Fast Fourier Analysis for Abelian Group Extensions

Daniel Rockmore*
Harvard University
Department of Mathematics

July 11, 1995

Abstract

Let G be a finite group and f any complex-valued function defined on G and ρ an irreducible complex matrix representation of G. The Fourier transform of f at ρ is defined to be the matrix $\sum_{s \in G} f(s)\rho(s)$. The Fourier transforms of f at all the irreducible representations of G determine f via the Fourier inversion formula $f(s) = \frac{1}{|G|}\sum_{\rho} d_{\rho} trace(\widehat{f}(\rho)\rho(s^{-1}))$. Direct computation of all Fourier transforms of f involves on the order of $|G|^2$ operations as does direct computation of Fourier inversion. Here fast algorithms are obtained for both operations in the case in which G contains some nontrivial normal subgroup K such that G/K is abelian. Consequently, fast algorithms for computing convolutions on G in this situation are also determined. Under the simplifiying assumption of exponent for matrix multiplication equal to 2 (it is 2.38 as of this writing) it is shown that the number of operations needed to compute all Fourier transforms on G is $O(\frac{|G|}{|K|}T(K) + |G| \log(\frac{|G|}{|K|}))$ where $T(K)$ is the number of operations needed to compute Fourier transforms on K. An analogous result is obtained for Fourier inversion and more careful estimates made for arbitrary exponents of matrix multiplication. In particular, in the case in which G is metabelian (an abelian extension of an abelian subgroup) the assumptions hold and the first large class of noncommutative groups is obtained for which Fourier inversion and the computation of all Fourier transforms may be performed in $O(|G| \log(|G|))$ operations.

*Supported by IBM and NSF Graduate Fellowships.

1. Introduction

Let G be a finite group, f a complex-valued function on G and ρ any complex irreducible matrix representation of G. Then the **Fourier transform** of f with respect to ρ is the matrix

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

The Fourier transforms at all the irreducible representations of G determine f via the **Fourier inversion** formula

$$f(s) = \frac{1}{|G|} \sum_{\rho} d_{\rho} trace(\widehat{f}(\rho)\rho(s^{-1})$$

where the sum is over all irreducible representations ρ of G.

Let $T(G)$ denote the number of operations needed to compute all Fourier transforms on G and $I(G)$ the number needed to perform Fourier inversion. Direct computation of all Fourier transforms would yield $T(G) = |G|^2$. Similarly, naive computation of Fourier inversion gives $I(G) = |G|^2$. For groups of large order this cost is prohibitive. Fast algorithms must be derived.

Motivated partly by their manifold applications, much is known for the case in which G is abelian. A large class of algorithms, collectively known as the "Fast Fourier Transform" ("FFT") have been developed. It is now a well known result that for abelian groups both $T(G)$ and $I(G)$ are $O(|G| \log(|G|))$. Over the years these fast algorithms have yielded a diversity of important applications. Aho, Hopcroft, and Ullman give a thorough treatment of the basic algorithm and its uses ([AHU], Chapters 7 and 8).

The problem for general finite groups has only recently begun to attract attention. In the most general setting, in joint work with Diaconis ([DR]), prompted by the study of spectral analysis for finite groups ([D1], [D2]), an algorithm for efficient computation of Fourier transforms for arbitrary groups has already been presented. There it is shown that great savings may be achieved by iterating a simple recurrence which allows the Fourier transforms of f to be written in terms of Fourier transforms on a subgroup of G. This is in fact the idea at the heart of many of the fast algorithms in the abelian case.

The efficacy of the methods in [DR] depends strongly on the exponent of matrix multiplication and the existence of representations of G which "split" when restricted down a given tower of subgroups of G. In particular, for the symmetric group S_n this situation exists and practical implementations for analyzing ranked data are possible and useful ([R1]).

Fourier analysis is also useful in the study of random walks on groups. ([D2], [DS]). A probability measure defined on the group gives rise to a random walk on the group. Explicit knowledge of the matrices for the Fourier transforms determined by the measure are useful in studying the rate at which the walk becomes random.

A problem left open in much of the work on fast algorithms for Fourier analysis is the efficient computation of Fourier inversion. In the case of abelian groups the group structure on the dual $\widehat{G}$ allows the techniques for computing Fourier transforms to be applied. In general there is no such structure and new methods must be developed.

In this paper the case of solvable groups is considered, and more generally, the case of abelian group extensions. Here, modifications of the ideas in [DR] combined with the use of induced and projective representations yield fast algorithms for both Fourier inversion and computing Fourier transforms. One main result is

Theorem 1 *Let G be a finite group containing K as a normal subgroup with quotient G/K an abelian group. Let $\{\eta_i\}_{i\in I}$ be a complete set of representatives for the orbits of the irreducible representations of K under the action of G. Denote the orbit of η as $\Delta(\eta)$. Let H_{η_i} be a subgroup of G (containing K) such that η_i extends to a representation of H_{η_i} but no further. Then $T(G)$ is bounded by*

$$\frac{|G|}{|K|}\cdot T(K)+\sum_{i\in I}\left(\frac{|H_{\eta_i}|}{|H|}\left(\frac{|G|}{|H_{\eta_i}|}\right)^2\cdot d^{\alpha}_{\eta_i}+O\left(\frac{|G|}{|K|}\cdot|\Delta(\eta_i)|\cdot\left(d^{\alpha}_{\eta_i}+d^2_{\eta_i}\log\left(\frac{|H_{\eta_i}|}{|K|}\right)\right)\right)\right)$$

where the O-notation indicates a universal constant determined by the by the FFT on cyclic groups. Furthermore, $I(G)$ is bounded by the same sum but with $I(K)$ replacing $T(K)$.

In Theorem 1 α denotes the exponent of matrix multiplication, (currently 2.38, see [W]). It is useful to point out that under the simplifying assumption of $\alpha = 2$ the bounds becomes

$$T(G) = \frac{|G|}{|K|}\cdot T(K) + O(|G|\cdot\log(\frac{|G|}{|K|}))$$

and

$$I(G) = \frac{|G|}{|K|}\cdot I(K) + O(|G|\cdot\log(\frac{|G|}{|K|}))$$

with the O-notation as in Theorem 1.

In brief, the argument for Theorem 1 relies on understanding how the irreducible representations of G can be constructed, given the representations of some normal subgroup K such that G/K is abelian. G acts on the representations of K, partitioning them into orbits. The representations in any single orbit are then extended maximally using the theory of projective representations. When they can be extended no further, the representation is induced to G and necessarily gives an irreducible representation of G. Careful rewriting of the Fourier transform and analysis of the structure of the resulting matrices gives the above result.

Iterating these ideas gives efficient methods for treating any solvable group (a sequence of abelian extensions). In the simplest nontrivial instance of this case, the so-called metabelian groups, (abelian extensions of abelian groups) this paper gives the first example of a class of non-abelian groups such that

$$T(G) = I(G) = O(|G|\log(|G|)).$$

Group convolutions may be computed using Fourier transforms (see section 2C). Both naive methods using Fourier transforms and direct computation require on the order of $|G|^2$ operations. To obtain an advantage by using the method of Fourier transforms fast algorithms for both computing Fourier transforms and performing Fourier inversion are necessary. Thus as an immediate consequence of Theorem 1 fast group convolution algorithms for abelian extensions are determined (Theorem 7).

Previous fundamental work was done by Beth in the first serious treatment of noncommutative Fourier analysis ([B]). Beth sketches ideas for achieving savings in the case in which G contains some nontrivial normal subgroup with more careful analysis being given to the case of normal subgroups of prime index. Also, novel applications to the theory of coding, VLSI design and vision are given.

Independently, Beth has also made a strong start on fast algorithms for abelian extensions by deriving some speedup in the case of prime extensions. In [B] a result like Theorem 1 given here (specialized to this situation) is obtained.

Clausen also derives speedups for the general computation of Fourier transforms as well as in the specific case of S_n ([C1], [C2]). This work is done in the context of VLSI design and the study of Wedderburn transforms. Also, it has recently been brought to my attention that simultaneous independent work by Clausen studies the case of Fourier analysis for metabelian groups. Again, in the setting of VLSI design, he obtains the $|G|\log(|G|)$ result ([C3]).

Collectively, this work makes a good start towards obtaining general results on more efficient computation of Fourier inversion. Work on this problem for general finite groups is in progress and will appear presently ([R3]).

The organization of this paper is as follows: Section 2 provides necessary group theoretic background, in particular giving a highly constructive treatment of induced representations. Section 3 explains the basic idea of the algorithm, providing a "road map" for the rigorous derivations of the main results in section 4. The methods of section 4 are illustrated with two examples in section 5. There the algorithm is applied to metabelian groups G such that G/K is of order p or p^2 for some prime p. In the former case the example of the Heisenberg group is carried along in detail as a concrete illustration. The ideas in section 4 suggest algorithms for any group in which the irreducible representations are given as induced representations. As a start, in section 6, these ideas are explored briefly in the case of semidirect products by abelian subgroups.

2. Background

A. Representation Theory

Here the most basic concepts of representation theory are presented briefly. Proofs of all unproved assertions may be found in Serre ([S]).

Let G be a finite group. Recall that a representation ρ of G is a map assigining matrices to group elements such that $\rho(st) = \rho(s)\rho(t)$ for all s and t in G. Thus, ρ is a homomorphism from G to $GL(V)$ with V a vector space of dimension d_ρ, the *dimension* or *degree* of ρ. Two representations are *equivalent* or *isomorphic* if they differ only by a change of basis. The representation is *irreducible* if and only if for any subspace $W \subseteq V$, if $\rho(s)W \subseteq W$ for all $s \in G$ then either $W = \{0\}$ or $W = V$. A basic fact due to Schur is

Theorem *(**Schur's lemma**, [S], Proposition 4) Let ρ be an irreducible representation of G and let A be an invertible matrix such that*

$$A \cdot \rho(s) = \rho(s) \cdot A$$

for all $s \in G$. Then A is a scalar multiple of the identity.

The number of inequivalent irreducible representations of G is finite with the degrees satisfying

$$\sum_{\rho} d_{\rho}^2 = |G|$$

where the sum is over all irreducible representations of G.

Any representation ρ of G may be decomposed as the direct sum of irreducible representations. Thus, suppose ρ has the decomposition

$$\rho = \rho_1 \oplus \ldots \oplus \rho_m \qquad (1.1)$$

where the ρ_i are irreducible (though not necessarily distinct) representations of G. In the language of matrices, this says that with respect to some basis, $\rho(s)$ is block diagonal with the representations $\rho_i(s)$ appearing on the diagonal. If exactly j of the ρ_i are equivalent, then we say that the irreducible ρ_i appears in ρ with *multiplicity* j.

To determinine the decomposition of a representation the notion of the *character* of a representation is important. If ρ is a representation of G then the character of ρ, written χ_{ρ} is defined as

$$\chi_{\rho}(s) = trace(\rho(s))$$

for all $s \in G$. As the trace is invariant under conjugation, the character is an invariant for any class of equivalent representations. If a representation is one-dimensional then it is equal to its character.

