Explicit Multiregister Measurements for Hidden Subgroup Problems; or, Fourier Sampling Strikes Back

Cristopher Moore
University of New Mexico

Alexander Russell
University of Connecticut

Abstract

We present an explicit measurement in the Fourier basis that solves an important case of the Hidden Subgroup Problem, including the case to which Graph Isomorphism reduces. This entangled measurement uses $k = \log_2 |G|$ registers, and each of the 2^k subsets of the registers contributes some information.

1 Introduction: The Hidden Subgroup Problem

Many problems of interest in quantum computing can be expressed as, or reduced to, an instance of the *Hidden Subgroup Problem* (HSP). We are given a group G and a function f with the promise that, for some subgroup $H \subseteq G$, f is invariant precisely under translation by H: that is, $f(g_1) = f(g_2)$ if any only if $g_1 = g_2 h$ for some $h \in H$. We then wish to determine the subgroup H. Every known efficient algorithm for this problem—and, indeed, almost every quantum algorithm that provides an exponential speedup over the best known classical algorithm—uses the approach of *Fourier sampling* [3]. By preparing a uniform superposition over the elements of G, querying the function f, and then measuring the value of f, we obtain a uniform superposition over one of the (left) cosets of H,

$$|cH\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle$$

where c is a uniformly random element of G. Alternately, we can view this as a mixed state over the left cosets, with density matrix

$$\rho = \frac{1}{|G|} \sum_{c \in G} |cH\rangle \langle cH| \quad .$$

We then carry out the quantum Fourier transform on $|cH\rangle$, or equivalently ρ, and measure the result.

For example, in Simon's problem [31], $G = \mathbb{Z}_2^n$ and there is some y such that $f(x) = f(x+y)$ for all x; in this case $H = \{0, y\}$ and we wish to identify y. In Shor's factoring algorithm [30] G is the group $\mathbb{Z}_n^*$ where n is the number we wish to factor, $f(x) = r^x \bmod n$ for a random $r < n$, and H is the subgroup of $\mathbb{Z}_n^*$ whose index is the multiplicative order of r. In both these algorithms, G is abelian, and it is not hard to see that for any abelian group a polynomial number[1] of experiments of this type allow us to determine H. In essence, each experiment yields a random element of the dual space $H^\perp$ perpendicular to H's characteristic function, and as soon as these elements span $H^\perp$ we can determine a set of generators for H by linear algebra.

While the *nonabelian* hidden subgroup problem appears to be much more difficult, solving it would provide enormous benefits. In particular, solving the HSP for the symmetric group S_n would provide an efficient quantum algorithm for the Graph Automorphism and Graph Isomorphism problems (see e.g. [18] for a review). Let G_1, G_2 be two rigid, connected graphs of size n, and let $H \subset S_{2n}$ be the automorphism group of their disjoint union. If $G_1 \cong G_2$, then $H = \{1, m\}$ is of order 2, consisting of the identity and an involution m composed of n disjoint transpositions; if $G_1 \ncong G_2$, then H is the trivial subgroup consisting

[1] Throughout the paper, the terms polynomial, subexponential, etc. refer to a function of $\log |G|$.

only of the identity. Thus even distinguishing subgroups of order 2 from the trivial subgroup would be sufficient to solve this case of Graph Isomorphism. Other important motivations include the relationship between the HSP on the dihedral group and hidden shift problems [4] and cryptographically important cases of the Shortest Lattice Vector problem [27].

So far, explicit polynomial-time quantum algorithms for the HSP are known only for a few families of nonabelian groups [8, 10, 13, 15, 17, 25, 28]. However, the basic idea of Fourier sampling can certainly be extended to the nonabelian case. Fourier basis functions are homomorphisms $\phi : G \to \mathbb{C}$ such as the familiar $\phi_k(x) = \mathrm{e}^{2\pi ikx/n}$ when G is the cyclic group $\mathbb{Z}_n$. In the nonabelian case, one instead considers *representations* of G, namely homomorphisms $\sigma : G \to \mathsf{U}(V)$ where $\mathsf{U}(V)$ is the group of unitary matrices acting on some vector space V of dimension d_σ. The *irreducible* representations are those which are not isomorphic to direct sums of representations on lower-dimensional subspaces, and we denote the set of irreducibles by $\widehat{G}$. We refer the reader to [9] for an introduction. We denote the set of functions $\psi : G \to \mathbb{C}$ with $\|\psi\|^2 = 1$, i.e., the Hilbert space of a group-valued register, as $\mathbb{C}[G]$; then the quantum Fourier transform consists of transforming vectors in $\mathbb{C}[G]$ from the basis $\{|g\rangle \mid g \in G\}$ to the basis $|\sigma, i, j\rangle$ where σ is the isomorphism type, or "name," of an irreducible representation and $1 \leq i, j \leq d_\sigma$ index a row and column (in a chosen basis for V). This transformation can be carried out efficiently for a wide variety of groups [2, 14, 24].

