Limitations of Fourier Sampling for Hidden Subgroup Problems

Adrian Chiu   Adrian She

January 2 2019

1 Introduction

Many problems that are known to have algorithms exhibiting an exponential quantum speedup over the best known classical algorithms are special instances of the Hidden Subgroup Problem (HSP). For the HSP, we are firstly given a function $f\colon G \to S$ from a finite group G to a set S that is constant on left cosets and distinct on different cosets of a hidden subgroup $H \leq G$. In other words, we have that $f(x) = f(y)$ if and only if $x^{-1}y \in H$. Such a function f is said to **hide** H. We will assume we can query f efficiently as a black box. Given f, we would like find a set of generators for H.

For example, in Simon's problem, the group in question is $G = (\mathbb{Z}/2\mathbb{Z})^n$, the set of n-bit strings under bitwise addition, and the hidden subgroup is $H = \{0, s\}$ for some unknown $s \in G$. Number-theoretical applications such as Shor's algorithm for factoring integers and the discrete logarithm problem also reduce to instances of the HSP for Abelian groups. In fact, the HSP has an efficient quantum solution over any finite Abelian group [HRT00, Wan10], since every finite Abelian group is a direct product of cyclic groups (groups of the form $\mathbb{Z}/d\mathbb{Z}$ for some d with the operation being addition modulo d.) Algorithms for these problems crucially use the quantum Fourier transform, also known as Fourier sampling.

The HSP for non-Abelian groups has many attractive applications. In particular, we are interested in the graph isomorphism and graph automorphism problems, which reduce to the HSP for the symmetric group S_n. We will outline this reduction as well as a method to define a Fourier transform over any finite group in Section 2. Since one can compute the Fourier transform over a symmetric group in quantum polynomial time (in n) [Bea97], the Fourier sampling method provides a promising approach for solving the HSP over S_n. However, the Fourier sampling method fails for S_n for various reasons. For example, we will show in Section 3 that the Fourier sampling algorithm cannot distinguish between conjugate subgroups, and this limits the types of subgroups recoverable for groups such as S_n, which have many conjugate subgroups.

2 Background

2.1 Representation Theory

To generalize the familiar Fourier transform to arbitrary groups G, we will need to use some representation theory that we will outline here. While representation theory can be used in more general settings than what we will present, we will always assume that G is a finite group and V is a finite-dimensional vector space over the complex numbers $\mathbb{C}$. The proofs of the following facts can be found in representation theory textbooks such as [Sag13] or [SS96].

Definition 1. A **representation** of G is a homomorphism $\rho\colon G \to GL(V)$ where $GL(V)$ is the set of all invertible linear transformations $V \to V$. The representation is said to be **unitary** if it is a homomorphism $\rho : G \to U(V)$ where $U(V)$ is the set of unitary transformations on V. The **dimension** of the representation ρ is the dimension of the vector space V.

Example 1. The **trivial representation** defined by $\rho(g) = 1$ for any $g \in G$ is a one-dimensional representation of any group.

Example 2. Let $0 \leq k \leq n-1$ and let $G = \mathbb{Z}/n\mathbb{Z}$ be the cyclic group of order n. The map $\phi(x) = e^{2\pi kxi/n}$ for any $x \in G$ is a one-dimensional representation of G.

Example 3. Let V be a vector space of dimension $|G|$ and label its standard basis vectors by $|g\rangle$ for $g \in G$. Define $\rho : G \mapsto GL(V)$ by $\rho(g)\,|x\rangle = |gx\rangle$ for any basis vector $|x\rangle \in V$. This is a $|G|$-dimensional representation of G known as the **(left) regular representation**, and $\rho(g)$ is a $|G| \times |G|$ permutation matrix in the basis $\{|g\rangle\}_{g \in G}$.

For example, if $G = \mathbb{Z}/2\mathbb{Z} = \{0, 1\}$, its regular representation can be realized by the matrices $\rho(0) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ and $\rho(1) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

Given a group G, we are interested in determining all possible representations of the group. The notion of irreducible representations is useful for this task.

