Normal Subgroup Reconstruction and Quantum Computation Using Group Representations

Sean Hallgren[*]
Computer Science Division
University of California
Berkeley, CA 94720
hallgren@cs.berkeley.edu

Alexander Russell[†]
Department of Computer
Science and Engineering
University of Connecticut
Storrs, CT 06269
acr@cse.uconn.edu

Amnon Ta-Shma
Computer Science Division
University of California
Berkeley, CA 94720
amnon@cs.berkeley.edu

ABSTRACT

The Hidden Subgroup Problem is the foundation of many quantum algorithms. An efficient solution is known for the problem over Abelian groups and this was used in Simon's algorithm and Shor's Factoring and Discrete Log algorithms. The non-Abelian case is open; an efficient solution would give rise to an efficient quantum algorithm for Graph Isomorphism. We fully analyze a natural generalization of the Abelian case solution to the non-Abelian case, and give an efficient solution to the problem for normal subgroups. We show, however, that this immediate generalization of the Abelian algorithm does not efficiently solve Graph Isomorphism.

1. INTRODUCTION

Peter Shor's seminal article [16] presented efficient quantum algorithms for computing integer factorizations and discrete logarithms, problems thought to be intractable for classical computation models. A primary ingredient in these algorithms is a solution to the *hidden subgroup problem* for certain Abelian groups; indeed Discrete-Log directly reduces to the hidden subgroup problem. Formally, the hidden subgroup problem is the following:

Definition 1. Hidden Subgroup Problem. (HSP) *Given an efficiently computable function $f : G \to S$, from a finite group G to a set S, that is constant on (left) cosets of some subgroup H and takes distinct values on distinct cosets, find a set of generators for H.*

The general paradigm is the following:

[*]Supported by an NDSEG Fellowship, a GAANN Fellowship, and NSF Grant CCR9800024

[†]Supported by NSF NYI Grant CCR-9457799 and a David and Lucile Packard Fellowship for Science

To appear in the Thirty-Second Annual ACM Symposium on Theory of Computing (STOC), May 2000, Portland, Oregon.

Algorithm 1. *(Algorithm for the Abelian HSP).*

1. *Compute $\sum_{g\in G} |g, f(g)\rangle$ and measure the second register $f(g)$. The resulting super-position is $\sum_{h\in H} |ch\rangle \otimes |f(ch)\rangle$ for some coset cH of H. Furthermore, the distribution of c over G is uniform.*
2. *Compute the Fourier transform of the coset state, resulting in*
$$\sum_{\rho\in\hat{G}} \sqrt{\frac{1}{|G|}}\sqrt{\frac{1}{|H|}}\sum_{h\in H}\rho(ch)|\rho\rangle,$$
where $\hat{G}$ denotes the set of homomorphisms $\{\rho : G \to \mathbb{C}\}$.
3. *Measure the first register and observe a representation ρ.*

A key fact about this procedure is that the resulting distribution over ρ does not depend on which coset cH arises after the first stage. Thus, we can repeat the same experiment many times, each time inducing the same distribution over ρ.

It is well known that an efficient solution to the HSP for the symmetric group S_n gives, in particular, an efficient quantum algorithm for Graph Isomorphism. It is also known how to efficiently compute the Fourier transform over many non-Abelian groups, most notably over S_n [2]. Nevertheless, until this work, there was no general understanding of the HSP over non-Abelian groups. In this paper we study the generalization of Algorithm 1 to non-Abelian groups. Namely, we study the following algorithm:

Algorithm 2. *(Potential Algorithm for the General HSP).*

1. *Compute $\sum_{g\in G} |g, f(g)\rangle$ and measure the second register $f(g)$. The resulting super-position is $\sum_{h\in H} |ch\rangle \otimes |f(ch)\rangle$ for some coset cH of H. Furthermore, c is uniformly distributed over G.*
2. *Compute the Fourier transform of the coset state, which is*
$$\sum_{\rho\in\hat{G}} \sqrt{\frac{d_\rho}{|G|}}\sqrt{\frac{1}{|H|}}\left(\sum_{h\in H}\rho(ch)\right)_{i,j}|\rho, i, j\rangle,$$

where $\hat{G}$ denotes the set of irreducible representations of G.

3. *Measure the first register and observe a representation ρ.*

For more details about the Fourier transform over non-Abelian groups, see Section 2.

As before, we wish the resulting distribution to be *independent* of the actual coset cH (and so depend only on the subgroup H). This is guaranteed by measuring only the name of the representation ρ, leaving the matrix indices (the values i and j) unobserved. The question we study is whether this procedure retains enough information to determine H, or, more precisely, whether $O(\log(|G|))$ samples of this distribution are enough to determine H with high probability. Our analysis of Algorithm 2 depends on the following theorem, which we believe is interesting on its own and is one of the main technical contributions of the paper:

THEOREM 1. *The probability of measuring the representation ρ in algorithm 2 is proportional to the dimension of ρ and number of times ρ appears in the induced representation $\mathrm{Ind}_H^G \mathbf{1}_H$, where $\mathbf{1}_H$ denotes the trivial representation on H.*

We apply this to solve the HSP for *normal* subgroups:

THEOREM 2. *(A solution to the Normal HSP). Let H be a normal subgroup of G. With high probability, H is uniquely determined by observing $m = O(\log |G|)$ independent trials of Algorithm 2.*

Our reconstruction result is information theoretic, and applies to any normal subgroup H of any group G without reference to the specific way that the representations ρ are expressed. We proceed at this level of abstraction because there is no known canonical presentation for the representation of a finite group G. In the same vein, there is no general method for computing the Fourier transform over an arbitrary group.

A corollary of Theorem 1 is that conjugate subgroups H_1 and H_2 (where $H_2 = gH_1g^{-1}$ for some $g \in G$) produce exactly the same distribution over ρ and hence cannot be distinguished by this process. In particular, the hidden subgroup problem cannot be solved by Algorithm 2 for a group G with two distinct conjugate subgroups H_1, H_2; the symmetric group S_n is such a group.

In light of this, one may ask whether Algorithm 2 can distinguish between a coset cH of a non-trivial subgroup H and a coset $cH_e = \{c\}$ of the trivial subgroup $H_e = \{e\}$, as even this would be enough for solving Graph Isomorphism. However, even for this weaker problem we show:

THEOREM 3. *For the symmetric group S_n, Algorithm 2 does not distinguish (even information theoretically) the case that the hidden subgroup is the trivial subgroup from the case that the hidden subgroup is non-trivial. (Specifically, the distributions induced on ρ in these two cases have exponentially small total variation distance.)*

1.1 Related Work.

Simon's algorithm [17] implicitly involves distinguishing the trivial subgroup from an order 2 subgroup over the group $\mathbb{Z}_2^n$. He shows that a classical probabilistic oracle machine would require exponentially many oracle queries to successfully distinguish the two cases with probability greater than 1/2. Shor [16] generalizes Simon's algorithm to solve integer factorization and the discrete log problem. In addition to solving a special case of the HSP, he also solves specific cases when the underlying group is not even known. Boneh and Lipton [3] handle a case when a periodic function is not fixed on a coset. Hales and Hallgren [11] generalize the results for the case when the underlying Abelian group is unknown, but an estimate is known for the cardinality of its cyclic factors. Kitaev [13] gave an algorithm for the Abelian stabilizer problem, which is a special case of the HSP. The efficient algorithm for general Abelian HSP solution is folklore.