Let $L(G)$ denote the space of complex-valued functions on G. There is a natural inner product on $L(G)$ given by

$$\langle f_1, f_2 \rangle = \frac{1}{|G|} \sum_{s \in G} f_1(s)\overline{f_2(s)}.$$

The irreducible characters are orthonormal with respect to this inner product. In terms of characters equation (1.1) becomes

$$\chi_{\rho} = \chi_{\rho_1} + \ldots + \chi_{\rho_m}.$$

Orthonormality implies the following useful result:

Theorem *([S], Theorem 3) Let ρ be an irreducible representation and η be any representation of G then the multiplicity of ρ in η is equal to $\langle \chi_{\rho}, \chi_{\eta} \rangle$. In particular, η is irreducible if and only if $\langle \chi_{\eta}, \chi_{\eta} \rangle = 1$.*

If ρ and η are representations of G we will sometimes write $\langle \rho, \eta \rangle$ for $\langle \chi_\rho, \chi_\eta \rangle$. An easy consequence of the orthogonality is that distinct irreducible representations have distinct characters.

B. Constructing Representations

Let $X = \{x_1, \ldots, x_n\}$ be some finite set on which G acts. That is for all $s, t \in G$ and $x \in X$, sx is an element of X such that $(st)x = s(tx)$. Associated to this action is the *permutation representation of G on X*. To define this, let $V(X)$ be the complex vector space generated by X,

$$V(X) = \mathbb{C}x_1 \oplus \ldots \oplus \mathbb{C}x_n.$$

G acts on $V(X)$ in the obvious way. An element s of G sends the basis vector e_{x_i} to e_{sx_i}. Thus, the corresponding matrix assigned to s, $\rho_X(s)$, has the form,

$$(\rho_X(s))_{i,j} = \begin{cases} 1 & \text{if } sx_j = x_i \\ 0 & \text{otherwise} \end{cases}.$$

The most natural finite set on which G acts is itself, by left multiplication. Here the associated permutation representation is called the *regular representation of G*. The associated matrices with respect to the standard basis $\{e_s \mid s \in G\}$ have the following properties:

Proposition 1 *Let ρ_G be the regular representation of G with respect to the standard basis of G. Then*

(i) The matrices $\{\rho_G(s) \mid s \in G\}$ have a unique nonzero entry in each row and column. This entry is equal to 1.

(ii) If s and t are distinct elements of G then all nonzero entries of $\rho_G(s)$ and $\rho_G(t)$ occur in distinct positions. (Another way to say this is that the matrix $\sum_{s \in G} \rho_G(s)$ is the $|G|$ by $|G|$ matrix of all ones.)

Proof: (i) is clear from the definition. As for (ii), note that $s \cdot t = s' \cdot t$ if and only if $s = s'$,

♠

Let H be a subgroup of G. Representations of H and G are related by the dual constructions of **induction** and **restriction**. A very explicit and concrete explananation of induced representations following the beautiful exposition of Coleman ([Co]) can be given.

Let η be a matrix representation of H. Then η gives rise to a representation of G called the *induced representation* of η to G. This is denoted $(\eta \uparrow G)$. To describe this, let $k = [G\colon H]$, the **index** of H in G, and fix a set of right coset representatives $\{s_1, ..., s_k\}$, with s_1 the identity, for the coset space G/H. Let $s \in G$, then $(\eta \uparrow G)(s)$ will be a k by k block matrix, with blocks of size d_η by d_η. Let $((\eta \uparrow G)(s))_{[i,j]}$ denote the i, j block of the matrix $(\eta \uparrow G)(s)$. Then

$$((\eta \uparrow G)(s))_{[i,j]} = \tilde{\eta}(s_i^{-1} s s_j)$$

where for all $t \in G$,

$$\tilde{\eta}(t) = \begin{cases} \eta(t) & \text{if } t \in H, \\ 0_{d_\eta \times d_\eta} & \text{otherwise.} \end{cases}$$

It is readily checked that this defines a representation of G ([Co], section 4). Also it is straightforward to show that induction is transitive ([Co], Theorem 4). That is, if

$$G \supseteq H \supseteq K,$$

and η is a representation of K then,

$$((\eta \uparrow H) \uparrow G) = (\eta \uparrow G).$$

The block structure of these matrices will prove to be of great importance. Several key properties follow immediately from the definition.

Proposition 2 *(i) $(\eta \uparrow G)(s)$ is a k by k block matrix with blocks of size d_η by d_η with at most one nonzero block in each row and column. Thus, the dimension of $(\eta \uparrow G)$ is $[G\colon H]\, d_\eta$.*

(ii) Suppose H is normal in G. If s and t are elements of G that lie in the same coset of G by H then $(\eta \uparrow G)(s)$ and $(\eta \uparrow G)(t)$ have the same structure (ie. nonzero blocks in the same position).

(iii) If H is normal in G then the matrices $\{(\eta \uparrow G)(s_i) \mid 1 \leq i \leq k\}$ have the same "structure" as the matrices for the regular representation on the group G/H. (That is, there are nonzero blocks in the induced representation in the same positions as the nonzero entries in the regular representation of G/H.) In particular, the matrices for the elements of H are block diagonal and elements not in H will have no blocks on the diagonal.

Proof: (i) is clear from the definition. To show (ii) and (iii) suppose now that H is a normal subgroup of G. s and t are in the same coset of G/H if and only if $s^{-1}t \in H$. Suppose $s_i^{-1}ss_j \in H$. Because H is normal in G $s_j^{-1}s^{-1}ts_j \in H$ as well. Consquently, the product,

$$s_i^{-1}ss_j \cdot s_j^{-1}s^{-1}ts_j = s_i^{-1}ts_j$$

is in H. This proves (ii).

For (iii), note that

$$s_i^{-1}ss_j \in H \Longleftrightarrow ss_j \in s_iH$$

which is if and only if

$$\overline{s}\,\overline{s_j} = \overline{s_i}$$

where the bar denotes the image in G/H. This proves (iii).

♠

Dual to the operation of induction is the process of *restriction*. That is, given a subgroup H of G and a representation ρ of G we obtain a representation of H by restricting ρ to H. This is denoted $(\rho \downarrow H)$. Clearly, this too is a transitive operation. Induction and restriction are related by **Frobenius reciprocity**.

Theorem *(**Frobenius reciprocity**, [Co], Theorem 6) Let η be a representation of H and ρ a representation of G. Then,*

$$\langle(\eta \uparrow G), \rho\rangle = \langle\eta, (\rho \downarrow H)\rangle.$$

When restricted to subgroups, irreducible representations need not remain irreducible. Frobenius reciprocity is a useful tool in determining the structure of a restricted representation. In general this can be a very difficult task. In the case of normal subgroups, however, much can be said. The study of the structure of restrictions of representations to normal subgroups is often known as **Clifford theory** or **Mackey's theory of little groups.**

Let H be a normal subgroup of G and η be a representation of H. Then given any $s \in G$ a (possibly) new representation of H may be defined, a **conjugate representation** of η, $\eta^{(s)}$ by

$$\eta^{(s)}(t) = \eta(s^{-1}ts).$$

Note that the conjugate of an irreducible representation is still irreducible.

Consequently, there is an action of G on the set of inequivalent irreducible representations of H. For any irreducible representation η of H let $\Delta(\eta)$ denote its orbit under this action. These are simply all the inequivalent conjugates of η. Abusing notation, $\Delta(\eta)$ will sometimes denote the direct sum of these representations as well. $\Delta(\eta)$ is often called the *star* of η. Let $Stab_G(\eta)$ be the isotropy subgoup of η under the action of G. Thus,

$$Stab_G(\eta) = \{s \in G \mid \eta^{(s)} \sim \eta\}$$

where "$\sim$" denotes equivalence of representations. In particular $Stab(\eta) \supseteq H$.

The language of conjugate representations and restrictions gives a nice way to explicitly restate part (iii) of Proposition 2.

Theorem 2 ***(Clifford Theory)*** *Let H be a normal subgroup of G.*

(i) ([Co], Theorem 10) Let η be a representation of H. Then

$$((\eta \uparrow G) \downarrow H) = [Stab(\eta){:}\, H] \cdot \Delta(\eta).$$

(ii) ([Co], Theorem 14) Let ρ an irreducible representation of G. Then

$$(\rho \downarrow H) = \frac{d_\rho}{[G{:}\, Stab(\eta)]d_\eta} \cdot \Delta(\eta)$$

where η is any irreducible representation of H which occurs in $(\rho \downarrow H)$.

Related to restriction and induction is the notion of **extension**. If ρ is an irreducible representation of G such that

$$(\rho \downarrow H) = \eta$$

then ρ is said to **extend** η. The most natural setting in which extensions arise is in the theory of **projective representations**. A thorough discussion of this is given in the first edition of Curtis and Reiner ([CR], section 51).

C. Fourier Transforms

Let G be a finite group, ρ a representation of G and $f \in L(G)$. Then the **Fourier transform of** f **at** ρ, denoted as $\widehat{f}(\rho)$, is the matrix given by

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

Let $H \subseteq G$ be a subgroup and $\{s_1, \ldots, s_k\}$ be a complete set of coset representatives for G/H. Then $\widehat{f}(\rho)$ may be rewritten as

$$\begin{aligned}
\widehat{f}(\rho) &= \sum_{i=1}^{k} \sum_{t \in H} f(s_i t)\rho(s_i)\rho(t) \\
&= \sum_{i=1}^{k} \rho(s_i) \sum_{t \in H} f(s_i t)\rho(t) \\
&= \sum_{i=1}^{k} \rho(s_i) \sum_{t \in H} f_i(t)\rho(t) \\
&= \sum_{i=1}^{k} \rho(s_i)\widehat{f_i}((\rho \downarrow H))
\end{aligned}$$

where $f_i \in L(H)$ is defined by

$$f_i(t) = f(s_i t)$$

for all $t \in H$ and i between 1 and k.

Via the **Fourier inversion** formula the Fourier transforms of a function $f \in L(G)$ at all the irreducible representations of G determine f,

$$f(s) = \frac{1}{|G|} \sum_{\rho} d_\rho trace(\widehat{f}(\rho)\rho(s^{-1})).$$

As usual, the Fourier transform converts convolution to multiplcation,

$$\widehat{f * g}(\rho) = \widehat{f}(\rho) \cdot \widehat{g}(\rho) \qquad (2.1)$$

for $f,\ g \in L(G)$.

The first finite groups for which the problem of efficient Fourier analysis was seriously studied were the cyclic groups. Let C_n denote the cyclic group of order n with generator s. Let ζ be a primitive n^{th} root of unity. Then the representations of C_n form a group (the so-called dual group) naturally isomorphic to C_n where

$$\zeta_j(s^k) = \zeta^{jk}$$

for $0 \leq j \leq n-1$.

Let $f \in L(C_n)$. The Fourier transform of f at ζ_j is given by

$$\widehat{f}(j) = \sum_{k=0}^{n-1} f(k)\zeta_j(s^k) = \sum_{k=0}^{n-1} f(k)\zeta^{jk}.$$

As all the representations are one-dimensional, Fourier inversion simply amounts to computing the Fourier transform twice

$$\begin{aligned} f(j) &= \frac{1}{|G|}\sum_{k=0}^{n-1}\widehat{f}(k)\zeta_k(s^j) \\ &= \frac{1}{|G|}\sum_{k=0}^{n-1}\widehat{f}(k)\zeta_{kj} \\ &= \frac{1}{|G|}\widehat{\widehat{f}}(j) \qquad (1.2) \end{aligned}$$

Thus for abelian groups computation of Fourier transforms and Fourier inversion are naturally of the same complexity.

Any of a number of different algorithms for computing all Fourier transforms on C_n go by the name of the "FFT" or "Fast Fourier Transform." Different approaches of course yield different expected running times. Diaconis ([D3]) gives an in-depth analysis of the expected and average running times of several different common versions of the FFT concluding,

Theorem *((2.8), [D3])* $T(C_n) = O(n \log n)$ *where the O-notation indicates some universal constant. Consequently, using (1.2), the same is true for* $I(C_n)$.

