetlin bases relative to K.

Proof. See Appendix B. □

Theorem 4.8. *Assume G is a finite group, and H, K are subgroups of G. Then the Fourier transform of any complex valued function on G may be computed at a complete set of irreducible representations, relative to a system of simultaneous Gel'fand-Tsetlin bases for both H and K, in no more than $C_1+C_2+C_3$ operations,* *where*

$$C_1 \le |G/K| \cdot C_K$$
$$C_2 \le \min\{|H| \cdot |G|, \quad |H\backslash G/K| \cdot |K| \cdot C_H\}$$
$$C_3 \le |H\backslash G/K| d_H d_K$$

where C_H and C_K denote the number of operations needed to compute Fourier transforms on H and K in the associated bases, respectively, and d_H and d_K are the maximum degrees of irreducible representations of H and K.

Proof. The proof of this theorem proceeds along the lines of Theorems 3.5 and 3.6. The variation comes from the need to consider Bratteli diagrams which possibly have multidimensional irreducible representations at the first level. This adds a level of indexing in the proof.

We work with a system of Gel'fand-Tsetlin bases relative to both the subgroups H and K. As we are not assuming that these subgroups form a chain of subgroups, we must use two different Bratteli diagrams to index the basis vectors. Thus basis vectors are in a one-to-one correspondence with both paths in a Bratteli diagram for the chain $G \ge H \ge 1$ as well as $G \ge K \ge 1$. This then gives a one-to-one correspondence between paths in the two Bratteli diagrams and we can index basis vectors in terms of either diagram according to our needs.

Let A be a complete set of double coset representatives for $H\backslash G/K$. Choose a subset Y of $H \times A$, of size $|G/K|$, such that

$$\{h \cdot a | (h, a) \in Y\}.$$

is a complete set of coset representatives of G/K. Then any element $g \in G$ may be written uniquely in the form

$$g = h \cdot a \cdot k, \quad \text{where } (h, a) \in Y, \text{ and } k \in K.$$

Let f be a complex valued function on G. Define a function f^0 on the Cartesian product $H \times A \times K$, by

$$f^0(h, a, k) = \begin{cases} f(h \cdot a \cdot k) & \text{if } (h, a) \in Y, \\ 0 & \text{otherwise.} \end{cases} \tag{4.3}$$

Then the Fourier transform of f at ρ may be written as

$$[\hat{f}(\rho)]_{rs} = \sum_{a\in A} \sum_{h\in H} \sum_{k\in K} \sum_{u,v} [\rho(h)]_{sv} [\rho(a)]_{vu} [\rho(k)]_{ur} f^0(h, a, k), \tag{4.4}$$

where $r = r_1 r_2, u = u_1 u_2$ are paths in the Bratteli diagram for $G \ge K \ge 1$ and $s = s_1 s_2, v = v_1 v_2$ are paths in the Bratteli diagram for $G \ge H \ge 1$, all starting at 1 and ending at ρ.

Now we use Equation (4.2) in the definition of Gel'fand-Tsetlin bases to see that only terms with $u_2 = r_2$ and $v_2 = s_2$ contribute to the sum (4.4). Thus (4.4) may be rewritten as

$$[\hat{f}(\rho)]_{rs} = \sum_{a\in A} \sum_{h\in H} \sum_{k\in K} \sum_{u_1,v_1} [\nu(h)]_{s_1 v_1} [\rho(a)]_{v_1 s_2, u_1 r_2} [\mu(k)]_{u_1 r_1} f^0(h, a, k), \tag{4.5}$$

where $\mu = \rho^{u_1}$, $\nu = \rho^{v_2}$, and the edges $s_1, s_2, r_1, r_2, u_1, v_1$ are connected in the relation shown in Figure 9. This diagram shows relations which must exist between the edges in order to have a nonzero contribution to the sum. However, we do not rule out additional relations that may exist for certain values of r, s, u_1, v_1.

We now evaluate Equation (4.5), by summing over k, h, a in three stages.

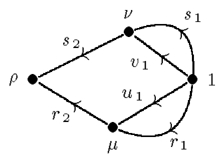

FIGURE 9. Paths in the Bratteli diagrams, general Fourier transform algorithm.

Stage 1. Compute

$$f^1(h, a, \mu, r_1, u_1) = \sum_{k \in k} f^0(h, a, k)[\mu(k)]_{u_1 r_1} \tag{4.6}$$

for all $(a, h) \in Y$, and for all edges u_1, r_1 from level 0 to level 1 in the Bratteli diagram of $G \geq K \geq 1$. For all other values of a and h we define f^1 to be zero.

Stage 2. Compute

$$f^2(a, \mu, r_1, u_1, \nu, s_1, v_1) = \sum_{h \in H} f^1(h, a, \mu, r_1, u_1)[\nu(h)]_{s_1 v_1} \tag{4.7}$$

for all $a \in A$, all edges u_1, r_1 starting at level 0 in the Bratteli diagram of $G \geq K \geq 1$, and for all edges s_1, v_1 starting at level 0 in the Bratteli diagram of $G \geq H \geq 1$.

Stage 3. First compute

$$f^3(a, \rho, r, s) = \sum_{u_1, v_1} f^2(a, \mu, r_1, u_1, \nu, s_1, v_1)[\rho(a)]_{v_1 s_2, u_1 r_2} \tag{4.8}$$

for all double coset representatives $a \in A$, for all paths r in the Bratteli diagram of $G \geq K \geq 1$, and s in the Bratteli diagram of $G \geq H \geq 1$, starting at 1 and ending at ρ, and for edges u_1, v_1 connected to r, s as shown in Figure 9. Note that $\mu = \rho^{r_1}$ and $\nu = \rho^{s_1}$.

Then compute

$$[\hat{f}(\rho)]_{rs} = \sum_{a \in A} f^3(a, \rho, r, s) \tag{4.9}$$

for all paths r in the Bratteli diagram of $G \geq K \geq 1$, and s in the Bratteli diagram of $G \geq H \geq 1$, starting at 1 and ending at ρ.

This is the basic algorithm for computing the Fourier transform of f. Stage 1 consists of $|Y| = |G/K|$ Fourier transforms on K, and thus requires $|G/K| C_K$ operations.

As in the proof of Theorems 3.5 and 3.6, we introduce two variations at Stage 2, which account for the two bounds on C_2 we stated. The first variation takes into account the fact that f^1 is zero when (h, a) does not lie in Y. Thus, in the sum (4.7), we need only sum over those h for which (h, a) lies in Y. This approach to the computation of (4.7) requires $|G|\,|H|$ operations, where we count the maximum of the number of multiplications and additions.

The second variation is to note that Stage 2 has the form of $|A| \cdot |K|$ Fourier transforms on H, and we may use a Fourier transform algorithm on H to compute these transforms, provided the algorithm works for the required sets of bases. This approach to the computation of (4.7) requires $|H\backslash G/K|\cdot|K|\cdot C_H$ operations.

The first part of Stage 3 requires as many operations as there are ways of choosing r, s, u_1, v_1 subject to the relations shown in Figure 9, and then choosing $a \in A$. To count this, we first choose ρ, μ, and ν. Then there are $\mathcal{M}(\rho,\mu)$ ways of choosing r_2, $\mathcal{M}(\rho,\nu)$ ways of choosing s_2, d_μ^2 ways of choosing s_1 and u_1 together, and d_ν^2 ways of choosing r_1 and v_1 together. Summing over all these possibilities for each $a \in A$ we get

$$|H\backslash G/K| \sum_{\rho,\mu,\nu} \mathcal{M}(\rho,\nu)\mathcal{M}(\rho,\mu)d_\mu^2 d_\nu^2$$

operations. The second part of Stage 3 consists of fewer than $|H\backslash G/K|\cdot|G|$ additions, but these do not add to the overall operation counts, because the first part of Stage 4 has exactly this many fewer additions than multiplications.

The bound C_3, for the number of operations required by Stage 3, now follows from the basic representation fact,

$$\sum_{\mu} \mathcal{M}(\rho,\mu)d_\mu = d_\rho.$$

$\square$

Remark 4.9. For fixed $a \in A$, the function f^3 computed at Stage 3 of the algorithm in Theorem 4.8 is actually the Fourier transform of the function f_a which is defined to be equal to f on the double coset HaK and zero elsewhere. Thus we see that our algorithm works by decomposing the function f according to the decomposition of G into double cosets, calculating the Fourier transform on each double coset and then reassembling the results.

Two special cases are worth singling out.

Corollary 4.10 ([40] Theorem 3.28)**.** *Assume G is a finite group and K is a subgroup of G. Then working in a system of Gel'fand-Tsetlin bases for G relative to K, the Fourier transform of a complex valued on G may be computed in no more than*

$$(2\,|K| + |K\backslash G/K|\,d_K^2)\cdot|G|$$

operations, where d_K is the maximum degree of an irreducible representation of K.

Corollary 4.11. *Assume G is a finite group and K is an abelian subgroup of G. Then working in a system of Gel'fand-Tsetlin bases for G relative to K, the Fourier transform of a complex valued on G may be computed in no more than*

$$(2\,|K| + |K\backslash G/K|)\cdot|G|$$

operations, using direct algorithms for computing Fourier transforms on K, or

$$(8\log|K| + 8\frac{|K|^2\,|K\backslash G/K|}{|G|}\log|K| + |K\backslash G/K|)\cdot|G|$$

operations using fast Fourier transform algorithms to compute Fourier transforms on K.

The results of Section 3 may be quickly rederived using Corollary 4.11. If we let $G = SL_2(F_q)$, and $K = N$, then $|G| = q^3 - q$, $|K| = q$, and $|K\backslash G/K| = 2(q-1)$. Thus, Theorem 3.5 follows from the first bound of Corollary 4.11, and Theorem 3.6 follows from the second bound of Corollary 4.11. However, we are now able to derive different algorithms for computing the Fourier transform on $SL_2(F_q)$, by using other subgroups of $SL_2(F_q)$ in place of N. In Section 5 we explore the algorithms that arise when we use the double coset space $N\backslash SL_2(F_q)/B$, where $B = T \cdot N$.

4.1. **Symmetric groups.** We may use Corollary 4.11 to construct new fast transforms on symmetric groups of low rank. We picked abelian subgroups of the symmetric groups S_4, S_5, S_6, S_7, S_8, and calculated the double cosets. The most meaningful way of comparing Fourier transform algorithms is to consider the complexity of the algorithm C_G divided by the order of the group. Corollary 4.11 shows that if G is a finite group and K is an abelian subgroup then we have

$$\frac{C_G}{|G|} \leq 2\,|K| + |K\backslash G/K|\,.$$

For $|K| \leq 43$ this bound is better than the one based on fast Fourier transforms on the abelian subgroup K.

Table 1 gives a comparison of the double coset based algorithms with two other algorithms for computing Fourier transforms on the symmetric group, due to Clausen [13, 14] and Maslen [38] respectively. These other algorithms use different Gel'fand-Tsetlin bases, so that the upper bounds are not directly comparable, and at any rate these are not exact operation counts. In addition, they do not take into account ease of programming, data flow, and most important, any real running times. Nevertheless, this may provide some insight into the relative utility of the different approaches.