As for computing the Fourier transform efficiently, Kitaev also shows how to efficiently compute the Fourier transform over any Abelian group. Beals [2] showed how to efficiently compute the Fourier transform over the symmetric group S_n.

Ettinger, Høyer and Knill [8] show that the HSP has polynomial query complexity. Ettinger and Høyer [5] give a solution for the HSP over the dihedral group D_n with polynomially many queries and exponential time. In [6] and [7] they address whether any measurement will distinguish subgroup states. Roetteler and Beth [15] give a solution to the HSP for a specific non-Abelian group.

Grigni, Schulman and Vazirani [10] independently showed that measuring the representation and the row of the matrix entry is not enough to solve graph isomorphism.

2. REPRESENTATION THEORY BACKGROUND

The main tool used by polynomial time quantum algorithms is the Fourier transform. To define the Fourier transform (over a group) we require the basic elements of representation theory, defined below.

Representation. A representation ρ of a group G is a homomorphism $\rho : G \to GL(V)$, where V is vector space over $\mathbb{C}$. Fixing a basis for V, each $\rho(g)$ may be realized as a $d \times d$ matrix over $\mathbb{C}$, where d is the dimension of V. As ρ is a homomorphism, for any $g, h \in G$, $\rho(gh) = \rho(g)\rho(h)$. The *dimension* d_ρ of the representation ρ is d, the dimension of V.

A representation provides a means for investigating a group by homomorphically mapping it into a family of matrices. With this realization, the group operation is matrix multiplication and tools from linear algebra can be applied to study the group. We shall be concerned with complex-valued functions on a group G; the representations of the group are relevant to this study, as they give rise to the Fourier transform for such functions.

Irreducibility. We say that a subspace W is an *invariant* subspace of a representation ρ if $\rho(g)W \subseteq W$ for all $g \in G$. We assume, without loss of generality, that every $\rho(g)$ is unitary, and, in particular, diagonalizable. Hence there are many subspaces fixed by an individual matrix $\rho(g)$. In order for W to be an invariant subspace for ρ, it must be simultaneously fixed under all $\rho(g)$.

The zero subspace and the subspace V are always invariant. If no nonzero proper subspaces are invariant, the representation is said to be *irreducible.*

Decomposition. When a representation *does* have a nonzero proper invariant subspace $V_1 \subset V$, it is always possible to find a complementary subspace V_2 (so that $V = V_1 \oplus V_2$) which is also invariant. Since $\rho(g)$ fixes V_1, we may let $\rho_1(g)$ be the linear map on V_1 given by $\rho(g)$. It is not hard to see that $\rho_1 : G \to \mathrm{GL}(V_1)$ is in fact a representation. Similarly, define $\rho_2(g)$ to be $\rho(g)$ restricted to V_2. Since $V = V_1 \oplus V_2$, the linear map $\rho(g)$ is completely determined by $\rho_1(g)$ and $\rho_2(g)$, and in this case we write $\rho = \rho_1 \oplus \rho_2$. In this case there is a basis for V so that each matrix $\rho(g)$ is block diagonal with two blocks.

Complete Reducibility. Repeating the process described above, any representation ρ may be written $\rho = \rho_1 \oplus \rho_2 \oplus \ldots \oplus \rho_k$, where each ρ_i is irreducible. In particular, there is a basis in which every matrix $\rho(g)$ is block diagonal, the ith block corresponding to the ith representation in the decomposition.

Characters. The *character* $\chi_\rho : G \to \mathbb{C}$ of a representation ρ is defined by $\chi_\rho(g) = \mathbf{tr}\,(\rho(g))$. It is basis independent, and, as it turns out, completely determines the representation ρ.

Orthogonality of Characters. For two functions f_1 and f_2 on a group, there is a natural inner product: $\langle f_1, f_2\rangle_G$ given by $\frac{1}{|G|}\sum_g f_1(g) f_2(g)^*$. The useful fact is the following: given the character χ_ρ of any representation ρ and the character χ_i of any irreducible representation ρ_i, the inner product $\langle \chi_\rho \,|\, \chi_i\rangle$ is precisely the number of times the representation ρ_i appears in the decomposition of ρ. Since each ρ is unitary, the inner product of two characters simplifies slightly:

$$\langle \chi_\rho \,|\, \chi_i\rangle_G = \frac{1}{|G|}\sum_{g\in G} \chi_\rho(g)\chi_i(g^{-1}).$$

Restriction. 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 $\mathrm{Res}_H\,\rho$. Note that even representations which are irreducible over G may be reducible when restricted to H.

Up to isomorphism, a finite group has a finite number of irreducible representations. For a group G, we let $\hat{G}$ denote this finite collection of irreducible representations. As mentioned above, any representation may be decomposed into a sum of representations in $\hat{G}$.

EXAMPLE 1. *Fix a group G and a representation ρ. Let $\rho_1, \ldots, \rho_k$ be the irreducible representations of G. Desiring to know how ρ decomposes in these ρ_i, we compute*

$$n_i = \langle \chi_\rho, \chi_i\rangle$$

for each $i = 1, \ldots, k$. Then $\rho = n_1\rho_1 \oplus \cdots \oplus n_k\rho_k$, and, after a unitary change of basis, the diagonal of the matrix $\rho(g)$ consists of n_1 copies of $\rho_1(g)$, followed by n_2 copies of $\rho_2(g)$, etc.

There are two representations possessed by every group:

The Trivial Representation. The trivial representation $\mathbf{1}_G$ maps every group element $g \in G$ to the 1 by 1 matrix (1). One feature of the trivial representation is that $\sum_g \mathbf{1}_G(g)$ is the 1×1 matrix $(|G|)$; this sum is the zero matrix for any other irreducible representation.

The Regular Representation. We take a vector space V with a basis vector e_g for every element $g \in G$. The regular representation $\mathrm{reg}_G : G \to \mathrm{GL}(V)$ is defined by $\mathrm{reg}_G(g) : e_x \mapsto e_{gx}$, for any $x \in G$. It has dimension $|G|$. With the basis above, for any $g \in G$, $\mathrm{reg}_G(g)$ is a permutation matrix.