Any abelian group is the product of cyclic groups, thus

Corollary ***(FFT)*** *Let G be any finite abelian group. Then*

$$T(G) = I(G) = O(|G|\log(|G|))$$

where the O-notation indicates some universal constant.

3. The Main Idea

A group G is an *extension of K by Q* if

$$(i)\;\; K \lhd G$$

$$(ii)\; G/K \cong Q.$$

In this paper the problem of computing Fourier transforms and performing Fourier inversion in the case in which G/K is abelian is considered. In this situation, the following structure theorem will be proved (Theorem 3).

This will allow the computations to be reduced to successive applications of the FFT where existing fast algorithms for these operations then speed the computation tremendously.

Theorem *Let G be an extension of K by Q such that Q is abelian. Then every irreducible representation of G is given by inducing an extension of an irreducible representation of K.*

In other words, given an irreducible representation ρ, of G there exists a subgroup H_ρ such that

$$K \subseteq H_\rho \subseteq G$$

and a representation η_ρ of H_ρ such that

(i) η_ρ is an extension of an irreducible representation of K

and

(ii) $\rho = (\eta_\rho \uparrow G)$.

Because G/K is abelian all subgroups between G and K are normal. An application of Clifford theory yields

$$(\rho \downarrow H_\rho) = \Delta(\eta_\rho).$$

Now consider the problem of computing $\widehat{f}(\rho)$ for some $f \in L(G)$. If ρ is constructed from η_ρ as in section 2B, Proposition 2 there implies

$$\begin{aligned} \widehat{f}(\rho) &= \sum_{s \in G} f(s)\rho(s) = \sum_{j=1}^{[G:H_\rho]} \rho(s_j) \sum_{h \in H_\rho} f_j(h)\rho(h) \\ &= \sum_{j=1}^{[G:H_\rho]} \rho(s_j) \begin{pmatrix} \widehat{f}_j(\eta_\rho^{(1)}) & 0 & \dots & 0 \\ 0 & \widehat{f}_j(\eta_\rho^{(2)}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \widehat{f}_j(\eta_\rho^{([G:H_\rho])}) \end{pmatrix}. \end{aligned}$$

where the s_j are coset representatives for G/H_ρ and the representations $\eta_\rho^{(i)}$ are the distinct conjugates in the star of η_ρ. Because ρ is induced from a normal subgroup the matrices

$$\{\rho(s_j) \mid 1 \leq j \leq [G\colon H_\rho]\}$$

have the same form as the matrices corresponding to the regular representation of the group G/H_ρ (Proposition 2). That is, these matrices have a block structure such that there is exactly one nonzero block in each row and column and furthermore, no two matrices $\rho(s_j)$ have nonzero blocks in the same position. This last observation will be crucial for performing fast inversion.

Suppose all restricted transforms

$$\{\widehat{f}_j(\eta_\rho^{(i)}) \mid 1 \leq i, j \leq [G\colon H_\rho]\}$$

had been computed previously and stored, and that furthermore, the matrices $\rho(s_i), (1 \leq i \leq [G:H])$ are given. For any k between 1 and $[G:H]$ let i_k be the unique column in row k that has a nonzero d_η by d_η block in $\rho(s_i)$. Then the product

$$\rho(s_i)\widehat{f}_i((\rho \downarrow H))$$

will be the block matrix with unique nonzero block

$$\rho(s_i)_{k,i_k} \cdot \widehat{f}_i(\eta_\rho^{(i_k)}) \qquad (3.1)$$

in row k.

Thus, d_η^α operations, where α is the exponent of matrix multiplication, are needed to form this block and consequently all of row k. So, in total an additional $[G:H_\rho] \cdot d_\eta^\alpha$ operations are needed to form $\rho(s_i)\widehat{f}_i((\rho \downarrow H))$. Finally, as there are $[G:H_\rho]$ values of s_i, in all an additional $[G:H_\rho]^2 \cdot d_\eta^\alpha$ operations would be needed to calculate $\widehat{f}(\rho)$.

Under the (strong) assumption of $\alpha = 2$ this yields

$$[G:H_\rho]^2 \cdot d_\eta^2 = d_{(\eta\uparrow G)}^2 = d_\rho^2$$

additional operations. In particular, if η is one dimensional then this is true independent of α. This is the case for all η_ρ when K is abelian.

For each irreducible representation ρ of G let T_ρ be the number of operations needed to calculate the collection of restricted transfoms

$$\{\widehat{f}_j(\eta_\rho^{(i)}) \mid 1 \leq i, j \leq [G\colon H_\rho]\}.$$

Then the total number of operations needed to compute all the Fourier transforms on G, $T(G)$ is

$$T(G) = \sum_{irreducible\ \rho} T_\rho + [G:H_\rho]^2 \cdot d_{\eta_\rho}^\alpha \qquad (3.2)$$

where η_ρ induces the irreducible representation ρ.

Again, under the simplifying assumption that $\alpha = 2$ or that K is abelian (3.2) becomes

$$\begin{aligned} T(G) &= \sum_\rho T_\rho + d_\rho^2 \\ &= |G| + \sum_\rho T_\rho . \end{aligned}$$

So the question becomes, can the calculations of the restricted transforms be scheduled in such a way so as to achieve savings. It will be shown that the computation essentially can be reduced to iteratively computing Fourier transforms over abelian groups (subgroups of G/K) until arriving at restricted transforms of K

$$\{\widehat{f}_j(\eta) \mid 1 \leq j \leq [G:K]\}$$

where η runs over all irreducible representations of K. Here methods depending on K are to be used.

In some sense, the earlier research in general FFT's had "predicted" the existence of fast algorithms for abelian extensions. For example, in [DR] it is shown that if G is a solvable group, then $T(G)$ is essentially $\mid G \mid \log(\mid G \mid)$. However, this assumes that a basis for the irreducible representations of G can be found with certain "splitting" properties with respect to a maximally refined tower of subgroups of G (Proposition 2, [DR]). In the present paper this prediction is realized, albeit by slightly different techniques. On the other hand up to now little has been said about computing Fourier inversion. The ideas presented here for computing Fourier transforms admit modifications allowing fast computation of Fourier inversion.

Computation of Fourier inversion relies on the observation that just as the Fourier transforms were easy to construct, dually, it is easy to take them apart. Equation (3.1) shows that the matrices $\widehat{f}(\rho)$ have block entries

$$\rho(s_i)_{k,i_k} \cdot \widehat{f}_i(\eta_\rho^{(i_k)})$$

while by Proposition 2, the blocks in the matrices $\rho(s_i)^{-1}$ will run through the collection $\{\rho(s_i)^{-1}_{k,i_k}\}$.

Thus, assuming that the matrices $\rho(s_i)^{-1}$ are given, appropriate multiplication of the blocks in $\widehat{f}(\rho)$ will permit the recovery of the restricted transforms

$$\{\widehat{f}_j(\eta_\rho^{(i)}) \mid 1 \leq i, j \leq [G: H_\rho]\}.$$

Assuming that these restricted transforms can be computed by iteratively computing Fourier transforms over abelian groups, iterating the process of Fourier inversion over these same groups will yield the collection

$$\{\widehat{f}_j(\eta) \mid 1 \le j \le [G:K]\}$$

of restricted transforms on K. As f is just the sum of the functions f_j on K, the problem has been reduced to inverting these $[G:K]$ functions on K where possibly new methods may be used.

The analysis proceeds by first considering the structure of the set of extensions of an irreducible representation of K when G/K is abelian. Assuming the existence of the extensions the computation of the Fourier transforms and Fourier inversion are then built in pieces. The construction of the extensions is carried out via a neat use of Schur's lemma that allows representations of arbitrary degree to be extended to cyclic extensions contained in their stabilizer. This is a construction that is essentially due to Clifford and is implicit in his original paper on projective representations ([Cl]).

Finally, careful consideration and counting of the irreducible representations which appear by this process give the main results (Theorems 5 and 6) on computations of Fourier transforms and Fourier inversion. These yield immediately fast algorithms for group convolution. Slightly more careful analysis is then given to the case of metabelian groups and a more precise result is obtained.

4. Fourier Analysis for Abelian Extensions

The starting point for this investigation is the following simple observation:

Proposition 3 *Let ρ be a representation of G and K a subgroup of G such that $(\rho \downarrow K)$ is irreducible. Then ρ is irreducible.*

Proof: Any direct sum decomposition of ρ implies some such for K.

♠

Dually, another way to say this is that extensions of irreducible representations are irreducible. This permits the simple analysis of the structure of the extensions of some representation of a subgroup.

Lemma 1 *Let G be an abelian extension of K. Let η be an irreducible representation of K that extends to an irreducible representation $\tilde{\eta}$ of G. For* *any* $\chi \in \widehat{(G/K)}$ *let* $\tilde{\chi}$ *denote the natural extension of* χ *to* G *by composition*

$$\tilde{\chi}(t) = \chi(\pi(t))$$

where π *is the projection map from* G *to* G/K*. Then for any* $\chi \in \widehat{G/K}$*, the map* $\tilde{\chi} \otimes \tilde{\eta}$ *defines an irreducible representation of* G *and these are all the irreducible representations of* G *which extend* η*.*

Proof: Recall that $\tilde{\chi} \otimes \tilde{\eta}$ is defined by

$$(\tilde{\chi} \otimes \tilde{\eta})(s) = \tilde{\chi}(s) \cdot \tilde{\eta}(s).$$

(In general it would be the matrix tensor product, but $\tilde{\chi}$ is one-dimensional.) As $\tilde{\chi}$ is trivial on K

$$((\tilde{\chi} \otimes \tilde{\eta}) \downarrow K) = \eta.$$

Thus the representations $(\tilde{\chi} \otimes \tilde{\eta})$ are all irreducible representations of G extending η. That these are all of them follows from the claim,

Claim: *Let all notation be as in the lemma, then*

$$(\eta \uparrow G) = \oplus_{\chi \in \widehat{G/K}} (\tilde{\chi} \otimes \tilde{\eta}).$$

Proof: By Frobenius reciprocity $(\eta \uparrow G)$ contains (possibly with multiplicity) all irreducible representations of G which restrict to η on K. Thus, by the lemma, $(\eta \uparrow G)$ contains each of the representations $\tilde{\chi} \otimes \tilde{\eta}$ at least once. Counting dimensions, it must be exactly once, with no others included.

♠

The particularly simple structure of these extensions gives the first step in building Fourier transforms and consequently, performing Fourier inversion.

Lemma 2 *Let* $K \lhd G$ *such that* G/K *is abelian and let* η *be an irreducible representation of* K *which extends to* G*. Assume* $m = [G : K]$ *and let* $\{s_1, \ldots, s_m\}$ *be coset representatives for* G/K *with* s_1 *the identity. Let* $f \in L(G)$ *and suppose that all restricted transforms* $\{\widehat{f}_i(\eta)\}$ *have been computed. Then the number of operations needed to compute the Fourier transforms at every extension of* η *to* G *is at most*

$$[G\colon K] \cdot d_\eta^\alpha + O(d_\eta^2 \cdot [G\colon K] \log([G\colon K])$$

where α *is the exponent of matrix multiplication and the* O*-notation indicates the universal constant determined by the FFT.*

Proof: Identify s_i with its image in G/K. Fix an extension $\tilde{\eta}$ of η to G. Consider now the computation of Fourier transforms at the extensions of η,

$$\begin{aligned}\widehat{f}(\tilde{\chi} \otimes \tilde{\eta}) &= \sum_{i=1}^{m} (\tilde{\chi} \otimes \tilde{\eta})(s_i)\widehat{f}_i(\eta) \\ &= \sum_{i=1}^{m} \chi(s_i)\tilde{\eta}(s_i)\widehat{f}_i(\eta) \\ &= \sum_{i=1}^{m} \chi(s_i) F_{s_i,\eta}\end{aligned}$$

where $F_{s_i,\eta}$ is defined to be the d_η by d_η matrix $\tilde{\eta}(s_i)\widehat{f}_i(\eta)$. Note that on the order of $[G:K] \cdot d_\eta^\alpha$ operations are needed to form the matrices $\{F_{s_i,\eta}\}$.

