

The Cooley–Tukey FFT and Group Theory

DAVID K. MASLEN AND DANIEL N. ROCKMORE

ABSTRACT. In 1965 J. Cooley and J. Tukey published an article detailing an efficient algorithm to compute the Discrete Fourier Transform, necessary for processing the newly available reams of digital time series produced by recently invented analog-to-digital converters. Since then, the Cooley–Tukey Fast Fourier Transform and its variants has been a staple of digital signal processing.

Among the many casts of the algorithm, a natural one is as an efficient algorithm for computing the Fourier expansion of a function on a finite abelian group. In this paper we survey some of our recent work on he “separation of variables” approach to computing a Fourier transform on an arbitrary finite group. This is a natural generalization of the Cooley–Tukey algorithm. In addition we touch on extensions of this idea to compact and noncompact groups.

Pure and Applied Mathematics: Two Sides of a Coin

The *Bulletin of the AMS* for November 1979 had a paper by L. Auslander and R. Tolimieri [3] with the delightful title “Is computing with the Finite Fourier Transform pure or applied mathematics?” This rhetorical question was answered by showing that in fact, the finite Fourier transform, and the family of efficient algorithms used to compute it, the Fast Fourier Transform (FFT), a pillar of the world of digital signal processing, were of interest to both pure and applied mathematicians.

Mathematics Subject Classification: 20C15; Secondary 65T10.

Keywords: generalized Fourier transform, Bratteli diagram, Gel’fand–Tsetlin basis, Cooley–Tukey algorithm.

This paper originally appeared in *Notices of the American Mathematical Society* 48:10 (2001), 1151–1160. Parts of the introduction are similar to the paper “The FFT: an algorithm the whole family can use”, which appeared in *Computing in Science and Engineering*, January 2000, pp. 62–67. Rockmore is supported in part by NSF PFF Award DMS-9553134, AFOSR F49620-00-1-0280, and DOJ 2000-DT-CX-K001. He would also like to thank the Santa Fe Institute and the Courant Institute for their hospitality during some of the writing.

Auslander had come of age as an applied mathematician at a time when pure and applied mathematicians still received much of the same training. The ends towards which these skills were then directed became a matter of taste. As Tolimieri retells it (private communication), Auslander had become distressed at the development of a separate discipline of applied mathematics which had grown apart from much of core mathematics. The effect of this development was detrimental on both sides. On the one hand applied mathematicians had fewer tools to bring to problems, and conversely, pure mathematicians were often ignoring the fertile bed of inspiration provided by real world problems. Auslander hoped their paper would help mend a growing perceived rift in the mathematical community by showing the ultimate unity of pure and applied mathematics.

We will show that investigation of finite and fast Fourier transforms continues to be a varied and interesting direction of mathematical research. Whereas Auslander and Tolimieri concentrated on relations to nilpotent harmonic analysis and theta functions, we emphasize connections between the famous *Cooley–Tukey FFT* and group representation theory. In this way we hope to provide further evidence of the rich interplay of ideas which can be found at the nexus of pure and applied mathematics.

1. Background

The finite Fourier transform or discrete Fourier transform (DFT) has several representation theoretic interpretations: either as an exact computation of the Fourier coefficients of a function on the cyclic group $\mathbb{Z}/n\mathbb{Z}$ or a function of band-limit n on the circle S^1 , or as an approximation to the Fourier transform of a function on the real line. For each of these points of view there is a natural group-theoretic generalization, and also a corresponding set of efficient algorithms for computing the quantities involved. These algorithms collectively make up the *Fast Fourier Transform* or *FFT*.

Formally, the DFT is a linear transformation mapping any complex vector of length n , $f = (f(0) \dots, f(n-1))^t \in \mathbb{C}^n$, to its *Fourier transform*, $\hat{f} \in \mathbb{C}^n$. The k^{th} component of \hat{f} , the *DFT of f at frequency k* , is

$$\hat{f}(k) = \sum_{j=0}^{n-1} f(j) e^{2\pi i j k / n} \quad (1-1)$$

where $i = \sqrt{-1}$, and the *inverse Fourier transform* is

$$f(j) = \frac{1}{n} \sum_{k=0}^{n-1} \hat{f}(k) e^{-2\pi i j k / n}. \quad (1-2)$$

Thus, with respect to the standard basis, the DFT can be expressed as the matrix-vector product $\hat{f} = \mathbb{F}_n \cdot f$ where \mathbb{F}_n is the *Fourier matrix* of order n , whose j, k entry is equal to $e^{2\pi i j k / n}$. Computing a DFT directly would require n^2

scalar operations. (For precision's sake: Our count of operations is the number of complex additions of the number of complex multiplications, whichever is greater.) Instead, the FFT is a family of algorithms for computing the DFT of any $f \in \mathbb{C}^n$ in $O(n \log n)$ operations. Since inversion can be framed as the DFT of the function $\check{f}(k) = \frac{1}{n} \widehat{f}(-k)$, the FFT also gives an efficient inverse Fourier transform.

One of the main practical implications of the FFT is that it allows any cyclically invariant linear operator to be applied to a vector in only $O(n \log n)$ scalar operations. Indeed, the DFT diagonalizes any group invariant operator, making possible the following algorithm: (1) compute the Fourier transform (DFT). (2) Multiply the DFT by the eigenvalues of the operator, which are also found using the Fourier transform. (3) Compute the inverse Fourier transform of the result. This technique is the basis of digital filtering and is also used for the efficient numerical solution of partial differential equations.

Some history. Since the Fourier matrix is effectively the character table of a cyclic group, it is not surprising that some of its earliest appearances are in number theory, the subject which gave birth to character theory. Consideration of the Fourier matrix goes back at least as far as to Gauss, who was interested in its connections to quadratic reciprocity. In particular, Gauss showed that for odd primes p and q ,

$$\left(\frac{p}{q}\right) \left(\frac{q}{p}\right) = \frac{\text{Trace } \mathbb{F}_{pq}}{\text{Trace } \mathbb{F}_p \text{ Trace } \mathbb{F}_q}, \quad (1-3)$$

where $\left(\frac{p}{q}\right)$ denotes the Legendre symbol. Gauss also established a formula for the quadratic Gauss sum $\text{Trace } \mathbb{F}_n$, which is discussed in detail in [3].