Several varieties of measurement in the Fourier basis have been proposed. *Weak Fourier sampling* consists of measuring just the name σ of the irreducible representation. *Strong Fourier sampling* consists of measuring the name σ and the column j in a basis of our choice. (As the state ρ is mixed uniformly over the left cosets, it is easy to show that measuring the row provides no information). As an intermediate notion, one can also consider measuring the column in a *random* basis for V.

Unfortunately, a series of negative results have shown that these types of measurement will not succeed in solving the Hidden Subgroup Problem in the cases we care most about—in particular, the case relevant to Graph Isomorphism [13, 10, 19]. Most recently, Moore, Russell and Schulman [22] showed that strong Fourier sampling requires an exponential number of experiments to distinguish the order-2 subgroups H defined above from each other or from the trivial subgroup.

However, there is still reason for hope. In the above description of Fourier sampling, f is queried just once, giving a coset state on a single group-valued register. One can also consider *multiregister* experiments, in which we carry out k queries of f, prepare k independent coset states, and then perform a joint measurement on the product state $\boldsymbol{\rho} = \rho^{\otimes k} = \rho \otimes \cdots \otimes \rho$. Note that this measurement does not generally consist of k independent measurements; rather, it is an *entangled* measurement, in which we measure vectors in $\mathbb{C}[G^k]$ along a basis whose basis vectors are not tensor products of k basis vectors in $\mathbb{C}[G]$. For instance, Ip [16] showed that the optimal measurement in the dihedral group is already entangled in the two-register case.

In one sense we already know that such a measurement can succeed. Ettinger, Høyer and Knill [6] showed that the density matrices $\boldsymbol{\rho}$ become nearly orthogonal for distinct subgroups for some $k = O(\log |G|)$. As a consequence, a measurement exists which determines the hidden subgroup with high probability. In [7] they make this result somewhat more constructive by giving an algorithm which solves the HSP by performing a brute-force search through the subgroup lattice of G; however, for groups of interest such as the symmetric groups, this algorithm takes exponential time. For the dihedral groups in particular, Kuperberg [20] devised a subexponential algorithm, which uses $2^{O(\sqrt{\log n})}$ time and registers, that works by performing entangled measurements on two registers at a time.

Regev [27] provided a beautiful kind of worst-case to average-case quantum reduction, by showing that the HSP for the dihedral group D_n can be reduced to uniformly random instances of the Subset Sum problem on $\mathbb{Z}_n$. Bacon, Childs, and van Dam [1] deepened this connection by determining the *optimal* multiregister measurement for the dihedral group, and showing that it consists of the so-called *pretty good measurement* (PGM); they used this to show a sharp threshold at $k = \log_2 n$ for the number of registers needed to solve the HSP. Moore and Russell [21] generalized their results to some extent, showing that the PGM is optimal for arbitrary groups G in the single-register case whenever we wish to distinguish the conjugates of some subgroup H from each other, and optimal in the multiregister case whenever (G, H) form a Gel'fand pair.

Whether a similar approach can be taken to the symmetric group S_n is a major open question. In particular, we would like to know whether there is a worst-case to average-case reduction analogous to Regev's, connecting the HSP to some Subset Sum-like problem and whether this would result in new subexponential-time quantum algorithms for Graph Isomorphism. We note that Moore and Russell [23] showed that performing strong Fourier sampling on two registers in S_n requires a superpolynomial number of experiments (specifically, $e^{\Omega(\sqrt{n}/\log n)}$) to distinguish order-2 subgroups $\{1, m\}$ from the identity, or from each other, and conjectured that $\Omega(n \log n)$ registers are necessary. On the other hand, they showed the variance over m in the observed probability distribution in the multiregister case [22] has a term for each subset of the registers, pointing towards an algorithm that finds a subset with particularly high variance, and thus gives a large amount of information about the hidden subgroup.

Our contribution. In this paper, as in Graph Isomorphism, we consider the special case of the HSP where we wish to distinguish the conjugates of some subgroup H from the trivial subgroup, where H has a "missing harmonic" (defined below). We give an explicit k-register measurement in the Fourier basis that distinguishes these two cases. Our approach relies on decomposing the tensor product of the representations observed in a given subset of the registers into a direct sum of irreducibles. Interestingly, each subset of the registers contributes a small amount of information, so that when $k \geq \log_2 |G|$ the measurement succeeds with constant probability. We hope that this may lead to worst-case to average-case quantum reductions involving generalizations of the Subset Sum problem.

2 A Sufficient Multiregister Experiment

We start by preparing independent coset states in k independent G-valued registers, giving the tensor product

$$\boldsymbol{\rho} = \rho^{\otimes k} = \left(\frac{1}{|G|} \sum_{c \in G} |cH\rangle \langle cH| \right)^{\otimes k} = \frac{1}{|G|^k} \sum_{\boldsymbol{c} \in G^k} \left|\boldsymbol{c}H^k\right\rangle \left\langle \boldsymbol{c}H^k\right| \quad . \tag{1}$$

Note that $\boldsymbol{\rho}$ can also be thought of as a random left coset of the product subgroup $H^k \in G^k$. Note also that $\boldsymbol{\rho}$ is the completely mixed state over $\mathbb{C}[G]^{\otimes k} = \mathbb{C}[G^k]$ if H is the trivial subgroup $\{1\}$.

