GENERIC QUANTUM FOURIER TRANSFORMS

Cristopher Moore
Department of Computer Science
University of New Mexico
moore@cs.unm.edu

Daniel Rockmore
Department of Mathematics
Dartmouth College
rockmore@cs.dartmouth.edu

Alexander Russell
Department of Computer Science and Engineering
University of Connecticut
acr@cse.uconn.edu

April 9, 2003

Abstract

The *quantum Fourier transform* (QFT) is the principal algorithmic tool underlying most efficient quantum algorithms. We present a generic framework for the construction of efficient quantum circuits for the QFT by "quantizing" the *separation of variables* technique that has been so successful in the study of classical Fourier transform computations. Specifically, this framework applies the existence of computable Bratteli diagrams, adapted factorizations, and Gel'fand-Tsetlin bases to offer efficient quantum circuits for the QFT over a wide variety a finite Abelian and non-Abelian groups, including all group families for which efficient QFTs are currently known and many new group families. Moreover, the method gives rise to the first subexponential-size quantum circuits for the QFT over the linear groups $\mathrm{GL}_k(q)$, $\mathrm{SL}_k(q)$, and the finite groups of Lie type, for any fixed prime power q.

1 Introduction

Peter Shor's spectacular application of the Fourier transform over the cyclic group $\mathbb{Z}_n$ in the seminal discovery of an efficient quantum factoring algorithm [25] has motivated broad interest in the problem of efficient quantum computation over arbitrary groups (see, e.g., [3, 9, 11, 13, 14, 20, 21, 27]). While this research effort has become quite ramified, two related themes have emerged: *(i.)* development of efficient *quantum Fourier transforms* and *(ii.)* development of efficient quantum algorithms for the *hidden subgroup problem*. The complexity of these two problems appears to relate intimately to the group in question: while quantum Fourier transforms and hidden subgroup problems over Abelian groups are well-understood, our understanding of these basic problems over non-Abelian groups remains embarrassingly sporadic. Aside from their natural appeal, this line of research been motivated by the direct relationship to the graph isomorphism problem: an efficient solution to the hidden subgroup problem over the (non-Abelian) symmetric groups would yield an efficient quantum algorithm for graph isomorphism.

Over the cyclic group $\mathbb{Z}_n$ the *quantum Fourier transform* refers to the transformation taking the state

$$\sum_{z\in\mathbb{Z}_n} f(z)\,|z\rangle \qquad \text{to the state} \qquad \sum_{\omega\in\mathbb{Z}_n} \hat{f}(\omega)\,|\omega\rangle\,,$$

where $f:\mathbb{Z}_n\to\mathbb{C}$ is a function with $\|f\|_2 = 1$ and $\hat{f}(\omega) = \sum_z f(z)e^{2\pi i\omega z/n}$ denotes the familiar discrete Fourier transform at ω. Over an arbitrary finite group G, this analogously refers to the transformation taking the state

$$\sum_{z\in G} f(z)\,|z\rangle \qquad \text{to the state} \qquad \sum_{\omega\in\hat{G}} \hat{f}(\omega)_{ij}\,|\omega,i,j\rangle\,,$$

where $f : G \to \mathbb{C}$, as before, is a function with $\|f\|_2 = 1$ and $\hat{f}(\omega)_{ij}$ denotes the i,jth entry of the Fourier transform at the representation ω. This is explained further in Section 2.

While there is no known explicit relationship between the quantum Fourier transform and the hidden subgroup problem over a group G, all known efficient hidden subgroup algorithms rely on an efficient quantum Fourier transform. Indeed, it is fair to say that the quantum Fourier transform is the only known non-trivial quantum algorithmic paradigm for such problems.

In this article we focus on the construction of efficient quantum Fourier transforms. Our research is motivated by dramatic progress over the last decade in the theory of efficient *classical* Fourier transforms (see, e.g., [4, 5, 8, 18, 22]). These developments have provided a collection of techniques which, taken together, yield a uniform framework for the efficient (classical) computation of Fourier transforms over a wide variety of important families of groups including, for example, the finite groups of Lie type (properly parametrized) and the symmetric groups.

We present here an adaptation to the quantum setting of a wide class of efficient classical Fourier transform algorithms; namely, those achieved by the "separation of variables" approach. This establishes the first generic quantitative relationship between efficient classical Fourier transforms and efficient circuits for the quantum Fourier transform.

Specifically, we define a broad class of *polynomially uniform* groups and show

Theorem 1 *If G is a polynomially uniform group with a subgroup tower $G = G_m > \cdots > \{1\}$ with adapted diameter D, maximum multiplicity M, and maximum index $I = \max_i [G_i : G_{i-1}]$, then there is a quantum circuit of size $\mathrm{poly}(I \times D \times M \times \log|G|)$ which computes the quantum Fourier transform over G.*

This quantifies the complexity of the quantum Fourier transform in exactly the same fashion as does Corollary 3.1 of [17] in the classical case. We extend this class further by showing that it is closed under a certain type of Abelian extension which may have exponential index.

Together, these results give efficient QFTs — namely, circuits of $\mathrm{polylog}(|G|)$ size — for many families of groups. These include *(i.)* the Clifford groups $\mathbb{CL}_n$; *(ii.)* the symmetric groups, recovering the algorithm of Beals [3]; *(iii.)* wreath products $G \wr S_n$ where $|G| = \mathrm{poly}(n)$; *(iv.)* metabelian groups, including metacyclic groups such as the dihedral and affine groups, recovering the algorithm of Høyer [13]; *(v.)* bounded extensions of Abelian groups such as the generalized quaternions, recovering the algorithm of Püschel et al. [21].

Our methods also give the first subexponential size quantum circuits for the linear groups $\mathrm{GL}_k(q)$, $\mathrm{SL}_k(q)$, $\mathrm{PGL}_k(q)$, and $\mathrm{PSL}_k(q)$ for fixed prime power q, various families of finite groups of Lie type, and the Chevalley and Weyl groups.

The paper is structured as follows. Sections 2 and 3 briefly summarize the representation theory of finite groups, the Bratteli diagram, and adapted bases. We give our algorithms in Section 4 along with a list of group families for which the provide efficient circuits for the QFT. We conclude with open problems in Section 5.

2 Representation theory background

Fourier analysis over a group G involves expressing arbitrary functions $f : G \to \mathbb{C}$ as linear combinations of specific functions on G which reflect the group's structure and symmetries. If G is Abelian, these are precisely the *characters* of G (the homomorphisms of G into $\mathbb{C}$). For a general group, they are the *irreducible matrix elements*, and the Fourier transform is the change of basis from the basis of delta functions to the basis of irreducible matrix elements.