An important fact about the regular representation is that it contains every irreducible representation of G. In particular, if $\rho_1, \ldots, \rho_k$ are the irreducible representations of G with dimensions $d_1, \ldots, d_k$, then

$$\rho_{\mathrm{reg}} = d_1\rho_1 \oplus \cdots \oplus d_k\rho_k,$$

that is, the regular representation contains each irreducible representation ρ_i exactly d_i times. Counting dimensions,

$$|G| = \sum_i d_i^2. \tag{1}$$

The main tool in quantum polynomial time algorithms is the Fourier transform.

DEFINITION 2. *Let $f : G \to \mathbb{C}$. The* Fourier transform *of f at the irreducible representation ρ is the $d_\rho \times d_\rho$ matrix*

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

We refer to the collection of matrices $\langle \hat{f}(\rho)\rangle_{\rho\in\hat{G}}$ as the *Fourier transform* of f. Thus f is mapped into $|\hat{G}|$ matrices of different dimensions. The total number of entries in these matrices is $\sum d_\rho^2 = |G|$, by equation 1 above. The Fourier transform is linear in f; with the constants used above (i.e. $\sqrt{d_\rho/|G|}$) it is in fact unitary, taking the $|G|$ complex numbers $\langle f(g)\rangle_{g\in G}$ to $|G|$ complex numbers organized into matrices.

A familiar case in computer science is when the group is cyclic of order n. Then the linear transformation (i.e., the Fourier transform) is a Vandermonde matrix with n-th roots of unity and the matrices are 1-by-1.

In the quantum setting we identify the superposition $\sum_{g\in G} f_g|g\rangle$ with the function $f : G \to \mathbb{C}$ defined by $f(g) = f_g$. In this notation, $\sum_{g\in G} f(g)|g\rangle$ is mapped under the Fourier transform to $\sum_{\rho\in\hat{G}, 1\le i,j\le d_\rho} \hat{f}(\rho)_{i,j}|\rho, i, j\rangle$. We remind the reader that $\hat{f}(\rho)_{i,j}$ is a complex number. When the first portion of this triple is measured, we observe $\rho \in \hat{G}$ with probability

$$\sum_{1\le i,j\le d_\rho} \left|\hat{f}(\rho)_{i,j}\right|^2 = \left\|\hat{f}(\rho)\right\|^2$$

where $\|A\|$ is the natural norm given by $\|A\|^2 = \mathbf{tr}\; A^*A$.

Let f be the indicator function of a left coset of H in G, i.e. for some $c \in G$,

$$f(g) = \begin{cases} \frac{1}{\sqrt{|H|}} & \text{if } g \in cH, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Our goal is to understand the Fourier transform of f. As mentioned above, the probability of observing ρ is $\left\|\hat{f}(\rho)\right\|^2 = \sum_{i,j} |(\hat{f}(\rho))_{i,j}|^2$. Our choice to measure only the representation ρ (and not the matrix indices) depends on the following key fact about the Fourier transform, also relevant to the Abelian solution:

CLAIM 1. *The probability of observing ρ is independent of the coset.*

PROOF. $\hat{f}(\rho) = \sum_{h\in H} \rho(ch) = \rho(c)\sum_{h\in H}\rho(h)$ and, since $\rho(c)$ is a unitary matrix,

$$\left\|\hat{f}(\rho)\right\|^2 = \left\|\rho(c)\sum_{h\in H}\rho(h)\right\|^2 = \left\|\sum_{h\in H}\rho(h)\right\|^2.$$

□

Given this, we may assume that our function f is positive on the subgroup H itself, and zero elsewhere.

3. THE PROBABILITY OF MEASURING ρ

The primary question is that of the the probability of observing ρ. We have seen that this is determined by $\sum_{h\in H}\rho(h)$ which, being a sum of the linear transformations $\rho(h)$, is a linear transformation. We being by showing that it is a projection:

LEMMA 1. *$\hat{f}(\rho)$ is a projection.*

With the right basis, then, $\hat{f}(\rho)$ will be diagonal, and the diagonal entries will consists of ones and zeros. The probability of observing a particular representation ρ will then correspond to the number of ones appearing on the diagonal (i.e., on the dimension of the image of $\hat{f}(\rho)$).

Given an irreducible representation ρ of G, we are interested in the sum of the matrices $\rho(h)$ for all $h \in H$. Since we only evaluate ρ on H, we can instead consider $\mathrm{Res}_H\rho$ without changing anything. As mentioned before, though ρ is irreducible (over G), $\mathrm{Res}_H\rho$ may not be irreducible on H. We may, however, decompose $\mathrm{Res}_H\rho$ into irreducible representations over H. Then the Fourier transform of f at ρ is comprised of blocks, each corresponding to a representation in the decomposition of $\mathrm{Res}_H\rho$. In particular, the matrix $\sum_{h\in H}\rho(h)$ is:

$$U\begin{bmatrix} \sum_{h\in H}\sigma_1(h) & 0 & \cdots & 0 \\ 0 & \sum_{h\in H}\sigma_2(h) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sum_{h\in H}\sigma_t(h)\end{bmatrix}U^\dagger \tag{2}$$

for some unitary transformation U and some irreducible representations σ_i of H (with possible repetitions). We know that the sum is nonzero only when the irreducible representation is trivial, in which case it is $|H|$. Then the probability of observing ρ is

$$\begin{aligned}\left\|(\hat{f}(\rho))\right\|^2 = \left\|\sqrt{\frac{d_\rho}{|G|}}\sum_{h\in H}\rho(h)\right\|^2 &= \frac{d_\rho}{|G|}\frac{1}{|H|}|H|^2\langle\chi_\rho,\chi_{\mathbf{1}_H}\rangle_H \\ &= \frac{|H|}{|G|}d_\rho\langle\chi_\rho,\chi_{\mathbf{1}_H}\rangle_H.\end{aligned}$$

We have proved:

THEOREM 4. $\left\|\hat{f}(\rho)\right\|^2 = \frac{|H|}{|G|}d_\rho\langle\chi_\rho,\chi_{\mathbf{1}_H}\rangle_H$.

Observe that one consequence of the theorem is that the probability of observing a representation ρ depends only on the character of ρ. It turns out that characters are class functions, i.e., $\chi(g) = \chi(hgh^{-1})$ for any character χ and $h, g \in G$. Hence conjugate subgroups (gHg^{-1} for some $g \in G$ is a conjugate subgroup of H) produce exactly the same distribution; this rules out using the paradigm of Algorithm 2 with representations names alone to solve the HSP for any group containing a non-normal subgroup.