In this section we give an explicit measurement in the Fourier basis which solves an important special case of the HSP, including the case relevant to Graph Isomorphism: namely, given a (non-normal) subgroup $H \subset G$, we wish to distinguish the conjugates of H from the trivial subgroup. Our measurement succeeds with constant probability whenever $k \geq \log_2 |G|$.

Missing harmonics. Recall that for any representation τ, the average of τ over a subgroup H is a projection operator denoted $\tau(H) = (1/|H|) \sum_{h \in H} \tau(h)$. Note that $\tau(H)$ is generally not of full rank, and indeed $\tau(H) = \mathbb{1}_{d_\tau}$ if and only if H is contained in the kernel of τ. Let us say that an irreducible representation η is a *missing harmonic* of H if $\eta(H) = 0$; this is then true for all of H's conjugates as well. For instance, if n is odd, then $\pi(H)$ for the sign representation π of S_{2n} where H is the order-2 subgroup defined for Graph Isomorphism above.

For simplicity, we focus on the case where H has some missing harmonic η; the idea is that if we ever observe it, then we know that the hidden subgroup must be trivial rather than a conjugate of H. The following lemma gives some sufficient conditions for H to have a missing harmonic.

Lemma 1. *If any of the following conditions hold, then H has a missing harmonic: 1) H is normal and nontrivial. 2) H intersects every coset of some proper normal subgroup $K \lhd G$. 3) $G = S_n$ and H is transitive. 4) $|H| > |G|/C$ where $C = \sum_{\tau \in \widehat{G}} d_\tau$.*

Proof. 1) Recall that if H is normal then for every $\tau \in \widehat{G}$, either $\tau(H) = \mathbb{1}_{d_\tau}$ or $\tau(H) = 0$. If H is not the trivial subgroup, then the latter must be true for at least one τ.

2) Recall that any irreducible representation of G/K gives an irreducible representation τ of G by composing it with the homomorphism $\phi : G \to G/K$. Since $\phi(H) = G/K$, we have $\tau(H) = 0$ for any such τ other than the trivial representation. (For instance, in Graph Isomorphism where n is odd, H is transverse to the alternating group A_n and τ is the sign representation.)

3) Let τ be the *standard representation*, corresponding to the Young diagram $(n-1, 1)$. This permutes the n vertices of an $(n-1)$-dimensional simplex centered at the origin, and so $\tau(H) = 0$ whenever H is transitive.

4) Recall that the *regular representation* Reg, namely $\mathbb{C}[G]$ under left multiplication by G, consists of d_τ copies of each $\tau \in \widehat{G}$. It is a simple exercise to show that $\mathbf{rk}\ \mathrm{Reg}(H)$ is the index $|G|/|H|$, and since

$$\mathbf{rk}\ \mathrm{Reg}(H) = \sum_{\tau \in \widehat{G}} d_\tau \mathbf{rk}\ \tau(H)$$

we have $\mathbf{rk}\ \mathrm{Reg}(H) \geq C$ if $\tau(H) \neq 0$ for all $\tau \in \widehat{G}$. □

Decomposing subsets of the registers. The state $\boldsymbol{\rho}$ is a density matrix defined on the Hilbert space $\mathbb{C}[G^k] = \mathbb{C}[G]^{\otimes k}$. Since it is completely mixed over left cosets of H^k, it commutes with left multiplication in G^k. It follows from Schur's lemma [20, 22, 23] that $\boldsymbol{\rho}$ is block-diagonal in the Fourier basis, where each block corresponds to one of the irreducible representations of G^k. These are tensor products of irreducible representations of G, $\boldsymbol{\sigma} = \sigma_1 \otimes \cdots \otimes \sigma_k$. To put it differently, the optimal measurement is consistent with first performing weak Fourier sampling on each of the k registers, observing the representation names $\sigma_1, \ldots, \sigma_k$.

The question is how to refine this measurement further, decomposing $\boldsymbol{\sigma}$ into smaller subspaces. (We abuse notation by identifying subspaces with the name of the representation that acts on them.) Happily, there is a natural way to do this that still respects the structure of G: specifically, we treat $\boldsymbol{\sigma}$ as a representation of G (rather than of G^k) by restricting to the *diagonal action*, where the element $g \in G$ acts by $\boldsymbol{\sigma}(g) = \sigma_1(g) \otimes \cdots \otimes \sigma_k(g)$. We can then further decompose $\boldsymbol{\sigma}$ into irreducible representations $\tau \in \widehat{G}$ under this action. If we observe a missing harmonic η under this decomposition, we know that the hidden subgroup is trivial. Unfortunately, in most cases of interest, the chances of observing η are exponentially small even if the hidden subgroup is trivial, so this direct approach does not work.