Another early appearance of the DFT occurs in the origins of representation theory in the work of Dedekind and Frobenius on the group determinant. For a finite group G , the group determinant Θ_G is defined as the homogeneous polynomial in the variables x_g (for each $g \in G$) given by the determinant of the matrix whose rows and columns are indexed by the elements of G with g, h -entry equal to $x_{gh^{-1}}$. Frobenius showed that when G is abelian, Θ_G admits the factorization

$$\Theta_G = \prod_{\chi \in \hat{G}} \left(\sum_{g \in G} \chi(g) x_g \right), \quad (1-4)$$

where \hat{G} is the set of characters of G . The linear form defined by the inner sum in (1-4) is a “generic” DFT at the frequency χ .

In the nonabelian case, Θ_G admits an analogous factorization in terms of irreducible polynomials of the form

$$\Theta_D(G) = \det \left(\sum_{g \in G} D(g) x_g \right)$$

where D is an *irreducible matrix representation* of G . The inner sum here is a generic Fourier transform over G . See [12] for a beautiful historical exposition of these ideas.

Gauss's interests ranged over all areas of mathematics and its applications, so it is perhaps not surprising that the first appearance of an FFT can also be traced back to him [10]. Gauss was interested in certain astronomical calculations, a recurrent area of application of the FFT, necessary for interpolation of asteroidal orbits from a finite set of equally-spaced observations. Surely the prospect of a huge laborious hand calculation was good motivation for the development of a fast algorithm. Making fewer hand calculations also implies less opportunity for error and hence increased numerical stability!

Gauss wanted to compute the Fourier coefficients, a_k, b_k of a function represented by a Fourier series of bandwidth n ,

$$f(x) = \sum_{k=0}^m a_k \cos 2\pi kx + \sum_{k=1}^m b_k \sin 2\pi kx, \quad (1-5)$$

where $m = (n-1)/2$ for n odd and $m = n/2$ for n even. He first observed that the Fourier coefficients can be computed by a DFT of length n using the values of f at equispaced sample points. Gauss then went on to show that if $n = n_1 n_2$, this DFT can in turn be reduced to first computing n_1 DFTs of length n_2 , using equispaced subsets of the sample points, i.e., a subsampled DFT, and then combining these shorter DFTs using various trigonometric identities. This is the basic idea underlying the *Cooley–Tukey FFT*.

Unfortunately, this reduction never appeared outside of Gauss's collected works. Similar ideas, usually for the case $n_1 = 2$ were rediscovered intermittently over the succeeding years. Notable among these is the doubling trick of Danielson and Lanczos (1942), performed in the service of x-ray crystallography, another frequent employer of FFT technology. Nevertheless, it was not until the publication of Cooley and Tukey's famous paper [7] that the algorithm gained any notice. The story of Cooley and Tukey's collaboration is an interesting one. Tukey arrived at the basic reduction while in a meeting of President Kennedy's Science Advisory Committee where among the topics of discussions were techniques for off-shore detection of nuclear tests in the Soviet Union. Ratification of a proposed United States/Soviet Union nuclear test ban depended upon the development of a method for detecting the tests without actually visiting the Soviet nuclear facilities. One idea required the analysis of seismological time series obtained from off-shore seismometers, the length and number of which would require fast algorithms for computing the DFT. Other possible applications to national security included the long-range acoustic detection of nuclear submarines.

R. Garwin of IBM was another of the participants at this meeting and when Tukey showed him this idea Garwin immediately saw a wide range of potential

applicability and quickly set to getting this algorithm implemented. Garwin was directed to Cooley, and, needing to hide the national security issues, told Cooley that he wanted the code for another problem of interest: the determination of the periodicities of the spin orientations in a 3-D crystal of He^3 . Cooley had other projects going on, and only after quite a lot of prodding did he sit down to program the “Cooley–Tukey” FFT. In short order, Cooley and Tukey prepared their paper, which, for a mathematics/computer science paper, was published almost instantaneously—in six months!. This publication, Garwin’s fervent proselytizing, as well as the new flood of data available from recently developed fast analog-to-digital converters, did much to help call attention to the existence of this apparently new fast and useful algorithm. In fact, the significance of and interest in the FFT was such that it is sometimes thought of as having given birth to the modern field of analysis of algorithms. See also [6] and the 1967 and 1969 special issues of the *IEEE Transactions in Audio Electronics* for more historical details.

The Fourier transform and finite groups. One natural group-theoretic interpretation of the Fourier transform is as a change of basis in the space of complex functions on $\mathbb{Z}/n\mathbb{Z}$. Given a complex function f on $\mathbb{Z}/n\mathbb{Z}$, we may expand f , in the basis of irreducible characters $\{\chi_k\}$, defined by $\chi_k(j) = e^{2\pi ijk/n}$. By (1–2) the coefficient of χ_k in the expansion is equal to the scaled Fourier coefficient $\frac{1}{n}\widehat{f}(-k)$, whereas the Fourier coefficient $\widehat{f}(k)$ is the inner product of the vector of function values of f with those of the character χ_k .

For an arbitrary finite group G there is an analogous definition. The characters of $\mathbb{Z}/n\mathbb{Z}$ are the simplest example of a *matrix representation*, which for any group G is a matrix-valued function $\rho(g)$ on G such that $\rho(ab) = \rho(a)\rho(b)$, and $\rho(e)$ is the identity matrix. Given a matrix representation ρ of dimension d_ρ , and a complex function f on G , the *Fourier transform of f at ρ* is defined as the matrix sum

$$\widehat{f}(\rho) = \sum_{x \in G} f(x)\rho(x). \quad (1-6)$$

Computing $\widehat{f}(\rho)$ is equivalent to the computation of the d_ρ^2 scalar Fourier transforms at each of the individual *matrix elements* ρ_{ij} ,

$$\widehat{f}(\rho_{ij}) = \sum_{x \in G} f(x)\rho_{ij}(x). \quad (1-7)$$

A set of matrix representations \mathcal{R} of G is called a *complete set of irreducible representations* if and only if the collection of matrix elements of the representations, relative to an arbitrary choice of basis for each matrix representation in the set, forms a basis for the space of complex functions on G . The Fourier transform of f with respect to \mathcal{R} is then defined as the collection of individual transforms, while *the Fourier transform on G* means any Fourier transform computed with respect to some complete set of irreducibles. In this case, the inverse

transform is given explicitly as

$$f(x) = \frac{1}{|G|} \sum_{\rho \in \mathcal{R}} d_\rho \text{Trace}(\widehat{f}(\rho)\rho(x^{-1})). \quad (1-8)$$

Equation (1-8) shows us a relation between the group Fourier transform and the expansion of a function in the basis of matrix elements. The coefficient of ρ_{ij} in the expansion of f is the Fourier transform of f at the dual representation $[\rho_{ji}(g^{-1})]$ scaled by the factor $d_\rho/|G|$.