3.1 Induced Representations

We have discussed the restriction of a representation ρ on G to a subgroup H of G. There is a dual operation, called *induction.* This extends to all of G a representation ρ defined on a subgroup H. We will only need to work with the representation induced from the trivial representation on H. Let $G/H \stackrel{\text{def}}{=} \{\alpha_1,\ldots,\alpha_t\}$ be a collection of representatives for the left cosets of H in G, so that $G = \alpha_1 H \cup\cdots\cup \alpha_t H$, this union being disjoint. Then the *induced* representation $\mathrm{Ind}_H^G\mathbf{1}_H : G \to \mathrm{GL}(W)$ is defined over the vector space W that has one basis vector $e_{[\alpha_i]}$ for each coset $\alpha_i H$. It is defined by linearly extending the rule

$$\mathrm{Ind}_H^G\mathbf{1}_H(g) : e_{[\alpha_i]} \to e_{[g\alpha_i]} = e_{[\alpha_j]}$$

where $g\alpha_i$ belongs to the coset $\alpha_j H$. Observe that this representation is a permutation representation. As suggested by the notation, the representation is independent of the choice of α_i.

EXAMPLE 2. $\mathrm{Ind}_{\{\mathbf{id}\}}^G\mathbf{1}_{\{\mathbf{id}\}} \cong \mathrm{reg}_G$.

LEMMA 2. *(A special case of* Frobenius reciprocity*) [12, §3.20]): Let $H \subset G$ and let $\rho : G \to GL(V)$ be an irreducible representation of G. Then*

$$\langle\chi_{\mathbf{1}_H} \,|\, \chi_\rho\rangle_H = \left\langle\chi_{\mathrm{Ind}_H^G\mathbf{1}_H}\,\middle|\,\chi_\rho\right\rangle_G,$$

the first inner product being computed over H and the second over G.

Hence the number of times that the trivial representation of H appears in $\mathrm{Res}_H\rho$ is the same as the number of times that ρ appears in $\mathrm{Ind}_H^G\mathbf{1}_H$. Theorem 4 shows that the probability of measuring the representation ρ in algorithm 2 is $\frac{|H|}{|G|}d_\rho\langle\chi_{\mathbf{1}}\,|\,\chi_\rho\rangle_H$. By Frobenius reciprocity, $\langle\chi_{\mathbf{1}}\,|\,\chi_\rho\rangle_H = \left\langle\chi_{\mathrm{Ind}_H^G\mathbf{1}_H}\,\middle|\,\chi_\rho\right\rangle_G$, proving Theorem 1.

4. A POSITIVE RESULT: NORMAL SUBGROUPS

We begin this section by showing that a polynomial number of samples suffices to reconstruct normal subgroups. The proof uses a Chernoff bound and Lemma 3, which we will prove in the next section. The lemma is the interesting part, and sections §4.1 and §4.2 are devoted to two proofs of the lemma, one from the perspective of restriction, the other from the perspective of induction.

For a normal subgroup H of G, let $\mathcal{D}_H$ denote the distribution induced on $\hat{G}$ by the quantum experiment of Algorithm 2. Lemma 3 below states that under this distribution $\sigma \in \hat{G}$ is sampled with probability $\frac{|H|}{|G|}d_\sigma^2$ if $H \subseteq \ker\sigma$, and zero otherwise[1]. We will see also that $\sum_{\substack{\sigma\in\hat{G},\\ H\subseteq\ker\sigma}} d_\sigma^2 = |G/H|$. Now, we are ready to prove:

THEOREM 5. *Let $\sigma_1,\ldots,\sigma_s$ be independent random variables distributed according to $\mathcal{D}_H$ with $s = c\log_2|G|$. Then*

$$\Pr[H \neq \bigcap_i \ker\sigma_i] \leq e^{-\frac{(c-2)^2}{2c}\log_2|G|}.$$

PROOF. Let $N_i = \bigcap_{j=1}^{i}\ker\sigma_j$ and $N_0 = G$. As an intersection of normal subgroups, each N_i is normal in G. Also, we know that $H \subseteq \ker\sigma_i$, for any i. Hence, $H \subseteq N_s \subset N_{s-1} \subseteq \cdots \subseteq N_0 = G$. We claim:

CLAIM 2. *If $N_i \neq H$ then $\Pr_{\sigma_{i+1}\in\mathcal{D}_H}(N_{i+1} = N_i) \leq \frac{1}{2}$*

PROOF. $N_{i+1} = N_i$ iff $N_i \subseteq \ker\sigma_{i+1}$. Now, we measure $\sigma\in\hat{G}$ with $H\subseteq\ker\sigma$ with probability $\frac{|H|}{|G|}d_\sigma^2$, by Lemma 3. Hence,

$$\begin{aligned}\Pr_{\sigma\in\mathcal{D}_H}(N_i\subseteq\ker\sigma) &= \sum_{\rho\in\hat{G}:N_i\subseteq\ker\rho}\frac{|H|}{|G|}d_\rho^2\\ &= \frac{|H|}{|G|}\cdot\frac{|G|}{|N_i|} = \frac{|H|}{|N_i|} \leq \frac{1}{2}\end{aligned}$$

□

For each $i = 1,\ldots,s$, let X_i be the indicator random variable taking value 1 if $N_i = H$ or $N_i \neq N_{i-1}$ and zero otherwise. The random variables $X_1,\ldots,X_s$ are not necessarily independent, but by the previous claim, $\Pr[X_i = 0 \mid \rho_1,\ldots,\rho_{i-1}] \leq \frac{1}{2}$. We can therefore define new independent random variables $Y_1,\ldots,Y_s$ with $\Pr(Y_i = 0) = \frac{1}{2}$ and such that $\sum Y_i \leq \sum X_i$ is always true. We now apply the Chernoff bound.

Chernoff bound [1, A.1] Let Y_i, $i = 1,\ldots,s$ be a collection of independent random variables, uniformly distributed in $\{0,1\}$. Then $\Pr\left[\sum_i Y_i < \frac{s-a}{2}\right] < e^{\frac{-a^2}{2s}}$.

Hence, $\Pr(\Sigma_{i=1}^s X_i \leq \log|G|) \leq e^{-\frac{(c-2)^2}{2c}\log_2|G|}$. However, whenever $\Sigma_{i=1}^s X_i \geq \log(|G|)$ we must have $N_s = H$, because each time $N_{i+1}\subsetneq N_i$ the size of N_{i+1} is reduced by at least half, and we can repeat this no more than $\log(|G|)$ times. Hence, except for probability at most $e^{-\frac{(c-2)^2}{2c}\log_2|G|}$, we have $N_s = H$, which completes the proof.

We now give two proofs for the facts stated at the beginning of this section.

4.1 Proof 1: Restricted representations

With this view, the main question, given the theorem of the last section, is whether or not the number of ones on the diagonal behaves in a useful way. We show that for normal subgroups, the Fourier transform at an irreducible representation is either (a multiple of) the identity or the zero matrix.