Instead, we focus on some subset $I \subseteq [k]$ of the registers. First, we can decompose $\boldsymbol{\sigma}$ into the tensor product of the registers inside and outside I, $\boldsymbol{\sigma} = \left(\bigotimes_{i \in I} \sigma_i\right) \otimes \left(\bigotimes_{i \notin I} \sigma_i\right)$. Now consider the decomposition of the registers in I into irreducible representations of G under the diagonal action (in which we left-multiply every register in I by g and leave the other registers fixed): we write $\bigotimes_{i \in I} \sigma_i = \tau_1 \oplus \cdots \oplus \tau_\ell$. For each nonempty I, this gives us a subspace

$$W^I_{\eta, \boldsymbol{\sigma}} = \left(\bigoplus_{i : \tau_i \cong \eta} \tau_i \right) \otimes \left(\bigotimes_{i \notin I} \sigma_i \right) \ ,$$

and we define $\Pi^I_{\eta, \boldsymbol{\sigma}}$ as the projection operator which projects onto this subspace. That is, $\Pi^I_{\eta, \boldsymbol{\sigma}}$ projects the registers in I into irreducible subspaces isomorphic to η, and leaves the other registers fixed. The following lemma shows that if η is a missing harmonic for H, then each of these projection operators annihilates $\boldsymbol{\rho}$.

Lemma 2. *Suppose that $\eta(H) = 0$. Then for all nonempty $I \subseteq [k]$ and all $\boldsymbol{\sigma}$, $\Pi^I_{\eta, \boldsymbol{\sigma}}\, \boldsymbol{\rho} = 0$.*

Proof. The state $\boldsymbol{\rho}$ is symmetric under right multiplication by any $\boldsymbol{h} \in H^k$. In particular, for each $h \in H$ it is symmetric under right multiplication by $\boldsymbol{h}$ such that $h_i = h$ for $i \in I$ and $h_i = 1$ for $i \notin I$. Let R^I be the average over all $h \in H$ of the operator that right-multiplies by this $\boldsymbol{h}$, i.e., that symmetrizes over the diagonal right H-action on the registers in I; then $\boldsymbol{\rho} = R^I \boldsymbol{\rho}$ and $\Pi^I_{\eta, \boldsymbol{\sigma}} R^I = 0$, so $\Pi^I_{\eta, \boldsymbol{\sigma}}\, \boldsymbol{\rho} = (\Pi^I_{\eta, \boldsymbol{\sigma}} R^I)\boldsymbol{\rho} = 0$. □

Now we patch these operators together to form our measurement. Let

$$W_{\eta,\boldsymbol{\sigma}} = \mathrm{span}_{I\subseteq[k],I\neq\emptyset} W^I_{\eta,\boldsymbol{\sigma}}$$

be the span of all these subspaces, and let $\Pi_{\eta,\boldsymbol{\sigma}}$ be the projection operator onto $W_{\eta,\boldsymbol{\sigma}}$. By Lemma 2, we know that $\Pi_{\eta,\boldsymbol{\sigma}}\,\boldsymbol{\rho} = 0$ whenever η is a missing harmonic for the hidden subgroup. Thus we can distinguish the conjugates of H from the trivial subgroup with a measurement operator that reports "trivial" if it observes the subspace $W_{\eta,\boldsymbol{\sigma}}$, and "don't know" if it observes the perpendicular subspace $W^{\perp}_{\eta,\boldsymbol{\sigma}}$. Since $\boldsymbol{\rho}$ is completely mixed if the hidden subgroup is trivial, the probability that our operator reports "trivial" in that case is $\dim W_{\eta,\boldsymbol{\sigma}}/d_{\boldsymbol{\sigma}}$. We wish to show that if $k \geq \log_2 |G|$, the expectation over $\boldsymbol{\sigma}$ of this fraction is at least $1/2$, so that our measurement distinguishes the trivial subgroup from conjugates of H with constant probability.

To calculate this expectation, it is convenient to work in the entire Hilbert space $\mathbb{C}[G^k]$ of the k registers, rather than conditioning on having observed the representation names $\boldsymbol{\sigma}$. Recall that the action of G on $\mathbb{C}[G]$ under (left, say) group multiplication yields the regular representation Reg, and that Reg contains d_σ copies of each $\sigma \in \widehat{G}$. It follows that the fraction of $\mathbb{C}[G]$, dimensionwise, consisting of copies of σ is $d_\sigma^2/|G|$. This fraction is also the probability that we observe the representation name σ in a given register when we perform weak Fourier sampling on the completely mixed state, and is called the *Plancherel distribution* P_{planch}. Similarly, $\mathbb{C}[G^k]$ can be thought of as the regular representation of G^k, in which case it contains $d_{\boldsymbol{\sigma}} = \prod_i d_{\sigma_i}$ copies of each $\boldsymbol{\sigma}$. Thus we have

$$\mathrm{Exp}_{\boldsymbol{\sigma}} \frac{\dim W_{\eta,\boldsymbol{\sigma}}}{d_{\boldsymbol{\sigma}}} = \sum_{\boldsymbol{\sigma}} P_{\mathrm{planch}}(\boldsymbol{\sigma}) \frac{\dim W_{\eta,\boldsymbol{\sigma}}}{d_{\boldsymbol{\sigma}}} = \frac{\sum_{\boldsymbol{\sigma}} d_{\boldsymbol{\sigma}} W_{\eta,\boldsymbol{\sigma}}}{|G|^k} = \frac{\dim W_\eta}{|G|^k} \ .$$