Viewing the Fourier transform on G as a simple matrix-vector multiplication leads to some simple bounds on the number of operations required to compute the transform. The computation clearly takes no more than the $|G|^2$ scalar operations required for any matrix-vector multiplication. On the other hand the column of the Fourier matrix corresponding to the trivial representation is all ones, so at least $|G| - 1$ additions are necessary. One main goal of this finite group FFT research is to discover algorithms which can significantly reduce the upper bound for various classes of groups, or even all finite groups.

The current state of affairs for finite group FFTs. Analysis of the Fourier transform shows that for G abelian, the number of operations required is bounded by $O(|G| \log |G|)$. For arbitrary groups G , upper bounds of $O(|G| \log |G|)$ remain the holy grail in group FFT research. In 1978, A. Willsky provided the first non-abelian example by showing that certain metabelian groups had an $O(|G| \log |G|)$ Fourier transform algorithm [20]. Implicit in the big- O notation is the idea that a family of groups is under consideration, with the size of the individual groups going to infinity.

Since Willsky's initial discovery much progress has been made. U. Baum has shown that the supersolvable groups admit an $O(|G| \log |G|)$ FFT, while others have shown that symmetric groups admit $O(|G| \log^2 |G|)$ FFTs (see Section 3). Other groups for which highly improved (but not $O(|G| \log^c |G|)$) algorithms have been discovered include the matrix groups over finite fields, and more generally, the Lie groups of finite type. See [15] for pointers to the literature. There is much work to be done finding new classes of groups which admit fast transforms, and improving on the above results. The ultimate goal is to settle or make progress on the following conjecture:

CONJECTURE 1. There exist constants c_1 and c_2 such that for any finite group G , there is a complete set of irreducible matrix representations for which the Fourier transform of any complex function on the G may be computed in fewer than $c_1|G| \log^{c_2} |G|$ scalar operations.

2. The Cooley–Tukey Algorithm

Cooley and Tukey showed [7] how the Fourier transform on the cyclic group $\mathbb{Z}/n\mathbb{Z}$, where $n = pq$ is composite, could be written in terms of Fourier transforms

on the subgroup $q\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p\mathbb{Z}$. The trick is to change variables, so that the one dimensional formula (1–1) is turned into a two dimensional formula, which can be computed in two stages. Define variables j_1, j_2, k_1, k_2 , through the equations

$$\begin{aligned} j &= j(j_1, j_2) = j_1q + j_2, & 0 \leq j_1 < p, & \quad 0 \leq j_2 < q, \\ k &= k(k_1, k_2) = k_2p + k_1, & 0 \leq k_1 < p, & \quad 0 \leq k_2 < q. \end{aligned} \tag{2-1}$$

It follows from these equations that (1–1) can be rewritten as

$$\widehat{f}(k_1, k_2) = \sum_{j_2=0}^{q-1} e^{2\pi i j_2(k_2p+k_1)/n} \sum_{j_1=0}^{p-1} e^{2\pi i j_1 k_1/p} f(j_1, j_2). \tag{2-2}$$

We now compute \widehat{f} in two stages:

- Stage 1: For each k_1 and j_2 compute the inner sum

$$\widetilde{f}(k_1, j_2) = \sum_{j_1=0}^{p-1} e^{2\pi i j_1 k_1/p} f(j_1, j_2). \tag{2-3}$$

This requires at most p^2q scalar operations.

- Stage 2: For each k_1, k_2 compute the outer sum

$$\widehat{f}(k_1, k_2) = \sum_{j_2=0}^{q-1} e^{2\pi i j_2(k_2p+k_1)/n} \widetilde{f}(k_1, j_2). \tag{2-4}$$

This requires an additional q^2p operations.

Thus, instead of $(pq)^2$ operations, the above algorithm uses $(pq)(p+q)$ operations.

Stage 1 has the form of a DFT on the subgroup $q\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p\mathbb{Z}$, embedded as the set of multiples of q , whereas stage 2 has the form of a DFT on a cyclic group of order q , so if n could be factored further, we could apply the same trick to these DFTs in turn. Thus, if N has the prime factorization $N = p_1 \cdots p_m$, then we recover Cooley and Tukey’s original m -stage algorithm which requires $N \sum_i p_i$ operations [7].

A group-theoretic interpretation. Auslander and Tolmieri’s paper [3] related the Cooley–Tukey algorithm to the Weil–Brezin map for the finite Heisenberg group. Here we present an alternate group-theoretic interpretation, originally due to Beth [4], that is more amenable to generalization.

The change of variables on the first line of (2–1) may be interpreted as the factorization of the group element j as the (group) product of $j_1q \in q\mathbb{Z}/n\mathbb{Z}$, with the coset representative j_2 . Thus, if we write $G = \mathbb{Z}/n\mathbb{Z}$, $H = q\mathbb{Z}/n\mathbb{Z}$, and let Y denote our set of coset representatives, the change of variables can be rewritten as

$$g = y \cdot h, \quad y \in Y, \quad h \in H \tag{2-5}$$

The second change of variables in (2–1) can be interpreted using the notion of restriction of representations. It is easy to see that restricting a representation

on a group G to a subgroup H yields a representation of that subgroup. In the case of $q\mathbb{Z}/n\mathbb{Z}$ this amounts to the observation that

$$e^{2\pi i j_1 q(k_2 p + k_1)/n} = e^{2\pi i j_1 k_1/p},$$

which is used to prove (2-2).