Let $a^{\eta}_{jk}(s_i)$ be the j,k entry of $F_{s_i,\eta}$. Then the j,k entry of the Fourier transform $\widehat{f}(\tilde{\chi} \otimes \tilde{\eta})$ is given by the sum

$$\sum_{i=1}^{m} \chi(s_i) a^{\eta}_{jk}(s_i) = \widehat{a^{\eta}_{jk}}(\chi).$$

To do this for all $\chi \in \widehat{G/K}$ for any fixed j and k requires $O([G\colon K]\log([G\colon K]))$ operations. Letting j and k vary over all possible values yields the theorem.

♠

To recover the restricted transforms $\{\widehat{f}_i(\eta)\}$ the above algorithm is run "backwards".

Lemma 3 *Let all notation be as in Lemma 2. Let $K \triangleleft G$ such that G/K is abelian. Let $f \in L(G)$ and let η be an irreducible representation of K which extends to G. Suppose that the matrices $\widehat{f}(\tilde{\eta})$ for all extensions $\tilde{\eta}$ of η are given. Then at most*

$$[G\colon K] \cdot d_\eta^\alpha + O(d_\eta^2 \cdot [G\colon K]\log([G\colon K])$$

operations are needed to recover the restricted transforms $\widehat{f}_i(\eta)$.

Proof: The previous lemma shows that the entries of the matrix $\widehat{f}(\tilde{\chi} \otimes \tilde{\eta})$ are the Fourier transforms

$$\widehat{a^{\eta}_{jk}}(\chi).$$

Thus, to recover the matrices $\{F_{s_i,\eta}\}$ these transforms must be inverted for all $\chi \in \widehat{G/K}$. This can be done in $O([G:K]\log([G:K]))$ operations. Now multiplication by the matrices $\{\tilde{\eta}(s_i)^{-1}\}$ gives the restricted transforms $\{\widehat{f}_i(\eta)\}$.

♠

All the previous work depends on having at least one extension. Thus, the problem of constructing extension is now considered.

Let K be a normal subgroup of G. As was shown in section 2, G acts on the representations of K, partitioning them into their respective stars. The stabilizer of any particular irreducible representation of K is the subgroup of elements of G which conjugate the representation to some equivalent irreducible representation. If η is an irreducible representation of K then $Stab_G(\eta)$ is its stabilizer. The following observation is basic.

Proposition 4 *Let $K \lhd G$, and let η be an irreducible representation of K. Then η can only be extended to subgroups that lie between $Stab_G(\eta)$ and K.*

Proof: Let $H \subseteq G$ be a subgroup of G such that η extends to a representation $\tilde{\eta}$ of H. Let χ be the character of η and $\tilde{\chi}$ denote the character of $\tilde{\eta}$. Note that

$$(\tilde{\chi} \downarrow K) = \chi.$$

As $\tilde{\chi}$ is a class function on H, then for any $s \in H$ and $t \in K$,

$$\tilde{\chi}((s^{-1}ts) = \tilde{\chi}(t) = \chi(t).$$

Thus, χ is stable under the action of H and so is η.

♠

The main tool in actually constructing an extension is the following theorem which is essentially due to Clifford ([Cl]).

Theorem 3 *Let $K \lhd G$ such that G/K is cyclic of order m. Let η be any irreducible representation of K such that η is stable under the action of G. That is*

$$\eta^{(g)} \sim \eta$$

for all $g \in G$. Then there are exactly m extensions of η to irreducible representations of G.

Proof: Let $s \in G$ generate G/K. Since $\eta^{(s)} \sim \eta$ there exists a matrix $C(s)$ such that

$$C(s)^{-1}\eta(h)C(s) = \eta(s^{-1}hs)$$

for all $h \in K$. Consequently,

$$C(s)^{-k}\eta(h)C(s)^{k} = \eta(s^{-k}hs^{k})$$

for all k.

In particular, for $k = m$

$$\begin{aligned} C(s)^{-m}\eta(h)C(s)^m &= \eta(s^{-m}hs^m) \\ &= \eta(s^m)^{-1}\eta(h)\eta(s^m) \end{aligned}$$

where the last line follows from $s^m \in K$. But this means that the matrix $\eta(s^m)C(s)^{-m}$ intertwines the irreducible representation η. That is

$$\eta(h) = \eta(s^m)C(s)^{-m}\eta(h)C(s)^m\eta(s^{-m})$$

for all $h \in K$. Schur's lemma now implies

$$\eta(s^m)C(s)^{-m} = \lambda \cdot I_{d_\eta \times d_\eta}$$

or

$$\eta(s^m) = \lambda \cdot C(s)^m$$

for some $\lambda \in \mathbb{C}$.

The above is to be used to obtain a representation of G extending η. The obvious thing to do would be to assign the matrix $C(s)$ to s and extend multiplicatively, hoping that this suffices. However, there is an extra constraint, that $\rho(s)^m = \rho(s^m) = \eta(s^m)$. This can satisfied by using the above relation.

Let a be any m^{th} root of λ. Now define a map η_a on G by

$$\eta_a(s^k h) = C_a(s)^k \eta(h)$$

where

$$C_s(s) = aC(s).$$

Note that

$$\begin{aligned} \eta_a(s^m) &= C_a(s)^m \\ &= a^m C(s)^m \\ &= \lambda C(s)^m = \eta(s^m). \end{aligned}$$

Claim: *η_a is a representation of G which extends η.*

Proof: Clearly, if η_a is a representation of G then it extends η. So all that is left to be shown is that η_a is a representation of G. In other words,

$$\eta_a(s^i h s^k h') = \eta_a(s^i h)\eta_a(s^k h').$$

By definition, the righthand side is

$$\begin{aligned}\eta_a(s^i h)\eta_a(s^k h') &= a^i C(s)^i \eta(h) a^k C(s)^k \eta(h') \\ &= a^{i+k} C(s)^i \eta(h) C(s)^k \eta(h')\end{aligned}$$

Now unravel the lefthand side.

$$\begin{aligned}\eta_a(s^i h s^k h') &= \eta_a(s^i s^k s^{-k} h s^k h') \\ &= \eta_a(s^{\overline{i+k}} s^{jm} s^{-k} h s^k h')\end{aligned}$$

where $\overline{i+k}$ is the unique integer, $0 \leq \overline{i+k} < m$ such that $\overline{i+k} + jm = i+k$. (This defines j.) Then

$$\begin{aligned}\eta_a(s^{\overline{i+k}} s^{jm} s^{-k} h s^k h') &= a^{\overline{i+k}} C(s)^{\overline{i+k}} \eta(s^m)^j \eta(s^{-k} h s^k) \eta(h') \\ &= a^{\overline{i+k}} C(s)^{\overline{i+k}} a^{jm} C(s)^{jm} C(s)^{-k} \eta(h) C(s)^k \eta(h') \\ &= a^{\overline{i+k}+jm} C(s)^{\overline{i+k}+jm-k} \eta(h) C(s)^k \eta(h') \\ &= a^{i+k} C(s)^i \eta(h) C(s)^k \eta(h')\end{aligned}$$

and the proof is complete.

♠

Note that in the case of $d_\eta = 1$, the "matrix" $C(s)$ is the constant 1. Then the map η_a is defined by

$$\eta_a(s^k h) = a^k \eta(h)$$

where recall that a is an m^{th} root of $\eta(s^m)$.

Of course, in general, for G an abelian extension of K, the stabilizer of any irreducible representation η need not be a cyclic extension. Then η may not be extendable all the way to the stabilizer. Any abelian extension can be written as a sequence of cycilc extensions. Thus, to build extensions of a representation η, first extensions to the biggest cyclic extension within the stabilizer could be constructed. Let $H \subseteq Stab(\eta)$ be this extension and let $\tilde{\eta}$ be an extension of η to H. As G/K is abelian, H is normal in G and there is an action of G on $\tilde{\eta}$. Any element of G stabilizing $\tilde{\eta}$ must stabilize η, thus

$$H \subseteq Stab_G(\tilde{\eta}) \subseteq Stab_G(\eta).$$

So, one of two things must occur:

(i) $Stab_G(\tilde{\eta}) \supsetneq H$

or

(ii) $Stab_G(\tilde{\eta}) = H$

Consider these in turn.

In case (i) again $Stab_G(\tilde{\eta})$ is an abelian extension of H so that $\tilde{\eta}$ may be extended to some cyclic extension H' within $Stab_G(\tilde{\eta})$. This is still an extension of η. So, by Lemma 1, η has $[H' : K]$ extensions to H'. Each such extension must come from an extension of η to H. Thus, the following is proved inductively:

Proposition 5 *Let K be a normal subgroup of G such that G/K is abelian. Let η an irreducible representation with $\tilde{\eta}$ and $\tilde{\eta}'$ extensions of η to some subgroup H between G and K. Then $\tilde{\eta}$ can be extended if and only if $\tilde{\eta}'$ can be extended.*

In case (ii) the following lemma is useful:

Lemma 4 *Let $K \lhd G$ and let η be an irreducible representation of K. Then $Stab_G(\eta) = K$ if and only if $(\eta \uparrow G)$ is irreducible.*

Proof: By Clifford theory

$$((\eta \uparrow G) \downarrow K) = c \cdot \Delta(\eta)$$

for some positive integer c. If $Stab_G(\eta) = K$ then $\mid \Delta(\eta) \mid = [G : K]$. Consequently, counting dimensions it must be that $c = 1$ if and only if $[G : Stab_G(\eta)] = [G : K]$, so if and only if

$$Stab_G(\eta) = K.$$

♠

Thus, any irreducible representation η of K is extended using Theorem 2 until it can be extended no further, whereupon it is induced to G and is irreducible. Note that Proposition 5 implies that if any extension of η to a subgroup H of G induces an irreducible representation of G then all extensions of η to H induce irreducible representations of G. By Frobenius reciprocity the resulting irreducible representations are independent of the sequence of extensions. Combining Proposition 5 with Frobenius reciprocity gives

Lemma 5 *Let $K \lhd G$ with G/K abelian.*

(i) Let η be an irreducible representation of K. If any extension of η to $H \subseteq G$ induces an irreducible representation of G then all extensions of η to H induce irreducible representations of G.

(ii) Let $\eta' \in \Delta(\eta)$. Then η' can be extended as far as η and the irreducible representations obtained by extending the maximal extensions of η' are the same as those from η. (The irreducible representations of G are determined by the different stars of K.)

Proof: (i) follows immediately from Proposition 5.