Definition 2. A representation $\rho\colon G \to GL(V)$ of a group G is **irreducible** if no nontrivial proper subspace W of the vector space V has $\rho(g)W \subseteq W$ for all $g \in G$ (the subspace W is not fixed by all $\rho(g)$.)

Example 4. Any one-dimensional representation of a group is irreducible since the only proper subspaces of any one-dimensional vector space is $\{0\}$.

Example 5. For any non-trivial group, the regular representation is not irreducible since if W is the one-dimensional subspace spanned by $\sum_{g \in G} |g\rangle$, then $\rho(g)W \subseteq W$ (i.e. it has a proper subspace which is fixed by all $\rho(g)$.)

Definition 3. Two representations ρ_1, ρ_2 of a group G are said to be **equivalent** if they are the same up to change of basis (there exists $A \in GL(V)$ such that $A\rho_1(g)A^{-1} = \rho_2(g)$ for every $g \in G$.)

The following are some facts about irreducible representations.

Theorem 1 (Maschke's Theorem, [Sag13, Theorem 1.5.3])**.** *Every representation ρ of a finite group G decomposes into a direct sum of irreducible representations, which can made equivalent to a unitary and block-diagonal matrix representation by a change of basis.*

Hence, without loss of generality, we will now assume all of our representations to be unitary.

Theorem 2 (Schur's Lemma, [Sag13, Corollary 1.6.8])**.** *Let V be a complex vector space and ρ be an irreducible representation of G on V. Suppose $M : V \to V$ is a linear transformation for which $\rho(g)M = M\rho(g)$ for all $g \in G$. Then $M = \lambda I$ for some $\lambda \in \mathbb{C}$, where I is the identity matrix.*

The **characters** of representations can be used to investigate them.

Definition 4. Let ρ be a matrix representation of a group G. Its **character** χ_ρ is a function $G \to \mathbb{C}$ defined by $\chi_\rho(g) = \operatorname{Tr}(\rho(g))$, where Tr is the trace of a matrix. We call a character χ_ρ of ρ **irreducible** if ρ is an irreducible representation.

Example 6. Let ρ be the regular representation of a group G. Then $\chi_\rho(e) = |G|$ for the identity element $e \in G$ and $\chi_\rho(g) = 0$ for any other element (since $gx = x$ if and only if $g = e$.) In fact, for any representation ρ of dimension d, $\chi_\rho(e) = d$.

Recall that elements g_1 and g_2 in a group are **conjugate** if there is exists some element h for which $g_1 = hg_2h^{-1}$. Since conjugacy of elements is an equivalence relation on the elements of the group, we call the resulting equivalence classes the **conjugacy classes** of the group. We now recall some results about conjugate elements in S_n since we will use them later. We will write an element $\sigma \in S_n$ by its cycle decomposition.

Example 7 ([DF04, Section 4.3])**.** In the symmetric group S_n, recall that the **cycle type** of an element $\sigma \in S_n$ written as a cycle decomposition is the length of its cycles sorted in non-increasing order. For example, the permutation $(145)(23) \in \sigma$ has cycle type $(3, 2)$. The permutations $\sigma, \tau \in S_n$ are conjugate if and only if they have the same cycle type. For instance, $(145)(23)$ and $(123)(45)$ are conjugate. Hence, for every integer partition $\lambda \vdash n$, there is unique conjugacy class in S_n consisting of permutations of that cycle type λ.

The relationship between conjugacy classes and representations of a finite group is explained in the following theorem.

Theorem 3. *Let G be a finite group and χ_ρ be a character of some representation.*

(a) ([Sag13, Proposition 1.8.5]) If g and h are conjugate in G, then $\chi_\rho(g) = \chi_\rho(h)$.

(b) ([Sag13, Corollary 1.9.4]) The character χ_ρ can be decomposed as $\chi_\rho = \sum_i m_i \chi_{\sigma_i}$ where each m_i is a non-negative integer and χ_{σ_i} is an irreducible character.

(c) ([Sag13, Proposition 1.10.1c]) The number of irreducible representations of G is equal to the number of conjugacy classes of G.

