GENERIC QUANTUM FOURIER TRANSFORMS

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

Daniel Rockmore
Dartmouth College
rockmore@cs.dartmouth.edu

Alexander Russell
University of Connecticut
acr@cse.uconn.edu

Abstract

The *quantum Fourier transform* (QFT) is the principal ingredient of most efficient quantum algorithms. We present a generic framework for the construction of efficient quantum circuits for the QFT by "quantizing" the highly successful *separation of variables* technique for the construction of efficient classical Fourier transforms. Specifically, we apply the existence of computable Bratteli diagrams, adapted factorizations, and Gel'fand-Tsetlin bases to provide efficient quantum circuits for the QFT over a wide variety of finite Abelian and non-Abelian groups, including all group families for which efficient QFTs are currently known and many new group families. Moreover, our method provides 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 seminal discovery of efficient quantum algorithms for factoring and discrete logarithm [25] relies crucially on the fact that the Fourier transform over the cyclic group $\mathbb{Z}_n$ can be carried out efficiently on a quantum computer, even when n is exponentially large. This 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 already 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 be intimately related to the structure of the group in question: while quantum Fourier transforms and hidden subgroup problems over Abelian groups are well-understood, for non-Abelian groups our understanding of these problems remains embarrassingly sporadic. Aside from their natural appeal, these lines of research are motivated by their 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 $\sum_{z\in\mathbb{Z}_n} f(z)\,|z\rangle$ to the state $\sum_{\omega\in\mathbb{Z}_n} \hat{f}(\omega)\,|\omega\rangle$, where $f:\mathbb{Z}_n\to\mathbb{C}$ is a function with $\|f\|_2 = 1$ and $\hat{f}(\omega) = \sum_z f(z)\,\mathrm{e}^{2\pi i\omega z/n}$ denotes the familiar discrete Fourier transform at the frequency ω. Over an arbitrary finite group G, this analogously refers to the transformation taking the state $\sum_{z\in G} f(z)\,|z\rangle$ to the state $\sum_{\rho\in\hat{G}} \hat{f}(\rho)_{ij}\,|\rho,i,j\rangle$, where $f:G\to\mathbb{C}$, as before, is a function with $\|f\|_2 = 1$ and $\hat{f}(\rho)_{ij}$ denotes the i,j entry of the Fourier transform at the representation ρ. 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—the so-called *transform and measure* approach—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, 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 computation of Fourier transforms over a wide variety of important families of groups. These include, for example, the finite groups of Lie type (properly parametrized) and the symmetric groups.

Our main result is an adaptation to the quantum setting of the most successful and general of these techniques, namely the "separation of variables" approach. While almost all efficient classical Fourier transforms are divide-and-conquer algorithms, which recursively perform the Fourier transform for a series of subgroups and combine the results according to their coset structure, the separation of variables approach uses the existence of *adapted bases* to streamline this process considerably.

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

THEOREM 1.1. *If G is a polynomially uniform group with a subgroup tower $G = G_m > G_{m-1} > \cdots > \{1\}$ 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 Corollary 3.1 of [17] does for the classical case. In fact, for many of the group families we study, the quantum and classical circuit complexities of the Fourier transform differ by a factor of $|G|$. We extend this class further by showing that it is closed under a certain type of Abelian extension which may have exponential index.

This framework allows us to give efficient QFTs—that is, circuits of $\mathrm{polylog}(|G|)$ size—for many new families of groups, as well as to place existing QFT algorithms in a uniform framework. These include

(i.) the Clifford groups $\mathbb{CL}_n$;

(ii.) symmetric groups, recovering Beals' algorithm [3];

(iii.) wreath products $G \wr S_n$ where $|G| = \mathrm{poly}(n)$;

(iv.) metabelian groups (semidirect products of two Abelian 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, 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 our techniques provide efficient circuits for the QFT. We conclude with open problems in Section 5.

2 Representation theory background

Fourier analysis over a group G consists of expressing arbitrary functions $f : G \to \mathbb{C}$ as linear combinations of basis functions which reflect the group's structure and symmetries. If G is Abelian, these are the *characters* of G, i.e., the homomorphisms of G into $\mathbb{C}$; for a general group, they are the *irreducible matrix elements*. Then 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 $\mathrm{U}(V)$ denotes the group of unitary linear operators on a finite-dimensional vector space V whose dimension we denote d_ρ. Once we fix an orthonormal basis for V, each $\rho(g)$ is a $d_\rho \times d_\rho$ unitary matrix and is called a *matrix representation* of G. Each of the d_ρ^2 functions $\rho_{ij}(g) = [\rho(g)]_{ij}$ is called a *matrix element* of ρ; note that while ρ is a homomorphism, in general ρ_{ij} is not.