[1]We remind the reader that a representation σ is a homomorphism $\sigma : G \to GL(V)$. The *kernel* of ρ is the set $\ker(\sigma) = \{g\in G|\sigma(g) = \mathbf{1}_V\}$ and is a normal subgroup of G, which we write $\ker\sigma \trianglelefteq G$.

LEMMA 3. *Let $H \trianglelefteq G$. If $H\subseteq\ker\rho$ then ρ is observed with probability $\frac{|H|}{|G|}d_\rho^2$. If $H\not\subseteq\ker\rho$ then ρ is observed with probability 0.*

PROOF. If $H\subseteq\ker\rho$ then $\rho(h)$ is the identity for every $h\in H$, so $\mathrm{Res}_H\rho$ must contain d_ρ copies of the trivial representation of H, and by theorem 1 the probability is $\frac{|H|}{|G|}d_\rho^2$. Now we will use a simple counting argument to show that no other representations can contain a trivial representation of H. First we compute how many times the trivial representation of H appears in the regular representation of G restricted to H:

$$\begin{aligned}\langle\chi_{\mathrm{reg}_G},\chi_{\mathbf{1}_H}\rangle_H &= \frac{1}{|H|}\sum_{h\in H}\chi_{\mathrm{reg}_G}(h)\cdot\chi_{\mathbf{1}_H}(h^{-1})\\ &= \frac{1}{|H|}\chi_{\mathrm{reg}_G}(e) = \frac{|G|}{|H|}.\end{aligned}$$

On the other hand, let $d_1,\ldots,d_l$ be the dimensions of those representations with H in their kernel. Each occurs in the restriction of the regular representation of G dimension many times, so we get a total of $\sum d_i^2$ copies of the trivial representation of H that we already counted. Using Equation 1 from Section 2 $\sum d_i^2 = \frac{|G|}{|H|}$, so there can be no more.
The final fact follows from group theory. Let groups G, G_1 and $H\trianglelefteq G$ be given. The number of homomorphisms of G that map H to the identity in G_1 is the same as the number of homomorphisms of the quotient group G/H to G_1.

□

4.2 Proof 2: Induced representations

We now provide an alternative proof for Lemma 3. Lemma 4 restates the result in Lemma 3 using induced representations.

LEMMA 4. *Let $H \trianglelefteq G$. Then*

$$\mathrm{Ind}_H^G\mathbf{1}_H = \bigoplus_{\sigma\in\hat{G}:H\subseteq\ker\sigma} d_\sigma\sigma.$$

PROOF. We may think of $\tau = \mathrm{Ind}_H^G\mathbf{1}_H$ as the natural permutation representation on the cosets of G/H (see Section 3.1). When $H\trianglelefteq G$ the coset $h\cdot gH$ is just the coset gH. Hence, for any $h\in H$, $\tau(h) = \mathbf{1}_V$ and $H\subseteq\ker\mathrm{Ind}_H^G\mathbf{1}_H$. In particular, if we express $\tau = \bigoplus_{\rho\in\hat{G}} m_\rho\rho$, then for any ρ with $m_\rho > 0$ it must be that $H\subseteq\ker\rho$.
On the other hand, let $\rho\in\hat{G}$ with $H\subseteq\ker\rho$. Then,

$$\begin{aligned}\langle\chi_\tau\,|\,\chi_\rho\rangle_G &= \langle\chi_{\mathbf{1}}\,|\,\chi_\rho\rangle_H = \frac{1}{|H|}\Sigma_{h\in H}\chi_\rho(h)\\ &= \chi_\rho(e) = d_\rho\end{aligned}$$

where the first equality is by Frobenius reciprocity, and the third is because $H\subseteq\ker\rho$. We conclude that $\mathrm{Ind}_H^G\mathbf{1}_H = \bigoplus_{\sigma\in\hat{G}:H\subseteq\ker\sigma} d_\sigma\sigma$, as desired. □

The dimension of $\mathrm{Ind}_H^G\mathbf{1}_H$ is $\frac{|G|}{|H|}$, which yields the following corollary:

COROLLARY 1. *Let $H\trianglelefteq G$.*

$$\Sigma_{\substack{\sigma\in\hat{G},\\ H\subseteq\ker\sigma}} d_\sigma^2 = |G/H|.$$

5. A NEGATIVE RESULT: DETERMINING TRIVIALITY IN S_N

In this section we show that a well known reduction of graph isomorphism to finding a hidden subgroup over S_n will not work using Algorithm 2. Some representation theory of S_n is needed and is in the appendix.
Graph Automorphism is the problem of determining if a graph G has a nontrivial automorphism, and is easier than Graph Isomorphism [14]. A natural special case occurs when the graph G consists of two disjoint connected rigid graphs G_1, G_2 (i.e., $\mathrm{Aut}(G_1) = \mathrm{Aut}(G_2) = \{e\}$). In this case there are two possibilities for the automorphism group of G:

CLAIM 3.

- *If* $G_1 \not\approx G_2$, *then* $\mathrm{Aut}(G) = \{e\}$.
- *If* $G_1 \approx G_2$ *then* $\mathrm{Aut}(G) = \{e, \sigma\}$ *where* $\sigma \in S_n$ *is a permutation with* $n/2$ *disjoint* 2*–cycles.*

PROOF. For the first part notice that any automorphism maps a connected component onto a connect component. In our case we have two connected components G_1 and G_2. However, G_1 and G_2 are not isomorphic and have no non-trivial automorphisms.
For the second part, let σ reflect an automorphism between G_1 and G_2. Now, suppose there was another non-trivial automorphism τ. Then $\sigma\tau$ is also an automorphism, and $\sigma\tau$ maps the connected component of G_1 onto G_1, and G_2 onto G_2. As G_1 and G_2 have no non-trivial automorphisms it follows that $\sigma\tau = 1$, $\tau = \sigma^{-1} = \sigma$. $\square$

Thus, if one knows how to solve the HSP for S_n, or if one knows how to distinguish between cosets of a trivial subgroup and the cosets of a non-trivial subgroup, one can give an efficient quantum algorithm for Graph Automorphism. In particular, one might try the following algorithm for reconstructing $H = \mathrm{Aut}(G)$.

ALGORITHM 3. **Input** *a graph* G *s.t. either* $\mathrm{Aut}(G) = \{e\}$ *or* $\mathrm{Aut}(G) = \{e, \sigma\}$.

1. *Compute* $\Sigma_{\pi \in S_n} |\pi, \pi(G)\rangle$ *and measure the second register* $\pi(G)$. *The resulting super-position is* $\Sigma_{h \in H} |ch\rangle \otimes |f(ch)\rangle$ *for some coset* cH *of* H. *Furthermore,* c *is uniformly distributed over* G.
2. *Compute the Fourier transform of the coset state, which is*
$$\sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G|}} \sqrt{\frac{1}{|H|}} \left(\sum_{h \in H} \rho(ch) \right)_{i,j} |\rho, i, j\rangle.$$
3. *Measure the first register and observe a representation* ρ.