In order to be precise we need the language of (finite) group representation theory (see, e.g., Serre [24] for an excellent introduction). A *representation* ρ of a finite group G is a homomorphism $\rho : G \to \mathrm{U}(V)$, where V is a (finite) d_ρ-dimensional vector space over $\mathbb{C}$ with an inner product and $\mathrm{U}(V)$ denotes the group of unitary linear operators on V. Fixing an orthonormal basis for V, each $\rho(g)$ may be realized as a $d_\rho \times d_\rho$ unitary matrix. When a basis has been selected in this way for V, we refer to ρ as a *matrix representation* of G; then each of the d_ρ^2 functions $\rho_{ij}(g) = [\rho(g)]_{ij}$ is called a *matrix element* (corresponding to ρ). As ρ is a homomorphism, for any $g, h \in G$, $\rho(gh) = \rho(g)\rho(h)$, implying that in general, $\rho_{ij}(gh) = \sum_{k=1}^{d_\rho} \rho_{ik}(g)\rho_{kj}(h)$.

A matrix representation ρ of G on V is *irreducible* if no subspace (other than the trivial $\{0\}$ subspace and V) is mapped into itself. This is equivalent to the statement that there is no change of basis that finds a simultaneously block diagonalization (of given shape) of all $\rho(g)$. Otherwise the representation is said to be *reducible*. The irreducible representations will play a role in the theory analogous to that of the characters of an Abelian group. Two

representations ρ and σ are *equivalent* if they differ only by a change of basis, so that for some fixed unitary matrix U, $\sigma(g) = U^{-1}\sigma(g)U$, for all $g \in G$. Up to equivalence, a finite group G has a finite number of irreducible representations equal to the number of its conjugacy classes. For a group G, we let $\hat{G}$ denote a collection of representations of G containing exactly one from each isomorphism class of irreducible representations.

Selecting bases B for the representations of $\hat{G}$ results in a set of (inequivalent irreducible) matrix representations; when we wish to be explicit about this selection of bases, we denote such a collection $\hat{G}_B$. The matrix elements of the matrix representations $\rho \in \hat{G}_B$ in fact form an orthonormal basis for the $|G|$-dimensional vector space of complex-valued functions on G. This implies the important relationship between the dimensions of the irreducible representations of G and $|G|$: $\sum_{\rho \in \hat{G}} d_\rho^2 = |G|$. Such a family gives rise to a general definition of Fourier transform.

Definition 1 *Let $f : G \to \mathbb{C}$; let $\rho : G \to U(V)$ be a matrix representation of G. The* Fourier transform of f at ρ, *denoted $\hat{f}(\rho)$, is the matrix*

$$\hat{f}(\rho) = \sqrt{\frac{d_\rho}{|G|}} \sum_{g \in G} f(g)\rho(g).$$

We typically restrict our attention to $\hat{f}(\rho)$, where ρ is irreducible.

We refer to the collection of matrices $\langle \hat{f}(\rho) \rangle_{\rho \in \hat{G}_B}$ as the *Fourier transform* of f. Thus f is mapped into $|\hat{G}|$ matrices of varying dimensions. The total number of entries in these matrices is $\sum d_\rho^2 = |G|$, by the equation mentioned above. The Fourier transform is linear in f; with the constants used above ($\sqrt{d_\rho / |G|}$) it is in fact unitary, taking the $|G|$ complex numbers $\langle f(g) \rangle_{g \in G}$ to $|G|$ complex numbers organized into matrices.

For two complex-valued functions f_1 and f_2 on a group G, there is a natural inner product $\langle f_1, f_2 \rangle$ given by $\frac{1}{|G|} \sum_g f_1(g) f_2(g)^*$. For any pair of matrix representations $\rho, \sigma \in \hat{G}_B$, the corresponding irreducible matrix elements are orthogonal according to the inner product: let ρ and σ be two elements of $\hat{G}$; then

$$\left\langle [\rho(\cdot)]_{ij}, [\sigma(\cdot)]_{kl} \right\rangle = \begin{cases} 0 & \text{if } \rho \ncong \sigma \\ \frac{1}{d_\rho} \delta_{ik} \delta_{jl} & \text{if } \rho = \sigma. \end{cases} \tag{1}$$

Computation of the Fourier transform (with respect to a given choice of $\hat{G}$) is equivalent to the change of basis from that of the point masses to the irreducible matrix elements determined by $\hat{G}$. This linear map (of the vector space of functions on G) is invertible, with (point-wise) inverse given by the *Fourier inversion formula*:

$$f(s) = \sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G|}} \operatorname{tr}\left(\rho(s) \hat{f}(\rho)^{-1} \right).$$

A reducible matrix representation $\rho : G \to U(V)$ may always be decomposed into irreducible representations; specifically, there is a basis of V in which each $\rho(g)$ is block diagonal where the ith block of $\rho(g)$ is precisely $\sigma_i(g)$ for some irreducible matrix representation σ_i. In this case we write $\rho = \bigoplus \sigma_i$. The number of times a given $\sigma \in \hat{G}$ appears in this decomposition is the *multiplicity* of σ in ρ. If the irreducible representation σ_i appears with multiplicity w_i in decomposition of ρ, we may write $\rho = \oplus^{w_1} \sigma_1 \ldots \oplus^{w_r} \sigma_r$.

A representation ρ of a group G is also automatically a representation of any subgroup H. We refer to this *restricted* representation on H as $\rho|_H$. Note that in general, representations that are irreducible over G may be reducible when restricted to H.

Note: The familiar *Discrete Fourier Transform* (DFT) corresponds to the case in which the group is cyclic. In this case the representations are all one dimensional, and if $G = \mathbb{Z}_n$, the linear transformation (i.e., the Fourier transform,) is an order n Vandermonde matrix using the n-th roots of unity.

3 Bratteli diagrams, Gel'fand-Tsetlin bases, and adapted diameters

The main ingredients for our algorithm are *(i.)* a tower of subgroups (or *chain*) which provides a means by which the Fourier transform on G can be built iteratively as an accumulation of Fourier transforms on increasingly larger subgroups and *(ii.)* a natural indexing scheme for the representations given by paths in the *Bratteli diagram* corresponding to the group tower and finally *(iii.)* a factorization of group elements in terms of a basic set of generators, which, when judiciously chosen, provide a factorization of the Fourier transform as a product of structured (direct sums of tensor products) and sparse matrices. The complexity of a corresponding efficient Fourier transform which uses these basic ingredients can then be derived in terms of basic representation-theoretic and combinatorial data.