A matrix representation ρ of G on V is called *irreducible* if the only subspaces it preserves are the trivial one, $\{0\}$, and V itself. This is equivalent to the statement that there is no change of basis that simultaneously gives a block diagonalization (of a given shape) of $\rho(g)$ for all 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 explicit bases for the representations of $\hat{G}$ results in a set of (inequivalent irreducible) matrix representations, whose *matrix elements* then form an orthonormal basis for the $|G|$-dimensional vector space of complex-valued functions on G. Since there must be enough matrix elements to span this space, this implies the following important relationship between $|G|$ and the dimensions of the irreducible representations:

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

We are now equipped to give the general definition of the Fourier transform over arbitrary groups. Marvelously, this definition possesses many of the properties of the Fourier transform over $\mathbb{Z}_n$ that we know and love; for instance, it transforms convolution into (matrix) product.

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.

The Fourier transform is linear in f; with the constants $\sqrt{d_\rho/|G|}$ we use here, it is in fact unitary, taking the $|G|$ complex numbers $\langle f(g) \rangle_{g \in G}$ to a total of $\sum d_\rho^2 = |G|$ complex numbers organized into $|\hat{G}|$ matrices with varying dimensions d_ρ.

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 $\frac{1}{|G|} \sum_g f_1(g) f_2(g)^*$. The orthonormality of the matrix elements can then be expressed as follows: for any pair of matrix representations $\rho, \sigma \in \hat{G}$,

$$\langle \rho_{ij}, \sigma_{kl} \rangle = \begin{cases} 0 & \text{if } \rho \ncong \sigma \ , \\ \frac{1}{d_\rho} \delta_{ik} \delta_{jl} & \text{if } \rho = \sigma \ . \end{cases} \tag{2.1}$$

This is one form of *Schur's lemma* [24]. We can use this orthonormality to invert the Fourier transform, giving 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 $\rho : G \to U(V)$ can always be decomposed into a direct product of irreducible representations. Specifically, there is a basis of V in which ρ is block diagonal, where the ith block of ρ is precisely σ_i for some irreducible matrix representation σ_i. In this case we write $\rho = \bigoplus_i \sigma_i$. The number of times a given $\sigma_i \in \hat{G}$ appears in this decomposition is the *multiplicity* of σ_i in ρ; denoting this multiplicity w_i, we will 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.

Remark. The familiar *Discrete Fourier Transform* (DFT) corresponds to the case $G = \mathbb{Z}_n$. In this case the representations are all one-dimensional, and the Fourier transform is an $n \times n$ Vandermonde matrix whose entries are nth roots of unity.

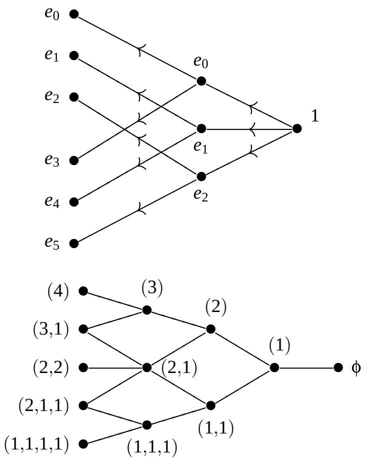

Figure 1: The Bratteli diagrams for the subgroup towers $\mathbb{Z}_6 > \mathbb{Z}_3 > 1$ (top) and $S_4 > S_3 > S_2 > 1$ (bottom). Cyclic groups of order n have representations indexed by the integers mod n, and (assuming $m|n$) the representation corresponding to j restricts to the representation corresponding to j mod 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.

3 Making divide-and-conquer feasible: Bratteli diagrams, Gel'fand-Tsetlin bases, and adapted diameters

The classic Cooley-Tukey Fast Fourier Transform relies on the fact that the cyclic group $\mathbb{Z}_{2^k}$ can be decomposed into a tower of subgroups:

$$\mathbb{Z}_{2^k} > \mathbb{Z}_{2^{k-1}} > \cdots > \mathbb{Z}_4 > \mathbb{Z}_2 > \mathbb{Z}_1 = \{1\}$$

The Cooley-Tukey algorithm works recursively, by calculating the FFT for each subgroup in the tower, and then combining the results from that subgroup's two cosets to form the FFT at the next level up.