As for (ii), observe first that if $\eta' \in \Delta(\eta)$ then $Stab_G(\eta') = Stab_G(\eta)$. This is clear for if G/K is abelian then $Stab_G(\eta)$ is normal in G. Thus,

$$Stab_G(\eta') = s \cdot Stab_G(\eta) \cdot s^{-1} = Stab_G(\eta)$$

where $\eta^{(s)} = \eta'$. Let $\tilde{\eta}$ be any extension of η to a subgroup H. Then $(\tilde{\eta}^{(s)} \downarrow K) = \eta'$ and η' is extendable to H. Frobenius reciprocity implies that conjugate representations have the same irreducible representations of G occurring in the representations which they induce in G. Combined with (i) this gives the second part of (ii).

♠

Now the basic representation theory fact may be proved:

Theorem 4 *Let K be a normal subgroup of G such that G/K is abelian. Then every irreducible representation of G is given by inducing an irreducible representation which is an extension of some irreducible representation of K.*

(This may be restated as:

Given a representation ρ of G there exists a subgroup H

$$K \subseteq H \subseteq G$$

and an irreducible representation η of H such that

(i) $(\eta \uparrow G) = \rho$ and

(ii) η is an extension of an irreducible representation of K.)

Proof: The proof is by induction on $[G : K]$. If $G = K$ then there is nothing to prove. So assume $[G : K] > 1$ and let ρ be an irreducible representation of G. In addition assume ρ is an extension of an irreducible representation of K. Let p be a prime such that $p \mid [G : K]$. Since G/K is abelian, there exists a subgroup H such that

$$K \subseteq H \subseteq G$$

and $[G : H] = p$. Let ψ be an irreducible representation of H which occurs in $(\rho \downarrow H)$. Either $Stab_{_G}(\psi)$ is G or it is H. By Lemma 4 if it is the latter, then $(\psi \uparrow G)$ is irreducible and $\rho = (\psi \uparrow G)$ and by induction ρ is induced from some extension of K.

So, suppose that $Stab_{_G}(\psi) = G$. By Theorem 2 ψ may be extended to G in p distinct ways. and $(\psi \uparrow G)$ is the direct sun of these extensions. Consequently, as ρ is irreducible, ρ must be one of these extensions and $(\rho \downarrow H) = \psi$.

By induction there exists a subgroup $H_{_\rho}$ and an irreducible representation η of $H_{_\rho}$ such that

$$K \subseteq H_{_\rho} \subseteq H \subseteq G$$

and $(\eta \uparrow H) = \psi$. Thus, $(\psi \downarrow H) = \Delta(\eta)$. H acts transitively on $\Delta(\eta)$ and $|\Delta(\eta)| = [H : H_{_\rho}]$. Consequently, $Stab_{_G}(\eta) \neq H$. However,

$$|\Delta(\eta)| = \frac{|G|}{|Stab_{_G}(\eta)|} = \frac{|H|}{|H_{_\rho}|}.$$

So,

$$\frac{|G|}{|H|} = p = \frac{|Stab_{_G}(\eta)|}{|H_{_\rho}|}.$$

Thus, the situation is

$$\begin{tikzcd} & G & \\ H \arrow[ur] & & Stab_{_G}(\eta) \arrow[ul] \\ & H_{_\rho} \arrow[ul] \arrow[ur] \arrow[d, dash] & \\ & K & \end{tikzcd}$$

By Theorem 2 every $\eta' \in \Delta(\eta)$ extends in p distinct ways to $Stab_{_G}(\eta)$. Each of these induces something irreducible in G. By Frobenius reciprocity one of these must be ρ.

♠

Corollary *Let K be a normal subgroup of G such that $G/K \cong \mathbb{Z}/p\mathbb{Z}$ for some prime p. Let ρ be an irreducible representation of G, then $(\rho \downarrow K)$ is multiplicity-free.*

Proof: In the above theorem it was determined that either $(\rho \downarrow K)$ remains irreducible or else is equal to the star of irreducible representation of K. In either case, every irreducible which occurs does so exactly once.

♠

The main theorems now follow easily.

Theorem 5 *Let $K \lhd G$ such that G/K is abelian. Let $\{\eta_i\}_{i\in I}$ be a complete set of representatives for the orbits of the irreducible representations of K under the action of G. Let H_{η_i} be a subgroup of G (containing K) such that η_i extends to H_{η_i} but no further. Then the number of operations needed to compute all Fourier transforms for any complex-valued function f defined on G is at most*

$$\frac{|G|}{|K|}\cdot T(K)+\sum_{i\in I} m_{\eta_i}(\frac{|G|}{|H_{\eta_i}|})^2 d_{\eta_i}^\alpha + |\Delta(\eta_i)|\,(\frac{|G|}{|K|}\cdot d_{\eta_i}^\alpha + d_{\eta_i}^2 O(\frac{|G|}{|K|}\log(\frac{|H_{\eta_i}|}{|K|})))$$

where m_{η_i} is the number of irreducible representations of G ultimately induced by extensions of η_i, α is the exponent of matrix multiplication, and the O-notation indicates the universal constant as determined by the FFT.

Proof: Fix a set of coset representatives $\{s_i\}$ for G/K and compute the restricted transforms $\{\widehat{f}_j(\eta)\}$ for all irreducible representations η of K. This requires $\frac{|G|}{|K|}T(K)$ operations.

Now for each irreducible representation η of K, extend η maximally to H_η. Using Lemma 3 to compute all restricted transforms at these extensions requires

$$|\Delta(\eta)|\cdot(\frac{|G|}{|K|}\cdot d_\eta^\alpha + d_\eta^2 O(\frac{|G|}{|K|}\log(\frac{|H_\eta|}{|K|})))$$

operations. Let $\rho = (\tilde{\eta} \uparrow G)$ be irreducible where $\tilde{\eta}$ is an extension of η to H_η. Then

$$\widehat{f}(\rho) = \sum_{j=1}^{[G:H_\eta]} \rho(s_j)\begin{pmatrix} \widehat{f}_j(\eta_\rho^{(1)}) & 0 & \dots & 0 \\ 0 & \widehat{f}_j(\eta_\rho^{(2)}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \widehat{f}_j(\eta_\rho^{([G:H_\eta])}) \end{pmatrix}.$$

As the matrices $\rho(s_j)$ are block matrices with the structure of the matrices of the permutation representation on G/H_η each multiplication takes $\frac{|G|}{|H_\eta|} \cdot d_\eta^\alpha$ operations. So in total,

$$(\frac{|G|}{|H_\eta|})^2 \cdot d_\eta^\alpha$$

operations are required.

By Theorem 3 all irreducible representations of G are obtained in this fashion. For each star $\Delta(\eta_i)$, m_{η_i} irreducible representations of G are obtained. Summing over all stars gives the result.

♠

Note that $m_{\eta_i} \leq [H_{\eta_i} : K]$. This yields the bound for Theorem 1 given in the introduction.

As stated at the beginning of this section. The algorithm outlined above can be run "backwards" to obtain the same result for Fourier inversion, up to replacing the $T(K)$ term by $I(K)$. Multiplication by the blocks in the matrices $\rho(s_i)^{-1}$ return the restricted transforms at the extensions which induce the irreducible representation. Iteration of Fourier inversion on abelian groups ultimately yields the restricted transforms on K. Finally, inversion of these restricted transforms gives back the function f.

Theorem 6 *With all notation as in Theorem 4, the number of operations needed to perform Fourier inversion for functions on G is*

$$\frac{|G|}{|K|} \cdot I(K) + \sum_{i \in I} m_{\eta_i} (\frac{|G|}{|H_{\eta_i}|})^2 d_{\eta_i}^\alpha + |\Delta(\eta_i)| \cdot (\frac{|G|}{|K|} d_{\eta_i}^\alpha + d_{\eta_i}^2 O(\frac{|G|}{|K|} \cdot \log(\frac{|H_{\eta_i}|}{|K|})))$$

Thus, having fast algorithms for performing Fourier inversion and computing Fourier transforms, a fast algorithm for computing convolutions can be given.

Theorem 7 *Let all notation be as in Theorems 4 and 5. Then if K is a normal subgroup of G such that G/K is abelian, then convolution of two functions in $L(G)$ may be computed in at most*

$$\frac{|G|}{|K|} \cdot (2T(K) + I(K)) + \sum_\rho d_\rho^\alpha + 3 \sum_{i \in I} m_{\eta_i} (\frac{|G|}{|H_{\eta_i}|})^2 d_{\eta_i}^\alpha + O(|\Delta(\eta_i)| \frac{|G|}{|K|} \cdot (d_{\eta_i}^\alpha + d_{\eta_i}^2 \log(\frac{|H_{\eta_i}|}{|K|})))$$

operations, where the constant indicated by the O-notation is universal and determined by the FFT.

Proof: Let f and g be in $L(G)$. By (2.1) the Fourier transform of their convolution at any irreducible representation ρ of G is simply the product to their individual Fourier transforms at ρ. Thus, compute the Fourier transforms of f and g at all irreducible representations of G, multiply the matrices to obtain the Fourier transform of $f * g$ at all the irreducible reporesentations of G and then invert. The asserted bound is but the sum of the number of operations needed to do each of these.

♠

Some crude estimates under the simplifying assumption of exponent for matrix multiplication 2 gives the following bounds:

Corollary *Let the notation be as in Theorem 4. Let K be a normal subgroup of G such that G/K is abelian. Then assuming that $\alpha = 2$, the number of operations needed to compute Fourier transforms on G is at most*

$$\frac{|G|}{|K|} \cdot T(K) + O(|G| \log(\frac{|G|}{|K|}))$$

while the number of operations needed to perform Fourier inversion is at most

$$\frac{|G|}{|K|} \cdot I(K) + O(|G| \log(\frac{|G|}{|K|}))$$

and the number of operations needed to convolve two functions on G is at most

$$\frac{|G|}{|K|} \cdot (2T(K) + I(K)) + |G| + O(|G| \log(\frac{|G|}{|K|})).$$

Proof: Consider the inner sum in Theorem 4 in two parts. The first expression is now

$$\sum_{i \in I} m_{\eta_i} (\frac{|G|}{|\Delta(\eta_i)|} \cdot d_{\eta_i})^2.$$

However $\frac{|G|}{|H_{\eta_i}|} \cdot d_{\eta_i}$ is just the degree of each irreducible induced by the extensions of irreducibles in $\Delta(\eta_i)$. Consequently, this sum is just a rewriting of the sum

$$\sum_{irreducible\ \rho} d_\rho^2$$

and so is equal to $|G|$.

Similarly, the sum

$$\sum_{i \in I} |\Delta(\eta_i)| \cdot d_{\eta_i}^2$$

is a rewriting of the sum of squares of the degrees of the irreducible representations of K so that it is equal to $|K|$. Lastly, replacing $\frac{|H_{\eta_i}|}{|K|}$ by $\frac{|G|}{|K|}$ and multiplying things out gives the result.

♠

Any solvable group can be written as a sequence of abelian extensions. Consequently,

Corollary *Suppose G is a solvable group and the exponent for matrix multiplicaition is 2. Then the number of operations needed to compute Fourier transforms on G or perform Fourier inversion on G is at most $O(|G|\log(|G|))$.*

If G is an abelian extension of an abelian subgroup (ie. metabelian), then Theorem 4 says that all irreducible representations of G are induced from one-dimensional representations. Since $1^\alpha = 1$ always,

Theorem 8 *let G be metabelian. Then*

$$T(G) = I(G) = O(|G|\log(|G|)).$$

Actually, in the case of metabelian groups a slight rearrangement of the sum in Theorem 5 gives a more precise result.