3.1 Bratteli diagrams and Gel'fand-Tsetlin bases

Much of Abelian Fourier analysis is simplified by the fact that in this case the dual (that is, the set of characters $\chi : G \to \mathbb{C}$) also forms a group isomorphic to the original group; furthermore, in this isomorphism lies a natural correspondence providing an indexing of the irreducible representations (i.e., matrix elements). However, in the general case there is no immediate indexing scheme for the dual $\hat{G}$ and the landscape is further complicated by the absence of a canonical basis for the now multidimensional representations. Indeed, for the goal of efficient Fourier analysis, not all bases are created alike! In particular, a fairly general methodology for the construction of group FFTs, the "separation of variables" approach [17, 18] relies on the use of *Gel'fand-Tsetlin* or *adapted* bases for efficient computation. These bases allow for a Fourier transform on G to be built from Fourier transforms on subgroups, a general technique whose efficiency improves as it is used through a tower of subgroups. This is in fact the main idea in the famous "Cooley-Tukey" (decimation-in-time) FFT.

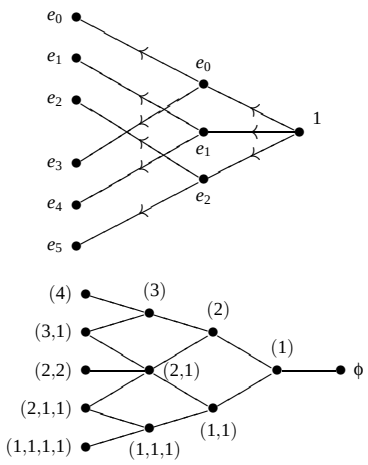

Figure 1: These are the Bratteli diagrams for the subgroup towers $\mathbb{Z}_6 > \mathbb{Z}_3 > 1$ (top) and $S_4 > S_3 > S_2 > 1$ (bottom). Cyclic groups of order n have representations indexed by the integers mod n, and (assuming $m|n$) then the representation corresponding to j restricts to the representation corresponding to $j \bmod m$. The lower diagram uses the well-known correspondence between irreducible representations of S_n and partitions of n. In this case restrictions from S_n to S_{n-1} are determined by those partitions obtained via the decrement of a part of the original partition.

A crucial ingredient of the general separation of variables approach is the incorporation of an indexing scheme that permits the computational to be organized efficiently. The same *Bratteli diagram* formalism is key to both the organization and manipulation of the calculation for a quantum FFT; we present it below.

Given a finite group G and let

$$G = G_m > G_{m-1} > \cdots > G_1 > G_0 = \{1\}$$

be a tower of subgroups of length m for G. The corresponding *Bratteli diagram*, denoted $\mathfrak{B}$, is a leveled directed multigraph whose nodes of level $i = 0, \ldots, m$ are in one-to-one correspondence with the (inequivalent) irreducible representations of G_i. For convenience, we refer to vertices in the diagram by the representation with which they are associated. The number of edges from an irreducible representation η of G_i to ρ of G_{i+1} is equal to the *multiplicity* of η in the restriction of ρ to G_i. Since there is a unique irreducible representation of the trivial group, a Bratteli diagram for a given tower is in fact a rooted tree.

Thus, the edges out of a node η of $\hat{G}_i$ represent a complete set of orthogonal embeddings of the corresponding representation space into the representations of G_{i+1} and conversely, the edges entering a given representation $\rho : G_{i+1} \to U(V_\rho)$ of G_{i+1} index a set of mutually orthogonal subspaces of V_ρ whose direct sum represents the decomposition of V_ρ under the (restricted) action of G_i. Thus, the paths from the root node to a vertex $\rho : G_i \to U(V_\rho)$ index a basis of V_ρ with the following property: for any $G_j < G_i$, there is a partition of the basis vectors into subsets, each of which spans an irreducible G_j-invariant subspace, so that the associated matrix representation is block diagonal

according to this partition when restricted to G_j and, moreover, that blocks for equivalent irreducible representations are actually equal. Such bases are said to be (subgroup-)adapted or *Gel'fand-Tsetlin*. Consequently, the number of paths to a node η is equal to d_η, and pairs of path with common endpoint η index an irreducible matrix element of η.

The block diagonal nature of the restriction (combined with the fact that blocks corresponding to equivalent representations are actually equal) allows the Fourier transform on $G = G_m$ to be expressed as a sum of Fourier transforms on G_{m-1}, each translated from a distinct coset: specifically, if $T \subset G$ is a transversal, i.e. a set of representatives for the left cosets of G_{m-1} in G_m, we define $f_\alpha : G_{m-1} \to \mathbb{C}$ by $f_\alpha(x) = f(\alpha x)$. Then

$$\hat{f}(\rho) = \sum_{\alpha \in T} \rho(\alpha) \sum_{x \in G_{m-1}} \rho(x) f(\alpha x) = \sum_{\alpha \in T} \rho(\alpha) \cdot \hat{f}_\alpha(\rho|_{G_{m-1}}). \tag{2}$$

3.2 Strong generating sets and adapted diameters

Adapted representations are only part of the story for the construction of efficient Fourier transform algorithms. In general, $\rho(\alpha)$ of Equation (2), the "twiddle factor", could be an arbitrary matrix of exponential size, so implementing it in (2) could be costly. Luckily, under fairly mild assumptions, the matrices $\rho(\alpha)$ can be factored into polylog($|G|$) sparse, highly structured matrices, and can therefore be implemented with polylog($|G|$) elementary quantum operators.

We say that S is a *strong generating set* for the tower of subgroups $\{G_i\}$ if $S \cap G_i$ generates G_i. Say that we have chosen a transversal T_i for each i indexing the cosets of G_{i-1} in G_i. Now define $D_i = \min\{\ell > 0 : \cup_{j \le \ell} (S \cap G_i)^j \supseteq T_i\}$, and define the *adapted diameter* $D = \sum_i D_i$. Then clearly any group element can be factored as a series of coset representatives, which in turn can be factored as a total of at most D elements of S.