Here

$$W_\eta = \mathrm{span}_{\boldsymbol{\sigma}} W_{\eta,\boldsymbol{\sigma}} = \mathrm{span}_{I\subseteq[k],I\neq\emptyset} W^I_\eta \ ,$$

where $W^I_\eta = \mathrm{span}_{\boldsymbol{\sigma}} W^I_{\eta,\boldsymbol{\sigma}}$ is the subspace of $\mathbb{C}[G^k]$ spanned by vectors which, if we decompose the Hilbert space of the registers in I into G-irreducibles, we observe the representation name η. To put it differently, W^I_η is spanned by vectors which, if we apply the diagonal G-action on the registers in I and leave the other registers fixed, are transformed in a way isomorphic to η, regardless of their G^k-representation $\boldsymbol{\sigma}$.

Now, recall that $\phi \otimes \mathrm{Reg} \cong d_\phi \,\mathrm{Reg}$ for any representation ϕ. In particular, $\mathrm{Reg}^{\otimes \ell}$ contains $|G|^{\ell-1} d_\sigma$ copies of each $\sigma \in \widehat{G}$. Thus if $|I| = \ell$, we have $W^I_\eta = |G|^{\ell-1} d_\eta \eta \otimes \mathrm{Reg}^{\otimes k-\ell}$ and so

$$\frac{\dim W^I_\eta}{|G|^k} = \frac{d_\eta^2}{|G|} \ .$$

We wish to lower bound the fraction of $\mathbb{C}[G^k]$ consisting of $W_\eta = \mathrm{span}_I W^I_\eta$. If the subspaces W^I_η for different I were orthogonal, their dimensions would simply add, giving $\dim W_\eta = (2^k - 1)\dim W^I_\eta$; however, it is easy to see (even for $G = \mathbb{Z}_2$ and $k = 2$) that this is not the case. Instead, it turns out that the subspaces W^I_η have a remarkable statistical property akin to pairwise independence.

Independent subspaces. We say that two subspaces W_1, W_2 are *independent* if the expected squared projection of a random vector $v \in W_1$ into W_2 is just what it would be if v were a random vector in the entire space. (This is a kind of statistical independence between the events that we observe W_1 and W_2, but if their projection operators do not commute we cannot consider these events simultaneously.) Formally:

Definition 1. *Let V be a vector space of dimension D, let $W_1, W_2 \subset V$ be subspaces of dimension d, and let Π_1, Π_2 project onto W_1 and W_2 respectively. Let w be chosen uniformly at random from the vectors in W_1 with norm 1. Then W_1 and W_2 are* independent *if $\mathrm{Exp}_w |\Pi_2 w|^2 = d/D$. A family of subspaces $W_1, \ldots, W_m$ is independent if W_i and W_j are independent for any distinct i, j.*

We start with the following general lemma. The idea is that if we have three representations U, V and W, and consider decomposing $U \otimes V$ and $V \otimes W$ into irreducible subspaces X and Y respectively under

the diagonal action, then the resulting subspaces $X \otimes W$ and $U \otimes Y$ are independent. Indeed, this lemma proves a slightly stronger property; namely, we get independence even if we tensor a *fixed* vector in X with a random vector in W, rather than choosing a random vector from all of $X \otimes W$.

Lemma 3. *Let U, V, and W be (not necessarily irreducible) G-representations and consider their tensor product $U \otimes V \otimes W$. Let $X \subset U \otimes V$ and $Y \subset V \otimes W$ be irreducible subspaces under the diagonal action, and let $\Pi_{U \otimes Y}$ project $U \otimes V \otimes W$ onto $U \otimes Y$. Let $x \in X$ be a fixed vector with $|x|^2 = 1$, and let w be chosen uniformly at random from the vectors in W with norm 1. Then*

$$\operatorname{Exp}_w \left\|\Pi_{U \otimes Y}(x \otimes w)\right\|^2 = \frac{\dim Y}{\dim V \dim W} = \frac{\dim(U \otimes Y)}{\dim(U \otimes V \otimes W)} \quad .$$

In particular, $X \otimes W$ and $U \otimes Y$ are independent subspaces of $U \otimes V \otimes W$.

Proof. In this proof and the next one we use the fact that the definition of independent subspaces is equivalent to one where w is chosen uniformly from an orthonormal basis $\{w_i\}$ for W_1, rather than from all vectors in W_1 with norm 1. This is because for any bilinear form A we have

$$\operatorname{Exp}_{w:|w|^2=1} w^\dagger A w = \frac{1}{\dim W_1} \mathbf{tr}\, A = \operatorname{Exp}_{w_i} w_i^\dagger A w_i$$

and here we take $A = \Pi_1 \Pi_2$.