The restriction relations between representations may be represented diagrammatically using a directed graded graph with three levels. At level zero there is a single vertex labeled 1, called the root vertex. The vertices at level one are labeled by the irreducible representations of $\mathbb{Z}/p\mathbb{Z}$, and the vertices at level two are labeled by the irreducible representations of $\mathbb{Z}/n\mathbb{Z}$. Edges are drawn from the root vertex to each of the vertices at level one, and 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 directed graph obtained is the *Bratteli diagram* for the chain of subgroups $\mathbb{Z}/n\mathbb{Z} > \mathbb{Z}p\mathbb{Z} > 1$. Figure 1 shows the situation for the chain $\mathbb{Z}/6\mathbb{Z} > 2\mathbb{Z}/6\mathbb{Z} \cong \mathbb{Z}/3\mathbb{Z} > 1$.




Figure 1. The Bratteli diagram for $\mathbb{Z}/6\mathbb{Z} > 2\mathbb{Z}/6\mathbb{Z} > 1$. The representation χ_k of $\mathbb{Z}/m\mathbb{Z}$ is defined by $\chi_k(l) = e^{2\pi i k l/m}$.

In this way the irreducible representations of $\mathbb{Z}/n\mathbb{Z}$ are indexed by paths (k_1, k_2) in the Bratteli diagram for $\mathbb{Z}/n\mathbb{Z} > \mathbb{Z}/p\mathbb{Z} > 1$. The DFT factorization (2-2) now becomes

$$\widehat{f}(k_1, k_2) = \sum_{y \in Y} \chi_{k_1, k_2}(y) \sum_{h \in H} f(y \cdot h) \chi_{k_1}(h). \tag{2-6}$$

The two-stage algorithm is now restated as first computing a set of sums that depend on only the first leg of the paths, and then combining these to compute the final sums that depend on the full paths.

In summary, the group elements have been indexed according to a particular factorization scheme, while the irreducible representations (the dual group) are now indexed by paths in a Bratteli diagram, describing the restriction of representations. This allows us to compute the Fourier transform in stages, using one fewer group element factor at each stage, but using paths of increasing length in the Bratteli diagram.

3. Fast Fourier Transforms on Symmetric Groups

A fair amount of attention has been devoted to developing efficient Fourier transform algorithms for the symmetric group. One motivation for developing these algorithms is the goal of analyzing data on the symmetric group using a spectral approach. In the simpler case of time series data on the cyclic group, this approach amounts to projecting the data vector onto the basis of complex exponentials.

The spectral approach to data analysis makes sense for a function defined on any kind of group, and such a general formulation is due to Diaconis (see [8], for example). The case of the symmetric group corresponds to considering *ranked data*. For instance, a group of people might be asked to rank a list of 4 restaurants in order of preference. Thus, each respondent chooses a permutation of the original ordered list of 4 objects, and counting the number of respondents choosing each permutation yields a function on S_4 . It turns out that the corresponding Fourier decomposition of this function naturally describes various coalition effects that may be useful in describing the data.

To get some feel for this notice that the Fourier transform at the matrix elements $\rho_{ij}(\pi)$ of the (reducible) defining representation count the number of people ranking restaurant i in position j . If instead ρ is the (reducible) permutation representation of S_n on unordered pairs $\{i, j\}$, then for each choice of $\{i, j\}$ and $\{k, l\}$ the individual Fourier transforms count the number of respondents ranking restaurants i and j in positions k and l . See [8] for a more thorough explanation.

The first FFT for symmetric groups (an $O(|G| \log^3 |G|)$ algorithm) was due M. Clausen. In what follows we summarize recent improvements on Clausen's result.

Example: Computing the Fourier transform on S_4 . The fast Fourier transform for S_4 is obtained by mimicking the group-theoretic approach to the Cooley–Tukey algorithm. More precisely, we shall rewrite the formula for the Fourier transform using two changes of variables: one using factorizations of group elements, and the other using paths in a Bratteli diagram. The former comes from the reduced word decomposition of $g \in S_4$, by which g may be uniquely expressed as

$$g = s_2^4 \cdot s_3^4 \cdot s_4^4 \cdot s_2^3 \cdot s_3^3 \cdot s_2^2, \quad (3-1)$$

where s_i^j is either e or the transposition $(i \ i-1)$, and $s_{i_1}^j = e$ implies that $s_{i_2}^j = e$ for $i_2 \leq i_1$. Thus any function on the group S_4 may be thought of as a function of the 6 variables $s_2^4, s_3^4, s_4^4, s_2^3, s_3^3, s_2^2$.

To index the matrix elements of S_4 paths in a Bratteli diagram are used, this time relative to the chain of subgroups $S_4 \geq S_3 \geq S_2 \geq S_1 \geq 1$. The irreducible representations of S_n are in one-to-one correspondence with partitions of the integer n , with restriction of representations corresponding to deleting a box in

the Young diagram. The corresponding Bratteli diagram is called Young’s lattice, and is shown in Figure 2. Paths in Young’s lattice from the empty partition ϕ




Figure 2. Young’s lattice up to level 4.

to β_4 , a partition of 4, index the basis vectors of the irreducible representation corresponding to β_4 . Matrix elements, however, are determined by specifying a pair of basis vectors, so to index the matrix elements, we must use pairs of paths in Young’s lattice, starting at ϕ and ending in the same partition of 4. Since there are no multiple edges in Young’s lattice, each path may be described by the sequence of partitions $\phi, \beta_1, \beta_2, \beta_3, \beta_4$, through which it passes.

Before we can state a formula for the Fourier transform, analogous to (2–2) and (2–6), we must choose bases for the irreducible representations of S_4 in order to define our matrix elements. Efficient algorithms are known only for special choices of bases, and our algorithm uses the representations in Young’s orthogonal form, which is equivalent to the following equation (3–2) for the Fourier transform in the new sets of variables.

$$\begin{aligned} & \widehat{f} \begin{pmatrix} \beta_4 & \beta_3 & \beta_2 & \beta_1 \\ \gamma_3 & \gamma_2 & \gamma_1 \end{pmatrix} \\ &= \sum_{g=s_2^4 s_3^4 s_4^4 s_3^3 s_2^2 \varphi_2, \varphi_1, \eta_1} \sum \left(P_{s_4^4}^4 \begin{pmatrix} \beta_4 & \beta_3 \\ \gamma_3 & \varphi_2 \end{pmatrix} P_{s_3^3}^3 \begin{pmatrix} \beta_3 & \beta_2 \\ \varphi_2 & \varphi_1 \end{pmatrix} P_{s_2^2}^2 \begin{pmatrix} \beta_2 & \beta_1 \\ \varphi_1 \end{pmatrix} \right. \\ & \quad \left. \times P_{s_3^3}^3 \begin{pmatrix} \gamma_3 & \varphi_2 \\ \gamma_2 & \eta_1 \end{pmatrix} P_{s_2^2}^2 \begin{pmatrix} \varphi_2 & \varphi_1 \\ \eta_1 \end{pmatrix} P_{s_2^2}^2 \begin{pmatrix} \gamma_2 & \eta_1 \\ \gamma_1 \end{pmatrix} f(g) \right). \quad (3-2) \end{aligned}$$

The functions $P_{s_i^i}^i$ in equation (3–2) are defined below, and for each i , the variables $\beta_i, \gamma_i, \varphi_i, \eta_i$ are partitions of i , satisfying the restriction relations described by Figure 3. A solid line between partitions means that the right partition is obtained from the left partition by removing a box.