Almost all efficient classical algorithms for the Fourier transform work in this way. However, in the non-Abelian case, making this divide-and-conquer approach concrete is far from trivial. Even if the group has a natural subgroup tower, we need to choose a set of bases for the representations which allows us to embed the cosets of each subgroup in the next one up in an efficient way. Furthermore, we need to choose a set of generators into which we can factor group elements efficiently, and our choice of bases should make the matrix representations for these generators sparse and highly structured, so that they can be multiplied together efficiently. (Finally, in the quantum setting, we will have to write the resulting transform as a product of elementary unitary operations.)

Luckily, there are principled ways to choose these bases and these generating sets. These techniques allow us to construct an efficient classical Fourier transform from the following ingredients:

(i.) a tower (or "chain") of subgroups, by which the Fourier transform on G can be built recursively as an accumulation of Fourier transforms on increasingly larger subgroups;

(ii.) a natural indexing scheme for the representations given by paths in the *Bratteli diagram* corresponding to that subgroup tower, which in turn provides a convenient basis for each representation; 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 the resulting algorithm can then be derived in terms of the basic representation-theoretic and combinatorial data of the subgroup tower, the Bratteli diagram, and the generating set. We describe the recipe by which these ingredients are made into efficient classical transforms in the next two sections.

3.1 Bratteli diagrams and Gel'fand-Tsetlin bases

Much of Abelian Fourier analysis is simplified by the fact that in this case the the set of characters $\hat{G} = \{\chi : G \to \mathbb{C}\}$, also called the *dual*, forms a group isomorphic to the original group G. Furthermore, in this isomorphism lies a natural correspondence which provides an indexing of the irreducible representations, and thus the matrix elements of the transform. 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, where efficient Fourier analysis is concerned, not all bases are created alike!

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. These bases allow us to carry out the recursive divide-and-conquer approach described above, building the transform efficiently at each level of the subgroup tower. To construct these bases, we need a natural indexing scheme for the representations, and for each of their matrix elements, Happily, such an indexing scheme is given by the *Bratteli diagram* formalism, which we now present. Given a finite group G, 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 at 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 σ of G_i to ρ of G_{i+1} is equal to the *multiplicity* of σ 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. Bratteli diagrams for the cyclic group $\mathbb{Z}_6$ and the symmetric group S_4 are shown in Figure 1.

We now describe how paths in the Bratteli diagram index the rows and columns of each representation, and thus provide a natural set of bases. Each edge, from a node $\sigma : G_i \to U(V_\sigma)$ of $\hat{G}_i$ to a node $\rho : G_{i+1} \to U(V_\rho)$ of $\hat{G}_{i+1}$, represents an embedding of V_σ into V_ρ. Thus the edges into ρ describe a decomposition of V_ρ into a direct sum of orthogonal subspaces V_σ, each of which are invariant under the (restricted) action of G_i; and conversely, the edges out from σ correspond to embeddings of V_σ into orthogonal subspaces V_ρ. Thus these edges describe how the subspaces acted on by the representations of G_{i+1} are decomposed into smaller subspaces acted on by representations of G_i, and conversely how the subspaces of G_i are embedded in the subspaces of G_{i+1}.

Since the only representation of the trivial group $\{1\}$ is one-dimensional, composing these edges into paths from the root to a given node $\rho \in \hat{G}_i$ gives a decomposition of V_ρ into a direct sum of orthogonal one-dimensional subspaces; but this is tantamount to providing a basis for V_ρ. Moreover, since paths from the root to ρ consist of paths from the root to various σ composed with paths from σ to ρ, where $\sigma \in \hat{G}_j$ for some $G_j < G_i$, this basis has 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. Therefore, in this basis the matrix representation ρ is block diagonal according to this partition when restricted to G_j and, moreover, the blocks corresponding to some σ which appears in ρ with multiplicity greater than 1 are actually equal. Such bases are said to be *G_j-adapted* or *Gel'fand-Tsetlin*.

Note that the number of paths to a node ρ is equal to d_ρ (so, for instance, the Bratteli diagram of an Abelian group is a directed tree). Furthermore, each ordered pair of paths with common endpoint ρ indexes an irreducible matrix element of ρ, since one path indexes a row and the other indexes a column.