Next, suppose F_G is the set of all functions from a group $G \mapsto \mathbb{C}$. One can define an inner product on two functions $\psi, \phi \in F_G$ by setting $\langle \psi, \phi \rangle_G = \frac{1}{|G|} \sum_{g \in G} \phi(g)\overline{\psi(g)}$. Under this inner product, characters satisfy certain orthogonality relations listed below.

Theorem 4. *Let ρ_1, ρ_2 be irreducible representations of a group G. Then the following hold:*

(a) ([Sag13, Theorem 1.9.3]) $\langle \chi_{\rho_1}, \chi_{\rho_2} \rangle = \begin{cases} 1 & \rho_1, \rho_2 \textit{ equivalent} \\ 0 & \textit{otherwise} \end{cases}$

(b) ([Sag13, Theorem 1.10.3]) If $g, h \in G$ and K is the conjugacy class of g in G, then

$$\sum_\chi \chi(g)\chi(h^{-1}) = \begin{cases} \frac{|G|}{|K|} & g, h \textit{ are conjugate} \\ 0 & \textit{otherwise} \end{cases},$$

where the sum is over all irreducible characters of G.

Corollary 1. *If χ_ρ is the character of a non-trivial irreducible representation of G, then $\sum_{g \in G} \chi(g) = 0$.*

Proof. If χ_ρ is non-trivial, it is orthogonal to the trivial and irreducible character $\chi(g) = 1$ for every $g \in G$ by Theorem 4(a). □

Theorem 5 ([Sag13, Proposition 1.10.1b])**.** *Let $\hat{G}$ be the set of irreducible representations of G. Then $\sum_{\rho \in \hat{G}} d_\rho^2 = |G|$ where d_ρ is the dimension of ρ.*

Finally, suppose $H \leq G$ is a subgroup. We can construct representations of H from G.

Definition 5. Let $H \leq G$ is a subgroup and ρ is a representation of G, the **restriction** ρ_H is a representation of H defined by $\rho_H(h) = \rho(h)$ for $h \in H$.

Note that although ρ may be an irreducible representation of G, the restriction ρ_H is in general not an irreducible representation of H.

2.2 Graph Isomorphism as an HSP

The graph isomorphism problem is the problem of determining whether two finite graphs are isomorphic. Two graphs G and H are said to be isomorphic if there exists a bijection $f\colon V(G) \to V(H)$ from the vertices of G to the vertices of H, such that there is an edge $(u, v) \in G$ if and only if there is an edge $(f(u), f(v)) \in H$. An automorphism of a graph G is an isomorphism between G and itself. These concepts are illustrated in Figure 1. The graph automorphism problem is the problem of finding (a generating set of) $\mathrm{Aut}(G)$, the automorphism group of a graph G. The graph isomorphism problem is polynomial time reducible to the graph automorphism problem. In particular, the graph automorphism problem is believed to be at least as hard as the graph isomorphism problem [SL17].

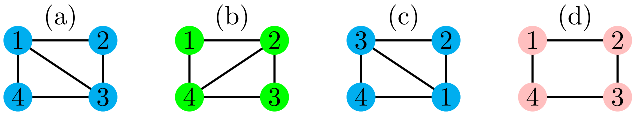

Figure 1: Graph (a) is isomorphic to graph (b) via the isomorphism (12)(34). However, (12)(34) is not an automorphism of graph (a). Instead, graph (c) illustrates an automorphism of graph (a) by the permutation (13). Neither (a), (b), or (c) is isomorphic to the graph (d) since they have a different number of edges.

2.2.1 Graph Automorphism Problem to HSP Reduction

Let G be a graph with n vertices. Note that $\mathrm{Aut}(G)$ embeds into S_n the symmetric group of n symbols; any automorphism of G is a permutation of its n vertices. Let $H = \mathrm{Aut}(G)$ be the hidden subgroup of S_n of interest. Consider the function $f\colon S_n \to X$, where X is the set of graphs created by permutations of G, by $\sigma \mapsto \sigma(G)$. We can think of this operation as creating a new graph $\sigma(G)$ by a permutating the rows and columns of the adjacency matrix of the original graph G according to the permutation σ.