Theorem 9 *Let K be a normal abelian subgroup of G such that G/K is abelian. Let m_d be the number of inequivalent irreducible representations of G of dimension d. Then both $T(G)$ and $I(G)$ are bounded by*

$$O(|G|\log(|K|) + |G| + \sum_{d|[G:K]} d^2 m_d \log(\frac{[G:K]}{d}))$$

where the O-notation indicates a universal constant determined by the FFT.

As always, having fixed a set of coset representatives, compute all restricted transforms on K. This requires $O(|G|\log(|K|))$ operations. Following the algorithm as presented in Theorem 5 all the representations of K (one-dimensional) are extended as far as possible and induced to G where they give irreducible representations. Representations which extend exactly as far as a subgroup of degree $\frac{[G:K]}{d}$ induce irreducible representations of degree d. On the order of $[G:K]\log(\frac{[G:K]}{d})$ additional operations are required to compute the restricted transforms at these extensions.

Let n_d be the number of irreducible representations of G whose extensions eventually induce irreducible representations of G of degree d. Then

$$\frac{[G:K]}{d} n_d = d m_d.$$

Hence, to compute all the restricted transforms for extensions which induce irreducible representations of degree d an additional

$$O(\frac{1}{[G:K]} \cdot d^2 m_d [G:K] \log(\frac{[G:K]}{d}))$$

are required. Only representations whose degree divides $[G:K]$ are possible.

If ρ is the irreducible representation induced by a one-dimensional representation of index d then only an additional d^2 are needed to compute the Fourier transform at ρ. Thus, summing over all ρ, at most an additional $|G|$ operations.

Similarly for recovering the restricted transforms, given the Fourier transform at all irreducible representations ρ at most an additional $|G|$ operations are needed. To go from these to the restricted transforms on K via Fourier inversion the above analysis holds. Finally, collecting terms yields the result.

♠

Many groups of interest occur as metabelian groups. For example the dihedral and generalized quaternion groups are metabelian. A theorem due to M. Hall says that all groups of square-free order are metabelian. The methods given above are readily applied for performing Fourier analysis. As illustration, in section 5 two examples are worked out.

It is worth mentioning here that in practice the exponent for matrix multiplication, α, is 3. There is active research in pursuing what is thought to be the theoretical limit of $\alpha = 2$. Currently, the best theoretical result is $\alpha = 2.38$ (see [CW]). It is straightforward to carry through the preceding analysis using assumptions like $\alpha = 2+\epsilon$ or assuming that $d^2 \log d$ operations are needed to multiply d by d matrices.

5. Examples

To illustrate the methods presented in section 4 two examples are worked out here. The cases in which G contains a normal abelian subgroup K such that either G/K is isomorphic to $\mathbb{Z}/p\mathbb{Z}$ or to an abelian group of order p^2 (so $\mathbb{Z}/p\mathbb{Z} \oplus \mathbb{Z}/p\mathbb{Z}$ or $\mathbb{Z}/p^2\mathbb{Z}$) for some prime p are considered. The first case illustrates the basic construction with a minimum of complexity. Here, to make things concrete the explicit example of the Heisenberg group is carried along.

Most generally the discussion in the first case will apply any time all irreducible representations of K either extend fully to G or else induce irreducible representations in G. As a consequence, this also gives the analysis for A_p, the affine group mod p.

The second case shows how things get more complex as the possibilities for extending and inducing representations of K increase. Here no explicit example is included.

Example 1: K abelian, $G/K \cong \mathbb{Z}/p\mathbb{Z}$.

In this case consider the action of G on $\widehat{K}$. As $\mathbb{Z}/p\mathbb{Z}$ has no nontrivial subgroups, the only possibilities for stabilizers for elements of $\widehat{K}$ are G and K. By Theorem 3, any representation fixed by G has p extensions to G. On the other hand, by Lemma 4, any representation with stabilizer K, or equivalently, any representation with p conjugates with respect to G, induces an irreducible representation of G.

In particular consider the case when G is the **Heisenberg group**, the subgroup of $GL_3(\mathbb{Z}/p\mathbb{Z})$ given by the matrices of the form

$$\begin{pmatrix} 1 & x & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix}$$

where x, y, z are any elements of $\mathbb{Z}/p\mathbb{Z}$. The Heisenberg group is of interest in many areas of mathematics and in particular plays a deep role in the theory of the classical fast Fourier transform (see [AT]). It is easy to check that the subset of matrices above with $x = 0$ form a normal subgroup K that is isomorphic to the abelian group $\mathbb{Z}/p\mathbb{Z} \oplus \mathbb{Z}/p\mathbb{Z}$ under the mapping

$$\begin{pmatrix} 1 & 0 & z \\ 0 & 1 & y \\ 0 & 0 & 1 \end{pmatrix} \mapsto (z, y).$$

Consequently, the representations of K are all one-dimensional, given by characters,

$$\chi_{r,s}\begin{pmatrix}1 & 0 & z\\ 0 & 1 & y\\ 0 & 0 & 1\end{pmatrix} = \chi_r(z)\cdot\chi_s(y) = e^{2\pi i r z/p}e^{2\pi i s y/p}$$

for any r or s with $0 \leq r, s < p$. Thus, $\chi_{r,s}$ is the tensor product of the characters χ_r and χ_s in $\widehat{\mathbb{Z}/p\mathbb{Z}}$.

G/K is isomorphic to $\mathbb{Z}/p\mathbb{Z}$ and the subgroup of matrices of the form

$$\begin{pmatrix}1 & x & 0\\ 0 & 1 & 0\\ 0 & 0 & 1\end{pmatrix}$$

give a complete set of coset representatives for G/K. Let this subgroup be denoted as A.

It is useful here to introduce more concise notation. Let $[z, y]$ denote the matrix

$$\begin{pmatrix}1 & 0 & z\\ 0 & 1 & y\\ 0 & 0 & 1\end{pmatrix}$$

in K and x the matrix

$$\begin{pmatrix}1 & x & 0\\ 0 & 1 & 0\\ 0 & 0 & 1\end{pmatrix}$$

in A.

Any character $\chi_{r,s}$ is determined by its values on the generators $[1, 0]$ and $[0, 1]$ of K. In order to determine the orbits within $\widehat{K}$ it is only necessary to consider the action of A on $\widehat{K}$. Note that

$$x^{-1}\cdot[z, y]\cdot x = [z - xy, y].$$

Consequently,

$$\chi_{r,s}^{(x)}([1,0]) = \chi_{r,s}([1,0]) = \chi_r(1)$$

and

$$\begin{aligned}\chi_{r,s}^{(x)}([0,1]) = \chi_{r,s}([-x,1]) &= \chi_r(-x)\chi_s(1)\\ &= \chi_{s-rx}(1).\end{aligned}$$

Thus,

$$\chi_{r,s} \stackrel{x}{\longmapsto} \chi_{r,s-rx}.$$

So $\widehat{K}$ decomposes as

(i) Characters by fixed by G: $\{\chi_{0,s}\}$, $0 \leq s < p$.

(ii) Characters fixed by K: Existing in orbits of size p given by the sets

$$O_j = \{\chi_{j,s} \mid 0 \leq s < p\}$$

with $j \neq 0$.

The following general theorem determines the number of operations needed to compute the Fourier transforms, after which the case of the Heisenberg group is worked out. Fourier inversion is treated in Theorem 12 below. Although these results could simply be stated as corollaries of Theorem 9, the "special" proofs supplied here give explicit recipes for performing these calculations.

Theorem 10 *Let K be an abelian normal subgroup of G such that $G/K \cong \mathbb{Z}/p\mathbb{Z}$ and let $f \in L(G)$. Then all Fourier transforms of f can be computed in at most*

$$m_1 p^2 + O(m_0 \log p + |G| \log(|K|))$$

operations where m_i is the number of irreducible representations of G of degree p^i and the O-notation indicates the universal constant as determined by the FFT.

Construction and Proof: Fix a generator $s \in G$ for G/K. For each j such that $0 \leq j \leq p-1$ define a function $f_j \in L(K)$ by

$$f_j(t) = f(s^j t) \quad (5.1)$$

for all $t \in K$. The first step is to calculate $\widehat{f_j}(\chi)$ for all j and each $\chi \in \widehat{K}$. As K is abelian this can be done for any f_j in

$$O(|K| \log(|K|))$$

operations. So, for all j

$$O(\frac{|G|}{|K|} |K| \log(|K|)) = O(|G| \log(|K|))$$

operations are needed.

Consider now the problem of computing the Fourier transforms at the irreducible representations of dimension p. Let ρ be any such representation. By Theorem 4, $\rho = (\chi \uparrow G)$ for some $\chi \in \widehat{K}$. Thus, by Clifford theory,

$$\begin{aligned}
\widehat{f}(\rho) &= \sum_{j=0}^{p-1} \rho(s)^j \sum_{t \in K} f_j(t) \begin{pmatrix} \chi^{(1)}(t) & 0 & \dots & 0 \\ 0 & \chi^{(2)}(t) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \chi^{(p)}(t) \end{pmatrix} \\
&= \sum_{j=0}^{p-1} \rho(s)^j \begin{pmatrix} \widehat{f}_j(\chi^{(1)}) & 0 & \dots & 0 \\ 0 & \widehat{f}_j(\chi^{(2)}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \widehat{f}_j(\chi^{(p)}) \end{pmatrix}
\end{aligned}$$

where $\{\chi^{(1)}, \dots, \chi^{(p)}\} = \Delta(\chi)$.

By Proposition 2 the matrices $\{\rho(s^j) \mid 0 \le j \le p-1\}$ have the structure of the matrices giving the regular representation of $G/K \cong \mathbb{Z}/p\mathbb{Z}$. That is, s may be chosen so that $\rho(s)$ has the form

$$\rho(s) = \begin{pmatrix} 0 & 0 & \dots & 0 & a^1_1 \\ a^1_2 & 0 & \dots & 0 & 0 \\ 0 & a^1_3 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & a^1_p & 0 \end{pmatrix}$$

where the a^1_i are nonzero and in general, $\rho(s)^j$ has p nonzero entries $\{a^j_1, \dots, a^j_p\}$. Thus,

$$\widehat{f}(\rho) = \begin{pmatrix} \widehat{f}_0(\chi^{(1)}) & a^{p-1}_1 \widehat{f_{p-1}}(\chi^{(2)}) & \dots & a^2_1 \widehat{f}_2(\chi^{(p-1)}) & a^1_1 \widehat{f}_1(\chi^{(p)}) \\ a^1_2 \widehat{f}_1(\chi^{(1)}) & \widehat{f}_0(\chi^{(2)}) & \dots & a^3_2 \widehat{f}_3(\chi^{(p-1)}) & a^2_2 \widehat{f}_2(\chi^{(p)}) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a^{p-1}_p \widehat{f_{p-1}}(\chi^{(1)}) & a^{p-2}_p \widehat{f_{p-2}}(\chi^{(2)}) & \dots & a^1_p \widehat{f}_1(\chi^{(p-1)}) & \widehat{f}_0(\chi^{(p)}) \end{pmatrix} \quad (5.2)$$

Consequently, the above sum may be performed in $d_\rho^2 = p^2$ additional operations. Doing this for each of the m_1 such representations yields $m_1 p^2$ operations in total.

The form of the matrix $\widehat{f}(\rho)$ is crucial. That is, up to scalars, the entries of $\widehat{f}(\rho)$ are just the restricted transforms at the characters which induce ρ. This is the key to performing Fourier inversion.