The relationship between (3–2) and Figure 3 is extremely close—we derived the diagram from the reduced word decomposition first, and then read the equation off the diagram. Each 2-cell in Figure 3 corresponds to a factor in the product of P functions in (3–2), and the labels on the boundary of each cell give the arguments of $P_{s_i^i}^i$. The sum in (3–2) is over those variables occurring in the interior of Figure 3. Thus, the variables describing the Fourier transformed function are exactly those appearing on the boundary of the figure.




Figure 3. Restriction relations for (3–2).

Equation (3–2) can be summarized by saying that we take the product over 2-cells, and sum on interior indices, in Figure 3. This suggests a generalization of the Cooley–Tukey algorithm, that corresponds to building up the diagram one cell at a time. At each stage multiply by the factor corresponding to a 2-cell, and form the diagram consisting of those 2-cells that have been considered so far. Then sum over any indices that are in the interior of the diagram for this stage, but were not in the interior for previous stages. At the end of this algorithm we have multiplied by the factors for each 2-cell, and summed over all the interior indices, and have therefore computed the Fourier transform.

The order in which the cells are added matters, of course. The order $s_2^2, s_3^3, s_3^3, s_2^4, s_3^4, s_4^4$ is known to be most efficient. Here is the algorithm in detail.

- Stage 0: Start with $f(s_2^4 s_3^4 s_4^4 s_2^3 s_3^3 s_2^2)$, for all reduced words.
- Stage 1: Multiply by $P_{s_2}^2$. Sum on s_2^2 .
- Stage 2: Multiply by $P_{s_3}^2$. Sum on s_3^2 .
- Stage 3: Multiply by $P_{s_3}^3$. Sum on η_1, s_3^3 .
- Stage 4: Multiply by $P_{s_4}^2$. Sum on s_4^2 .
- Stage 5: Multiply by $P_{s_3}^4$. Sum on φ_1, s_3^4 .
- Stage 6: Multiply by $P_{s_4}^3$. Sum on φ_2, s_4^4 .

The indices occurring in each stage of the algorithm are shown in Figure 4.

To count the number of additions and multiplications used by the algorithm, we must count the number of configurations in Young’s lattice corresponding to each of the diagrams in Figure 4. This yields a grand total of 130 additions and 130 multiplications for the Fourier transform on S_4 .

The generalization to higher order symmetric groups is straightforward. The reduced word decomposition gives the group element factorization and Young’s orthogonal form allows us to change variables, and the formula and algorithm for the Fourier transform can be read off a diagram generalizing Figure 3. The diagram for S_5 is shown, for example, in Figure 5.




Figure 4. Variables occurring at each stage of the fast Fourier transform for S_4




Figure 5. Restriction relations in the Fourier transform formula for S_5 .


We have computed the exact operation counts for symmetric groups S_n with $n \leq 50$, and a general formula seems hard to come by. (Presumably $n \leq 50$ would cover all cases where the algorithm might ever be implemented, but the same numbers arise in FFTs on homogeneous spaces, which have far fewer elements.)

However, bounds are easier to obtain:

THEOREM 3.1 ([13]). *The number of additions (or multiplications) required by the above algorithm (as generalized to $S_n > S_{n-1} > \dots > S_1$) is exactly*

$$n! \cdot \sum_{k=2}^n \frac{1}{k} \sum_{i=2}^k \frac{1}{(i-1)!} F_i$$

where F_i is the number of configurations in Young's lattice of the form



(3-3)

Furthermore, $F_i \leq 3(1 - \frac{1}{i})i!$, so the number of additions (multiplications) is bounded by $\frac{3}{4}n(n-1) \cdot n!$.

Why stop at S_n ? The algorithm for the FFT on S_n generalizes to any wreath product $S_n[G]$ with the symmetric group. The subgroup chain is replaced by the chain

$$S_n[G] > S_{n-1}[G] \times G > S_{n-1}[G] > \dots > S_2[G] > G \times G > G, \tag{3-4}$$

and the reduced word decomposition is replaced by the factorization

$$x = s_2^n \cdots s_n^n g^n s_2^{n-1} \cdots s_{n-1}^{n-1} g^{n-1} \cdots s_s^2 g^2 g^1. \tag{3-5}$$

Adapting the S_n argument along these lines gives the following new result.

THEOREM 3.2. *The number of operations needed to compute a Fourier transform on $S_n[G]$ is at most*

$$\left(\frac{3n(n-1)}{4} |G| d_G^2 + n(t_G + \frac{1}{4}|G|(h_G d_G^2 - |G|)) \right) |S_n[G]|$$

where h_G is the number of conjugacy classes in G , d_G is the maximal degree of an irreducible representation of G , and t_G is the number of operations required to compute a Fourier transform on G . If G is abelian, then the inner term $h_G d_G^2 - |G| = 0$.