Let us define $f(x) = \operatorname{Exp}_w \left\|\Pi_{U \otimes Y}(x \otimes w)\right\|^2$. Consider applying the diagonal G-action to U, V, and W. This transports x to gx and fixes the uniform distribution on W; but it also fixes $U \otimes Y$. Therefore we have $f(x) = f(gx)$. However, $f(x)$ is bilinear in x, and it is easy to show using Schur's lemma that any bilinear form defined on an irreducible subspace which is invariant under the G-action is a scalar. Thus $f(x)$ is constant for all x of norm 1.

In particular, let x take the form $u \otimes v$ where $|u|^2 = |v|^2 = 1$, and let Π_Y project $V \otimes W$ onto Y. Then

$$f(x) = \operatorname{Exp}_w \left\|\Pi_{U \otimes Y}(x \otimes w)\right\|^2 = \operatorname{Exp}_w \left\|\Pi_Y(v \otimes w)\right\|^2 \quad .$$

As in the previous paragraph, this is bilinear in v and is invariant under the G-action, and so it is constant for all v of norm 1. Since $\left\|\Pi_Y(v \otimes w)\right\|^2$ is also bilinear in w, we can choose orthonormal bases $\{v_i\}$ and $\{w_j\}$ for V and W respectively, and replace the expectation over w with the expectation over this basis. This gives

$$f(x) = \frac{1}{\dim V \dim W} \sum_{i,j} \left\|\Pi_Y(v_i \otimes w_j)\right\|^2 \quad .$$

But since $\{v_i \otimes w_j\}$ is an orthonormal basis for $V \otimes W$, this sum is $\dim Y$, giving the stated result. (Note that Y does not actually need to be irreducible for this argument to go through.)

Finally, to prove that $X \otimes W$ and $U \otimes Y$ are independent, we consider the expectation of $\left\|\Pi_{U \otimes Y}(w_1)\right\|^2$ with w_1 chosen uniformly from an orthonormal basis for $X \otimes W$ as discussed above. In particular, if this basis is of the form $\{x_i \otimes w_j\}$, this expectation is $\dim Y / (\dim V \dim W) = \dim(U \otimes Y) / \dim(U \otimes V \otimes W)$ for each fixed x_i and we are done. □

With Lemma 3, we prove the following.

Lemma 4. *Let $I, J \subseteq [k]$ be distinct and nonempty. Then W_η^I and W_η^J are independent: that is, if $w \in W_\eta^I$ is uniformly random with norm 1, or, equivalently, chosen uniformly from an orthonormal basis for W_η^I, then $\operatorname{Exp}_w \left\|\Pi_\eta^J w\right\|^2 = d_\eta^2 / |G|$.*

Proof. There are two cases; we address the simpler one first. If I and J are disjoint, consider an orthonormal basis for W_η^I whose basis vectors take the form $\{u_i \otimes v_j\}$ where $v_j \in \mathbb{C}[G^{|J|}]$ describes the registers in J

and $u_i \in \mathbb{C}[G^{k-|J|}]$ describes the others. Let Π_η project the registers in J onto copies of η, and recall that $\Pi_\eta^J = \Pi_\eta \otimes \mathbb{1}^{\otimes(k-|J|)}$. Then for any fixed u_i, taking the expectation over v_j gives an expectation

$$\mathrm{Exp}_{v_j} \left\|\Pi_\eta^J(u_i \otimes v_j)\right\|^2 = \mathrm{Exp}_{v_j} \left\|\Pi_\eta v_j\right\|^2 = d_\eta^2/|G| \ .$$

Now suppose that $I \cap J = K \neq \emptyset$. Without loss of generality, assume that $J \setminus I \neq \emptyset$. Then apply Lemma 3 where W describes the registers in $J \setminus I$, V describes those in $I \cap J$, and U describes the registers in $I \setminus J$ (note that $U = \mathbb{C}$, i.e., the identity of the tensor product, if $I \subset J$). Let X and Y consist of the subspaces of $U \otimes V$ and $V \otimes W$ isomorphic to η. Lemma 3 implies that $X \otimes W$ and $U \otimes Y$ are independent. Finally, note that W_η^I and W_η^J are the tensor products of $X \otimes W$ and $U \otimes Y$, respectively, with $\mathbb{C}[G^{k-|I\cup J|}]$. We choose an orthonormal basis for W_η^I of the form $\{u_i \otimes v_j\}$ where u_i describes the registers in $I \cup J$ and v_j describes the others; then since $X \otimes W$ and $U \otimes Y$ are independent, the expectation over u_i for any fixed v_j of $\Pi_\eta^J(u_i \otimes v_j)$ is

$$\frac{\dim Y}{\dim V \dim W} = \frac{|G|^{|J|-1} d_\eta^2}{|G|^J} = \frac{d_\eta^2}{|G|}$$

since $V \otimes W = \mathrm{Reg}^{\otimes|J|}$ contains $|G|^{|J|-1}$ copies of η. □

Finally, we lower bound the dimension of the span of a independent family of subspaces with the following lemma, and show $\dim W_\eta/|G|^k \geq 1/2$ whenever $k \geq \log_2 |G|$.