Now all that is left is to compute the Fourier transforms at the one-dimensional representations of G. By Theorem 3 the characters come in families of size p all of which extend some specific character of K. Applying Lemma 2, for each such family all the Fourier transforms may be computed in $O(p \log p)$ operations. In total then, an additional

$$O(\frac{m_0}{p} p \log p)$$

operations are accumulated. Simplifying and summing the terms gives the asserted result.

♠

It is instructive to note that in the case in which G is itself abelian the algorithm may be viewed as a generalization of the Cooley-Tukey method. In this case, $m_1 = 0$ and it may be assumed that the generator s of G/K is such that $s^p = 1$. Then $\lambda = 1$ and thus $F_{\lambda,\chi}(i) = \widehat{f}_i(\chi)$ for each i.

The constructions in Theorem 10 are readily applied to the Heisenberg group. Let $a \in A$ be the generator

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

Then $a^{-i} a a^{i+1} = 1$ in G. Thus for $\chi_{j,o} \in O_j$

$$(\chi_{j,0} \uparrow G)(a) = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix}$$

and

$$\widehat{f}(\chi_{j,0} \uparrow G) = \begin{pmatrix} \widehat{f}_0(\chi_{j,0}) & \widehat{f_{p-1}}(\chi_{j,1}) & \dots & \widehat{f}_2(\chi_{j,p-2}) & \widehat{f}_1(\chi_{j,p-1}) \\ \widehat{f}_1(\chi_{j,0}) & \widehat{f}_0(\chi_{j,1}) & \dots & \widehat{f}_3(\chi_{j,p-2}) & \widehat{f}_2(\chi_{j,p-1}) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \widehat{f_{p-1}}(\chi^{(j)}0) & \widehat{f_{p-2}}(\chi_{j,1}) & \dots & \widehat{f}_1(\chi_{j,p-2}) & \widehat{f}_0(\chi_{j,p-1}) \end{pmatrix}. \quad (5.3)$$

Thus, actually no additional operations are needed to form $\widehat{f}(\chi_{j,0} \uparrow G)$.

As for the characters $\chi_{0,s}$ fixed by G, the constructions of Theorem 3 show how to extend these to p one-dimensional representations of G. Since $a^p = 1$, the procedure is to take some p^{th} root of 1, λ and define a representation of G, $\chi^{\lambda}_{0,s}$ by

$$\begin{aligned} \chi^{\lambda}_{0,s}(a^j \cdot [z, y]) &= \lambda^j \chi_{0,s}([z, y]) \qquad (5.4) \\ &= \psi_{\lambda}(j)\chi_{0,s}([z, y]) \end{aligned}$$

where ψ_{λ} is the corresponding character of $\mathbb{Z}/p\mathbb{Z}$. Thus,

$$\widehat{f}(\chi^{\lambda}_{0,s}) = \widehat{F_s}(\lambda) \qquad (5.5)$$

where $F_s \in L(\mathbb{Z}/p\mathbb{Z})$ is defined by

$$F_s(j) = \widehat{f_j}(\chi_{0,s}).$$

Theorem 11 *Let G denote the Heisenberg group. Then*

$$T(G) = O(|\,G\,| \log(p^2) + p^2 \log p).$$

<u>Remark 1:</u> When G is the Heisenberg group, the fact that the coset representatives formed a subgroup of G exactly reflects that G is the semidirect product of $\mathbb{Z}/p\mathbb{Z} \oplus \mathbb{Z}/p\mathbb{Z}$ by $\mathbb{Z}/p\mathbb{Z}$. In another case of interest, the affine group mod p, A_p, this is also given as a semidirect product of abelian groups, this time $\mathbb{Z}/p\mathbb{Z}$ by $(\mathbb{Z}/p\mathbb{Z})^*$. In this case, the characters of $\mathbb{Z}/p\mathbb{Z}$ are split into two orbits under the action of $(\mathbb{Z}/p\mathbb{Z})^*$ The trivial character is fixed while the nontrivial characters comprise the other orbit. The trivial character extends in $p-1$ ways to A_p and a single irreducible representation of degree $p-1$ is induced by the other orbit. Thus,

$$T(A_p) = O(|\,A_p\,| \log p + (p-1)\log(p-1) + (p-1)^2)$$

In general, if G is the semidirect product of two abelian groups H and K with K normal in G, the characters of K will in fact extend all the way to their stabilizers (see [S], Proposition 25). This simplifies the analysis greatly and is very useful for computation. In the context of spectral analysis for data on finite groups, explicit examples of this type,

A_{p^n} $(= (\mathbb{Z}/p^n\mathbb{Z})^* \ltimes \mathbb{Z}/p^n\mathbb{Z})$ and $\mathbb{Z}/n\mathbb{Z} \ltimes (\mathbb{Z}/2\mathbb{Z})^n$ have been worked out in detail ([R2]).

Remark 2: It is useful to try to understand which is the dominant term in the expected running time for this algorithm. If K has order prime to p then in fact G is the direct product of K and G/K and the methods of abelian FFT's are used. Thus, assuming that p divides K and the above algorithm is used, the analysis in the proof of Theorem 10 shows that

$$m_0 \log p + m_1 p^2 \;\; < |G| \log p \;\; \leq |G| \log(|K|)$$

so computing the restrictions to K dominates the running time.

Now consider the problem of performing Fourier inversion.

Theorem 12 *With notation as in Theorem 10, given the matrices $\widehat{f}(\rho)$ and $\rho(s)$ for all irreducible representations ρ and all s in G, the number of operations needed to determine f has the same bound as the number needed to calculate the Fourier transforms.*

Proof: The first observation is that it is sufficient to determine the p restricted functions $f_j \in L(K)$ defined in equation (5.1). The other key observation is that, as shown by figure (5.2) in the proof of Theorem 10, the p-dimensional Fourier transforms actually contain as entries (up to scalars) the restricted transforms $\{\widehat{f_j}(\chi)\}$, where j is between 0 and $p-1$ and χ ranges over all characters of K which have stars of size p. Thus, at most $m_1 p^2$ operations are needed to determine the restricted transforms

$$\{\widehat{f_j}(\chi) \mid 0 \leq j \leq p-1, \chi \in \widehat{K}, \mid \Delta(\chi) \mid = p\}.$$

Note that the values needed to divide by to determine the restricted transforms are simply the entries in the representations of the chosen coset representatives.

Now consider the Fourier transforms at the one-dimensional representations of G. Again, these come in families of size p of Fourier transforms of characters that extend some specific character of K. Using Lemma 3 the restricted transforms at this character may be determined in $O(p \log p)$ operations.

Thus, on the order of an additional $\frac{m_0}{p} p \log p$ or $m_0 \log p$ operations are required to determine the restricted transforms

$$\{\widehat{f}_j(\chi) \mid 0 \leq j \leq p-1, \mid \Delta(\chi) \mid = 1\}.$$

As K is abelian, Fourier inversion for any function in $L(K)$ may be peformed in at most $O(\mid K \mid \log(\mid K \mid))$ operations. This must be done $\frac{|G|}{|K|}$ times, so $O(\mid G \mid \log(\mid K \mid))$ in all. Summing all the terms yields the result.

♠

Again, in the case in which G is the Heisenberg group, the application is straightforward. In this situation, figure (5.3) shows that the p-dimensional representations actually contain as entries the restricted transforms $\widehat{f}_j(\chi_{r,s})$ for $r \neq 0$ and thus, no additional operations are needed to recover these.

To recover the restricted transforms at the characters $\chi_{0,s}$ fixed by G, the extensions $\chi_{0,s}^{\lambda}$ (see (5.4)) must be decomposed. Using (5.5) this requires only that the Fourier transforms $\widehat{F_s}(\lambda)$ be inverted. Thus, a bound identical to Theorem 11 is obtained for performing inversion.

Theorem 13 *Let G denote the Heisenberg group. Then*

$$T(G) = O(\mid G \mid \log(p^2) + p^2 \log p).$$

<u>Example 2:</u> K abelian, G/K abelian of order p^2.

It is instructive to also consider the case of an extension of order p^2. That is, let K be an abelian normal subgroup of G such that either $G/K \cong \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$ or $G/K \cong \mathbb{Z}/p^2\mathbb{Z}$. The bounds achieved here are independent of the particular case. The complications here arise from the possibilty of having to extend characters more than once.

Suppose that G is an abelian extension of K of order p^2. Then $\widehat{K}$ is partitioned into orbits of 3 sizes:

(i) orbits of size p^2.
(ii) orbits of size p, and
(iii) orbits of size 1.

Let these characters be of types (i), (ii), and (iii) respectively. As in example 1 the possibilities for the orbits will organize both the analysis of the representations theory and the subsequent computations.

Lemma 4 says that characters of type (i) induce irreducible representations of dimension p^2. Theorem 3 shows that the characters of type (ii) extend to characters of the stabilizer of their orbit which further split into orbits of size p under the action of G. Consequently, as characters of a subgroup of index p in G that lie in an orbit of size p with respect to the action of G, they in turn induce an irreducible representation of dimension p in G - one for each such orbit, so p in total.

Finally consider the characters of type (iii). Pick <u>any</u> subgroup H of degree p over K. As K is also normal in H, H acts on $\widehat{K}$ as well and any character fixed by G is necessarily fixed by H. So, these characters extend to characters of H. G acts on the extensions and under this action one of two things may happen. Either the p extensions of any character become conjugate, or they remain fixed. If they remain fixed they extend further to characters of G. If they become conjugate, then as in case (ii) the orbit induces an irreducible representation of dimension p. The following diagram illustrates case these possibilities:

All this gives,

Theorem 14 *Let K be an abelian normal subgroup of G where either*

$$G/K \cong \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$$

or

$$G/K \cong \mathbb{Z}/p^2\mathbb{Z}.$$

Let ρ be a representation of G of dimension d_ρ. Then
(i) If $d_\rho = p^2$ then ρ is induced by a character of type (i).
(ii) If $d_\rho = p$ then either ρ is induced by a character of type (ii) or by a character of K that only extends to a subgroup of degree p.
(iii) If $d_\rho = 1$ then ρ is the repeated extension of a character of type (iii).

Taking a cue from the first example, the Fourier transforms of f are computed in stages. First all restricted transforms on K are computed. Then the computations of Fourier transforms at the three different types of representations are considered separately. Computations at the largest representation will require only negligible additional scalar multiplication. At the next largest representations (of dimension p) one "level" of computing Fourier transfroms on $\mathbb{Z}/p\mathbb{Z}$ will be required after isolating the restricted Fourier transforms which appear. Lastly, at the the characters of G the Fourier transforms will require computation of Fourier transforms over abelian groups of size p^2. This is made precise in the following theorem.

Theorem 15 *Let K be an abelian normal subgroup of G such that $[G\colon K] = p^2$ and let $f \in L(G)$. Then all Fourier transforms of f can be computed in at most*

$$m_2 p^4 + m_1 p^2 + O(|G| \log(|K|) + m_1 p^2 \log p + m_0 \log(p^2))$$

operations where m_i is the number of irreducible representations of G of degree p^i and the O-notation indicates the universal constant as determined by the FFT.

Proof: As before, first the restricted transforms

$$\{\widehat{f}_j(\chi) \mid \chi \in \widehat{K},\ 0 \le j \le p^2 - 1\}.$$

are computed. This requires $O(|G| \log(|K|))$ operations. Again following example 1 all the Fourier transforms at all the irreducible representations of degree p^2 may now be computed, each with at most an additional $(p^2)^2$ operations. Thus, in total $m_2 p^4$ more operations computes all Fourier transforms at all the p^2-dimensional representations.