Of course, to perform the QFT efficiently we would like $\rho(\gamma)$ to have a simple form for each $\gamma \in S$. Given a subgroup $K < G$, recall that the *centralizer* of K is the subgroup $Z(K) = \{g \in G : gk = kg \text{ for all } k \in K\}$. The following is implicit in the oft-cited lemma of Schur:

Lemma 1 *(Schur, [17, Lemma 5.1]) Let $K < G$, let $\gamma \in Z(K)$, and let ρ be a K-adapted representation of G. Suppose that $\rho|_K = \oplus^{m_1}\eta_1 \cdots \oplus^{m_r}\eta_r$. Then $\rho(\gamma)$ has the form*

$$(GL_{m_1}(\mathbb{C}) \otimes I_{d_1}) \oplus \cdots \oplus (GL_{m_r}(\mathbb{C}) \otimes I_{d_r}) \tag{3}$$

where I_k is the $k \times k$ identity matrix and $d_i = d_{\eta_i}$.

Since any unitary operator in $\mathrm{GL}_m(\mathbb{C})$ can be carried out with poly(m) elementary quantum gates [2], and since we can condition on the η_i to find out which subspace of ρ we are in, we can write $\rho(\gamma)$ as a series of poly(M) elementary quantum operations where $M = \max_i m_i$ in (3). Therefore, the total number of elementary quantum operators we need to implement $\rho(\alpha)$ is then $D \times \text{poly}(M)$.

Moreover, if γ is itself in a subgroup $H > K$, and ρ is adapted to both H and K, then $\rho(a)$ also possesses the block structure corresponding to $\rho|_H$. This places an upper bound on M of the maximum multiplicity with which representations of K appear in restrictions of representations of H. Thus we can minimize M by choosing generators γ inside subgroups as low on the tower as possible, which centralize subgroups as high on the tower as possible.

For instance, in the symmetric group S_n we take the tower to be $S_n > S_{n-1} > \cdots > \{1\}$, where S_i fixes all elements greater than i. Let S be the set of pairwise adjacent transpositions $(j, j+1)$; each of these is contained in S_{j+1} and centralizes S_{j-1}. The maximum multiplicity with which a representation of S_{j-1} appears in a representation of S_{j+1} is 2, corresponding to the two orders in which we can remove two cells from a Young diagram. Since the adapted diameter is easily seen to be $O(n^2)$, this means that the $\rho(\alpha)$ can be carried out in $O(n^2) = \text{polylog}(|S_n|)$ elementary quantum operations [3]. We will see that a similar situation obtains for a large class of groups.

4 Efficient quantum Fourier transforms

We describe our algorithm in this section. The algorithm performs the Fourier transform inductively on the tower of subgroups, using the structure of the Bratteli diagram to construct the transform at each level from the transform at the previous level.

Recall that for each level of our tower of subgroups $G = G_m > G_{m-1} > \cdots > G_0 = \{1\}$ we have chosen a transversal T_i for the left cosets of G_{i-1} in G_i. At the beginning of the computation, we represent each group element g as a product $\alpha = \alpha_m \cdots \alpha_1$ where $\alpha_i \in T_i$. This string becomes shorter as we work our way up the tower, and after having performed the Fourier transform for G_i the remaining string $\alpha = \alpha_m \cdots \alpha_{i+1}$ indexes the coset of G_i in G in which g lies.

At the end of the computation, we have a pair of paths in the Bratteli diagram, $s = s_1 \cdots s_m$ and $t = t_1 \cdots t_m$, which index the rows and columns of the representations ρ of G. These paths begin empty and grow as we work our way up the tower; after having performed the Fourier transform for G_i, the paths $p = p_1 \cdots p_i$ and $q = q_1 \cdots q_i$ of length i index the rows and columns of representations σ of G_i.

With a compact encoding, one could store α in the same registers as s and t, at each step replacing a coset representative α_i with a pair of edges s_i, t_i. However, our algorithm is simpler to describe if we double the number of qubits and store α and s,t in separate registers. Padding out α, s, and t to length m with zeroes, our computational basis consists of unit vectors of the form

$$|\alpha\rangle\,|s,t\rangle = \left|\alpha_m \cdots \alpha_{i+1}\, 0^i\right\rangle \otimes \left|s_1 \cdots s_i\, 0^{m-i}, s_1 \cdots s_i\, 0^{m-i}\right\rangle \quad .$$

Keep in mind the basis $\{|s,t\rangle\}$, where s and t have length i and end in the same representation, is just a permutation of our adapted Gel'fand-Tsetlin basis $\{|\sigma, j, k\rangle\}$ for $\hat{G}_i$, where σ ranges over the representations of G_i and $1 \le j,k \le d_\sigma$ index its rows and columns. Therefore, we will sometimes abuse notation by writing $\hat{f}(s,t)$ and $\hat{f}(\sigma)_{j,k}$ for the Fourier transform over G_i indexed in these two different ways.

Each stage of the algorithm consists of calculating the Fourier transform over G_{i+1} from that over G_i. By induction it suffices to consider the last stage, where we go from $H = G_{m-1}$ to $G = G_m$. Specifically, choose a transversal T of H in G such that every $g \in G$ can be written αh where $\alpha \in T$ and $h \in H$. For each $\alpha \in T$, define a function f_α on H as $f_\alpha(h) = f(\alpha h)$; this is the restriction of f to the coset αH, shifted into H.

After having performed the Fourier transform on H, our state will be

$$\sum_{\alpha \in T} |\alpha\rangle \otimes \sum_{s,t \text{ of length } m-1} \hat{f}_\alpha(s,t)\,|s,t\rangle \;=\; \sum_{\alpha \in T} |\alpha\rangle \otimes \sum_{(\sigma,j,k) \in \hat{H}} \hat{f}_\alpha(\sigma)_{j,k}\,|\sigma,j,k\rangle \quad . \tag{4}$$

Our goal is to transform this state into the Fourier basis of G, namely

$$|0\rangle \otimes \sum_{s,t \text{ of length } m} \hat{f}(s,t)\,|s,t\rangle \;=\; |0\rangle \otimes \sum_{(\rho,j,k) \in \hat{G}} \hat{f}(\rho)_{j,k}\,|\rho,j,k\rangle \quad . \tag{5}$$

where $|0\rangle$ occupies the register that held the coset representative α before.

This transformation is greatly simplified by the following two observations, which are common to nearly every algorithm for the FFT. First, as described in Equation (2) above, $\hat{f}$ can be written as a sum over contributions from f's values on each coset αH, giving

$$\hat{f}(\rho) = \sum_{\alpha \in T} \rho(\alpha) \cdot \hat{f}_\alpha(\rho) \quad . \tag{6}$$

Since f_α has support only in H, the matrix $\hat{f}_\alpha(\rho)$ is a direct sum of sub-matrices of the form $\hat{f}_\alpha(\sigma)$, summed over the σ appearing in ρ. In the quantum setting we accomplish this via an *embedding* operation which reverses the restriction to H,

$$|\sigma\rangle \to \sum_{\rho:\, \sigma \text{ appears in } \rho|_H} A_{\sigma,\rho}\,|\rho\rangle \tag{7}$$

where this "scale factor" is

$$A_{\sigma,\rho} = \sqrt{\frac{|H|}{|G|}\frac{d_\rho}{d_\sigma}} \quad .$$

(Note that $\sum_\rho |A_{\sigma,\rho}|^2 = 1$.)