In Table 1 we have

$$\begin{aligned} C_G^{\mathrm{dc}} &= (2\,|K| + |K\backslash\, G/K|) \cdot |G| \\ C_G^{\mathrm{ma}} &= \text{operation count for the algorithm of [38]} \\ C_G^{\mathrm{cl}} &= \text{operation count for the algorithm of [13, 14]} \\ &= \frac{(n+1)n(n-1)}{3}, \text{ for } G = S_n. \end{aligned}$$

4.2. **The rotation group.** One of the nicest examples of a double coset decomposition is the use of Euler angles to put coordinates on the group of rotations in three-space. The rotation group is a continuous group, but the definitions and techniques we have developed for finite groups apply here with very few changes. The algorithm for computing Fourier transforms using double cosets suggests itself as a standard separation of variables approach to the computation. In fact the starting point for our work in this paper was the observation that the obvious algorithm for computing Fourier transforms on the rotation group could be generalized to arbitrary finite groups, and to the special linear group in particular.

The rotation group is usually identified with the group of all 3×3 orthogonal matrices with determinant one, and is denoted $SO(3)$. The Euler angle decomposition writes any rotation as a product of rotations about the z and x axes. For ths,

G	K	$\lvert G\rvert$	$\lvert K\rvert$	$\lvert K\backslash G/K\rvert$	$\frac{C_G^{\mathrm{dc}}}{\lvert G\rvert}$	$\frac{C_G^{\mathrm{ma}}}{\lvert G\rvert}$	$\frac{C_G^{\mathrm{cl}}}{\lvert G\rvert}$
S_4	$\langle(1,2),(3,4)\rangle$	24	4	3	11	5.4	20
	$\langle(1,2,3,4)\rangle$		4	3	11		
S_5	$\langle(1,2,3),(4,5)\rangle$	120	6	5	17	9.1	40
	$\langle(1,2,3,4,5)\rangle$		5	8	18		
	$\langle(1,2),(4,5)\rangle$		4	11	19		
S_6	$\langle(1,2,3,4),(5,6)\rangle$	720	8	14	30	13.6	70
	$\langle(1,2,3),(4,5,6)\rangle$		8	14	30		
	$\langle(1,2),(3,4),(5,6)\rangle$		8	21	37		
	$\langle(1,2,3,4,5,6)\rangle$		6	24	36		
S_7	$\langle(1,2,3,4),(5,6,7)\rangle$	5040	12	40	64	19.1	112
	$\langle(1,2),(3,4),(5,6,7)\rangle$		12	46	70		
	$\langle(1,2,3,4,5),(6,7)\rangle$		10	56	76		
S_8	$\langle(1,2,3),(4,5,6),(7,8)\rangle$	40320	18	152	186	25.7	168
	$\langle(1,2,3,4),(5,6,7,8)\rangle$		16	172	204		
	$\langle(1,2,3,4),(5,6),(7,8)\rangle$		16	189	221		
	$\langle(1,2,3),(4,5,6,7,8)\rangle$		15	192	222		
	$\langle(1,2),(3,4),(5,6),(7,8)\rangle$		16	282	314		

TABLE 1. Operation counts for Fourier transforms on symmetric groups.

define for any real angle θ the matrices

$$r_2(\theta) = \begin{pmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \qquad r_3(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & \sin\theta \\ 0 & -\sin\theta & \cos\theta \end{pmatrix}. \tag{4.10}$$

Any $g \in SO(3)$ may be expressed uniquely in the form

$$g = r_2(\varphi_2) r_3(\theta) r_2(\varphi_1) \tag{4.11}$$

for $\varphi_1, \varphi_2 \in [0, 2\pi)$ and $\theta \in (0, \pi)$ or $\varphi_1 = \varphi_2 = 0$ and $\theta \in \{0, \pi\}$. This defines a coordinate system on a dense open subset of the rotation group, and the coordinates, $\varphi_2, \theta, \varphi_1$, of any rotation are called its Euler angles (see e.g. [52], p. 106).

The double coset decomposition corresponding to (4.11) is obtained as follows. Let $SO(2) = \{r_2(\theta) \mid \theta \in \mathbf{R}\}$. This is the subgroup of $SO(3)$ consisting of the rotations about the z-axis. Then by Equation (4.11), the set of group elements $\{r_3(\theta) \mid \theta \in [0, \pi]\}$ is a compete set of double coset representatives for $SO(2)\backslash SO(3)/SO(2)$.

It is well-known that the irreducible representations of $SO(3)$ may be realized by the action of $SO(3)$ on the vector space of harmonic polynomials of degree l in three variables [52]. As such they are usually indexed by nonnegative integers l and we denote them as T_l for $l \geq 0$. The restriction of T_l to $SO(2)$ is equivalent to the sum of one-dimensional representations $r_2(\varphi) \mapsto e^{im\varphi}$ of $SO(2)$ with $|m| \leq l$. Thus, in a manner completely analogous to Equation (3.2) the irreducible representations of $SO(3)$ may be written in a system of Gel'fand-Tsetlin bases relative to $SO(2)$ as functions of the Euler angles of the form

$$T^l_{m_1,m_2}(\theta, \varphi_1, \varphi_2) = P^l_{m_1,m_2}(\theta) e^{im_1\varphi_1} e^{im_2\varphi_2} \tag{4.12}$$

where $|m_1|, |m_2| \leq l$ index representations of $SO(2)$ occurring in the restriction of T_l to $SO(2)$. This is simply a continuous analogue of Lemma 3.4.

When the system of Gel'fand-Tsetlin bases is chosen carefully, the functions $P^l_{m_1,m_2}$ can be written in terms of Jacobi polynomials as

$$(4.13) \quad P^l_{m_1,m_2}(\theta) = \left[\frac{(l-m_1)!(l+m_1)!}{(l-m_2)!(l+m_2)!}\right]^{\frac{1}{2}} \times \\ \times (\sin\frac{\theta}{2})^{m_1-m_2}(\cos\frac{\theta}{2})^{m_1+m_2} P^{(m_1-m_2,m_1+m_2)}_{l-m_1}(\cos\theta)$$

For a proof of this formula see [52].

With this notation we may define the Fourier transform of a continuous function on $SO(3)$ as follows.

Definition 4.12 (Fourier transform on $SO(3)$)**.** Assume f is a continuous complex valued function on $SO(3)$. Then the Fourier transform of f is the collection of matrices $\hat{f}(T_l)_{m_1,m_2}$ defined by

(4.14)
$$\hat{f}(T_l)_{m_1,m_2} = \frac{1}{8\pi^2}\int_{SO(3)} f(\varphi_2,\theta,\varphi_1)P^l_{m_1,m_2}(\theta)e^{i(m_1\varphi_1+m_2\varphi_2)}\sin\theta d\varphi_1 d\varphi_2 d\theta.$$

To state an analogue of our finite group algorithms for the rotation group, we must first place conditions on the function f that reduces the calculation of its Fourier transform to a finite computation. One simple way to do this is to require that the function be band-limited, in the sense of the following definition. This will allow us to formulate a quadrature rule which will compute the Fourier transform from values of f on a finite grid.

Definition 4.13 (Band-limited functions on $SO(3)$)**.** A continuous function f on $SO(3)$ is band-limited with band-limit b if $\hat{f}(T_l)_{m_1,m_2} = 0$ for all $l > b$.

Given a function on $SO(3)$ with band-limit b, the grid we use to sample functions on $SO(3)$ is an equi-angular grid in the Euler angle coordinate system. Let $\varphi^j_1, \varphi^j_2, \theta^k$ be defined by

$$\varphi^j_1 = \varphi^j_2 = \frac{j}{2b+1} \quad \text{and} \quad \theta^k = \frac{(k+\frac{1}{2})\pi}{2b+1}, \text{ for } 0 \leq j,k \leq 2b.$$

For this grid and band-limit we have the following simple quadrature rule. (For a proof, see e.g. [40] Theorem 4.5.)

Lemma 4.14. *Let f be a band-limited function of $SO(3)$ of band-limit b. Then there is a sequence of $2b+1$ numbers $w_b(0),\ldots,w_b(2b)$, independent of f, such that for $l \leq b$ the Fourier transform $\hat{f}(T_l)_{m_1,m_2}$ may be computed by the finite sum*

(4.15)
$$\hat{f}(T_l)_{m_1,m_2} = \frac{1}{(2b+1)^2}\sum_{k,j_1,j_2=0}^{2b} w_b(k)P^l_{m_1,m_2}(\theta^k)e^{im_2\varphi_2^{j_2}}e^{im_1\varphi_1^{j_1}}f(\varphi_2^{j_2},\theta^k,\varphi_1^{j_1}).$$

This is an analogue of Equations (4.5) and (3.4), and yields here an efficient algorithm for its evaluation (just as in the case of $SL_2(F_q)$).

Theorem 4.15 ([24] [40]). *The Fourier transform of a band-limited function, f, on $SO(3)$ with band-limit b may be computed in $O(b^3(\log b)^2)$ scalar operations, given the function values of f on the set*

$$\{r_2(\varphi_2^{j_2})r_3(\theta^k)r_2(\varphi_1^{j_1})|0 \leq j_1, j_2, k\}$$

Proof sketch. By using (4.15) the Fourier transform may be computed in three stages as follows.

Stage 1. Compute

$$f^1(j_2, k; m_1) = \frac{1}{2b+1} \sum_{j_1=0}^{2b} e^{im_1\varphi_1^{j_1}} f(\varphi_2^{j_2}, \theta^k, \varphi_1^{j_1})$$

for $0 \leq j_2, k \leq 2b$, and $|m_1| \leq b$.

Stage 2.

$$f^2(k; m_1, m_2) = \frac{1}{2b+1} \sum_{j_2=0}^{2b} e^{im_2\varphi_2^{j_2}} f^1(j_2, k; m_1)$$

for $0 \leq k \leq 2b$, and $|m_1|, |m_2| \leq b$.

Stage 3.

$$\hat{f}(T_l)_{m_1,m_2} = \sum_{k=0}^{2b} w_b(k) P^l_{m_1,m_2}(\theta^k) f^2(k; m_1, m_2)$$

for $|m_1|, |m_2| \leq l \leq b$.

A direct approach to computing the sums at the three stages results in an operation count of $(5b+1)(2b+1)^3$ scalar operations for the entire calculation. However, as in the case of finite groups we may compute the first two stages by means of fast Fourier transforms on the abelian subgroup $SO(2)$, which in this discretized form amounts to computing Fourier transforms on the cyclic group $\mathbf{Z}/(2b+1)\mathbf{Z}$. Thus, the first two stages together take $O(b^3 \log b)$ operations to complete.

To obtain an efficient method of computing Stage 3, we must use the special form of the functions $P^l_{m_1,m_2}(\theta)$. According to Equation (4.13), these may be written in terms of Jacobi polynomials, so the sums at Stage 3, each have the form of a discrete orthogonal polynomial transform at $2b+1$ points. Such polynomial transforms may be computed in $O(b(\log b)^2$ scalar operations using the fast polynomial transform techniques [24, 25]. Therefore Stage 3 may be computed in $O(b^3(\log b)^2)$ scalar operations. □

It is interesting to compare the result for the rotation groups with the result for finite special linear groups. The algorithms in each case have the same overall structure, and in each case a direct approach to the three stages gives a bound of $O(N^{\frac{4}{3}})$ on the number of operations needed to compute a Fourier transform on N points. However, the matrix elements of the rotation group at the double coset representatives have a form that lets us compute Stage 3 efficiently, while it is not clear how to compute Stage 3 for the special linear group in fewer than $O(N^{\frac{4}{3}})$ operations. This is currently the bottleneck to obtaining an $O(q^3 \log q)$ algorithm for the special linear groups via this approach.