Following the divide and conquer approach, the Fourier transform on $G = G_m$ can be written a sum of Fourier transforms on G_{m-1}, each of which is translated from a different 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

$$\begin{aligned} \hat{f}(\rho) &= \sum_{\alpha \in T} \rho(\alpha) \sum_{x \in G_{m-1}} \rho(x) f(\alpha x) \\ &= \sum_{\alpha \in T} \rho(\alpha) \cdot \hat{f}_\alpha(\rho|_{G_{m-1}}). \end{aligned} \tag{3.2}$$

These matrices $\rho(\alpha)$ are called the "twiddle factors". Note that the number of terms in this sum is $|T| = [G_m : G_{m-1}] = |G_m|/|G_{m-1}|$, the *index* of G_{m-1} in G_m. As we will see below, the recursion of (3.2) will be greatly simplified by the fact that in the adapted basis, the restricted representations $\rho|_{G_j}$ become block diagonal, where the blocks are simply the matrices σ.

3.2 Strong generating sets and adapted diameters

Adapted representations are only part of the story for the construction of efficient Fourier transforms. In general, the twiddle factors $\rho(\alpha)$ in Equation (3.2) could be an arbitrary matrices of exponential size, so an algorithm which simply performs the sum in (3.2) could be costly. Luckily, under fairly mild assumptions, these twiddle factors can be factored into polylog$(|G|)$ sparse, highly structured matrices, and can therefore be implemented with polylog$(|G|)$ elementary quantum operations.

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

$$D_i = \min\{\ell > 0 : \cup_{j \le \ell} (S \cap G_i)^j \supseteq T_i\} \ ,$$

i.e., the length of words over $S \cap G_i$ we need to generate every representative in T, 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 3.1. *(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*

$$(GL_{m_1}(\mathbb{C}) \otimes I_{d_1}) \oplus \cdots \oplus (GL_{m_r}(\mathbb{C}) \otimes I_{d_r}) \tag{3.3}$$

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 (3.3). Therefore, the total number of elementary quantum operators we need to implement $\rho(\alpha)$ is $D \times \text{poly}(M)$.

Moreover, if γ is itself in a subgroup $H > K$, and ρ is adapted to both H and K, then $\rho(\gamma)$ also possesses the block structure corresponding to $\rho|_H$. This places an upper bound on M, namely the maximum multiplicity with which representations of K appear in restrictions of representations of H. Thus we can minimize M by choosing generators γ which (1) are inside subgroups as low on the tower as possible, and (2) centralize subgroups as high on the tower as possible.

For instance, for the symmetric group S_n we take the tower to be

$$S_n > S_{n-1} > \cdots > \{1\} \ ,$$

where S_i fixes all elements of $\{1, \ldots, n\}$ 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. In this case the adapted basis defined by the Bratteli diagram is exactly the *Young orthogonal basis*, in which each block of $\rho((j, j+1))$ differs from the identity only by a 2×2 minor. Since the adapted diameter is easily seen to be $O(n^2)$, this means that the twiddle factors $\rho(\alpha)$ can be carried out in $O(n^2) = \text{polylog}(|S_n|)$ elementary quantum operations [3]. We will see in the next section that a similar situation obtains for a large class of groups.

4 Efficient quantum Fourier transforms

We describe our algorithm in this section. As in the classical case, we perform 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. (This is how Coppersmith's circuit for the QFT over $\mathbb{Z}_{2^k}$ works; see below.) 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

$$|\alpha\rangle\,|s,t\rangle = \left|\alpha_m \cdots \alpha_{i+1} 0^i\right\rangle \otimes \left|s_1 \cdots s_i 0^{m-i}, t_1 \cdots t_i 0^{m-i}\right\rangle \quad .$$

Keep in mind that 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 $1 \leq j, k \leq d_\sigma$ 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 $h \in H$. As in (3.2), for each $\alpha \in T$ we define a function f_α on H as $f_\alpha(h) = f(\alpha h)$; this is the restriction of f to the coset αH, translated into H. After having performed the Fourier transform on H, our state will be

$$\begin{aligned} & \sum_{\alpha \in T} |\alpha\rangle \otimes \sum_{s,t \text{ of length } m-1} \hat{f}_\alpha(s,t)\,|s,t\rangle \\ (4.4) \qquad = & \sum_{\alpha \in T} |\alpha\rangle \otimes \sum_{(\sigma,j,k) \in \hat{H}} \hat{f}_\alpha(\sigma)_{j,k}\,|\sigma, j, k\rangle \quad . \end{aligned}$$