Thus the algorithm consists of *(i.)* embedding the σ in the appropriate ρ, *(ii.)* applying the "twiddle factor" $\rho(\alpha)$, and *(iii.)* summing over the cosets. However, in general, doing these things efficiently is no simple matter. First, a given σ might appear in a given ρ with an arbitrary change of basis; the twiddle $\rho(\alpha)$ could be an arbitrary unitary

matrix of exponential size; and summing over an exponential number of cosets will take exponential time unless parallelized in some way.

It is here that the Bratteli diagram proves to be extremely helpful. It allows us to implement the twiddle factors $\rho(\alpha)$ efficiently when coupled with a strong generating set as discussed in Section 3.2 by providing an adapted basis. It simplifies the embedding operation as well: first note that $\hat{f}_\alpha(s,t)$ is nonzero only when s and t end in the same representation σ of G_t, i.e. in the same vertex of the diagram. Moreover, recall that the Bratteli diagram indexes an adapted basis in which $\rho|_H$ is block-diagonal with the σ_j as its blocks. This means that the σ appear in the ρ in an extremely simple way: namely, where s and t are extended by appending the same edge e to both.

Let adopt some notation. Given a path s in the Bratteli diagram of length $m-1$ or m, denote the representation in which it ends by $\sigma[s]$ or $\rho[s]$ respectively, and if $s = s_1 \cdots s_{m-1}$, denote $s_1 \cdots s_{m-1} e$ as se. We will index the edges of each vertex $\{1,\ldots,k\}$ where it has out-degree k. It will be convenient to carry out this embedding only if the register containing the coset representative is zero, and leave other basis vectors in $(T \cup \{0\}) \otimes \hat{H}$ fixed. Then (7) becomes

$$U : \begin{cases} |0\rangle |s,t\rangle \to |0\rangle \sum_e A_{\sigma[s],\rho[se]} |se,te\rangle \\ |\alpha\rangle |s,t\rangle \to |\alpha\rangle |s,t\rangle \text{ for all } \alpha \in T \end{cases} \tag{8}$$

where the sum is over all outgoing edges e of $\sigma[s] = \sigma[s]$.

Note that we have not defined U on the entire space; in particular, since we are moving probability from $\hat{H}$ to $\hat{G}$, basis vectors $|0\rangle |se,te\rangle \in (T \cup \{0\}) \otimes \hat{G}$ cannot stay fixed. As we will see below, it does not matter precisely how U behaves on the rest of the state space, as long as its behavior on $\hat{H}$ is as described in (8). This can be accomplished simply by putting the m'th registers of s and t in the superposition $\sum_e A_{\sigma[s],\rho[se]} |e\rangle \otimes |e\rangle$, and for a large class of extensions we can prepare this superposition efficiently.

4.1 Extensions of subexponential index

In this section we generalize Beals' QFT for the symmetric group [3] to a large class of groups. First we show that the Fourier transform can be extended from H to G, modulo some reasonable uniformity conditions on G.

Definition 2 *For a group G and a tower of subgroups G_i, let $\mathfrak{B}$ be the corresponding Bratteli diagram, let T_i be a set of coset representatives at each level, and let S be a strong set of generators for G. Then we say that G is* polynomially uniform *(with respect to $\{G_i\}$, $\mathfrak{B}$, $\{T_i\}$, and S) if the following functions are computable by a classical algorithm in* $\text{polylog}(|G|)$ *time:*

1. *Given two paths s,t in $\mathfrak{B}$, whether $\rho[s] = \rho[t]$;*
2. *Given a path s in $\mathfrak{B}$, the dimension and the out-degree of $\rho[s]$;*
3. *Given a coset representative $\alpha_i \in T_i$, a factorization of α as a word of* $\text{polylog}(|G|)$ *length in $(S \cap G_i)^*$.*

Lemma 2 *If G is polynomially uniform with respect to a tower of subgroups where $G = G_m$ and $H = G_{m-1}$ and a strong generating set S with adapted diameter D and maximum multiplicity M, then the Fourier transform of G can obtained from the state* (4) *using* $\text{poly}([G:H] \times D \times M \times \log|G|)$ *elementary quantum operations.*

Proof. First, to carry out the embedding transformation U, we use the classical algorithm to compute the list of edges e and $d_{\rho[se]}$ conditional on s, and thus compute the $A_{\sigma,\rho}$ (say, to n digits in $\text{poly}(n)$ time). Note that σ appears in at most $[G:H]$ many ρ. We then carry out a series of $[G:H]$ conditional rotations, each of which rotates the appropriate amplitude from $|0\rangle |s,t\rangle$ to $|0\rangle |se,te\rangle$. Thus U, and therefore U^{-1}, can be carried out in $O([G:H])$ quantum operations.

To apply the twiddle factor and sum over the cosets as in (6), we use a technique of Beals [3] and carry out the following for-loop. For each $\alpha \in T$, we do the following three things: left multiply $\hat{f}(\rho)$ by $\rho(\alpha)^{-1}$; add $\hat{f}_\alpha(\rho)$ to $\hat{f}(\rho)$; and left multiply $\hat{f}(\rho)$ by $\rho(\alpha)$. This loop clearly produces $\sum_{\alpha \in T} \rho(\alpha) \cdot \hat{f}(\rho)$, so we just need to show that each of these three steps can be carried out efficiently.

Recall that $\hat{f}(\rho)$ is given in the $|s,t\rangle$ basis, where s and t index the row and column of ρ respectively. To left multiply $\hat{f}(\rho)$ by $\rho(\alpha)$, we apply $\rho(\alpha)$ to the s register and leave the t register unchanged. Since G is polynomially uniform, a classical algorithm can factor α as the product of D generators $\gamma_i \in S$, and provide a factorization of each

$\rho(\gamma_i)$ as the product of poly(M) many elementary quantum operations, in polylog$(|G|)$ time. This implements $\rho(\alpha)$ and $\rho(\alpha)^{-1}$ in $D \times \text{poly}(M) + \text{polylog}(|G|)$ operations.