Lemma 5. *Let V have dimension D, and let $W_1, \ldots, W_m \subset V$ be a independent family of subspaces of dimension d, and let $W = \mathrm{span}_i W_i$. Then*

$$\frac{\dim W}{D} \geq 1 - \frac{1}{1 + md/(D-d)} \ .$$

Proof. Let Π_i project onto W_i for each $1 \leq i \leq m$, and consider the operator $M = \sum_{i=1}^m \Pi_m$. Since M is positive and symmetric, it can be diagonalized, and so has nonzero eigenvalues $\lambda_1, \ldots, \lambda_t > 0$ where its rank is $t = \dim W$. Its trace is

$$\sum_{\ell=1}^t \lambda_\ell = md \ .$$

Now suppose that $\{e_i\}$ and $\{e_j\}$ are orthonormal bases for W_i and W_j where $i \neq j$. Since W_i and W_j are independent, we have

$$\sum_{e_i, e_j} |\langle e_i \,|\, e_j \rangle|^2 = d^2/D \ .$$

Then the Frobenius norm of M is

$$\begin{aligned} \|M\|^2 = \sum_{\ell=1}^t \lambda_\ell^2 = \mathbf{tr}\, M^\dagger M = \sum_{i,j} \mathbf{tr}\, \Pi_i \Pi_j &= md + \sum_{i \neq j} \sum_{e_i, e_j} \mathbf{tr}\ |e_i\rangle \langle e_i \,|\, e_j\rangle \langle e_j| \\ &= md + \sum_{i \neq j} \sum_{e_i, e_j} |\langle e_i \,|\, e_j \rangle|^2 = md + m(m-1)d^2/D \ . \end{aligned} \tag{2}$$

On the other hand, we have

$$\sum_{\ell=1}^t \lambda_\ell^2 \geq \left(\sum_{\ell=1}^t \lambda_\ell \right) / \, t = (md)^2/t \ .$$

Combining this with (2) and simplifying gives

$$\frac{t}{D} \geq 1 - \frac{D-d}{md + D - d} = 1 - \frac{1}{1 + md/(D-d)} \ .$$

□

Applying this to the independent family $\{W_\eta^I \mid I \subseteq [k], I \neq \emptyset\}$ gives the following corollary.

Corollary 1. *For any $k \geq \log_2 |G|$, we have* $\dim W_\eta / |G|^k \geq \frac{1}{2}$.

Proof. We have $V = \mathbb{C}[G^k]$, $D = |G|^k$, $m = 2^k - 1 \geq |G| - 1$, and $d/D = d_\eta^2/|G| \geq 1/|G|$. Thus $md/(D-d) \geq 1$, and Lemma 5 completes the proof. □

Implementing the measurement: the representation kickback trick. Fixing, for the moment, a subset I of the registers in the experiment above, we now show how to efficiently carry out the von Neumann measurement associated with the subspace $W_{\eta,\boldsymbol{\sigma}}^I$ (that projects onto the subspace or its orthogonal complement). Focusing on the registers in I, it suffices to consider the space $V = \bigotimes_{i \in I} \sigma_i$, decompose $V = \bigoplus_{\tau \in \widehat{G}} a_\tau \tau$ into irreducible representations of G, and implement the measurement associated with the projection operator Π_η that projects onto the space spanned by the a_η copies of η in this direct sum above. Our approach is essentially the same as the "summand extraction" of Kuperberg [20].

To carry out this measurement, we introduce a new G-valued control register, in which we initially prepare $|G\rangle$, the uniform superposition over G. Treating our state now as an element of $\mathbb{C}[G] \otimes V$, we apply the *controlled G-action* operator: $\mathcal{M} : |g\rangle \otimes |\phi\rangle \mapsto |g\rangle \otimes \boldsymbol{\sigma}_I(g^{-1}) |\phi\rangle$, where $\boldsymbol{\sigma}_I(g^{-1})$ is the unitary operator $\bigotimes_{i \in I} \sigma_i(g^{-1})$. Finally, we compute the quantum Fourier transform on the control register and carry out the measurement (on the control register only) corresponding to the operators Π_η and $\mathbb{1} - \Pi_\eta$, where $\Pi_\eta : \mathbb{C}[G] \to \mathbb{C}[G]$ is the operator that projects onto the irreducible subspaces isomorphic to η.

To see why this works, we return our attention to $\mathbb{C}[G] \otimes V$; consider the following two G-actions on this space: let $D_h : |g\rangle \otimes |\phi\rangle \mapsto |hg\rangle \otimes \boldsymbol{\sigma}_I(h) |\phi\rangle$ apply the group action to *both* the control register and V, and let $L_h : |g\rangle \otimes |\phi\rangle \mapsto |hg\rangle \otimes |\phi\rangle$ apply the group action only on the control register. Then observe that $\mathcal{M} \circ D_h = L_h \circ \mathcal{M}$: evidently, any subspace of $\mathbb{C}[G] \otimes V$ that is invariant under D_h is (unitarily) transformed by $\mathcal{M}$ to a subspace that is invariant under L_h. Observe now that if $W \subset V$ is an invariant subspace of V (under $\sigma_I(h)$) that is isomorphic to η, then $|G\rangle \otimes W$ is isomorphic to η under the action D_h, as $|G\rangle$ is invariant under left multiplication by G. Thus such a space is carried to a L_h-invariant space by $\mathcal{M}$, still isomorphic to η.