We show that even for this particular case of Graph Isomorphism (and Graph Automorphism) the algorithm fails.

THEOREM 6. *Let* p_ρ *be the probability of sampling* ρ *in Algorithm 3 when* $G_1 \not\approx G_2$, *and* q_ρ *when* $G_1 \approx G_2$ *and* G_1, G_2 *are connected and rigid. Then* $|p - q|_1 \leq 2^{-\Omega(n)}$.

PROOF. When $G_1 \not\approx G_2$, $H = \mathrm{Aut}(G) = \{e\}$. Then $\mathrm{Ind}_H^{S_n} 1$ is the regular representation, and so $\left\langle \chi_\rho \,\middle|\, \chi_{\mathrm{Ind}_A^{S_n}} 1 \right\rangle$, the multiplicity of ρ in the regular representation, is d_ρ. Hence, $p_\rho = \frac{d_\rho^2}{n!}$.
When $G_1 \approx G_2$ and G_1, G_2 are connected and rigid, $A = H = \{e, \sigma\}$. By Theorem 4

$$q_\rho \;=\; \frac{|H|}{|G|} d_\rho \langle \chi_{\mathbf{1}} \mid \chi_\rho \rangle_H$$

H has only two elements, e and σ, hence

$$\langle \chi_{\mathbf{1}} \mid \chi_\rho \rangle_H = \frac{1}{2}(\chi_\rho(e) + \chi_\rho(\sigma)) = \frac{1}{2}(d_\rho + \chi_\rho(\sigma)).$$

That is, $q_\rho = \frac{d_\rho}{n!}(d_\rho + \chi_\rho(\sigma))$ and so,

$$|p_\rho - q_\rho| \;=\; \frac{d_\rho}{n!} \, |\chi_\rho(\sigma)|$$

We will soon prove:

LEMMA 5. $|\chi_\rho(\sigma)| \leq 2^{O(n)} \sqrt{n}^{\frac{n}{2}}$.

Therefore,

$$\begin{aligned}
|p - q|_1 &= \Sigma_\rho \, |p_\rho - q_\rho| \\
&\leq \Sigma_\rho \frac{d_\rho}{n!} 2^{O(n)} \sqrt{n}^{n/2} \\
&\leq \Sigma_\rho \frac{\sqrt{n!}}{n!} 2^{O(n)} \sqrt{n}^{n/2} \\
&\leq \frac{2^{O(n)} \sqrt{n}^{n/2}}{\sqrt{n!}} \\
&= 2^{O(n)} \frac{\sqrt{\sqrt{n}^n}}{n!} \\
&\leq 2^{O(n)} \frac{1}{\sqrt{(n/2)!}} \lll 2^{-\Omega(n)}
\end{aligned}$$

where the second inequality is due to the fact that the every irreducible representation ρ has dimension $d_\rho \leq \sqrt{n!}$, for $\Sigma_\rho d_\rho^2 = n!$. The third inequality follows from the fact that the number of irreducible representation is the partition number and is about $2^{\sqrt{n}}$ (see Equation (3) in the appendix). $\square$

PROOF OF LEMMA 5. We use the Murnaghan-Nakayama rule. We refer the reader to the appendix for more background about representations of S_n and an explanation of the rule.

THEOREM 7 (THE MURNAGAN-NAKAYAMA RULE). *Let* c *be a permutation with cycle structure* $(c_1, \ldots, c_t)$, $c_1 \geq \ldots \geq c_t$. *Then*

$$\chi_\lambda(\mathbf{c}) = \sum_{s_1, \ldots, s_t} (-1)^{v(s_1)} \cdots (-1)^{v(s_t)},$$

where each s_i *is a skew hook of length* c_i *of the diagram* λ *after* $s_1, \ldots, s_{i-1}$ *have been removed, and* $v(s_i)$ *denotes the number of vertical steps in* s_i.

A sequence $s_1, \ldots, s_t$ where each s_i is a skew hook of length c_i of the diagram λ after $s_1, \ldots, s_{i-1}$ have been removed, is called an *ordered decomposition* of λ into c. The (unordered) decomposition of $s_1, \ldots, s_t$ is the set $\{s_1, \ldots, s_t\}$.

We first claim that the number of unordered decompositions is at most 4^n. To see that notice that a decomposition is in particular a domino covering. For each cell in the shape, there is a unique neighbor (up,down left or right) such that the two are covered together by a domino. In particular, one way to completely specify such a covering is to write, for each cell in the shape, which neighbor (up,down,left or right) is covered with this cell. If the shape has n cells, this takes $2n$ bits, and therefore the number of such coverings is at most 2^{2n}.
We now bound the number of ordered decomposition that correspond to the same unordered decomposition Λ. Λ contains $n/2$ dominos (skew-hooks of length 2). An ordered decomposition with unordered decomposition Λ, must consecutively pick a domino from Λ that lies on the boundary of the current shape. However,

CLAIM 4. *For any shape λ with m cells, the number of dominos from Λ that lie on the boundary of λ is at most $O(\sqrt{m})$*

PROOF. Let $B(\lambda)$ be the set of cells that belong to a domino that lie on the boundary of λ. We prove by induction that if λ covers m cells then $|B(\lambda)| \leq 4\sqrt{m}$. It then follows that the number of dominos of Λ that lie on the boundary of λ is at most $2\sqrt{m}$.
The base case ($m = 1$) is clearly true. Now, let λ be a shape with m cells. If a cell $(i,j) \in B(\lambda)$ then (i,j) is one of the last two cells on the i'th row, and one of the last two cells on the j'th column. Now, λ covers m cells, hence either the first row or the first column contains at least $\sqrt{m}$ cells. W.l.o.g. let us assume the first row contains at least $\sqrt{m}$ cells. Let λ' be the shape λ without the first row. λ' contains $m' \leq m-\sqrt{m}$ cells. By induction, $|B(\lambda')| \leq 4\sqrt{m'}$. Now, $|B(\lambda)| \leq |B(\lambda')| + 2 \leq 4\sqrt{m-\sqrt{m}} + 2 \leq 4\sqrt{m}$ and the induction follows. $\square$

Hence, the number of ordered decompositions that correspond to Λ is at most $(2\sqrt{n})^{n/2}$. By the Murnaghan-Nakayama rule $\chi_\rho(\sigma) \leq 4^n(2\sqrt{n})^{n/2}$.

6. REFERENCES

[1] Noga Alon and Joel H. Spencer. *The Probabilistic Method.* John Wiley & Sons, Inc., 1992.