This takes care of the contributions due to characters in orbits of size p^2. Turning to those that lie in orbits of size p, extend each of these characters to their respective p extensions within their respective stabilizers.

The function f has p restrictions to any such stabilizer. As in example 1 all the Fourier transforms of any of these restrictions at each of the p extensions of any fixed $\chi \in \widehat{K}$ may be computed in $O(p \log p)$ operations. The same is true for those characters of K fixed by G which may only be extended as far as degree p. It is important to note that whether or not the extensions are conjugate, in either case, all restricted transforms at each of the p extensions must be computed. Thus, in total, on the order of an additional $m_1 p^2 \log p$ operations are needed. (The extra factor of p comes from the p restrictions to the various subgroups of index p.) Finally, summing over the restrictions, at most another $m_1 p^2$ operations are needed.

Lastly the Fourier transforms at the characters of G are computed. Theorem 4 shows that these come in families of size p^2 extending specific characters of K, all of whose Fourier transforms may be computed in $O(p^2 \log(p^2))$ operations. So, in total, on the order of $\frac{m_0}{p^2} p^2 \log(p^2)$ operations are needed. Simplifying and summing yields the theorem.

♠

Again as in example 1, the algorithm may be reversed to perform Fourier inversion in at worst, the same number of operations.

Theorem 16 *With the notation as in Theorem 12 above, given the matrices $\widehat{f}(\rho)$ and $\rho(s)$ for all irreducible representations ρ and all $s \in G$, the number of operations needed to compute Fourier inversion for f has the same bound as the number needed to calculate Fourier transforms.*

Proof: As in the proof of Theorem 11, first note that up to scalars, the m_2 representations of degree p^2 contain the restricted transforms

$$\{\widehat{f}_j(\chi) \mid \chi \in \widehat{K}, \ |\Delta(\chi)| = p^2, \ 0 \le j \le p^2 - 1\}.$$

Thus, at most $m_2 p^4$ operations are needed to determine these.

Similarly, as in (5.2), up to scalars, the Fourier transforms of degree p contain the restricted transforms at extensions of characters of K that could extend no further than a subgroup of degree p. Again for any single representation of this type, at most p^2 operations are needed to recover these restricted transforms, namely dividing through by the scalars which occur as entries in the matrices in the induced representations. So, in total m_1p^2 operations are needed. The m_1p^2 restricted Fourier transforms which occur in these transforms of degree p come in families of size p as extensions of fixed characters of K. Thus, using Lemma 3 an additional

$$\frac{(m_1p^2)}{p} p \log p$$

or $m_1p^2 \log p$ operations are needed to determine the restricted transforms at these characters.

All that is left is to work back down from the one-dimensional representations of G to the restricted transforms at the characters of K which extended to G. Again, using Lemma 3, for each family of p^2 extensions $O(p^2 \log(p^2))$ operations are needed, so $\frac{m_0}{p^2} p^2 \log(p^2)$ in all. Now the p^2 functions on K can be inverted in $O(|G| \log(|K|))$ additional operations.

♠

6. Strategies for General Semidirect Products by Abelian Groups

When a representation is induced from a subgroup that is not normal the situation can be more complex. To illustrate, in this section the case in which G is the semidirect product by an abelian group is considered. That is, G contains a normal abelian subgroup K and another subgroup H such that

(i) $H \cap K = \{1\}$

(ii) $G = H \cdot K$.

This is denoted as $G = H \ltimes K$.

After a brief review of the representation theory as given in Serre ([S], Section 8.2), ideas are given and compared for calculating Fourier transforms. This work is preliminary in nature but serves as a starting point for the study of fast computations for these groups.

Let $G = H \ltimes K$ where K is abelian. The representation theory of G is determined by the equivalent actions of G and H on $\widehat{K}$. Let

$$\widehat{K} = \cup_{i=1}^{r} \Delta(\chi_i)$$

be the decomposition of $\widehat{K}$ into orbits under this action. Define

$$H_i = Stab_H(\chi_i)$$

and

$$G_i = H_i \ltimes K = Stab_G(\chi_i).$$

Let ρ be any representation of H_i. Then ρ extends by the identity on K to a representation $\tilde{\rho}$ of G_i. Similarly, because H_i stabilizes χ_i, χ_i extends by the identity to G_i. Call this extension $\tilde{\chi}_i$. Then the tensor product $\tilde{\rho} \otimes \tilde{\chi}_i$ defines an irreducible representation of G_i.

$$\tilde{\rho} \otimes \tilde{\chi}_i(ts) = \rho(t)\chi_i(s)$$

for all $t \in H_i$ and $s \in K$. Let

$$\Theta_{\rho,i} = (\tilde{\rho} \otimes \tilde{\chi}_i \uparrow G).$$

Note that if H is abelian then this means that χ_i has been extended to its stabilizer. An application of Lemma 4 then proves that representations $\{\Theta_{\rho,i}\}$ give all the irreducible representations of G. For arbitrary H this is still true ([S], Proposition 25).

Consider now the problem of computing all Fourier transforms for such groups. If G_i is normal in G then the ideas of section 4 may be used. However it need not be the case that G_i is normal in G. Consequently, the restriction of $\Theta_{\rho,i}$ to G_i may not be diagonal and in general, the structure of the matrix of the induced representations will not be so well-behaved. However, $(\Theta_{\rho,i} \downarrow K)$ will be diagonal and in fact, it is easy to see that

$$(\Theta_{\rho,i} \downarrow K) = d_\rho \Delta(\chi_i).$$

This suggests two possible approaches to computing Fourier transforms of a function f defined on G.

<u>Method 1:</u>

(a) Compute all Fourier transforms for the $|K|$ restrictions of f to H.

(b) Compute every Fourier transform of f as

$$\hat{f}(\Theta_{\rho,i}) = \sum_{s \in K} \Theta_{\rho,i}(s)\hat{f}_s(\Theta_{\rho,i} \downarrow H).$$

Step (a) above can be performed in $|K|\ T(H)$ operations. Note that here, even though H need not be abelian, fast algorithms for computing Fourier transforms on H may exist. In particular, methods suggested in [DR] may prove useful.

Now, for each i and ρ step (b) must be performed. As $(\Theta_{\rho,i} \downarrow K)$ is diagonal the matrix multiplication here takes $d^2_{\Theta_{\rho,i}}$ operations. So, in total an additional $|K||G|$ operations will be needed. This is recorded as

Proposition 6 *Let $G = H \ltimes K$ where K is a normal abelian subgroup of G. Then Method 1 takes at most $O(|K|(|G| + T(H)))$ operations to compute all Fourier transforms on G.*

Method 2:

(a) Compute all Fourier transforms of the $|H|$ different restrictions of f to K.

(b) Compute all Fourier transforms as

$$\hat{f}(\Theta_{\rho,i}) = \sum_{t \in H} \Theta_{\rho,i}(t) \hat{f}_t(\Theta_{\rho,i} \downarrow K).$$

Since $T(K)$ is at most $O(|K| \log(|K|))$ operations, step (a) takes at most $O(|G| \log(|K|))$ operations. Again, these matrices are all diagonal, so that an additional $d^2_{\Theta_{\rho,i}}$ operations for each matrix multiplication. So, in total there are an additional $|K||G|$ operations. In total,

Proposition 7 *Let $G = H \ltimes K$ where K is a normal abelian subgroup of G. Then Method 2 takes at most $O(|G|(|H| + \log(|K|))$ operations to compute all Fourier transforms on G.*

It is instructive to make some comparisons here. In cases in which $T(H) = O(|H| \log(|H|))$ Method 1 yields a running time of

$$|G|(|K| + \log(|H|))$$

operations, versus method 2 which gives an estimate of

$$|G|(|H| + \log(|K|))$$

operations.

In a particular case of interest, $B_n = S_n \ltimes (\mathbb{Z}/2\mathbb{Z})^n$, use of the ideas in [DR] gives an algorithm for $T(S_n)$ which has a running time crudely bounded by $O((n!)^{3/2})$. Then Method 1 estimates a running time of on the order of

$$|(\mathbb{Z}/2\mathbb{Z})^n|(|B_n| + T(S_n)) = 2^n(2^n n! + (n!)^{3/2})$$

operations, while method 2 gives

$$| B_n | (| S_n | + \log((\mathbb{Z}/2\mathbb{Z})^n)) = 2^n n!(n! + n \log 2).$$

Acknowledgements

I would like to thank Persi Diaconis, my thesis advisor, for originally suggesting that I work on noncommutative FFT's and for reading early versions of this paper. Most of all, I would like to thank him for his continuing encouragement and friendship. I would also like to thank Joseph Bernstein for helpful conversations and for discovering an error in an early proof which led to a much cleaner and fuller development of the material here.

References

[AHU] Aho, A., Hopcroft, J., and Ullman, J. (1976). *The Design and Analysis of Computer Algorithms.* Addison Wesley, Reading, Mass.

[AT] Auslander, L. and Tolmieri, R. (1979). Is Computing with the Fast Fourier Transform Pure or Applied Mathematics? *Bulletin (New Series) of the AMS Vol. 1* 847-897.

[B] Beth, T. (1984). *Vehrfahren der Schnellen Fourier-Transform.* Teubner, Stuttgart.

[C1] Clausen, M. (1988). Fast Generalized Fourier Transforms. To appear in *Theoretical Computer Science*

[C2] Clausen, M. (1988). Spectral Transforms for Symmetric Groups. *Proceedings of Third International Workshop on Spectral Techniques.*

[C3] Clausen, M. (1988). Fast Fourier Transforms for Metabelian Groups. To appear in *Siam J. Comp.*

[Cl] Clifford, A.H., (1937). Representations Induced in an Invariant Subgroup. *Annals of Mathematics Vol. 38*, 533-550.

[Co] Coleman, A.J. (1966). *Induced Representations with Applications to* S_n *and* GL_n., Queen's Papers in Pure and Applied Mathematics, No. 4.

[CW] Coppersmith, D. and Winograd, S. (1987). Matrix Multiplication via Arithmetic Progressions. *Proceedings of the* 19^{th} *ACM Symposium on the Theory of Computing* (STOC), 1-6.

[CR] Curtis, C. and Reiner, I., (1962). *Representation Theory of Finite Groups and Associative Algebras.*, Wiley, New York.

[D1] Diaconis, P. (1987). Spectral Analysis for ranked data. To appear in *Ann. Statist.*

[D2] Diaconis, P. (1988). *Group Representations in Probability and Statistics.* Institute of Mathematical Statistics, Hayward, CA.

[D3] Diaconis, P. (1980). Average Running Time for the Fast Fourier Transform. *J. of Algorithms 1*, 187-208.

[DR] Diaconis, P. and Rockmore, D. (1988). Efficient Computation of Fourier Transforms on Finite Groups, *To appear.*

[DS] Diaconis P. and Shashahani, M. (1986). Products of Random Matrices as They Arise in the Study of Random Walks on Groups. *Contemporary Mathematics* **50**, 183-195.

[R1] Rockmore, D. (1988). Computation of Fourier Transforms on the Symmetric Group. Submitted to *1989 Conference in Computers and Mathematics.*

[R2] Rockmore, D. (1988). Fast Fourier Analysis for A_{p^n} and $(\mathbb{Z}/2\mathbb{Z})^n \ltimes \mathbb{Z}/n\mathbb{Z}$. *Unpublished manuscript.*

[R3] Rockmore, D. (1988). Efficient Computation of Fourier Inversion for Finite Groups. *In preparation.*

[W] Winograd, S. (1978). On Computing the discrete Fourier transform. *Math. Comp. 32*, 175-199.