The functions $P_{s_i}^i$ defining Young’s orthogonal form are defined as follows: For any two boxes b_1 and b_2 in a Young diagram, we define the axial distance from b_1 to b_2 to be $d(b_1, b_2)$, where $d(b_1, b_2) = \text{row}(b_1) - \text{row}(b_2) + \text{column}(b_1) - \text{column}(b_2)$. Now suppose $\beta_i, \beta_{i-1}, \alpha_{i-1}, \alpha_{i-2}$ are partitions and that $\alpha_{i-1}, \beta_{i-1}$ are obtained from β_i by removing a box, and are obtained from α_{i-2} by adding a box. Then the skew diagrams of $\beta_i - \beta_{i-1}$ and $\beta_{i-1} - \alpha_{i-2}$ each consist of a single box, and P^i is given by

$$P_e^i \begin{pmatrix} \beta_i & \beta_{i-1} \\ \alpha_{i-1} & \alpha_{i-2} \end{pmatrix} = \begin{cases} 1 & \text{if } \alpha_{i-1} = \beta_{i-1}, \\ 0 & \text{if } \alpha_{i-1} \neq \beta_{i-1}. \end{cases}$$

$$P_{(i \ i-1)}^i \begin{pmatrix} \beta_i & \beta_{i-1} \\ \alpha_{i-1} & \alpha_{i-2} \end{pmatrix} = \begin{cases} d(\beta_i - \beta_{i-1}, \beta_{i-1} - \alpha_{i-2})^{-1} & \text{if } \alpha_{i-1} = \beta_{i-1}, \\ \sqrt{1 - d(\beta_i - \beta_{i-1}, \beta_{i-1} - \alpha_{i-2})^{-2}} & \text{if } \alpha_{i-1} \neq \beta_{i-1}. \end{cases} \tag{3-6}$$

For a proof of this formula, in slightly different notation, see [11], Chapter 3.

4. Generalization to Other Groups

The FFT described for symmetric groups suggests a general approach to computing Fourier transforms on finite groups. Here is the recipe.

- (i) Choose a chain of subgroups

$$G = G_m \geq G_{m-1} \geq \dots \geq G_1 \geq G_0 = 1 \tag{4-1}$$

- for the group. This determines the Bratteli diagram that we will use to index the matrix elements of G . In the general case, this Bratteli diagram may have multiple edges, so a path is no longer determined by the nodes it visits.
- (ii) Choose a factorization $g = g_n \cdot g_{n-1} \cdots g_1$ of each group element g . Choose the g_i so that they lie in as small a subgroup G_k as possible, and commute with as large a subgroup G_l as possible.
 - (iii) Choose a system of Gel'fand–Tsetlin bases [9] for the irreducible representations of G relative to the chain (4–1). These are bases that are indexed by paths in the Bratteli diagram, that behave well under restriction of representations. Relative to such a basis, the representation matrices of g_i will be block diagonal whenever g_i lies in a subgroup from the chain, and block scalar whenever g_i commutes with all elements of a subgroup from the chain.
 - (iv) Now write the Fourier transform in coordinates, as a function of the pairs of paths in the Bratteli diagram with a common endpoint, and with the original function written as a function of g_1, \dots, g_n . This will be a sum of products indexed by edges in the Bratteli diagram which lie in some configuration generalizing (3). This configuration of edges specifies the way in which the nonzero elements of the representation matrices appear in the formula for the Fourier transform in coordinates.
 - (v) The algorithm proceeds by building up the product piece by piece, and summing on as many partially indexed variables as possible.

Further considerations and generalizations. The efficiency of the above approach, both in theory, in terms of algorithmic complexity, and practice, in terms of execution time, depends on both the choice of factorization and the Gel'fand–Tsetlin bases. In particular, very interesting work of L. Auslander, R. Johnson and J. Johnson [2] shows how in the abelian case, different factorizations correspond to different well-known FFTs, each well suited for execution on a different computer architecture. This work shows how to relate the 2-cocycle of a group extension to construction of the important “twiddle factor” matrix in the factorization of the Fourier matrix. It marks the first appearances of group cohomology in signal processing and derives an interesting connection between group theory and the design of retargetable software.

The analogous questions for nonabelian groups and other important signal processing transform algorithms, that is, the problem of finding architecture-optimized factorizations, is currently being investigated by the SPIRAL project at Carnegie Mellon [19].

Another abelian idea: the “chirp-z” FFT. The use of subgroups depends upon the existence of a nontrivial subgroup. Thus, for a reduction in the case of a cyclic group of prime order, a new idea is necessary. In this case, C. Rader’s “chirp-z transform” (the “chirp” here refers to radar chirp—the generation of an extremely short electromagnetic pulse, i.e., something approaching the ideal delta function) may be used [16].

The chirp-z transform proceeds by turning computation of the DFT into computation of convolution on a different, albeit related, group. Let p be a prime. Since $\mathbb{Z}/p\mathbb{Z}$ is also a finite field, there exists a generator g of $\mathbb{Z}/p\mathbb{Z}^\times$, a cyclic group (under multiplication) of order $p - 1$. Thus, for any $f : \mathbb{Z}/p\mathbb{Z} \rightarrow \mathbb{C}$ and nonzero frequency index g^{-b} , we can write $\hat{f}(g^{-b})$ as

$$\hat{f}(g^{-b}) = f(0) + \sum_{a=0}^{p-2} f(g^a) e^{2\pi i g^{a-b}/p}. \quad (4-2)$$

The summation in (4-2) has the form of a convolution on $\mathbb{Z}/(p-1)\mathbb{Z}$, of the sequence $f'(a) = f(g^a)$, with the function $z(a) = \exp^{2\pi i g^a/p}$, so that \hat{f} may be almost entirely computed using Fourier transforms of length $p - 1$ for which Cooley–Tukey-like ideas may be used. It is an interesting open question to discover if the chirp-z transform has a nonabelian generalization.

Modular FFTs. A significant application of the abelian FFT is in the efficient computation of Fourier transforms for functions on cyclic groups defined over finite fields. These are necessary for the efficient encoding and decoding of various polynomial error correcting codes. Many abelian codes, e.g., the Golay codes used in deep-space communication, are defined as \mathbf{F}_p -valued functions on a group $\mathbb{Z}/m\mathbb{Z}$ with the property that $\hat{f}(k) = 0$ for $k \in S$ some specified set of indices S , where now the Fourier transform is defined in terms of a primitive $(p - 1)^{st}$ root of unity.

These sorts of *spectral constraints* define *cyclic codes*, and they may immediately be generalized to any finite group. Recently, this has been done in the construction of codes over $SL_2(\mathbf{F}_p)$, using connections between expander graphs and linear codes discovered by M. Sipser and D. Spielman. For further discussion of this and other applications see [17].