Notice that f is constant on cosets of H in S_n, since for $\sigma, \tau \in S_n$, we have that $\sigma(G) = \tau(G)$ if and only if $\tau^{-1}\sigma \in \mathrm{Aut}(G)$. Hence, we have created an instance of the hidden subgroup problem where $f : S_n \mapsto X$ hides the automorphism group of graph G.

2.2.2 Graph Isomorphism Problem to Graph Automorphism Problem Reduction

We give the reduction due to Jozsa in [Joz01]. Consider two connected graphs A and B both with n vertices, their disjoint union $C = A \sqcup B$, and the automorphism group $\mathrm{Aut}(C)$ which embeds into S_{2n}. We only consider the case where both A and B have n vertices as it is a necessary condition for them to be isomorphic. Any automorphism of C must either permute the vertices of A and B separately, or swap the their vertices entirely since no vertex in A is connected to a vertex in B.

Therefore if $H = S_n \times S_n$ (and identify it in S_{2n}), and $\sigma \in S_{2n}$ a permutation which swaps $\{1, 2, \ldots, n\}$ with $\{n+1, n+2, \ldots, 2n\}$, then $\mathrm{Aut}(C)$ is a subset of $H \cup \sigma(H)$. Without loss of generality, we can choose $\sigma = (1\ n+1)(2\ n+2)\ldots(n\ 2n)$, which is an involution (a permutation with $\sigma^2 = e$.)

If A and B are not isomorphic, $\mathrm{Aut}(C)$ lies entirely in H. Otherwise, if A and B are isomorphic, then for every $\tau \in \mathrm{Aut}(C) \cap H$, and the composition $\sigma\tau \in \sigma(H)$ is also an automorphism of C. In this case, half of $\mathrm{Aut}(C)$ lies in H, and the other half lies in $\sigma(H)$. We can easily check if an element τ of $\mathrm{Aut}(C)$ lies in H or $\sigma(H)$ by evaluating $\tau(1)$, so if we randomly sample elements of $\mathrm{Aut}(C)$, with high probability we can determine if A and B are isomorphic. In the general case where G and H may not be connected, it suffices to check there is some isomorphism between pairs of their connected components.

By a result of Erdös and Renyi in [ER63], for almost all graphs we have that $\mathrm{Aut}(G) = \{e\}$, the trivial group, in the sense that the probability that a random n-vertex graph has a trivial automorphism group tends to one as $n \to \infty$. Hence, we do not lose much generality by considering the **rigid graph isomorphism** problem, where we assume that the graphs to be tested have trivial automorphism group. In this case, if G, H are graphs of this type, it suffices to check if $\mathrm{Aut}(G \sqcup H) = \{e\}$ or if $\mathrm{Aut}(G \sqcup H) \simeq \mathbb{Z}/2\mathbb{Z}$ to check if G and H are isomorphic or not by the previous remarks.

2.3 Quantum Fourier Transform over Arbitrary Groups

Definition 6. Let $f\colon G \to \mathbb{C}$, ρ an irreducible representation of G of dimension d_ρ. The Fourier transform of f with respect to ρ is defined to be a $d_\rho \times d_\rho$ matrix $\hat{f}(\rho) = \sqrt{\dfrac{d_\rho}{|G|}} \displaystyle\sum_{g \in G} f(g)\rho(g)$.

Let $\hat{G}$ be the set of distinct irreducible representations (up to equivalence) of G. Then the **Fourier transform** of f is said to be the collection of matrices $\{\, \hat{f}(\rho) \mid \rho \in \hat{G} \,\}$.

By Theorem 5, the Fourier transform of f is an operation from $\mathbb{C}^{|G|}$, since the function f is specified by $|G|$ values, to $\mathbb{C}^{|G|}$, since the set of resulting matrices given by the Fourier transform are described by

$\sum_{\rho \in \hat{G}} d_\rho^2 = |G|$ numbers. Furthermore, the Fourier transform is also a unitary operation in the sense that $\langle f, g \rangle = \sum_{\rho \in \bar{G}} \left\langle \hat{f}(\rho), \hat{g}(\rho) \right\rangle_F$, where $\langle A, B \rangle_F = \mathrm{Tr}(A^* B) = \sum_{i,j} \overline{a_{ij}} b_{ij}$ is the Frobenius inner product on matrices.