The step "add $\hat{f}_\alpha(\rho)$ to $\hat{f}(\rho)$" is slightly more mysterious, and indeed it does not even sound unitary at first. However, as Beals points out, at each point in the loop we are adding $\hat{f}_\alpha(\rho)$, which is the Fourier transform of a function with support only on H, to $\sum_{\beta<\alpha} \rho(\alpha^{-1}\beta)\hat{f}_\beta(\rho)$, which is the Fourier transform of a function with support only *outside* H. Thus these two states are orthogonal, and adding two orthogonal vectors can be done unitarily by rotating one vector into the other while fixing the subspace perpendicular to both. Let V_α be the operation that exchanges $|\alpha\rangle|s,t\rangle$ with $|0\rangle|s,t\rangle$ and leaves $|\beta\rangle|s,t\rangle$ fixed for all $\beta \leq \alpha, 0$; then Beals showed that this step can be written $U^{-1}V_\alpha U$ where U is the embedding operator defined in (8). We showed earlier that U can be carried out in $O([G:H])$ quantum operations, and V is a simply a Boolean operation on the α register. Finally, the for-loop runs $|T| = [G:H]$ times, so we're done. □

Proof of Theorem 1. This follows immediately from the fact that the depth of the Bratteli diagram is at most $\log|G|$. □

As noted above, for many groups, the maximum index $I = \max_i[G_i : G_{i-1}]$, the adapted diameter D, and the maximum multiplicity M are all polylog$(|G|)$. In this case, Theorem 1 gives circuits for the QFT of polylog$(|G|)$ size. This includes the following three families of groups:

The symmetric groups S_n. As stated above, we take the tower $S_n > S_{n-1} > \cdots > \{1\}$ where S_i fixes all elements greater than i. The maximum index is then $n = o(\log|S_n|)$. The generators are the adjacent transpositions; the adapted diameter is $O(n^2)$ and the maximum multiplicity is 2. The adapted basis is precisely the Young orthogonal basis.

Wreath products $G = H \wr S_n$ for H of size poly(n)**.** These groups arise naturally as automorphism groups of graphs obtained by composition [12]. As in [23] the tower is

$$H \wr S_n > H \times (H \wr S_{n-1}) > H \wr S_{n-1} > \cdots > \{1\} \ .$$

The maximum index is $\max(n, |H|)$, the generators are the adjacent transpositions and an arbitrary set of $\log|H|$ generators for each factor of H, the adapted diameter is $O(n^2 \log|H|)$, and the maximum multiplicity is $O(|H|)$. Then note that $|H| = \text{polylog}(|G|)$. See [17] for details and [15] for discussion on wreath products.

The Clifford groups. The Clifford groups $\mathbb{CL}_n$ are generated by $x_1, \ldots, x_n$ where $x_i^2 = 1$ and $x_i x_j = -x_i x_j$ for all $i \neq j$ [26]. We take the tower $\mathbb{CL}_n > \mathbb{CL}_{n-1} > \cdots > \{1\}$ which has maximum index 2, and the generators $\big\{\{x_1\}, \{x_1x_2\}, \ldots, \{x_{n-1}x_n\}\big\}$. The adapted diameter is $O(n)$, and since each $x_i x_{i+1}$ centralizes $\mathbb{CL}_{i-1}$, the maximum multiplicity is 4.

In addition to giving polylog$(|G|)$-size circuits for these groups, this technique also gives the first subexponential-size circuits for the following classical groups:

The linear groups $\mathrm{GL}_n(q)$, $\mathrm{SL}_n(q)$, $\mathrm{PGL}_n(q)$, and $\mathrm{PSL}_n(q)$; the finite groups of Lie type; the Chevalley and Weyl groups. The case of $\mathrm{GL}_n(q)$ is emblematic of all these families. We have a natural tower:

$$\mathrm{GL}_n(q) > \mathrm{P}_n(q) > \mathrm{GL}_{n-1}(q) \times \mathrm{GL}_1(q) > \mathrm{GL}_{k-1}(q) > \{1\} \ .$$

Here $\mathrm{P}_k(q)$ is the so-called maximal parabolic subgroup of the form shown in Figure 2, where $A \in \mathrm{GL}_{k-1}(q), v \in \mathbb{F}_q^{k-1}$, and $c \in \mathbb{F}_q^\times$. Our generators are block-diagonal with an arbitrary element of $\mathrm{GL}_2(q)$ in the $i, i-1$ block and all other diagonal elements equal to 1. The adapted diameter is $O(n^2)$, the maximum index is q^{n-1}, and the maximum multiplicity is $q^{O(n)}$. Analogous factorizations arise in the case of the finite groups of Lie type as well as the finite unitary groups [18].

$$\left(\begin{array}{c|c} A & v \\ \hline 0\ldots0 & c \end{array}\right)$$

Figure 2: P_k

Theorem 1 then implies a quantum circuit of size $q^{O(n)}$ for the QFT over these groups. Since $|G| = O(q^{n^2})$ we can write this as $|G|^{O(1/n)}$, which is $\exp\big(O(\sqrt{\log|G|})\big)$ if q is fixed. Note that the best-known classical algorithm for these groups [17] has complexity $|G|\, q^{\Theta(n)} = G^{1+\Theta(1/n)}$; therefore, we argue that this quantum speedup is the most we could expect relative to the existing classical algorithm. Note, for instance, that for the group families above for which we obtain circuits of size polylog$(|G|)$, there are classical algorithms of complexity $|G|\,\text{polylog}(|G|)$. In both cases it appears that the natural quantum speedup is to remove a factor of $|G|$ (modulo polylogarithmic terms).

4.2 Extensions of exponential index and Coppersmith-type circuits

The reader familiar with Coppersmith's circuit [7] for the QFT over $G = \mathbb{Z}_{2^n}$, where $H = \mathbb{Z}_{2^{n-1}}$, will recall that the Hadamard gate embeds a character $\sigma \in \hat{H}$ in two characters $\rho \in \hat{G}$, applies part of the twiddle factor, and sums over the two cosets of H, all in one operation. This is in contrast to Beals' technique, which sums over the cosets serially. Indeed, if the index $[G:H]$ is exponential — for instance, if G is an extension of H by $\mathbb{Z}_p$ where p is exponentially large — then Beals' technique takes exponential time.

For a certain type of extension, we can construct circuits analogous to Coppersmith's, which use quantum parallelism to embed σ in the ρ, sum over all p cosets simultaneously, and apply the twiddle factor as well. Recall that G is a *split extension* or *semidirect product* of H by T, written $T \ltimes H$, if $H \lhd G$ and there is a transverse subgroup $T < G$ so that $T \cong G/H$.