[2] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 48–53, El Paso, Texas, 4–6 May 1997.

[3] Dan Boneh and Richard J. Lipton. Quantum cryptanalysis of hidden linear functions (extended abstract). In Don Coppersmith, editor, *Advances in Cryptology—CRYPTO '95*, volume 963 of *Lecture Notes in Computer Science*, pages 424–437. Springer-Verlag, 27–31 August 1995.

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

[5] Mark Ettinger and Peter Høyer. On quantum algorithms for noncommutative hidden subgroups. In *Symposium on Theoretical Aspects in Computer Science*, University of Trier, 4–6 March 1999.

[6] Mark Ettinger and Peter Høyer. Quantum state detection via elimination. Technical report, quant-ph/9905099, 1999.

[7] Mark Ettinger and Peter Høyer. A quantum observable for the graph isomorphism problem. Technical report, quant-ph/9901029, 1999.

[8] Mark Ettinger and Peter Høyer and Emanuel Knill. Hidden subgroup states are almost orthogonal. Technical report, quant-ph/9901034, 1999.

[9] M. Goldmann and A. Russell. The computational complexity of solving systems of equations over finite groups. In *Fourteenth Annual IEEE Conference on Computational Complexity*, Atlanta, Georgia, 4–6 May 1999.

[10] Michaelangelo Grigni, Leonard Schulman, and Umesh Vazirani. Unpublished, 1997.

[11] Lisa Hales and Sean Hallgren. Quantum fourier sampling simplified. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 330–338, Atlanta, Georgia, 1–4 May 1999.

[12] Joe Harris and William Fulton. *Representation Theory.* Number 129 in Graduate Texts in Mathematics. Springer-Verlag, New York, NY, 1991.

[13] Alexey Yu. Kitaev. Quantum measurements and the abelian stabilizer problem. Technical report, quant-ph/9511026, 1995.

[14] Johannes Köbler, Uwe Schöning, and Jacobo Torán. *The graph isomorphism problem: its structural complexity.* Birkhäuser Boston Inc., Boston, MA, 1993.

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

[16] 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.

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

APPENDIX

A. THE IRREDUCIBLE REPRESENTATIONS OF S_N

The irreducible representations of S_n may be placed into one-to-one correspondence with the partitions of n. A *partition* of n is a sequence $(\lambda_1, \ldots, \lambda_k)$ of positive integers, with $\lambda_1 \geq \ldots \geq \lambda_k$ for which $\sum \lambda_i = n$. The number of distinct partitions of n (also called the *partition number* of n and denoted $p(n)$) is a very well-studied function. Though no explicit formula is known,

$$p(n) \sim \frac{1}{\alpha n} e^{\beta\sqrt{n}} \tag{3}$$

where $\alpha = 4\sqrt{3}$, $\beta = \pi\sqrt{2/3}$ and the notation $f \sim g$ means that $\lim_{n\to\infty} f(n)/g(n) = 1$. It is customary to identify the partition $\lambda = (\lambda_1, \dots, \lambda_k)$ with a diagram consisting of k rows of boxes, the ith row containing λ_i boxes. We will let λ stand for both the partition and the associated diagram. For example, the diagram corresponding to the partition $\lambda = (6, 4, 3, 3, 2)$ is shown in figure 1.

The irreducible representation associated with λ is denoted ρ_λ. There is an explicit formula for the dimension of ρ_λ. This involves the notion of a *hook*: for a cell (i, j) of a Young tableau λ, the (i, j)-hook $h_{i,j}$ is the collection of all cells of λ which are beneath (i, j) (but in the same column) or to the right of (i, j) (but in the same row), including the cell (i, j). The *length* of the hook $\ell(h)$ is the number of cells appearing in the hook. With this notation, the dimension of ρ_λ may be expressed:

$$d_\lambda = \frac{n!}{\prod_{i,j} \ell(h_{i,j})},$$

this product being taken over all hooks h of λ.

Fixing a partition λ, the function $\chi_\lambda(\cdot)$ is invariant on conjugacy classes of S_n (that is, depends only on the conjugacy class of its argument). Conjugacy classes in S_n consist of all elements with like cycle structure. Let $\mathbf{c} = (c_1, \dots, c_t)$, with $c_1 \geq \cdots \geq c_t$, denote the class of permutations π which can be written as a disjoint product of cycles $\pi = C_1C_2\cdots C_t$ with each C_i having length c_i. Though no explicit formula for $\chi_\lambda(\mathbf{c})$ is known, the remarkable *Murnaghan-Nakayama rule* will be sufficient for our needs.

Definition 3. *A skew hook s of a Young diagram λ is a connected collection of boundary boxes such that their removal from λ results in a (smaller) diagram. The set of skew hooks of a diagram λ may be placed in a natural one-to-one correspondence with the set of hooks as indicated in Figure 2.*

Theorem 8 (The Murnaghan-Nakayama Rule). *Let c be a permutation with a cycle structure $(c_1, \dots, c_t)$, $c_1 \geq \ldots \geq c_t$.*

$$\chi_\lambda(\mathbf{c}) = \sum_{s_1,\dots,s_t} (-1)^{v(s_1)} \cdots (-1)^{v(s_t)},$$

where each s_i is a skew hook of length c_i of the diagram λ after $s_1, \dots, s_{i-1}$ have been removed, and $v(s_i)$ denotes the number of vertical steps in s_i.

A sequence $s_1, \dots, s_t$ where each s_i is a skew hook of length c_i of the diagram λ after $s_1, \dots, s_{i-1}$ have been removed, is called an *ordered decomposition* of λ into c. The (unordered) decomposition of $s_1, \dots, s_t$ is the set $\{s_1, \dots, s_t\}$.

Example 3. *For $\lambda = (2, 2, 1, 1)$ there are three possible ordered decomposition; the first is $(2, 2, 1, 1)$, $(2, 2)$, $(1, 1)$;*

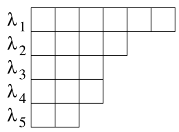

Figure 1: The diagram for $\lambda = (6, 4, 3, 3, 2)$.

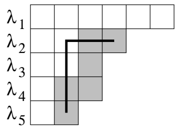

Figure 2: A hook with it's associated skew hook in $\lambda = (6, 4, 3, 3, 2)$.

the second is $(2, 2, 1, 1)$, $(2, 2)$, $(2, 0)$; the third is $(2, 2, 1, 1)$, $(1, 1, 1, 1)$, $(1, 1)$. Notice that as sets there are only two possible decompositions, the first and the third coincide. The first decomposition has three vertical steps, the second has one vertical step, and the third three vertical steps. In all those cases the sign is -1, so the rule says that $\chi_\rho(\sigma) = -3$.

B. THE HSP FOR HAMILTONIAN GROUPS

B.1 Abelian HSP

We now show how the above gives an efficient procedure for solving the hsp for abelian groups.