Hence, the quantum Fourier transform over an arbitrary finite group may be expressed as an operation on quantum states since it is unitary. As such, we can identify a function f with the superposition $\sum_{g \in G} f(g) |g\rangle$ and write the Fourier transform as the transformation:

$$\sum_{g \in G} f(g) |g\rangle \mapsto \sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G|}} \sum_{1 \le i,j \le d_\rho} \left(\sum_{g \in G} f(g) \rho(g) \right)_{i,j} |\rho, i, j\rangle . \tag{1}$$

Note that in general, we will treat the representation ρ, the row i, and the column j as separate registers.

The standard algorithm for the HSP is referred to as **Fourier sampling**, and is outlined below. Let G be a finite group, H a subgroup of G, and f a function which hides H.

1. Prepare a uniform superposition over the elements of G, with a second register initialized to zero:

$$|\psi_0\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle \otimes |0\rangle .$$

2. Query f and XOR it with the second register, resulting in the state:

$$|\psi_1\rangle = \frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle \otimes |f(g)\rangle .$$

3. Measure the second register. If we measure $|f(ch)\rangle = |f(c)\rangle$ for some coset cH of H, the first register is in a uniform superposition over the coset cH:

$$|\psi_2\rangle = \frac{1}{\sqrt{|H|}} \sum_{h \in H} |ch\rangle \otimes |f(c)\rangle .$$

   Let $\mathbb{1}_{cH}(g)$ denote the indicator function of the coset cH. Then we can write the first register as: $\frac{1}{\sqrt{|H|}} \sum_{g \in G} \mathbb{1}_{cH}(g) |g\rangle$.

4. Let $\hat{G}$ be the set of irreducible representations of G. Perform a Fourier transform on the first register with respect to the indicator function to obtain:

$$\begin{aligned}\frac{1}{\sqrt{|H|}} \sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G|}} \sum_{1 \le i,j, \le d_\rho} (\widehat{\mathbb{1}_{cH}(\rho)})_{i,j} |\rho, i, j\rangle &= \frac{1}{\sqrt{|H|}} \sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G|}} \sum_{1 \le i,j, \le d_\rho} \left(\sum_{g \in G} \mathbb{1}_{cH}(g) \rho(g) \right)_{i,j} |\rho, i, j\rangle \\ &= \sum_{\rho \in \hat{G}} \sqrt{\frac{d_\rho}{|G||H|}} \sum_{1 \le i,j, \le d_\rho} \left(\sum_{h \in H} \rho(ch) \right)_{i,j} |\rho, i, j\rangle ,\end{aligned}$$

   since 1_{cH} was an indicator function of some coset.

5. Make a measurement on the state obtained in Step 4. We may either measure the first register only to obtain a representation ρ, which is known as **weak Fourier sampling**. Otherwise, we can measure ρ and then subsequently measure the row and column registers i and j. This is known as **strong Fourier sampling**.

We then repeat Steps 1-5 a number of times to obtain a set of representations and possibly some row and column indices. It is hoped that H can be reconstructed from the representations observed by this algorithm. Recall that the kernel of a representation, $\mathrm{Ker}\,\rho$, is the set of elements $g \in G$ for which $\rho(g) = I$, the identity. In the case where H is promised to be a normal subgroup of G, if we choose $s = c\log_2|G|$ for appropriate c, then with high probability $H = \bigcap_{i=1}^{s} \operatorname{Ker}\rho_i$, where ρ_i are the representations observed during Fourier sampling [HRT00, Theorem 5].