4.3. **Icosahedral grids.** For applications in meteorology, the use of sets of points, invariant under the rotational symmetries of an icosahedron, have been suggested for sampling data on the two-sphere (see e.g. [31, 32]). A typical problem in such applications is to compute the spherical harmonic expansion of a function from its values at the invariant set of points. This simplifies the form of rotationally invariant operators such as the Laplacian or gradient. In this setting it is natural to ask whether the symmetries of the sample points can be used to speed up the computation. We shall now show that this is indeed the case, assuming a choice of the correct spherical harmonic basis.

For the remainder of this section, we let G denote the finite subgroup of $SO(3)$ consisting of rotations that leave a given icosahedron invariant. The group G has 60 elements: 12 rotations by $\frac{2\pi}{5}$, 12 rotations by $\frac{4\pi}{5}$, 20 rotations by $\frac{2\pi}{3}$, 15 rotations by π, and the identity element. In fact, G is isomorphic to $PSL_2(F_5)$, and therefore has 5 irreducible representations of degrees $1, 3, 3, 4, 5$ respectively. We shall denote these by $\rho_1, \rho_3^+, \rho_3^-, \rho_4, \rho_5$, where ρ_3^+ is the standard representation given by mapping each group element to its matrix in the usual coordinates on $\mathbf{R}^3$.

Definition 4.16 (Icosahedral grid)**.** Let X denote any finite subset of the two-sphere. Then the icosahedral grid associated with X is the set of points

$$G \cdot X = \{g \cdot x | g \in G, x \in X\}.$$

Let l be any non-negative integer. Then the space V_l of spherical harmonics of degree l is defined to be the space of harmonic complex polynomials $f(x_0, x_1, x_2)$ of three real variables of total degree l. We view these polynomials as defining functions on the two-sphere $x_0^2 + x_1^2 + x_2^2 = 1$, and consider V_l a subspace of the space of continuous functions on the sphere. If we pick any basis for V_l, then we may write the natural action $(g \cdot f) = f(g^{-1}x)$ of $SO(3)$ on functions in this basis, and obtain a matrix representation of $SO(3)$ equivalent to T_l, as defined in Section 4.2.

The usual choice of bases for the spherical harmonics forms a system Gel'fand-Tsetlin bases for the chain of subgroups $SO(3) > SO(2) > 1$. This gives rise to the matrix representations T_l. This choice of bases is particularly suited to working in spherical polar coordinates. When computing on icosahedral grids, however, a system of Gel'fand-Tsetlin bases relative to the chain of subgroups $SO(3) > G > 1$ is more convenient. This is because working with an icosahedral grid amounts to using a double coset decomposition of the form $G \backslash SO(3)/SO(2)$, under the correspondence of the two-sphere with the coset space $SO(3)/SO(2)$.

We now describe the Bratteli diagram of the chain $SO(3) > G > 1$, which is defined in terms of the restrictions of the representations V_l to G.

Theorem 4.17. *Define rational functions $R_i(z)$, $i = 1, 2, 4$, and $R_3^{\pm}$ by*

$$\begin{aligned} R_1(z) &= 1 + \frac{z^6(1+z-z^3)}{(1+z)(1-z^3)(1-z^5)}, \\ R_3^+(z) &= \frac{z(1-z^2+z^4)}{(1-z^2)(1-z^5)}, \\ R_3^-(z) &= \frac{z^3}{(1-z^2)(1-z^5)}, \\ R_4(z) &= \frac{z^3(1+z)}{(1-z^3)(1-z^5)}, \qquad \textit{and} \\ R_5(z) &= \frac{z^2}{(1-z^2)(1-z^3)}. \end{aligned}$$

Then the multiplicity of the representation ρ_i, $i = 1, 4, 5$, or $\rho_3^{\pm}$ in the restriction of T_l to G is the coefficient of z^l in the formal power series for $R_i(z)$, or $R_3^{\pm}(z)$ respectively.

Proof. The proof is by character theory. First, compute the generating function for the characters of the symmetric l^{th} powers of ρ_3^+, using the method given in [50] Chapter 9.1. Then note that the restriction of T_l to G is the quotient of $\mathrm{Sym}^l\,\rho_3^+$ by a subrepresentation equivalent to $\mathrm{Sym}^{l-2}\,\rho_3^+$. This means that the generating function for the characters of the restricted representations $T_l|G$ is $1 - z^2$ times the generating function for the symmetric powers. The result now follows by taking inner products with each of the irreducible characters of G. □

Since the dimension of T_l is $2l + 1$, which increases without bound, but the dimension of a representation of G is at most 5, it is clear that the Bratteli diagram must contain vertices with arbitrarily large numbers of edges between them. In Figure 10 we show the portion of the Bratteli diagram for $l \leq 8$. The edges between levels zero and one may be identified with paths of length three in a Bratteli diagram for $PSL_2(F_5)$ (see Figure 14).

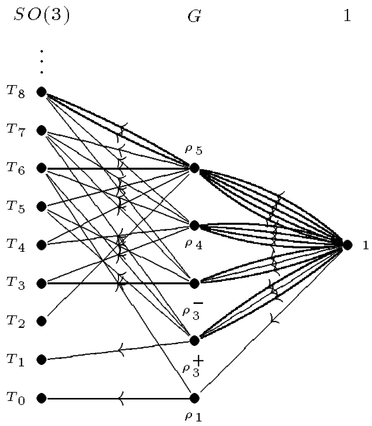

FIGURE 10. The Bratteli diagram for $SO(3) > G > 1$, $0 \leq l \leq 8$.
$G =$ the icosahedral group.

Definition 4.18 (Icosahedral spherical harmonics)**.** We call a system of Gel'fand-Tsetlin bases relative to $SO(3) > G > 1$ a complete set of icosahedral spherical harmonics, and call the basis vectors icosahedral spherical harmonics. We denote the icosahedral spherical harmonics $Y^l_{r_1 r_2}$, where $r_1 r_2$ is a path in the Bratteli diagram for $SO(3) > G > 1$ ending at T_l.

Since the Bratteli diagram of $SO(3) > G > 1$ has multiple edges, there are many different complete sets of icosahedral spherical harmonics unrelated by scalar multiples. Some of this ambiguity may be resolved by including subgroups of G in the chain. The most convenient subgroups for this purpose are dihedral group $\mathbf{D}_{10}$ of order 10 about a fixed axis through a vertex, and the group of rotations $\mathbf{C}_5$ about the same axis. These correspond to the Borel subgroup and upper unipotent matrices in $PSL_2(F_5)$.

Theorem 4.19. *Let X be a finite subset of the two-sphere, let b be a nonnegative integer, and let f be a complex valued function on the icosahedral grid $G \cdot X$. Then the collection of sums*

$$\hat{f}^l_{r_1 r_2} = \sum_{g \in G, x \in X} f(g \cdot x) Y^l_{r_1 r_2}(g \cdot x) \tag{4.16}$$

may be computed for all $0 \le l \le b$ and all paths $r_1 r_2$ ending at T_l, in no more than

$$|X| \cdot \left[C_G + 5(b+1)^2\right]$$

scalar operations, where C_G is the number of operations required to compute a Fourier transform on G. If the icosahedral spherical harmonics form a system of Gel'fand-Tsetlin bases relative to the entire chain $SO(3) > G > \mathbf{D}_{10} > \mathbf{C}_5 > 1$, then the number of operations required is

$$|X| \cdot \left[580 + 5(b+1)^2\right]$$

Proof. Let $r_1 r_2$ be a path in the Bratteli diagram ending at T_l, and let ρ^{r_1} denote the representation of G at the tip of r_1. Then we have

$$Y^l_{r_1 r_2}(g \cdot x) = \sum_{s_1} \left[\rho^{r_1}(g)\right]_{r_1, s_1} Y^l_{s_1 r_2}(x),$$

for any $g \in G$, $x \in X$, where s_1 ranges over edges ending at ρ^{r_1}. Now choose a subset $Y \subseteq G \times X$ of size $|G \cdot X|$, such that

$$\{g \cdot x | (g, x) \in Y\} = G \cdot X,$$

and define the function f^0 on $G \times X$ by

$$f^0(g, x) = \begin{cases} g \cdot x & \text{if } (g, x) \in Y \\ 0 & \text{otherwise.} \end{cases}$$

Then we may write

$$\hat{f}^l_{r_1 r_2} = \sum_{x \in X} \sum_{g \in G} \sum_{s_1} \left[\rho^{r_1}(g)\right]_{r_1, s_1} Y^l_{s_1 r_2}(x) f^0(g, x). \tag{4.17}$$

We evaluate Equation (4.17) as follows

Stage 1. Compute

$$f^1(x, r_1, s_1) = \sum_{g \in G} \left[\rho^{r_1}(g)\right]_{r_1, s_1} f^0(g, x),$$

for all $x \in X$, and for all edges r_1, s_1 from level 0 to ρ^{r_1} at level one in the Bratteli diagram.

Stage 2. Compute

$$\hat{f}^l_{r_1 r_2} = \sum_{x \in X} \sum_{s_1} Y^l_{s_1 r_2}(x) f^1(x, r_1, s_1)$$

for all $0 \le l \le b$ and paths $r_1 r_2$ from level 0 to T_l, where s_1 is an edge from level 0 ending at ρ^{r_1}.

Stage 1 of the computation may be performed using $|X|$ Fourier transforms on G. This requires $|X| C_G$ operations, where C_G depends on the particular bases chosen for the representations of G. In the case where the icosahedral spherical harmonics are relative to the chain $SO(3) > G > \mathbf{D}_{10} > \mathbf{C}_5 > 1$, Theorem 5.5 allows us to bound C_G by 580.

Stage 2 is computed by summing first on s_1, and then on x. This requires

$$|X|\sum_{l=0}^{b}\sum_{\rho}\mathcal{M}(T_l,\rho)d_\rho^2 \leq 5\sum_{l=0}^{b}(2l+1)$$

scalar operations, where ρ ranges over irreducible representations of G, so that d_ρ is bounded by 5. Here $\mathcal{M}(T_l,\rho)$ denotes the multiplicity of ρ in $T_l|G$, so that $\sum_\rho \mathcal{M}(T_l,\rho) = 2l+1$. □

Remark 4.20. A naive approach to the computation of the sums (4.16) in general requires $60\,|X|\,(b+1)^2$ scalar operations, so for b larger than 2, the result of Theorem 4.19 is an improvement over the naive approach. The coefficient 5 in front of $(b+1)^2$ can be improved.

Remark 4.21. It is interesting to note that the double coset approach occurs twice in the computation of (4.16). First in the overall structure of the algorithm, and second in the computation of Fourier transforms on $G \cong PSL_2(F_5)$.

In order to relate the sums (4.16) to the expansion of functions in icosahedral spherical harmonics requires some sampling theory. By elementary linear algebra it follows that for large enough X there is a choice of weighting function w on $G\cdot X$ so that the result of computing (4.16) using $w \cdot f$ instead of f will give the coefficients of f in the icosahedral spherical harmonic basis. However, for certain sets X, the sampling properties may be significantly better. There is a large literature on group invariant cubature formula (see e.g. [34, 51, 43, 30, 28]).

For the purpose of applying invariant differential operators to a function on the sphere, the icosahedral spherical harmonics are in principle as useful as the spherical harmonics, because they diagonalize any such operator. In practice the major obstacle to their use is the problem of finding a set of icosahedral spherical harmonics in the first place. We have only described the first step in this direction, computing the restriction multiplicities. The problem of explicitly decomposing the representations of $SO(3)$ under G remains to be solved.