5. FFTs for Compact Groups

The DFT and FFT also have a natural extension to continuous compact groups. The terminology “discrete Fourier transform” derives from the algorithm having been originally designed to compute the (possibly approximate) Fourier transform of a continuous signal from a discrete collection of sample values.

Under the simplifying assumption of periodicity a continuous function may be interpreted as a function on the unit circle, and compact abelian group, S^1 . Any such function f has a *Fourier expansion* defined as

$$f(e^{2\pi it}) = \sum_{l \in \mathbb{Z}} \hat{f}(l) e^{-2\pi i lt} \quad (5-1)$$

where

$$\hat{f}(l) = \int_0^1 f(e^{2\pi it}) e^{2\pi i lt} dt. \quad (5-2)$$

If $\widehat{f}(l) = 0$ for $|l| \geq N$, then f is *band-limited* with *band-limit* N and the DFT (1–1) is in fact a *quadrature rule* or *sampling theorem* for f . In other words, the DFT of the function

$$\frac{1}{2N-1} f(e^{2\pi it})$$

on the group of $(2N-1)$ -st roots of unity computes exactly the Fourier coefficients of the band-limited function. The FFT then efficiently computes these Fourier coefficients.

The first nonabelian FFT for a compact group was a fast spherical harmonic expansion algorithm discovered by J. Driscoll and D. Healy. Several ingredients were required: (1) A notion of “band-limit” for functions on S^2 ; (2) A sampling theory for such functions; and (3) A fast algorithm for the computation.

The spherical harmonics are naturally indexed according to their order (the common degree of a set of homogeneous polynomials on S^2). With respect to the usual coordinates of latitude and longitude, the spherical harmonics separate as a product of exponentials and associated Legendre functions, each of which separately has a sampling theory. Finally, by using the usual FFT for the exponential part, and a new fast algorithm (based on three-term recurrences) for the Legendre part, an FFT for S^2 is formed.

These ideas generalize nicely. Keep in mind that the representation theory of compact groups is much like that of finite groups: there is a countable complete set of irreducible representations and any square-integrable function (with respect to Haar measure) has an expansion in terms of the corresponding matrix elements. There is a natural definition of band-limited in the compact case, encompassing those functions whose Fourier expansion has only a finite number of terms. The simplest version of the theory is as follows:

DEFINITION 5.1. Let \mathcal{R} denote a complete set of irreducible representations of a compact group G . A *system of band-limits on G* is a decomposition of $\mathcal{R} = \bigcup_{b \geq 0} \mathcal{R}_b$ such that

- (i) \mathcal{R}_b is finite for all $b \geq 0$;
- (ii) $b_1 \leq b_2$ implies that $\mathcal{R}_{b_1} \subseteq \mathcal{R}_{b_2}$;
- (iii) $\mathcal{R}_{b_1} \otimes \mathcal{R}_{b_2} \subseteq \text{span}_{\mathbf{Z}} \mathcal{R}_{b_1+b_2}$.

Let $\{\mathcal{R}_b\}_{b \geq 0}$ be a system of band-limits on G and $f \in L^2(G)$. Then, f is *band-limited with band-limit b* if the Fourier coefficients are zero for all matrix elements in ρ for all $\rho \notin \mathcal{R}_b$.

The case of $G = S^1$ provides the classical example. If $\mathcal{R}_b = \{\chi_j : |j| \leq b\}$ where $\chi_j(z) = z^j$, then $\chi_j \otimes \chi_k = \chi_{j+k}$ and the corresponding notion of band-limited (as per Definition 1) coincides with the usual notion.

For a nonabelian example, consider $G = \text{SO}(3)$. In this case the irreducible representations of G are indexed by the nonnegative integers with V_λ the unique

irreducible of dimension $2\lambda + 1$. Let $\mathcal{R}_b = \{V_\lambda : \lambda \leq b\}$. The Clebsch-Gordan relations

$$V_{\lambda_1} \otimes V_{\lambda_2} = \sum_{j=|\lambda_1-\lambda_2|}^{\lambda_1+\lambda_2} V_j \quad (5-3)$$

imply that this is a system of band-limits for $\text{SO}(3)$. When restricted to the quotient $S^2 \cong \text{SO}(3)/\text{SO}(2)$, band-limits are described in terms of the highest order spherical harmonics that appear in a given expansion.

This notion of band-limit permits the construction of a sampling theory [14]. For example, in the case of the classical groups, a system of band-limits \mathcal{R}_b^n is chosen with respect to a particular norm on the dual of the associated Cartan subalgebra. Such a norm $\|\cdot\|$ (assuming that it is invariant under taking duals, and $\|\alpha\| \leq \|\beta\| + \|\gamma\|$, for α occurring in $\beta \otimes \gamma$) defines a notion of band-limit given by all α with norm less than a fixed b . This generalizes the definition above. The associated sampling sets X_b^n are contained in certain one-parameter subgroups. These sampling sets permit a separation of variables analogous to that used in the Driscoll–Healy FFT. Once again, the special functions satisfy certain three-term recurrences which admit a similar efficient divide-and-conquer computational approach (see [15] and references therein.) one may derive efficient algorithms for all the classical groups, $U(n)$, $\text{SU}(n)$, $\text{Sp}(n)$.

THEOREM 5.2. *Assume $n \geq 2$.*

- (i) *For $U(n)$, $T_{X_b^n}(\mathcal{R}_b^n) \leq O(b^{\dim U(n)+3n-3})$.*
- (ii) *For $\text{SU}(n)$, $T_{X_b^n}(\mathcal{R}_b^n) \leq O(b^{\dim \text{SU}(n)+3n-2})$.*
- (iii) *For $\text{Sp}(n)$, $T_{X_b^n}(\mathcal{R}_b^n) \leq O(b^{\dim \text{Sp}(n)+6n-6})$.*

Here $T_{X_b^n}(\mathcal{R}_b^n)$ denotes the number of operations needed for the particular sample set X_b^n and representations \mathcal{R}_b^n for the associated group.

6. Further and Related Work

Noncompact groups. Much of modern signal processing relies on the understanding and implementation of Fourier analysis for $L^2(\mathbf{R})$, i.e., the noncompact abelian group \mathbf{R} . Nonabelian, noncompact examples have begun to attract much attention.

In this area some of the most exciting work is being done by G. Chirikjian and his collaborators. They have been exploring applications of convolution on the group of rigid motions of Euclidean space to such diverse areas as robotics, polymer modeling and pattern matching. See [5] for details and pointers to the literature.