Definition 3 *Suppose G is a split extension of H by T, and let S be a set of at most $\log_2 |T|$ generators for T, and suppose that G is polynomially uniform with respect to a tower of subgroups where $G = G_m$ and $H = G_{m-1}$ and a Bratteli diagram $\mathfrak{B}$. Then G is a* homothetic extension *of H by T if*

1. *Given $\sigma \in \hat{H}$ and $\gamma \in S$, define $\sigma^\gamma(h) = \sigma(\gamma^{-1} h \gamma)$. Then for every $\sigma \in \hat{H}$, either $\sigma^\gamma = \sigma$, or the orbit of q distinct representations σ^{γ^j}, for $0 \le j < q$ where q divides the order of γ, appears among the representations of H given by $\mathfrak{B}$.*

2. *For each $\gamma \in S$, there is a classical algorithm which runs in* polylog$(|G|)$ *time which, given a path s in $\mathfrak{B}$ indexing a row of $\sigma[s]$ and an integer j, returns the size q of σ's orbit under conjugation by γ, and returns a path s^{γ^j} that indexes the same row of $\sigma[s^{\gamma^j}] = \sigma^{\gamma^j}$.*

Theorem 2 *If G is a homothetic extension of H by an Abelian group, then the Fourier transform of G can be obtained from the state* (4) *using* polylog$(|G|)$ *elementary quantum operations .*

Proof. It is easy to show that a homothetic extension of H by $A \times B$ consists of a homothetic extension of H by A, followed by a homothetic extension by B. Therefore it suffices to prove the lemma for homothetic extensions by cyclic groups of prime power order, so let T be generated by γ of order p^z.

We recall some representation theory from [6, 22]. Given $\sigma \in \hat{H}$, the *stabilizer* of σ is $K = \{x \in T : \sigma^x \cong \sigma\}$, and for a homothetic extension we can replace $\sigma^x \cong \sigma$ with $\sigma^x = \sigma$. Then K is the subgroup of T of order p^ℓ generated by γ^q where $q = p^{z-\ell}$, and σ's orbit under conjugation by γ is of size q.

The representations ρ in which σ appears can be obtained in two steps. First, we extend σ to $K \ltimes H$ by multiplying σ by one of the p^ℓ characters of K. This yields $\tau_b \in \widehat{K \ltimes H}$ where $\tau_b(\gamma^{qj} h) = \chi_b(j)\sigma(h)$ and $\chi_b(\gamma^{qj}) = \omega_{p^\ell}^{bj}$. Since $d_{\tau_b} = d_\sigma$, we have $A_{\sigma,\tau_b} = \sqrt{1/p^\ell}$ and σ embeds in a uniform superposition over the τ_b, so we append a uniform superposition of edges $1 \le e \le p^\ell$ where $b = e - 1$. Combining this with the twiddle factor χ_b gives the unitary transformation

$$\left|\gamma^{qj+k}\right\rangle |s,t\rangle \to \left|\gamma^k\right\rangle \otimes \frac{1}{\sqrt{p^\ell}} \sum_{e=1}^{p^\ell} \omega_{p^\ell}^{(e-1)j} |se, te\rangle \quad . \tag{9}$$

Here we write the power of γ in two registers $0 \le j < p^\ell$ and $0 \le k < q$. Then this operation Fourier transforms the first register over $\mathbb{Z}_{p^\ell}$ and transfers the result to the mth register of s and t. This transform can be carried out with $O(\log p^\ell \log\log p^\ell) = O(\log|G| \log\log|G|)$ elementary operations [10, 16]. Note that p^ℓ takes at most $\log|G|$ different values, and can be obtained from the classical algorithm which computes q.

If $K = T$, then the $\rho \in \hat{G}$ containing σ are simply the extensions τ_b and we're done. If $K < T$, i.e. if $q > 1$, we carry out a second step as follows. Each τ_b appears in a single induced representation ρ_b whose restriction to $K \ltimes H$ is the direct product of all the representations in σ's orbit, times χ_b: that is, $\rho_b|_H = \chi_b \oplus_{i=0}^{q-1} \sigma^{\gamma^i}$. The twiddle factor $\rho_b(\gamma^k)$ is then a permutation matrix which cycles these p blocks k times, with an additional phase change $\omega_{p^z}^{bk}$. This gives the unitary transformation

$$\left|\gamma^k\right\rangle |se, te\rangle \to \omega_{p^z}^{(e-1)k} |0\rangle \left|s^{\gamma^k} e, te\right\rangle \quad . \tag{10}$$

Since s^{γ^k} can be calculated by the classical algorithm in polylog($|G|$) time, and since it is easy to implement $\omega_{p^z}^{bk}$ with phase shifts $\omega_{p^z}^{2^y b}$ for $0 < y < \log_2 k$ conditioned on the binary digit sequence of bk, we can perform this operation in polylog($|G|$) quantum steps. Composing (9) and (10) transforms the state (4) to the Fourier transform (5) over G. $\square$

Closure under homothetic extensions and the metacyclic groups. Theorem 2 shows that the set of groups for which circuits of polylog($|G|$) size exist is closed under homothetic extensions by Abelian groups. It also generalizes the efficient quantum Fourier transform of Høyer [13] for the *metacyclic* groups $\mathbb{Z}_q \ltimes \mathbb{Z}_p$, since these are homothetic extensions of $\mathbb{Z}_p$ by $\mathbb{Z}_q$. Note that the metacyclic groups include the dihedral groups (where $q = 2$) and the affine groups (where $q = p - 1$) as special cases.

The general case. In general, Abelian extensions can be slightly more complicated; consider extensions by $\mathbb{Z}_p$. If σ^γ is isomorphic to σ, rather than equal to it, γ induces an additional twiddle factor $C(\gamma)$ which changes σ's basis [22]. This occurs, for instance, if γ^p is an element of H other than the identity, in which case the cyclic group generated by γ is not transverse to H and the extension is not split. In this case $C(\gamma)$ is a p'th root of $\sigma(\gamma^p)$.