5. The special linear group revisited

The special linear group was treated in Section 3 using the double coset space $N\backslash SL_2(F_q)/N$. The subgroup N is abelian so the analysis in that section was particularly simple. However, given the results of Section 4, it is more efficient to compute the Fourier transform using the double coset space $N\backslash SL_2(F_q)/B$, where $B = T \cdot N$ is the Borel subgroup; in that situation we can use efficient Fourier transform algorithms on B to aid the computation. After a few slight improvements specific to $SL_2(F_q)$, we obtain our best algorithm for computing Fourier transforms on $SL_2(F_q)$ in Theorem 5.3.

Recall that the trace of a representation is called it's character. In the case of abelian groups, the irreducible representations are one-dimensional, and therefore coincide with the characters. For abelian groups we will therefore use the terms "irreducible representation" and character interchangeably.

5.1. **The Fourier transform on the Borel subgroup.** A complete set of irreducible representations of $B = T \cdot N$ may be constructed using the action of the group T on the space $\hat{N}$ of characters of N. For any character $\zeta \in \hat{N}$, and any $t \in T$ we define the character ζ^t by

$$\zeta^t(n) = \zeta(t^{-1} \cdot n \cdot t), \quad \text{for all } n \in N.$$

In the case when q is odd there are two nontrivial orbits of T on $\hat{N}$ and when q is even there is just one nontrivial orbit. This action corresponds to the action of T on the vector space dual of F_q over its prime field.

Pick any orbit X of T on $\hat{N}$, and pick any character ψ, of T. Then we define a matrix representation $\nu^{X,\psi}$ of B whose rows and columns are indexed by elements of X, by

$$\left[\nu^{X,\psi}(t \cdot n)\right]_{\zeta_1,\zeta_2} = \delta_{\zeta_1,\zeta_2^t}\psi(t)\zeta_2(n), \quad \zeta_1,\zeta_2 \in X, t \in T, n \in N. \tag{5.1}$$

Lemma 5.1. *The representations defined by Equation* (5.1) *are irreducible and include all the irreducible representations of B, though for certain values of X and ψ the representations may be equivalent. A complete set of inequivalent representations may be obtained by restricting the values of ψ to the following.*

$$\begin{aligned} q \text{ odd}: \quad & X = \{1\}, \psi \text{ arbitrary, or } X \neq \{1\}, \psi = 1 \text{ or } \psi = \psi^- \\ q \text{ even}: \quad & X = \{1\}, \psi \text{ arbitrary, or } X \neq \{1\}, \psi = 1. \end{aligned} \tag{5.2}$$

When q is odd, the character ψ^- on T is chosen to be any fixed character such that $\psi^-(-1) = -1$.

Proof. We start by verifying that Equation (5.1) does indeed define a representation of B. To do this, we need only check the commutation relations of the matrices $\nu^{X,\psi}(t)$ and $\nu^{X,\psi}(n)$, since it is clear that these already define representations on T and N respectively. For any $t \in T$ and $n \in N$, we have

$$\begin{aligned} \left[\nu^{X,\psi}(t^{-1})\nu^{X,\psi}(n)\nu^{X,\psi}(t)\right]_{\zeta_1,\zeta_4} &= \sum_{\zeta_2,\zeta_3 \in X} \delta_{\zeta_1,\zeta_2^{t^{-1}}}\psi(t^{-1})\delta_{\zeta_2,\zeta_3}\zeta_3(n)\delta_{\zeta_3,\zeta_4^t}\psi(t) \\ &= \delta_{\zeta_1,\zeta_4}(\zeta_4^t)(n) \\ &= \delta_{\zeta_1,\zeta_4}\zeta_4(t^{-1}nt) = \left[\nu^{X,\psi}(t^{-1}nt)\right]_{\zeta_1,\zeta_4}. \end{aligned}$$

We complete the proof by elementary character theory (see e.g. [50] Chapter 2). The character of $\nu^{X,\psi}$ is

$$\operatorname{Tr}\nu^{X,\psi}(t \cdot n) = \begin{cases} \psi(t)\sum_{\zeta \in X}\zeta(n) & \text{if } t^2 = 1, \\ 0 & \text{otherwise} \end{cases}$$

The orthogonality of the characters $\zeta(n)$ then implies that the inner product of the character of $\nu^{X,\psi}$ with itself is 1, so $\nu^{X,\psi}$ is irreducible. To show that the given representations are distinct, either look at their characters, or consider the restrictions to N and T. The completeness of the set of representations follows by summing the squares of the degrees. $\square$

Theorem 5.2 ([14], [49]). *Working in a system of Gel'fand-Tsetlin bases, relative to N, for the irreducible representations of B, the Fourier transform of a complex valued function can be computed using no more than*

$$(q-1)C_N + \frac{1}{2}(q-1)^2 + C_T + 2(q-1)^2$$

scalar operations, when q is odd, and no more than

$$(q-1)C_N + C_T + 2(q-1)^2$$

scalar operations, when q is even. C_N and C_T denote the number of operations required to compute Fourier transforms on N and T respectively.

Proof. First note that any character of N occurs at most once in the restriction of any given irreducible representation of B. Therefore the system of Gel'fand-Tsetlin bases for the irreducible representations of B, relative to N, is uniquely determined up to a choice of nonzero scalar multiples. Furthermore, the choice of basis for the one-dimensional representations does not affect the Fourier transform matrices of a given function. This means that it takes at most $2(q-1)^2$ scalar multiplications to change the Fourier transform matrices for a function from one Gel'fand-Tsetlin basis to another.

We now consider the case where the irreducible representations of B are given in the form of Equation (5.1), where X and ψ satisfy (5.2). When q is odd we choose a set T^+ of $\frac{q-1}{2}$ elements of T such that $T = T^+ \cup (-T^+)$. we also fix a character ψ^- of T such that $\psi^-(-1) = -1$. Then we may compute the transform as follows.

Stage 1. Compute

$$f^1(t,\zeta) = \sum_{n \in N} \zeta(n) f(t \cdot n) \tag{5.3}$$

for all characters ζ of N, and for all $t \in T$. Do this using $q-1$ Fourier transforms on N, one for each value of t.

Stage 2. Compute

$$f^2(t,\zeta,\epsilon) = f^1(t,\zeta) + \epsilon f^1(-t,\zeta) \tag{5.4}$$

for all $t \in T^+$, for all nontrivial characters ζ_1 of N, and for $\epsilon = \pm 1$.

Stage 3. Compute

$$\left[\hat{f}(\nu^{X,\psi})\right]_{\zeta_1,\zeta_2} = \psi(t) f^2(t,\zeta_2,\psi(-1)) \tag{5.5}$$

for all nontrivial orbits X, $\zeta_1, \zeta_2 \in X$, and $\psi \in \{1, \psi^-\}$. Here, t is the unique element of T^+ such that $\zeta_1 = \zeta_2^t$. When $\psi = 1$ this requires no computation. Also compute

$$\left[\hat{f}(\nu^{\{1\},\psi})\right] = \sum_{t \in T} \psi(t) f^1(t,1)$$

for all characters ψ of T. Do this using a Fourier transform on T.

Stage 1 of this algorithm requires $(q-1)C_N$ operations. Stage 2 requires $(q-1)^2$ additions. Stage 3 requires $\frac{1}{2}(q-1)^2$ multiplications, and a further C_T operations (maximum of multiplications and additions for the Fourier transform on T. Taking into account the multiplications required to change to a different Gel'fand-Tsetlin basis, we see that for q odd our algorithm takes fewer than $(q-1)C_N + (q-1)^2 + \frac{1}{2}(q-1)^2 + C_T$ additions and $(q-1)C_N + \frac{1}{2}(q-1)^2 + C_T + 2(q-1)^2$ multiplications.

The case when q is even is much simpler. In that case there is just 1 nontrivial orbit of T on the characters of N. We then compute the transform in two stages.

Stage 1. Compute

$$f^1(t,\zeta) = \sum_{n \in N} \zeta(n) f(t \cdot n)$$

for all characters ζ of N, and for all $t \in T$. Do this using $q-1$ Fourier transforms on N, one for each value of t.

Stage 2. Compute

$$\left[\hat{f}(\nu^{\{1\},\psi})\right] = \sum_{t\in T} \psi(t) f^1(t,1)$$

for all characters ψ of T. Do this using a Fourier transform on T. The Fourier transform at the representation corresponding to the nontrivial orbit X of T on the characters of N is

$$\left[\hat{f}(\nu^{X,1})\right]_{\zeta_1,\zeta_2} = f^1(t,\zeta_2),$$

where $\zeta_1, \zeta_2 \in X$, and t is the unique element of T such that $\zeta_1 = \zeta_2^t$.

□

5.2. **The special linear group.** Theorems 4.8 and 5.2 immediately imply an improved result for the computation of Fourier transforms on $SL_2(F_q)$. Combining the bounds of these theorems we find that for q odd it takes either $(q^3-q)(3q+\frac{5}{2}-\frac{7}{2q})$ scalar operations or

$$(q^2+\frac{5}{2}q+\frac{5}{2})(q-1)^2 + 16q(q^2+\frac{1}{2}q+\frac{1}{2})\log q + 8(q^2-1)\log(q-1)$$

scalar operations to perform such a transform, where the second result comes from the use of fast Fourier transforms to compute Fourier transforms on N and T.

A better result is possible, if we apply Theorem 4.8 and Theorem 5.2 to functions supported on each double coset, as appropriate for that double coset.

Theorem 5.3. *Let q be an odd prime power. Working in a system of Gel'fand-Tsetlin bases, relative to $SL_2(F_q) > B > N > 1$, for the irreducible representations of $SL_2(F_q)$, the Fourier transform of a complex valued function can be computed using no more than*

$$\frac{5}{2}(q+1)^2(q-\frac{2}{5})(q-1)$$

scalar operations. Also, this transform may be computed in no more than

$$\frac{1}{2}(q^2-1)q(q+3) + 16(q+\frac{1}{2})q(q-1)\log q + 8(q^2-1)\log(q-1)$$

scalar operations.

Proof. We start by recalling the decomposition

$$SL_2(F_q) = B \cup (N\cdot w\cdot B)$$

of $SL_2(F_q)$ into double cosets. Let f be a complex valued function on $SL_2(F_q)$. Then we may write $f = f_1 + f_w$, where f_1 is zero on $N\cdot w\cdot B$, and f_w is zero on B. To compute the Fourier transform of f we compute the Fourier transform of f_1 using Theorem 5.2 and the Fourier transform of f_w using the approach of Theorem 4.8. We then add the results.

In Appendix A, we show that the Bratteli diagram for the chain of subgroups $SL_2(F_q) > B > N > 1$ has no multiple edges. This means that a path in this diagram is determined by a sequence of irreducible representations $1\zeta\nu\rho$, where ζ, ν, ρ are representations of N, B, and $SL_2(F_q)$ respectively.

The full algorithm may now be described as follows.

Stage 1. Compute

$$f^1(g,\zeta_1,\zeta_2) = \sum_{b\in B} f(g\cdot b)\,[\nu_1(b)]_{\zeta_2,\zeta_1} \tag{5.6}$$

for all irreducible representations ν_1 of B, given in the system of Gel'fand-Tsetlin bases, for all characters ζ_1,ζ_2 of N occurring in the restriction of ν_1 to N, and for all $g \in \{1\}\cup(N\cdot w)$. Do this using $q+1$ Fourier transforms on B, one for each value of g.

Stage 2. Compute

$$f^2(\nu_1,\zeta_1,\zeta_2,\zeta_3) = \sum_{n\in N} f^1(n\cdot w,\nu_1,\zeta_1,\zeta_2)\zeta_3(n) \tag{5.7}$$

for all irreducible representations ν_1 of B, and all characters ζ_1,ζ_2,ζ_3 of N, where ζ_1,ζ_2 occur in the restriction of ν_1 to N. Do this using $q(q-1)$ Fourier transforms on N, one for each possible value of (ν_1,ζ_1,ζ_2).