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

$$\begin{aligned} & |0\rangle \otimes \sum_{s,t \text{ of length } m} \hat{f}(s,t)\,|s,t\rangle \\ (4.5) \qquad = & |0\rangle \otimes \sum_{(\rho,j,k) \in \hat{G}} \hat{f}(\rho)_{j,k}\,|\rho, j, k\rangle \quad . \end{aligned}$$

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

As described in Equation (3.2) above, $\hat{f}$ can be written as a sum over contributions from f's values on each coset αH, giving

$$(4.6) \qquad \hat{f}(\rho) = \sum_{\alpha \in T} \rho(\alpha) \cdot \hat{f}_\alpha(\rho|_H) \quad .$$

Recall that the matrix $\hat{f}_\alpha(\rho|_H)$ is a direct sum of submatrices of the form $\hat{f}_\alpha(\sigma)$, summed over the σ appearing in ρ. We will construct $\hat{f}(\rho|_H)$ via an *embedding* operation which reverses the restriction to H,

$$(4.7) \qquad |\sigma\rangle \to \sum_{\rho:\,\sigma \text{ appears in } \rho|_H} A_{\sigma,\rho}\,|\rho\rangle$$

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, as discussed above, 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, if $[G : H]$ is exponentially large, summing over the cosets will take exponential time unless parallelized in some way.

The Bratteli diagram, and the adapted basis it provides, allow us to accomplish (i) and (ii) above with a minimum of trouble. For (i), the embedding operation, 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. The only change of basis required is to literally pick the matrix elements of σ up and place them in the appropriate place in ρ, and we discuss below how to do this unitarily.

Similarly, when coupled with a strong generating set of small adapted diameter as discussed in Section 3.2, the adapted basis allows us to carry (ii) out efficiently by writing $\rho(\alpha)$ as a product of a small number of $\rho(\gamma)$, each of which has the block-diagonal structure given by Lemma 3.1.

For (iii), summing over the cosets, for now we simply take the time to sum over all the cosets serially, paying a cost of $[G_i : G_{i-1}]$ per level as reflected in Theorem 1.1. This makes sense for subgroup towers where the index of each subgroup in the one above it is polynomial, such as the tower for S_n above, and we focus on that case in Section 4.1. However, in Section 4.2 we will see that even when the index of some level of the tower is exponentially large, in some cases we can use the parallelism of quantum mechanics to sum over all the cosets simultaneously, and still achieve an efficient QFT.

We adopt the following 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 k is its out-degree. 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 (4.7) becomes

$$(4.8) \qquad U : \begin{cases} |0\rangle\,|s,t\rangle \to |0\rangle \sum_e A_{\sigma[s],\rho[se]}\,|se,te\rangle \\ |\alpha\rangle\,|s,t\rangle \to |\alpha\rangle\,|s,t\rangle \text{ for all } \alpha \in T \end{cases}$$

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

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 (4.8). This can be accomplished simply by putting the mth 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.

We shall focus on group towers for which the Bratteli diagram data can be effectively computed:

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

4.1 Extensions of small index We begin by focusing on groups and towers which are fairly refined, i.e., with polynomial indexes at each level.

LEMMA 4.1. *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* (4.4) *using* 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 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 (4.6), 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 (4.8). 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, and we are done. □

Proof of Theorem 1.1. This follows by induction as the depth of the Bratteli diagram is at most $\log|G|$. □

For many families of 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 1.1 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 $S_n > S_{n-1} > \cdots > \{1\}$, so the $I = n = o(\log|S_n|)$. The generators are the adjacent transpositions, so $D = O(n^2)$ and $M = 2$. The adapted basis is precisely the Young orthogonal basis.

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

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

Then $I = \max(n, |H|)$, the generators are the adjacent transpositions and an arbitrary set of $\log |H|$ generators for each factor of H, $D = O(n^2 \log |H|)$, and $M = O(|H|)$. Then note that $|H| = \mathrm{polylog}(|G|)$. 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

$$\mathbb{CL}_n > \mathbb{CL}_{n-1} > \cdots > \{1\}$$

for which $I = 2$, and the generators $\{x_1, x_1x_2, \ldots, x_{n-1}x_n\}$. Then $D = O(n)$, and since each $x_i x_{i+1}$ centralizes $\mathbb{CL}_{i-1}$ we have $M = 4$.