To date, the techniques used here are approximate in nature and interesting open problems abound. Possibilities include the formulation of natural sampling, band-limiting and time-frequency theories. The exploration of other special cases such as semisimple Lie groups (see [1], for a beautifully written succinct survey

of the Harish-Chandra theory) would be one natural place to start. A sampling and band-limiting theory would be the first step towards developing a computational theory, i.e., FFT. “Fast Fourier transforms on semisimple Lie groups” has a nice ring to it!

Approximate techniques. The techniques in this paper are all exact, in the sense that if computed in exact arithmetic, they yield exactly correct answers. Of course, in any actual implementation, errors are introduced and the utility of an algorithm will depend highly on its numerical stability.

There are also “approximate methods”, approximate in the sense that they guarantee a certain specified approximation to the exact answer that depends on the running time of the algorithm. For computing Fourier transforms at nonequispaced frequencies, as well as spherical harmonic expansions, the fast *multipole method* due to V. Rokhlin and L. Greengard is a recent and very important approximate technique. Multipole-based approaches efficiently compute these quantities approximately, in such a way that the running time increases by a factor of $\log(1/\varepsilon)$, where ε denotes the precision of the approximation. M. Mohlenkamp has applied quasi-classical frequency estimates to the approximate computation of various special function transforms.

Quantum computing. Another related and active area of research involves connections with quantum computing. One of the first great triumphs of the quantum computing model is P. Shor’s fast algorithm for integer factorization on a quantum computer [18]. At the heart of Shor’s algorithm is a subroutine which computes (on a quantum computer) the DFT of a binary vector representing an integer. The implementation of this transform as a sequence of one- and two-bit quantum gates, is the *quantum FFT*, is effectively the Cooley–Tukey FFT realized as a particular factorization of the Fourier matrix into a product of matrices composed as tensor products of certain two by two unitary matrices, each of which is a “local unitary transform”. Extensions of these ideas to the more general group transforms mentioned above are a current important area of research of great interest in computer science.

So, these are some of the things that go into the computation of the finite Fourier transform. It is a tapestry of mathematics both pure and applied, woven from algebra and analysis, complexity theory and scientific computing. It is on the one hand a focused problem, but like any good problem, its “solution” does not end a story, but rather initiates an exploration of unexpected connections and new challenges.

References

- [1] J. Arthur, “Harmonic analysis and group representations”, *Notices Amer. Math. Soc.* **47**:1 (2000), 26–34.

- [2] L. Auslander, J. R. Johnson, R. W. Johnson, “Multidimensional Cooley–Tukey algorithms revisited”, *Adv. Appl. Math.* **17**:4 (1996), 477–519.
- [3] L. Auslander and R. Tolimieri, “Is computing with the finite Fourier transform pure or applied mathematics?”, *Bull. Amer. Math. Soc.* (N.S.) **1**:6 (1979), 847–897.
- [4] T. Beth, *Verfahren der schnellen Fourier–Transformation*, Teubner, Stuttgart, 1984.
- [5] G. S. Chirikjian and A. B. Kyatkin, *Engineering applications of noncommutative harmonic analysis*, CRC Press, Boca Raton (FL), 2000.
- [6] J. W. Cooley, The re-discovery of the fast Fourier transform algorithm, *Mikrochimica Acta* **III** (1987), 33–45.
- [7] J. W. Cooley and J. W. Tukey, “An algorithm for machine calculation of complex Fourier series”, *Math. Comp.* **19** (1965), 297–301.
- [8] P. Diaconis, *Group representations in probability and statistics*, IMS, Hayward (CA), 1988.
- [9] I. Gel’fand and M. Tsetlin, “Finite dimensional representations of the group of unimodular matrices”, *Dokl. Akad. Nauk SSSR* **71** (1950), 825–828 (in Russian).
- [10] M. T. Heideman, D. H. Johnson and C. S. Burrus, “Gauss and the history of the fast Fourier transform”, *Archive for History of Exact Sciences* **34**:3 (1985), 265–277.
- [11] G. James and A. Kerber, *The representation theory of the symmetric group*, Encyclopedia of Mathematics and its Applications **16**, Addison-Wesley, Reading (MA), 1981.
- [12] T. Y. Lam, “Representations of finite groups: a hundred years”, parts I and II, *Notices Amer. Math. Soc.* **45**:3 (1998), 361–372 and **45**:4 (1998), 465–474.
- [13] D. K. Maslen, “The efficient computation of Fourier transforms on the symmetric group”, *Math. Comp.* **67**(223) (1998), 1121–1147.
- [14] D. K. Maslen, “Efficient computation of Fourier transforms on compact groups”, *J. Fourier Anal. Appl.* **4**:1 (1998), 19–52.
- [15] D. K. Maslen and D. N. Rockmore, “Generalized FFTs: a survey of some recent results”, pp. 183–237 in *Groups and computation*, II (New Brunswick, NJ, 1995), DIMACS Ser. Discrete Math. Theoret. Comput. Sci. **28**, Amer. Math. Soc., Providence (RI), 1997.
- [16] C. Rader, “Discrete Fourier transforms when the number of data samples is prime”, *IEEE Proc.* **56** (1968), 1107–1108.
- [17] D. N. Rockmore, “Some applications of generalized FFTs” (an appendix with D. Healy), pp. 329–369 in *Proceedings of the DIMACS Workshop on Groups and Computation* (June 7–10, 1995), edited by L. Finkelstein and W. Kantor, 1997.
- [18] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer”, *SIAM J. Computing* **26** (1997), 1484–1509.
- [19] “SPIRAL: automatic generation of platform-adapted code for DSP algorithms”, <http://www.ece.cmu.edu/~spiral/>.
- [20] A. Willsky, “On the algebraic structure of certain partially observable finite-state Markov processes”, *Inform. Contr.* **38** (1978), 179–212.

DAVID K. MASLEN
SUSQUEHANNA INTERNATIONAL GROUP LLP
401 CITY AVENUE, SUITE 220
BALA CYNWYD, PA 19004
david@maslen.net

DANIEL N. ROCKMORE
DEPARTMENT OF MATHEMATICS
DARTMOUTH COLLEGE
HANOVER, NH 03755
rockmore@cs.dartmouth.edu