Stage 3. Compute

$$[\hat{f}_w(\rho)]_{r,s} = \sum_{\zeta_2} [\rho(w)]_{r,s}\, f^2(\nu_1,\zeta_1,\zeta_2,\zeta_3) \tag{5.8}$$

for all irreducible representations ρ of $SL_2(F_q)$, and paths $r = 1\zeta_1\nu_1\rho$, $s = 1\zeta_3\nu_2\rho$ in the Bratteli diagram, starting at 1 and ending at ρ.

Stage 4. Compute

$$[\hat{f}(\rho)]_{r,s} = [\hat{f}_1(\rho)]_{r,s} + [\hat{f}_w(\rho)]_{r,s}, \tag{5.9}$$

where

$$[\hat{f}_1(\rho)]_{r,s} = \begin{cases} f^1(1,\nu_1,\zeta_1,\zeta_2) & \text{if } \nu_1 = \nu_2,\\ 0 & \text{if } \nu_1 \neq \nu_2. \end{cases}$$

The relations between the indices occurring in the algorithm are summarized in Figure 11, which shows the paths connecting the indices in the Bratteli diagram.

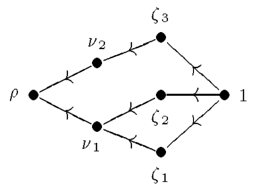

FIGURE 11. Relations between the indices, Fourier transform on $SL_2(F_q)$.

We now calculate the number of operations required by the algorithm. Stage 1 consists of $q+1$ Fourier transforms on B, and thus requires $(q+1)C_B$ operations. Stage 2 consists of $|B|$ Fourier transforms on N, and thus requires $C_N\,|B|$ operations.

Stage 3 requires as many operations as there are ways of choosing $\rho,\nu_1,\nu_2,\zeta_1,\zeta_2,\zeta_3$ subject to the relations indicated in Figure 11. By the proof of Theorem 4.8, this is

$$\sum_{\rho,\nu_1} \mathcal{M}(\rho,\nu_1)d_\rho d_{\nu_1}^2 = \sum_{\nu} d_{(\mathrm{Ind}_B^{SL_2(F_q)}\nu)} d_\nu^2 = (q+1)\sum_{\nu} d_\nu^3$$

where ν ranges over irreducible representations of B, and we have used Frobenius reciprocity to simplify the result; see [50] Chapter 7.2. The sum on the right hand side is easily found to be $\frac{1}{2}(q-1)(q^2-2q+3)$, so the number of operations for Stage 3 of the algorithm is $(q^2-1)(q^2-2q+3)$. This is the only place where the condition that q be odd is used. Stage 4 consists of fewer than $|SL_2(F_q)|$ additions, but these do not contribute to the overall operation count, because Stage 3 has exactly this many fewer additions than multiplications.

Adding up the operation counts for each of the stages we obtain

$$(q+1)C_B + (q^2-q)C_N + \frac{1}{2}(q^2-1)(q^2-2q+3)$$

Lemma 2.5 and Theorem 5.2 now imply the desired bounds. The first bound arises if we perform the Fourier transforms on N and T by a naive algorithm. The second arises when we use fast Fourier transform techniques to compute Fourier transforms on these groups. □

Theorem 5.4. *Let q be any power of 2. Working in a system of Gel'fand-Tsetlin bases, relative to $SL_2(F_q) > B > N > 1$, for the irreducible representations of $SL_2(F_q)$, the Fourier transform of a complex valued function can be computed using no more than*

$$q^2(q^2-1) + 4(q+\frac{1}{2})q(q-1)\log q + 8(q^2-1)\log(q-1)$$

scalar operations.

Proof. Following the proof of Theorem 5.3 word for word, we find that the number of operations required to compute the Fourier transform of a complex valued function f is no more than

$$(q+1)C_B + (q^2-q)C_N + (q+1)\sum_\nu d_\nu^3,$$

where ν ranges over irreducible representations of B. When q is even we have

$$\sum_\nu d_\nu^3 = (q-1)(q^2-2q+2)$$

We now use Theorem 5.2, Lemma 2.5, and the fact (cf. [55]) that when q is a power of two, the Fourier transform on N may be computed in under $2q\log q$ scalar operations, by a fast Hadamard-Walsh transform. The stated bound assumes that the Fourier transforms on T are computed using under $8(q-1)\log(q-1)$ scalar operations. For small values of q a naive algorithm may be more efficient. □

5.3. **The projective special linear group.** The projective special linear group is defined to be

$$PSL_2(F_q) = SL_2(F_q)/Z,$$

where Z is the center of $SL_2(F_q)$. For q odd, $Z = \{+1,-1\}$, but for q even, $Z = 1$. Thus when q is even, the special linear and projective special linear groups coincide. For the rest of this section we therefore assume that q is odd.

From the point of view of computational harmonic analysis the projective linear and special linear groups are virtually identical. The main changes are that T must be replaced by $\tilde{T} = T/Z$, and that the Bratteli diagram of $PSL_2(F_q)$, shown in

Figure 14, is simpler than that of $SL_2(F_q)$. The subgroup N embeds into $PSL_2(F_q)$ in a one-to-one fashion, and the Bruhat decomposition

$$PSL_2(F_q) = \tilde{B} \cup (N \cdot w \cdot \tilde{B}), \quad \tilde{B} = \tilde{T} \cdot N,$$

still holds provided we replace B by it's image $\tilde{B}$.

Given these changes, the proofs of Theorems 3.5, 3.6, 5.2, and 5.3 remain virtually unchanged. The final result for $PSL_2(F_q)$, q odd is given in Theorem 5.5 which follows.

Theorem 5.5. *Let q be an odd prime power. Working in a system of Gel'fand-Tsetlin bases, relative to $PSL_2(F_q) > \tilde{B} > N > 1$, for the irreducible representations of $PSL_2(F_q)$, the Fourier transform of a complex valued function can be computed using no more than*

$$(q-1)(q^3 + \frac{1}{2}q^2 + \frac{5}{4}q + \frac{5}{4})$$

scalar operations. Also, this transform may be computed using no more than

$$\frac{1}{4}(q^4 - 1) + 8(q + \frac{1}{2})q(q-1)\log q + 4(q^2-1)\log(\frac{q-1}{2})$$

scalar operations.

Proof. Following through the proof for $SL_2(F_q)$ word for word, we find that the transform requires no more than

$$\frac{1}{2}(q+1)C_{\tilde{B}} + \frac{1}{2}(q^2-q)C_N + \frac{1}{4}(q^2-1)(q^2-2q+3)$$

scalar operations. If we adapt the proof of Theorem 5.2 to $\tilde{B}$, we find that

$$C_{\tilde{B}} \le (q-1)C_N + C_{\tilde{T}} + (q-1)^2.$$

□

6. Conclusion

We have shown that the rather crude decomposition of a group into double cosets can lead to efficient algorithms for the computation of Fourier transforms. Better results can sometimes be obtained by using more refined decompositions, which factor group elements into more factors. This is treated in [42], but the techniques, though related, are considerably more complicated.

There are a number of questions raised by the current work. First, it would be interesting to find more examples of groups with small abelian subgroups and a small double coset space. Also, there is a lot of work to be done relating to icosahedral spherical harmonics; finding explicit formulae for a set of icosahedral spherical harmonics, and determining their other properties.

Another open problem is the problem of finding an algorithm for computing Fourier transforms on $SL_2(F_q)$ in $O(q^3(\log q)^c)$ operations, preferably with $c = 1$. It is apparent from the proof of Theorem 5.3 that for large q, the slowest part of the computation of Fourier transforms on $SL_2(F_q)$ occurs at Stage 3 of the algorithm. Indeed, Stages 1 and 2 may each be performed in $O(q^3 \log q)$ scalar operations, using fast Fourier transform techniques for abelian groups.

In the case of rotation groups, Stage 3 of the computation (corresponding to Stages 3 and 4 in Theorem 5.3) was performed efficiently using a fast orthogonal polynomial transform. This suggests that a more efficient algorithm for $SL_2(F_q)$

might follow from a better understanding of the matrix elements $[\rho(w)]_{rs}$ as special functions. It is also possible that a different set of double cosets, e.g., $T\backslash SL_2(F_q)/T$, might lead to a simplification at Stage 3.

Acknowledgements

We thank R. Leduc for alerting us to several references on Bratteli diagrams and their applications. We also thank W. Verkely for suggesting the problem of computing on icosahedral grids.

Appendix A. Representations of the special linear group.

The algorithms we have presented require certain values of the matrix elements of the group in question as precomputed data. In the case of the special linear group there are explicit formulae for these matrix elements, which we shall present in this section. First, we must be clear about what choices are necessary to describe the values of matrix elements on the subgroups of $SL_2(F_q)$.

A.1. The representations of the subgroups of $SL_2(F_q)$.

The characters of N. The characters of N may be indexed by elements of the dual space of F_q over its prime field $F_p = \mathbf{Z}/p\mathbf{Z}$. To obtain explicit values for the characters, we identify N with F_q, via $\begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix} \mapsto x$, and fix a basis $\{e_1, \ldots, e_d\}$ for F_q over $\mathbf{Z}/p\mathbf{Z}$, where $d = \dim_{F_p} F_q$. Notice that this is equivalent to fixing a root for a splitting polynomial for the extension, which in turn is equivalent to fixing a multiplicative generator for the (cyclic) multiplicative group $F_q^\times$. The dual of N, $\widehat{N}$, consisting of the characters of N, is isomorphic to N. With the above choices, a correspondence can be given via the map that associates to any row vector of length d, $(a_1, \ldots, a_d)$, each $a_i \in F_p$, a character $\zeta_{(a_1,\ldots,a_n)}$ of F_q according to the formula

$$\zeta_{(a_1,\ldots a_d)}(b_1 e_1 + \cdots + b_d e_d) = e^{2\pi i(a_1 b_1 + \cdots + a_d b_d)/p}$$

where $b_1, \ldots, b_d \in \mathbf{Z}/p\mathbf{Z}$ and $a_1 b_1 + \cdots a_d b_d \in \mathbf{Z}/p\mathbf{Z}$ is interpreted as an integer between 0 and $p-1$.

The characters of T. We identify T with the group $F_q^\times$ of nonzero elements of F_q under multiplication, associating the matrix t with it's upper left entry. Thus T is a cyclic group, and in order to obtain character values, we must choose a generator. Once we have a generator α the characters are constructed as in Section 2. Specifically, for each integer k with $0 \le k \le q-1$, define a character ψ_k by

$$\psi_k(\alpha^j) = e^{2\pi i jk/(q-1)}, \quad j = 0, 1, \ldots, q-1.$$

The representations of B. We have already given formulae for the representations of B in terms of characters of T and N (see Equation (5.1)). But these formulae still depend on calculations in F_q.