In addition to giving $\mathrm{polylog}(|G|)$-size circuits for these groups, our techniques give the first subexponential-size circuits for the following classical groups:

The linear groups $\mathbf{GL}_n(q)$, $\mathbf{SL}_n(q)$, $\mathbf{PGL}_n(q)$, and $\mathbf{PSL}_n(q)$; 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:

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

Here $\mathrm{P}_k(q)$ is the so-called *maximal parabolic subgroup*, consisting of elements of the form

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

where $A \in \mathrm{GL}_{k-1}(q), \vec{v} \in \mathbb{F}_q^{k-1}$, and $c \in \mathbb{F}_q^{\times}$, so $I = q^{n-1}$. Our generators are the block-diagonal matrices with an arbitrary element of $\mathrm{GL}_2(q)$ in the $i, i-1$ block and all other diagonal elements equal to 1. Then $D = O(n^2)$ and $M = q^{O(n)}$. Analogous factorizations arise for the finite groups of Lie type as well as the finite unitary groups [18].

Theorem 1.1 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. On the other hand, he best-known classical FFT for these groups [17] has complexity $|G|\, q^{\Theta(n)} = |G|^{1+\Theta(1/n)}$. Note for the group families above for which we obtain circuits of size $\mathrm{polylog}(|G|)$, there are classical algorithms of complexity $|G|\, \mathrm{polylog}(|G|)$. In both cases, it seems that the natural quantum speedup is a factor of $|G|$, modulo polylogarithmic terms; of course, we would like to know if this is the best possible.

4.2 Extensions of large index; 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 both cosets of H, all in one operation. This is in contrast to the technique of the previous section, which sums over the cosets serially—and which takes exponential time if, for instance, G is an extension of H by $\mathbb{Z}_p$ where p is exponentially large.

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 the 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, 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 (i.) 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 $0 \leq j < q$ where q divides the order of γ, appears among the representations of H given by $\mathfrak{B}$; (ii.) For each $\gamma \in S$, there is a classical algorithm which runs in* $\mathrm{polylog}(|G|)$ *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 4.1. *If G is a homothetic extension of H by an Abelian group, then the Fourier transform of G can be obtained from the state* (4.4) *using* $\mathrm{polylog}(|G|)$ *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 without loss of generality we 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 $\chi_b(\gamma^{qj}) = \omega^{bj}_{p^\ell}$. 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

$$\text{(4.9)} \quad \left|\gamma^{qj+k}\right\rangle |s,t\rangle \to \left|\gamma^k\right\rangle \otimes \frac{1}{\sqrt{p^\ell}} \sum_{e=1}^{p^\ell} \omega_{p^\ell}^{(e-1)j} \, |se,te\rangle \quad .$$

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

$$\text{(4.10)} \qquad \left|\gamma^k\right\rangle |se,te\rangle \to \omega_{p^z}^{(e-1)k} \, |0\rangle \left|s^{\gamma^k} e, te\right\rangle \quad .$$

Since s^{γ^k} can be calculated by the classical algorithm in $\mathrm{polylog}(|G|)$ 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 $\mathrm{polylog}(|G|)$ quantum steps. Composing (4.9) and (4.10) transforms the state (4.4) to the Fourier transform (4.5) over G. □

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 (4.10) 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 (4.9) 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).

Closure under homothetic extensions and metacyclic groups. Theorem 4.1 shows that the set of groups for which circuits of $\mathrm{polylog}(|G|)$ 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 quaternionic groups. The generalized quaternion group 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 4.1 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—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.

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; then $C(\gamma)$ is a pth root of $\sigma(\gamma^p)$.

5 Conclusion and open problems

We have shown that a general technique for constructing efficient classical fast Fourier transforms on groups—separation of variables using an adapted basis—can be carried over to the quantum context, producing circuits of $\mathrm{polylog}(|G|)$ size for a wide variety of groups, and of subexponential size for classical linear groups mod q.

While separation of variables is one of the most general techniques for classical FFTs, it is not the only one. It is possible to use the Bratteli diagram in a more precise fashion, looking for redundancy and sparsity on the level of individual matrix 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 explore 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 and Dartmouth College, 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.