Relation to Coppersmith's circuit. Let γ be a generator of $G = \mathbb{Z}_{2^n}$. Then G is an extension of $H = \mathbb{Z}_{2^{n-1}}$ with transversal $\{1, \gamma\}$. Since $\gamma^2 \neq 1$, γ induces an additional phase shift $C(\gamma) = \sqrt{\chi_b(\gamma^2)} = \omega_{2^n}^b$. (Similarly, the additional phase shift in (10) is due to the fact that $\mathbb{Z}_{p^z}$ is not a split extension of $\mathbb{Z}_{p^\ell}$.) In Coppersmith's circuit, $C(\gamma)$ appears as a set of phase shift gates conditional on the low-order bit of j. Finally, the Hadamard gate in Coppersmith's circuit is precisely the operation (9) in the case $p = 2$, $\ell = 1$ and $q = 1$, and where we use the same qubit register for e (the high-order bit of the frequency) as for α (the low-order bit of the time).

The quaternionic groups. Another example is the generalized quaternion group, which is an extension of $H = \mathbb{Z}_{2n}$ by $\mathbb{Z}_2$ where γ^2 is the element of order 2 in H. Then $C(\gamma) = \sqrt{\sigma(\gamma^2)} = 1$ or i. Püschel, Rötteler and Beth [21] gave an efficient quantum Fourier transform for these groups in the case where n is a power of 2. Of course, these groups are extensions of Abelian groups with bounded index, so Lemma 2 already provides an efficient QFT for them.

Metabelian groups. Even if an extension is neither homothetic nor of polynomial index, we can still construct an efficient QFT if we can apply arbitrary powers of $C(\gamma)$ in polynomial time. This is true, for instance, if $C(\gamma)$ is of polynomial size, which is true whenever all the representations of H are of polynomial size. This includes the *metabelian* groups, i.e. split extensions of Abelian groups by Abelian groups, since all the representations of H are one-dimensional. We discuss this further in the full paper.

5 Conclusion and open problems

The separation of variables is in essence a coarse scale use of a factorization of the dual, using blockwise redundancy as well as sparseness. It is possible to use the Bratteli diagram indexing and accompanying path factorizations in a more precise fashion, effectively looking for redundancy and sparsity on the level of individual elements. This finer analysis is responsible for the fastest known classical FFTs for the groups $\mathrm{SL}_2(q)$, as well as S_n and its wreath products [19]. It would be interesting to investigate the possibility of adapting these techniques to the quantum setting.

Acknowledgements

The authors gratefully acknowledge the support of the National Science Foundation under grants CCR-0220264, EIA-0218443, and CCR-0093065 and the hospitality of the Mathematical Sciences Research Institute, where portions of this research were completed.

References

[1] ACM, editor. *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing: Hersonissos, Crete, Greece, July 6–8, 2001*, New York, NY, USA, 2001. ACM Press.

[2] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Revew A*, pages 3457–, 1995.

[3] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In ACM, editor, *Proceedings of the twenty-ninth annual ACM Symposium on the Theory of Computing: El Paso, Texas, May 4–6, 1997*, pages 48–53, New York, NY, USA, 1997. ACM Press.

[4] Thomas Beth. On the computational complexity of the general discrete Fourier transform. *Theoret. Comput. Sci.*, 51(3):331–339, 1987.

[5] Michael Clausen. Fast generalized Fourier transforms. *Theoret. Comput. Sci.*, 67(1):55–63, 1989.

[6] A. H. Clifford. Representations induced in an invariant subgroup. *Annals of Mathematics*, 38:533–550, 1937.

[7] D. Coppersmith. An approximate fourier transform useful in quantum factoring. Technical Report RC19642, IBM, 1994. Quantum Physics e-Print Archive, quant-ph/0201067.

[8] Persi Diaconis and Daniel Rockmore. Efficient computation of the Fourier transform on finite groups. *J. Amer. Math. Soc.*, 3(2):297–332, 1990.

[9] Michelangelo Grigni, Leonard Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In ACM [1], pages 68–74.

[10] L. Hales and S. Hallgren. An improved quantum Fourier transform algorithm and applications. In IEEE, editor, *41st Annual Symposium on Foundations of Computer Science: proceedings: 12–14 November, 2000, Redondo Beach, California*, pages 515–525, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 2000. IEEE Computer Society Press. IEEE Computer Society Order Number PR00850.

[11] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. In ACM, editor, *Proceedings of the thirty second annual ACM Symposium on Theory of Computing: Portland, Oregon, May 21–23, [2000]*, pages 627–635, New York, NY, USA, 2000. ACM Press.

[12] Frank Harary. *Graph Theory*. Addison-Wesley, 1969.

[13] Peter Høyer. Efficient quantum transforms. Technical Report quant-ph/9702028, Quantum Physics e-Print Archive, 1997.

[14] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 263–270, Heraklion, Crete Island, Greece, 4-6 July 2001. ACM.

[15] Adalbert Kerber. *Representations of Permutations Groups I, II*, volume 240 and 495 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1971 and 1975.

[16] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. Technical Report quant-ph/9511026, Quantum Physics e-Print Archive, 1995.

[17] David K. Maslen and Daniel N. Rockmore. Adapted diameters and the efficient computation of Fourier transforms on finite groups. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 253–262, San Francisco, California, 22–24 January 1995.

[18] David K. Maslen and Daniel N. Rockmore. Separation of variables and the computation of Fourier transforms on finite groups, I. *Journal of the American Math Society*, 10(1):169–214, 1997.

[19] David K. Maslen and Daniel N. Rockmore. The Cooley-Tukey FFT and group theory. *Notices Amer. Math. Soc.*, 48(10):1151–1160, 2001.

[20] Cristopher Moore, Daniel Rockmore, Alexander Russell, and Leonard Schulman. The hidden subgroup problem in affine groups: Basis selection in Fourier sampling. Technical Report quant-ph/0211124, Quantum Physics e-Print Archive, 2003.

[21] Markus Püschel, Martin Rötteler, and Thomas Beth. Fast quantum fourier transforms for a class of non-abelian groups. In *Proceedings of Applied Algebra Algebraic Algorithms, and Error-Correcting Codes (AAECC-13)*, volume 1719 of *Lecture Notes in Computer Science*, pages 148–159. Springer-Verlag, 1999.

[22] Daniel Rockmore. Fast Fourier analysis for Abelian group extensions. *Advances in Applied Mathematics*, 11:164–204, 1990.

[23] Daniel Rockmore. Fast fourier transforms for wreath products. *J. Applied and Computational Harmonic Analysis*, 2:279–292, 1995.

[24] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Springer-Verlag, 1977.

[25] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997.

[26] Barry Simon. *Representations of Finite and Compact Groups*, volume 10 of *Graduate Studies in Mathematics*. American Mathematical Society, 1996.

[27] John Watrous. Quantum algorithms for solvable groups. In ACM [1], pages 60–67.