In particular, we need to consider an action of T on the characters of N which may be described as follows. Any element $\alpha^j \in T$ defines a linear operator on F_q via multiplication. Having chosen a particular basis for F_q over F_p, this assigns to α^j a particular matrix $[\alpha^j] \in GL_d(F_p)$. With this notation we now define the action of α^j on $\widehat{N}$ as

$$\alpha^j \zeta_{a_1,\ldots,a_d} = \zeta_{(a_1,\ldots,a_d)[\alpha^{-2j}]},$$

The orbit decomposition of $\widehat{N}$ under this action is easy to understand. This action always fixes the trivial character, which is thus always in an orbit alone. For $j \neq k$, and fixed $(a_1,\ldots,a_d) \in (\mathbf{Z}/p\mathbf{Z})^d$ we have $\zeta_{(a_1,\ldots,a_d)[\alpha^j]} \neq \zeta_{(a_1,\ldots,a_d)[\alpha^k]}$. Thus, when q is even, since the map $\alpha^j \mapsto \alpha^{2j}$ is an automorphism of T (the Frobenius automorphism) we get only one other orbit in $\widehat{N}$ which we denote by X. However, when q is odd the squares make up a subgroup of index 2 in F_q^X so that we get two distinct nontrivial orbits in $\widehat{N}$ which we denote by X' and X''.

When q is odd, we must also choose the representation ψ^-, which may be chosen to be ψ_j for any odd j. For definiteness we can choose $\psi^- = \psi_1$ but when $\frac{q-1}{2}$ is odd, it is more convenient to choose $\psi^- = \psi_{\frac{q-1}{2}} = \mathrm{sgn}$.

It is useful to describe here the change of basis matrices from $\nu^{X,\psi}$ to $\nu^{X,\gamma}$, where X is any nontrivial orbit of the action of T on $\hat{N}$, and ψ, γ are characters of T such that $\psi(-1) = \gamma(-1)$. These matrices are diagonal, so that the explicit change of basis may be written as

$$\left[\nu^{X,\gamma}\right]_{\zeta_1,\zeta_2} = f(\zeta_1)^{-1} f(\zeta_2) \nu^{X,\psi}, \quad \zeta_1, \zeta_2 \in X.$$

This corresponds to multiplication of the basis vectors $\nu^{X,\psi}$ by the function $f(\zeta)$. To define a suitable f, first make a choice of a character ζ_0 in the orbit X. Then we may set

$$f(\zeta) = \psi(t)\gamma(t)^{-1}, \quad \text{where } t \in T, \text{ and } \zeta = \zeta_0^t. \tag{A.1}$$

The condition that $\psi(-1) = \gamma(-1)$ implies that f is well-defined by Equation (A.1).

A.2. **The representations of $SL_2(F_q)$.** The representations of $SL_2(F_q)$ fall naturally into two classes. The principal series representations are those representations whose restriction to N contains the trivial representation. The discrete series representations are those representations whose restriction does not contain the trivial representation. Our discussion derives much from the presentation of [44].

A.2.1. *The discrete series.* The discrete series representations are defined using the action of $F_q^\times$ on the characters of N given by

$$(\beta\zeta)(n) = \zeta(\beta n), \quad \beta \in F_q^\times, \zeta \in \hat{N}, n \in N.$$

Of course this determines an action of T on $\widehat{N}$ by setting $t\zeta = \beta\zeta$, where β is the top left entry of t. This is related to the conjugation action of T defined in Section 5.

Let φ be a character of the cyclic group $F_{q^2}^\times$. Then we define the representation D_φ to be the matrix representation of $SL_2(F_q)$ with rows and columns indexed by the nontrivial characters of $\hat{N}$, such that

$$\begin{aligned} \left[D_\varphi(t \cdot n)\right]_{\zeta_1,\zeta_2} &= \delta_{\zeta_1,\zeta_2^t} \varphi(t) \zeta_2(n) \\ \left[D_\varphi(w)\right]_{\zeta_1,\zeta_2} &= \frac{1}{q} \sum_{u^{q+1}\zeta_2 = \zeta_1} \zeta_2(u + u^q)\varphi(u), \end{aligned}$$

for any $n \in N$, $t \in T$. Here we have used the identification of T with $F_q^\times \subseteq F_{q^2}^\times$ to restrict φ to T when necessary.

Let $C = \{c \in F_{q^2}^\times | c^{q+1} = 1\}$ denote the subgroup of $F_{q^2}^\times$ of order $q+1$. In the case when q is odd, it is possible that $\varphi^2|C = 1_C$ but $\varphi|C \neq 1_C$. In this case D_φ is the matrix direct sum of representations which are obtained by restricting the range of the indices ζ_1, ζ_2 to the same orbit of the action $(t, \zeta) \mapsto \zeta^t$. We denote the representations corresponding to the orbits X', X'' by D'_φ, D''_φ, respectively, though we will sometimes omit the reference to φ.

Theorem A.1 ([44] pp. 35–40). 1. *The representations D_φ, D'_φ, D''_φ are well defined.*
2. *When $\varphi^2|C \neq 1_C$, the representation D_φ is irreducible. When $\varphi^2|C = 1_C$ and $\varphi|C \neq 1_C$, the representations D'_φ, D''_φ are irreducible.*
3. *$D_\varphi, D'_\varphi, D''_\varphi$ are determined up to equivalence by the restriction $\varphi|C$.*
4. *When q is odd, the representations $D_\varphi, D'_\varphi, D''_\varphi$, include all the discrete series representations of $SL_2(F_q)$.*
5. *When q is even, the representations D_φ, $\varphi^2|C \neq 1_C$, include all the discrete series representations of $SL_2(F_q)$.*
6. *$D_\varphi|B = \nu^{X',\varphi|T} \oplus \nu^{X'',\varphi|T}$, $D'_\varphi|B = \nu^{X',\varphi|B}$, and $D''_\varphi|B = \nu^{X'',\varphi|B}$.*

To obtain a set of discrete series representations in a system of Gel'fand-Tsetlin bases we must ensure that the only matrix representations occurring in the restrictions to B are $\nu^{X',1}, \nu^{X'',1}, \nu^{X',\psi^-}$ and ν^{X'',ψ^-}, when q is odd, and $\nu^{X,1}$ when ν is even. This means choosing a complete set of coset representatives for $\widehat{F_{q^2}^\times}/(\widehat{F_{q^2}^\times/C})$ whose restrictions to T are each either 1 or ψ^-. This is always possible, because the intersection of T and C inside $F_{q^2}^\times$ is exactly $Z = \{x | x^2 = 1\}$. Thus, if ψ is any character of T and $\bar{\varphi}$ is any character of C, and $\psi|Z = \bar{\varphi}|Z$, then there is a unique character of TC that restricts to ψ on T and $\bar{\varphi}$ on C. This character may then be extended to $F_{q^2}^\times$.

For example, let q be odd, let η be a generator of $F_{q^2}^\times$, and let $\alpha = \eta^{q+1}$ be the generator for $F_q^\times$. Define $\varphi_k(\eta^j) = e^{2\pi i jk/(q^2-1)}$. Let $\psi^- = \psi_1$. Then the characters

$$\varphi_{(q-1)l}, \quad \varphi_{(q-1)l+1}, \quad l = 0, 1, \ldots, \frac{q+1}{2}$$

restrict to the $q+1$ distinct characters of C, and to either 1 or ψ^- on T. We may therefore use this choice of the characters φ to obtain the discrete series in a system of Gel'fand-Tsetlin bases.

A.2.2. *The principal series.* The principal series representations are usually defined by their action on the space of functions on the projective line over F_q. Any character of N, may be extended to the projective line by setting it equal to zero at infinity. A basis for the space of functions may then be constructed by taking all characters of N, and one additional function, a point mass at infinity, which we denote by e_∞. It turns out to be slightly easier to specify the principal series representations by their action on this basis, rather than giving formulae directly for the matrices.

Let ψ be any character of T. Then the representation U_ψ is defined on the basis $\hat{N} \cup \{e_\infty\}$ by

$$(A.2)\qquad \begin{array}{lcl} U_\psi(t\cdot n)\zeta &=& \psi^{-1}(t)\zeta^t \\ U_\psi(t\cdot n)e_\infty &=& \psi(t)e_\infty \\ U_\psi(w)\zeta &=& e_\infty + \sum_{\chi\in\hat{N}} \frac{1}{q}\left[\sum_{x\in F_q^\times} \psi(x)\zeta(x)\chi(x^{-1})\right]\chi \\ U_\psi(w)e_\infty &=& \frac{\psi(-1)}{q}\sum_{\chi\in\hat{N}}\chi \end{array}$$

We consider U_ψ as a matrix representation, with rows and columns indexed by $\hat{N}\cup\{e_\infty\}$.

Lemma A.2. *Let ψ be a character of T.*

1. *The representation U_ψ is well defined.*
2. *If $\psi^2 \neq 1$, then U_ψ is irreducible.*
3. *When q is odd,*
$$U_\psi|B = \nu^{X',\psi} \oplus \nu^{X'',\psi} \oplus \nu^{\{1\},\psi^{-1}} \oplus \nu^{\{1\},\psi},$$
where $\nu^{\{1\},\psi^{-1}}$ occurs in the diagonal entry of the trivial representation of N, and $\nu^{\{1\},\psi}$ occurs in the diagonal entry corresponding to e_∞.
4. *When q is even,*
$$U_\psi|B = \nu^{X,\psi} \oplus \nu^{\{1\},\psi^{-1}} \oplus \nu^{\{1\},\psi}.$$

When $\psi^2 = 1$, the representation U_ψ is no longer irreducible. When $\psi = 1$ the representation U_ψ leaves $e_\infty + 1_N$ fixed, where 1_N denotes the trivial character of N. The complementary subspace spanned by $e_\infty - \frac{1}{q}1_N$ and the nontrivial characters of N is an invariant subspace of U_1, and we denote the matrix representation obtained on this subspace in the given basis by V.

When q is odd, there is a unique nontrivial character sgn of T satisfying $\mathrm{sgn}^2 = 1$. Then define the vectors e', e'', by

$$(A.3)\qquad \begin{aligned} e' &= e_\infty + \frac{1}{q}\sum_{\zeta\in X'}(\zeta(x_0^2) - \zeta(x_0))1_N \\ e'' &= e_\infty + \frac{1}{q}\sum_{\zeta\in X''}(\zeta(x_0^2) - \zeta(x_0))1_N, \end{aligned}$$

where x_0 is any non-square in $F_q^\times$. Note that e' and e'' are independent of the choice of x_0. The subspaces spanned by $X'\cup\{e'\}$ and $X''\cup\{e''\}$ are complementary subspaces, each invariant under U_ψ. We define U', U'' to be the matrix representations obtained by expressing U_{sgn} on each of these subspaces, in the given bases.

For the sake of completeness we include the evaluation of the inner sums in (A.3). These may be realized as well-known Gauss sums and thus have the following simple expression.

Lemma A.3 ([37], Theorem 5.15, p. 199)**.** *Let q be odd, and let X be a nontrivial orbit of the action $(t,\zeta)\mapsto\zeta^t$ on the characters of F^q. Let χ denote the character of F_q defined by $\chi(\beta) = e^{2\pi i\,\mathrm{Tr}_{F_p}(\beta)/p}$, where F_p denotes the prime field, and $\mathrm{Tr}_{F_p}(\beta)$ is the trace of the F_p-linear map of multiplication by β. Then for any non-square x_0 in F_q, we have*

$$\frac{1}{q}\sum_{\zeta\in X}(\zeta(x_0^2) - \zeta(x_0)) = \begin{cases} (-1)^{d-1}\sqrt{q} & \text{if } p\equiv 1 \bmod 4 \\ (-1)^{d-1}i^d\sqrt{q} & \text{if } p\equiv 3 \bmod 4. \end{cases}$$

if $\chi \in X$ and minus this if $\chi \notin X$.