$G = \mathbb{Z}_n$. Any abelian group is isomorphic to a product of cyclic groups, so we start with cyclic groups $G = \mathbb{Z}_n$. G has n one-dimensional irreducible representations, $\rho_k : G \to \mathbb{C}$ defined by $\rho_k(x) = \omega^{kx}$ where ω is a primitive n root of unity, say $\omega = e^{2\pi i/n}$. When we sample $\rho_k \in \hat{G}$ we see the value k. The kernel is

$$\ker(\rho_k) = \left\{x \ : \ w^{kx} = 1\right\} = \{x \ : \ n|kx\} = \{x \ : \ m|x\}$$

where $m = \frac{n}{gcd(n,k)}$. Clearly, given n and k we can easily compute m. Now, to find H we take s samples. Suppose that we sampled $k_1, \dots, k_s$ and computed $m_1, \dots, m_s$. Then, w.h.p., $H = \bigcap_i \ker\sigma_i = \{x \in G \ : \ m_1|x, \dots, m_s|x\} = \{x \in G \ : \ l.c.m(m_1, \dots, m_s)|x\}$. We can easily compute $g = l.c.m(m_1, \dots, m_s)$. We then know that w.h.p. H is generated by g.

$G = \mathbb{Z}_2^n$. Next, we do the easy case of $G = \mathbb{Z}_2^n$. G has 2^n one-dimensional irreducible representations $\rho_{(a_1,\dots,a_n)} : G \to \mathbb{C}$, $a_1, \dots, a_n \in \{0, 1\}$, defined by $\rho_{(a_1,\dots,a_n)}(x_1, \dots, x_n) = (-1)^{\Sigma a_i x_i}$. The kernel is $\ker(\rho_{(a_1,\dots,a_n)}) = \left\{(x_1, \dots, x_n) \ : \ (-1)^{\Sigma a_i x_i} = 1\right\} = \{(x_1, \dots, x_n) \ : \ \Sigma a_i x_i = 0 \pmod 2\}$. To find H we do s samples and get s linear equations over the field $\mathbb{Z}_2$. We solve the system of linear equations for those $(x_1, \dots, x_n)$ that satisfy all of the s equations. This is Simon's algorithm [17].

$G = \mathbb{Z}_k \times \mathbb{Z}_l$. We now deal with a product of two cyclic groups: $G = \mathbb{Z}_k \times \mathbb{Z}_l$. Before getting into the solution, the reader might want to play with the following example: $G = \mathbb{Z}_8 \times \mathbb{Z}_8$ and the subgroup H is generated by $\{(4, 0), (0, 4), (2, 2)\}$.

G has $k \cdot l$ one-dimensional irreducible representations $\rho_{a,b} : G \to \mathbb{C}$ defined by $\rho_{a,b}(x_1, x_2) = \omega_k^{ax_1}\omega_l^{bx_2}$, where ω_k is a primitive k root of unity, and ω_l is a primitive

l root of unity. The kernel is

$$\begin{aligned}\ker(\rho_{a,b}) &= \left\{(x_1,x_2) \ : \ e^{2\pi i a x_1/k + 2\pi i b x_2/l} = 1\right\} \\ &= \left\{(x_1,x_2) \ : \ e^{2\pi i (l a x_1 + k b x_2)/(kl)} = 1\right\} \\ &= \{(x_1,x_2) \ : \ l a x_1 + k b x_2 = 0 \ (\mod kl)\}\end{aligned}$$

To find H we do s samples and get s linear equations over the *ring* $\mathbb{Z}_{kl}$. Even though $\mathbb{Z}_{kl}$ is a ring (and not a field) we can still diagonlize the system of equations, and sample a random solution [9], which, w.h.p., is a random element of H. Doing that $O(\log(|G|))$ times we get, w.h.p., a set of generators for H.

Any product of cyclic groups . The general case where G is a product of cyclic group is similar.

We note that even though any abelian group is isomorphic to a product of cyclic group, it is not always easy to find this isomorphism. E.g., it might not be easy to find a representations of the subgroup H as a product of cyclic groups.

B.2 The HSP for the Hamiltonian Group

A group is Hamiltonian if every subgroup is normal. Every abelian group is Hamiltonian. The only non-abelian, Hamiltonian group is $G = \mathbb{Z}_2^k \oplus B \oplus Q$ for some abelian group B with exponent b coprime with 2. The irreducible representations of G are

$$\hat{G} = \left\{\rho_{(a_1,\dots,a_k)} \otimes \rho_b \otimes \rho_q \mid \rho_{(a_1,\dots,a_k)} \in \hat{Z_2^k}, \rho_b \in \hat{B}, \rho_q \in \hat{Q}\right\}.$$

A and B are abelian and we know $\hat{A}, \hat{B}$. The quaternin group $Q = \{\pm 1, \pm i, \pm j, \pm k\}$. has a normal subgroup $A = \{1,-1\}$. Q/A has four cosets, $\{\{1,-1\},\{i,-i\},\{j,-j\},\{k,-k\}\}$. Furthermore, Q/A is isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2$ with $\{1,-1\}$ mapping to $(0,0)$, $\{i,-i\}$ mapping to $(1,1)$, $\{j,-j\}$ mapping to $(0,1)$ and $\{k,-k\}$ mapping to $(1,0)$. Hence, Q/A is abelian. Therefore, Q has four irreducible one dimensional representations $\rho_{(c_1,c_2)(x)}$ that given $x \in Q$ give $(-1)^{c_1 x_1 + c_2 x_2}$ where the coset of x in Q/A is isomorphic to $(x_1,x_2) \in \mathbb{Z}_2 \times \mathbb{Z}_2$. The last irreducible representation ρ of Q is obtained by realizing Q as $SU(2)$, the group of 2×2 matrices with determinant 1. We associate 1 with $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, i with $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, j with $\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$, and k with $\begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}$. $-x$ is associated with minus the corresponding matrix.

Now, suppose we measure $\rho = \rho_a \otimes \rho_b \otimes \rho_c$. If ρ_c is one dimensional, we are like in the abelian case, and

$$\begin{aligned}ker(\rho) &= \{x = (x_1,\dots,q) : \sum_{i=1}^{l} d_i x_i = 0 \ (\mod N), \\ &\quad \text{and the coset of } q \text{ in } Q/A \text{ is} \\ &\quad \text{isomorphic to } (x_{l-1}, x_l)\}\end{aligned}$$

where d_i and N are known integers. If ρ_c is two dimensional, then there are eight possibilities for the matrix $\rho(c)$. We group all equations with a particular matrix together, and then we are back to the abelian case. We solve each systems of equations separately, and we find a generating set for each of them. Finally, we find the intersections of the nine sets of equations.