Of course, while this measurement can be applied efficiently for any fixed subset I, it is unclear how to efficiently apply the measurement corresponding to W_η, which would, in the case of S_n, solve Graph Isomorphism.

Acknowledgments

We are grateful to Dorit Aharonov, Andrew Childs, Tracy Conrad, Gabor Ivanyos, Greg Kuperberg, Sally Milius, Rosemary Moore, Martin Rötteler, Leonard Schulman, Pranab Sen, Douglas Strain and Umesh Vazirani for helpful discussions.

References

[1] David Bacon, Andrew Childs, and Wim van Dam. Optimal measurements for the dihedral hidden subgroup problem. Preprint, quant-ph/0501044 (2005).

[2] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. *Proc. 29th ACM Symposium on Theory of Computing*, pages 48–53, 1997.

[3] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory (preliminary abstract). *Proc. 25th ACM Symposium on Theory of Computing*, pages 11–20, 1993.

[4] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 489–498, 2003.

[5] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. Preprint, quant-ph/9807029 (1998).

[6] Mark Ettinger and Peter Høyer and Emmanuel Knill. Hidden subgroup states are almost orthogonal. Preprint, quant-ph/9901034.

[7] Mark Ettinger and Peter Høyer and Emmanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, to appear.

[8] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. *Proc. 35th ACM Symposium on Theory of Computing*, 2003.

[9] William Fulton and Joe Harris. *Representation Theory: A First Course*. Number 129 in Graduate Texts in Mathematics. Springer-Verlag, 1991.

[10] Michelangelo Grigni, Leonard J. Schulman, Monica Vazirani, and Umesh Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. *Proc. 33rd ACM Symposium on Theory of Computing*, pages 68–74, 2001.

[11] Lisa Hales and Sean Hallgren. Quantum fourier sampling simplified. *Proc. 31st ACM Symposium on Theory of Computing*, 1999.

[12] Lisa Hales and Sean Hallgren. An improved quantum Fourier transform algorithm and applications. *Proc. 41st Symposium on Foundations of Computer Science*, 2000.

[13] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. *Proc. 32nd ACM Symposium on Theory of Computing*, pages 627–635, 2000.

[14] Peter Høyer. Efficient quantum transforms. Preprint, quant-ph/9702028 (1997).

[15] Yoshifumi Inui and François Le Gall. An efficient algorithm for the hidden subgroup problem over a class of semi-direct product groups. Proc. EQIS 2004.

[16] Lawrence Ip. Shor's algorithm is optimal. Preprint, 2004.

[17] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. *Int. J. Found. Comput. Sci.* 14(5): 723–740, 2003.

[18] Richard Jozsa. Quantum factoring, discrete logarithms and the hidden subgroup problem. Preprint, quant-ph/0012084 (2000).

[19] J. Kempe and A. Shalev, The hidden subgroup problem and permutation group theory. Preprint, quant-ph/0406046 (2004).

[20] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. Preprint, quant-ph/0302112 (2003).

[21] C. Moore and A. Russell. For distinguishing conjugate hidden subgroups, the pretty good measurement is as good as it gets. Preprint, quant-ph/0501177

[22] Cristopher Moore, Alexander Russell, and Leonard Schulman. The symmetric group defies strong Fourier sampling: Part I. Preprint, quant-ph-0501056.

[23] Cristopher Moore and Alexander Russell. The symmetric group defies strong Fourier sampling: Part II. Preprint, quant-ph-0501066.

[24] Cristopher Moore, Daniel Rockmore, and Alexander Russell. Generic quantum Fourier transforms. *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms*, pages 778–787, 2004.

[25] C. Moore, D. Rockmore, A. Russell, and L. Schulman, The value of basis selection in Fourier sampling: hidden subgroup problems for affine groups. *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms*, pages 1113–1122, 2004.

[26] Jaikumar Radhakrishnan, Martin Rötteler, and Pranab Sen. On the Power of Random Bases in Fourier Sampling: Hidden Subgroup Problem in the Heisenberg Groups. *Proc. ICALP*, 2005.

[27] O. Regev, Quantum computation and lattice problems. *Proc. 43rd Symposium on Foundations of Computer Science*, pages 520–530, 2002.

[28] Martin Rötteler and Thomas Beth. Polynomial-time solution to the hidden subgroup problem for a class of non-abelian groups. Preprint, quant-ph/9812070 (1998).

[29] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Number 42 in Graduate Texts in Mathematics. Springer-Verlag, 1977.

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

[31] Daniel R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, 1997.