Theorem A.4. *The principal series representations of $SL_2(F_q)$, up to equivalence are U_ψ, $\psi \in \hat{T}, \psi^2 \neq 1$, and $1, V$, when q is even, and these representation together with U', U'', when q is odd. The matrix representations V, U', U'' are always irreducible, and their restrictions to B are*

$$\begin{aligned} V|B &= \nu^{\{1\},1} \oplus \nu^{X',1} \oplus \nu^{X'',1} \\ U'|B &= \nu^{\{1\},\mathrm{sgn}} \oplus \nu^{X',\mathrm{sgn}} \qquad \textit{and} \\ U''|B &= \nu^{\{1\},\mathrm{sgn}} \oplus \nu^{X'',\mathrm{sgn}} \end{aligned}$$

Note that when q is odd, $\mathrm{sgn}\,|Z$ is 1 when $\frac{q-1}{2}$ is odd, and -1 when $\frac{q-1}{2}$ is even.

We have constructed matrix representations for the principal series of $SL_2(F_q)$, which restrict down to representations of the form $\nu^{X,\psi}$ on B. To have a Gel'fand-Tsetlin basis, however, we must ensure that ψ is either 1 or ψ^-. This requires a further change of basis that we now describe.

It suffices to rescale the basis vectors $\zeta \in \hat{N} \setminus \{1_N\}$, so that for any nontrivial orbit X of T on $\hat{N}$ the representation $\nu^{X,\psi}$ is mapped onto either $\nu^{X,1}$ or ν^{X,ψ^-}. This is achieved if we multiply the character ζ by the function $f(\zeta)$ defined in Equation (A.1) with $\gamma = 1$ or $\gamma = \psi^-$.

We now summarize the information we have gained about the representations of $SL_2(F_q)$. Table 2 gives a list of the representations and their degrees. Figures 13 and 12 show the Bratteli diagrams of $SL_2(F_q) > B > N > 1$ in the cases q even and odd respectively. Formulae for the representations of $PSL_2(F_q)$ follow immediately from the formulae for $SL_2(F_q)$. We also included the corresponding Bratteli diagram for $PSL_2(F_q)$, q odd, in Figure 14.

Rep	U_ψ	1	V	U'	U''	D_φ	D'	D''
Dim	$q+1$	1	q	$\frac{q+1}{2}$	$\frac{q+1}{2}$	$q-1$	$\frac{q-1}{2}$	$\frac{q-1}{2}$
Number	$\lfloor\frac{q-2}{2}\rfloor$	1	1	0 or 1	0 or 1	$\lfloor\frac{q}{2}\rfloor$	0 or 1	0 or 1

TABLE 2. The representation of $SL_2(F_q)$. "0 or 1" means 0 if q odd, 1 if q even.

APPENDIX B. THE CONSTRUCTION OF GEL'FAND-TSETLIN BASES

We now sketch the arguments proving the existence of Gel'fand-Tsetlin bases.

Proof of Lemma 4.6. The argument is by induction on the number of subgroups in the chain. Clearly the lemma holds for $n = 0$ and $n = 1$. Now suppose that $G = K_n \geq K_{n-1} \geq \cdots \geq K_0 = 1$ is a chain of subgroups. By induction we may assume that we have constructed a system of Gel'fand-Tsetlin bases for the chain $K_{n-1} \geq \cdots \geq K_0$. Let $\mathcal{R}_{n-1}$ denote the complete set of inequivalent irreducible matrix representations of K_{n-1} corresponding to this system.

Suppose we are given a complete set $\mathcal{R}_n$ of inequivalent irreducible matrix representations of G. Consider any $\rho \in \mathcal{R}_n$. Then $\rho|K_{n-1}$ is a matrix representation of K_{n-1} which is equivalent to a matrix direct sum of representations from $\mathcal{R}_{n-1}$. Thus, there is a change of basis matrix E_ρ, such that

$$\text{(B.1)} \qquad E_\rho \rho(k_{n-1}) E_\rho^{-1} = \mu_1(k_{n-1}) \oplus \cdots \oplus \mu_l(k_{n-1}), \quad \text{for all } k_{n-1} \in K_{n-1},$$

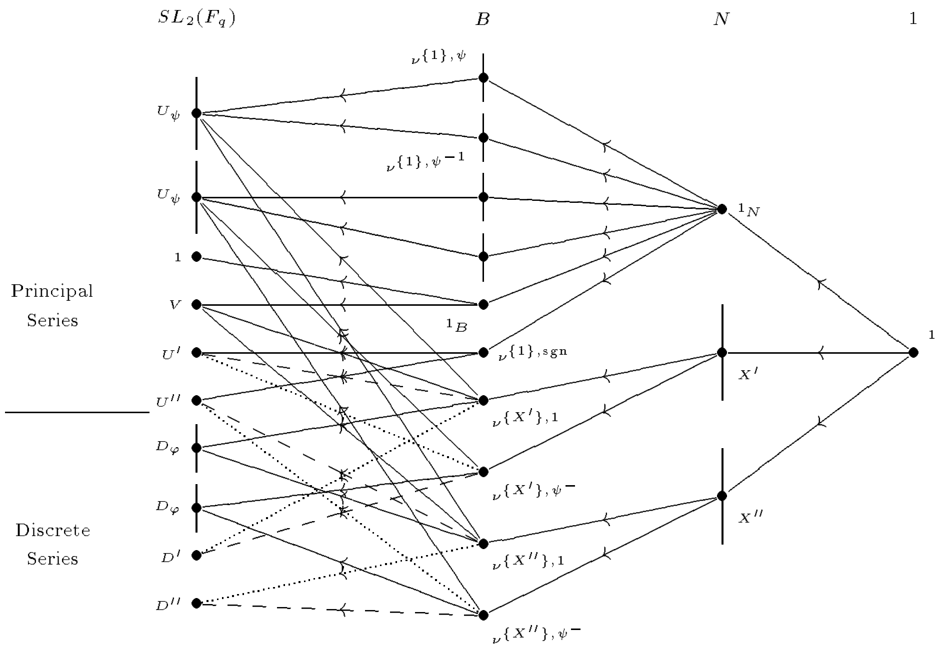

FIGURE 12. The Bratteli diagram for $SL_2(F_q) > B > N > 1$, q odd. Dashed lines are present when $\frac{q-1}{2}$ is even; dotted lines when $\frac{q-1}{2}$ odd.

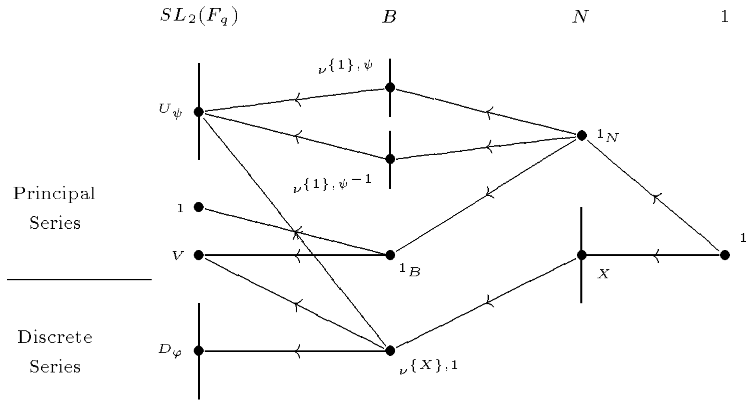

FIGURE 13. The Bratteli diagram for $SL_2(F_q) > B > N > 1$, q even. Dashed lines are present when $\frac{q-1}{2}$ is even; dotted lines when $\frac{q-1}{2}$ odd.

where $\mu_1, \ldots, \mu_l \in \mathcal{R}_{n-1}$, possibly with repeats among the μ_i. This allows us to collect the standard basis vectors of $\mathbf{C}^{d_\rho}$ into l 'blocks' according to the structure of the block diagonal matrices $\mu_1 \oplus \cdots \oplus \mu_l$

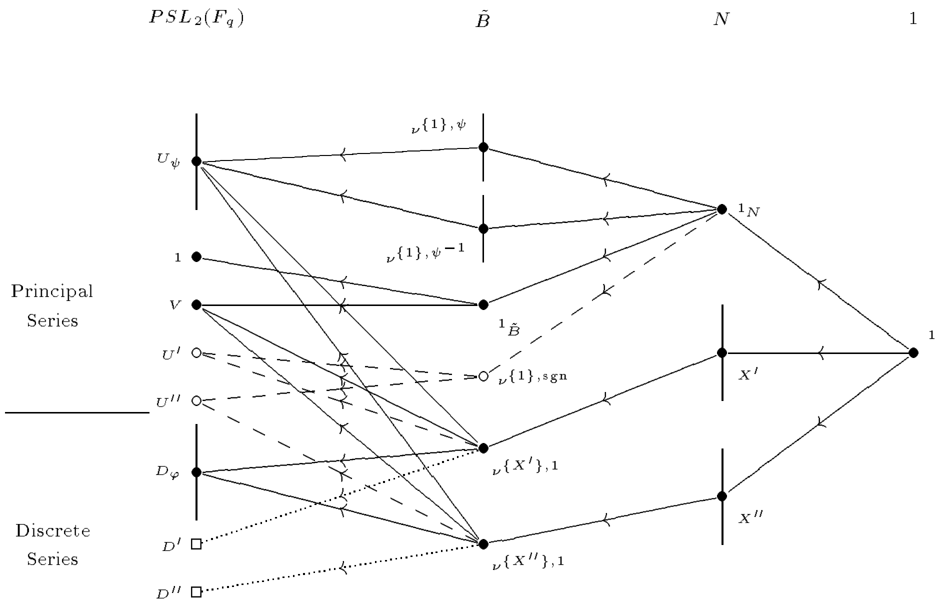

FIGURE 14. The Bratteli diagram for $PSL_2(F_q) > B > N > 1$, q odd. Dashed lines and hollow circles are present when $\frac{q-1}{2}$ is even; dotted lines and squares when $\frac{q-1}{2}$ odd. Here, $\psi|Z = 1$ and $\varphi|Z = 1$.

We now index the standard basis vectors of $\mathbf{C}^{d_\rho}$ by paths in the Bratteli diagram for $K_n \geq \cdots \geq K_0$, as follows. By construction of the Bratteli diagram, the blocks $\mu_1, \ldots, \mu_l$ in the matrix direct sum may be put into one-to-one correspondence with the edges between level $n-1$ and level n that end at ρ. For any such edge r_n, let μ_{r_n} denote the block in the matrix sum corresponding to the edge r_n. Then, if $r_1 \cdots r_{n-1} r_n$ is a path in the Bratteli diagram starting at level 0, and ending at ρ, we let $e_{r_1 \cdots r_n}$ denote the standard basis vector in the block for μ_{r_n} which is indexed by $r_1 \cdots r_{n-1}$ inside that block, using the indexing of the rows or columns of μ_{r_n} by paths in the Bratteli diagram of $K_{n-1} \geq \cdots \geq K_0$.

The vectors of a system of Gel'fand-Tsetlin bases for the set of matrix representations $\mathcal{R}_n$ are now the basis vectors we chose that put the restrictions of the representations of $\mathcal{R}$ in the block diagonal form (B.1). Thus, if $r_1 \cdots r_n$ is any path in the Bratteli diagram, then $E_{\rho^{r_n}}^{-1} e_{r_1 \cdots r_n}$ is the Gel'fand basis vector corresponding to that path, where ρ^{r_n} is the representation in $\mathcal{R}$ corresponding to the vertex at the tip of r_n.

When the Bratteli diagram has no multiple edges, the subspaces corresponding to the decomposition of $\rho|K_{n-1}$ into irreducibles are uniquely determined (cf. [50] Chapter 2.6). Any set of Gel'fand-Tsetlin basis vectors for ρ will lie in the union of these subspaces, and define a set of bases for these subspaces which are part of a system of Gel'fand-Tsetlin bases for K_{n-1}. But by induction the Gel'fand-Tsetlin bases for each of these subspaces as representations of K_{n-1} are uniquely determined up to scalar multiples, since the Bratteli diagram of $K_{n-1} \geq \cdots \geq K_0$ is also without multiple edges. This proves the uniqueness of the system of Gel'fand-Tsetlin bases up to scalar multiples. □

Proof of Lemma 4.7. In case when $H \geq K$ or $K \geq H$, Lemma 4.7 is essentially equivalent to the existence of a system of Gel'fand-Tsetlin bases relative either $G \geq H \geq K \geq 1$, or $G \geq K \geq H \geq 1$.

For example, when $H \geq K$, we may first construct a set of Gel'fand-Tsetlin bases for the chain $G \geq H \geq K \geq 1$. Now, construct a bijection between the paths in the Bratteli diagram of $G \geq H \geq K \geq 1$ and $G \geq K \geq 1$, by associating each path between the K-level and G-level in the first Bratteli diagram with an edge between the K-level and G-level in the second Bratteli diagram. In a similar fashion, we can construct a bijection between the paths in the Bratteli diagram of $G \geq H \geq K \geq 1$ and paths in the diagram of $G \geq H \geq 1$ by associating paths between the 1-level and H-level in the first diagram with edges between the 1-level and H-level in the second diagram.

Thus, the system of Gel'fand-Tsetlin bases for $G \geq H \geq K \geq 1$, may be indexed by paths in any of the three Bratteli diagrams. In each case we it forms a system of Gel'fand-Tsetlin bases for the corresponding chain of subgroups.

The case when all elements of H commute with all elements of K must be treated differently. We start with complete sets of inequivalent irreducible matrix representations of H and K, denoted $\mathcal{R}_H$, and $\mathcal{R}_K$ respectively. Next we form the set of tensor product representations

$$\mathcal{R}_{HK} = \{\mu \otimes \nu | \mu \in \mathcal{R}_H, \nu \in \mathcal{R}_K, \text{ and } \mu(c) \otimes \nu(c^{-1}) = I, \text{ for all } c \in H \cap K.\}$$

Then $\mathcal{R}_{HK}$ is a complete set of inequivalent irreducible matrix representations of HK which is in a system of Gel'fand-Tsetlin bases for $HK \geq K \geq 1$ and $HK \geq K \geq 1$. When can now use the argument of the proof of Lemma 4.6 to construct simultaneous Gel'fand-Tsetlin bases for the two chains $G \geq HK \geq K \geq 1$, and $G \geq HK \geq H \geq 1$. The result now follows by the argument given for the first part of the lemma. □

The arguments used in the above proofs are essentially due to Gel'fand and Tsetlin [29], though the most general formulation is due to Želobenko [56].

We have the existence of systems of Gel'fand-Tsetlin bases, but this still falls short of a practical algorithm for finding such bases. The difficult part is the construction of the change of basis matrix E, and the construction of a complete set of irreducible representations in the first place. We refer the reader to [3, 2, 27] for discussions of various aspects of these problems.

References

[1] L. Auslander and R. Tolmieri, *Is computing with the fast Fourier transform pure or applied mathematics?*, Bull. Amer. Math. Soc. (N.S.) **1** (1979), 847–897.

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

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

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

[5] U. Baum and M. Clausen, *Computing irreducible representations of supersolvable groups*, Math. Comp. **63** (1994), 351–359.

[6] ———, *Some lower and upper complexity bounds for generalized Fourier transforms and their inverses*, SIAM J. Comput. **20**(3) (1991), 451–459.

[7] U. Baum, M. Clausen, and B. Tietz, *Improved upper complexity bounds for the discrete Fourier transform*, AAECC **2** (1991), 35–43.
[8] T. Beth, *Verfahren der schnellen Fourier–Transformation*, Teubner Studienbücher, Stuttgart, 1984.
[9] L. Bluestein, *A linear filtering approach to the computation of the discrete Fourier transform*, IEEE Trans. **AU-18** (1970), 451–455.
[10] O. Bratteli, *Inductive limits of finite dimensional C^*-algebras*, Trans. Amer. Math. Soc. **171** (1972), 195–234.
[11] P. Bürgisser, M. Clausen, A. Shokrollahi, *Algebraic Complexity Theory*, Springer-Verlag, Berlin, 1996.
[12] M. Clausen, *Fast Fourier transforms for metabelian groups*, SIAM J. Comput. **18** (1989), 584–593.
[13] ——— *Fast generalized Fourier transforms*, Theoret. Comput. Sci. **67** (1989), 55–63.
[14] ———, *Beiträge zum Entwurf schneller Spektraltransformationen*, Habilitationsschrift, Fakultät für Informatik der Universität Karlsruhe (TH), 1988.
[15] M. Clausen and U. Baum, *Fast Fourier transforms*, Wissenschaftsverlag, Mannheim, 1993.
[16] ———, *Fast Fourier transforms for symmetric groups, theory and implementation*, Math. Comp. **61**(204) (1993), 833–847.
[17] J. W. Cooley and J. W. Tukey, *An algorithm for machine calculation of complex Fourier series*, Math. Comp. **19** (1965), 297–301.
[18] P. Diaconis, *Applications of group theory to statistical problems*, ICM-90 Sattel. Conf. Proc., Springer-Verlag, Tokyo, 1991, 1037–1048.
[19] ———, *A generalization of spectral analysis with applications to ranked data*, Ann. Stat. **17** (1989), 949–979.
[20] ———, *Group representations in probability and statistics*, IMS, Hayward, CA, 1988.
[21] ———, *Average running time of the fast Fourier transform*, J. Algorithms **1** (1980), 187–208.
[22] P. Diaconis and D. Rockmore, *Efficient computation of the Fourier transform on finite groups*, J. Amer. Math. Soc. **3**(2) (1990), 297–332.
[23] ———, *Efficient computation of isotypic projections for the symmetric group*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. **11**, L. Finkelstein and W. Kantor (eds.), 1993, 87–104.
[24] J. Driscoll and D. Healy, *Computing Fourier transforms and convolutions on the 2-sphere*, (extended abstract) Proc. 34^{th} IEEE FOCS, (1989) 344–349; Adv. in Appl. Math., **15** (1994), 202–250.
[25] J. R. Driscoll, D. Healy, and D. Rockmore. *Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs*, SIAM J. Comput. **26**(4) (1997), 1066–1099.
[26] D. Elliott and K. 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] K. Gatermann, *Linear representations of finite groups and the ideal theoretical construction of G-invariant cubature formulas*, Numerical Integration, Nato Adv. Sci. Inst. Ser. C Math. Phys. Sci. **357**, Kluwer Acad. Publ., Dordrecht, 1992, 25–35.
[29] I. Gel'fand and M. Tsetlin, *Finite dimensional representations of the group of unimodular matrices*, Dokl. Akad. Nauk SSSR **71** (1950), 825–828 (Russian).
[30] J. Goethals and J. Seidel, *Cubature formulae, polytopes, and spherical designs*, The Geometric vein, Springer, New York, 1981, 203–218.
[31] R. Heikes and D. A. Randall, *Numerical integration of the shallow water equations on a twisted icosahedral grid. Part I: Basic design and results of tests*, Month. Weath. Rev. **123** (1995), 1862–1880.
[32] ———, *Numerical integration of the shallow water equations on a twisted icosahedral grid. Part II: A detailed description of the grid and an analysis of numerical accuracy*, Month. Weath. Rev. **123** (1995), 1881–1887.
[33] G. James and A. Kerber, *The representation theory of the symmetric group*, Encyclopedia of Mathematics and its Applications, vol. 16, Addison-Wesley, Reading MA, 1981.
[34] S. Konjaev, *Gaussian quadratures for the sphere that are invariant with respect to the icosahedral group with inversion*, Mat. Zametski **25** no. 4 (1979), 629–634,638.

[35] J. Lafferty and D. Rockmore, *Spectral techniques for expander codes*, Proc. Symp. Theory of Computing (1997), 160–167.
[36] ______ *Fast Fourier analysis for SL_2 over a finite field and related numerical experiments*, J. Exp. Math. **1** (1992), 115–139.
[37] R. Lidl and H. Neiderreiter, *Finite fields*, Encyclopedia of Mathematics and its Applications, vol. 20, Addison-Wesley, Reading MA, 1983.
[38] D. Maslen, *The efficient computation of Fourier transforms on the symmetric group*, preprint IHES/M/96/52, and Math. Comp. (to appear).
[39] ______, *Efficient computation of Fourier transforms on compact groups*, Max-Planck preprint MPI/95-8 (1995).
[40] D. Maslen and D. Rockmore *Generalized FFTs - a survey of some recent results*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Groups and Computation, II, L. Finkelstein and W. Kantor (eds.), (1997), 183–237.
[41] ______, *Separation of variables and the efficient computation of Fourier transforms on finite groups, I*, J. Amer. Math. Soc. **10** (1) (1997), 169–214.
[42] ______, *Separation of variables and the efficient computation of Fourier transforms on finite groups, II*, (in preparation).
[43] I. Mysovskih, *Cubature formulas that are invariant with respect to groups of transformations*, Metody Vychisl. no. 11 (1978), 2–21, 187.
[44] I. Piatetski-Shapiro, *Complex representations of $GL[2,K]$ for finite fields K*, Contemporary Mathematics, Vol. 16, American Math. Soc., Providence, RI, 1983.
[45] L. Rabiner, R. Schafer, and C. Rader, *The chirp-z transform and its applications*, Bell System Tech. J. **48** (1969), 1249-1292.
[46] C. Rader, *Discrete Fourier transforms when the number of data samples is prime*, IEEE Proc. **56** (1968), 1107–1108.
[47] D. Rockmore, *Applications of generalized FFTs*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Groups and Computation, II, L. Finkelstein and W. Kantor (eds.), (1997), 329–369.
[48] ______ *Fast Fourier transforms for wreath products*, Appl. and Comp. Harm. Analysis **4** (1995), 279–292.
[49] ______ *Fast Fourier analysis for abelian group extensions*, Adv. in Appl. Math. **11** (1990), 164–204.
[50] J. P. Serre, *Linear representations of finite groups*, Springer-Verlag, New York, 1977.
[51] S. Sobolev, *The formlas of mechanical cubature on the surface of a sphere*, Sibirsk Mat. Zh. **3** (1962), 769–796.
[52] N. Vilenkin, *Special functions and the theory of group representations*, Translations of Mathematical Monographs, **22**, A.M.S., Providence RI, 1968.
[53] A. Willsky, *On the algebraic structure of certain partially observable finite-state Markov processes*, Inform. Contr. **38**, (1978), 179–212.
[54] S. Winograd, *Arithmetic complexity of computations*, SIAM CBMS-NSF Series **33** SIAM, Philadelphia PA, 1980.
[55] F. Yates, *The design and analysis of factorial experiments*, Imp. Bur. Soil Sci. Tech. Comm. **35** (1937).
[56] D. Želobenko, *Compact Lie groups and their representations*, Translations of Mathematical Monographs, **40**, A.M.S., Providence RI, 1973.

Department of Mathematics, Universiteit Utrecht, 3584 CD, Utrecht, Netherlands
Current address: Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098 SJ, Amsterdam, THE NETHERLANDS
E-mail address: maslen@cwi.nl

Dept. of Mathematics, Dartmouth College, Hanover, NH 03755
E-mail address: rockmore@cs.dartmouth.edu