Note that if G is Abelian (for instance, in the factoring problem), all irreducible representations of an Abelian group are 1-dimensional and furthermore there are exactly $|G|$ of them. Therefore the Fourier transform in Step 4 can be simplified to $\sum_{\rho\in\hat{G}} \sqrt{\frac{1}{|G||H|}} \sum_{h\in H} \rho(ch)\, |\rho\rangle$. Taking concretely the case where $G = \mathbb{Z}/n\mathbb{Z}$ and $H = \{e\}$, each coset is just a single element $\{x\}$. Hence, since we have calculated all irreducible representations of G in Example 2, we get that the state $|x\rangle$ becomes the state $\frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} e^{\frac{2\pi ikx}{n}} |k\rangle$ after the transform. This is exactly the Fourier transform used in Shor's algorithm.

3 Weak Fourier Sampling

3.1 Analysis of the Algorithm

We now analyze the weak Fourier sampling algorithm given in Section 2.3.

Theorem 6 ([HRT00, Theorem 4])**.** *Suppose f hides a subgroup H. Then, the probability of measuring a representation ρ in Step 5 of the weak Fourier sampling algorithm is* $\frac{|H|}{|G|} d_\rho \langle \chi_{\rho_H}, \chi_{\mathbb{1}_H}\rangle_H = \frac{d_\rho}{|G|} \sum_{h\in H} \chi_\rho(h)$ *(by the definition of the inner product), where d_ρ is the dimension of ρ.*

Proof. Let $\|A\|$ be the Frobenius norm of matrix A given by $\|A\|^2 = \operatorname{Tr}(A^*A)$. As ρ is a homomorphism and $\rho(c)$ is a unitary matrix:

$$\sum_{1\le i,j\le d_\rho} \left| \sqrt{\frac{d_\rho}{|G||H|}} \left(\sum_{h\in H} \rho(ch) \right)_{i,j} \right|^2 = \frac{d_\rho}{|G||H|} \left\| \sum_{h\in H} \rho(ch) \right\|^2 = \frac{d_\rho}{|G||H|} \left\| \rho(c) \sum_{h\in H} \rho(h) \right\|^2 = \frac{d_\rho}{|G||H|} \left\| \sum_{h\in H} \rho(h) \right\|^2 .$$

Note that this calculation also shows that the probability of measuring $\rho \in \hat{G}$ in the first register is independent of the coset state to which the original superposition collapses in Step 3 of the algorithm.

Therefore we wish to calculate the matrix $\sum_{h\in H} \rho(h)$. We consider the restriction ρ_H as we only evaluate ρ on H. ρ_H may not be irreducible over H, but it decomposes into irreducible representations over H by Theorem 1. Then for some unitary matrix U and irreducible representations σ_i of H, we have that,

$$\sum_{h\in H} \rho_H(h) = U \begin{bmatrix} \sum_{h\in H} \sigma_1(h) & 0 & \cdots & 0 \\ 0 & \sum_{h\in H} \sigma_2(h) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sum_{h\in H} \sigma_s(h). \end{bmatrix} U^\dagger.$$

Let $M_i = \sum_{h\in H} \sigma_i(h)$. Notice that for any $h' \in H$, and for the irreducible representation $\sigma_i \in \hat{H}$ we have:

$$\sigma_i(h')M_i = \sigma_i(h') \sum_{h\in H} \sigma_i(h) = \sum_{h\in H} \sigma_i(h'h) = \sum_{h\in H} \sigma_i(hh') = M_i\sigma_i(h'),$$

as we can replace $h'h$ with hh' because we are summing over all elements of the group H. Then by Schur's Lemma (Theorem 2), $M_i = \sum_{h\in H} \sigma_i(h) = \lambda I$ for some $\lambda \in \mathbb{C}$. Taking the trace on both sides, we get:

$$d_\rho\lambda = \operatorname{Tr}\left(\sum_{h\in H} \sigma_i(h) \right) = \sum_{h\in H} \operatorname{Tr}(\sigma_i(h)) = \sum_{h\in H} \chi_{\sigma_i}(h) = |H|\langle \chi_{\sigma_i}, \chi_{\mathbb{1}}\rangle_H = \begin{cases} |H| & \sigma_i \text{ trivial} \\ 0 & \text{otherwise} \end{cases}$$

by orthogonality of characters (Corollary 1).

Therefore $M_i = 0$ if σ_i is nontrivial. If σ_i is trivial, then $M_i = |H|$. Finally, we have:

$$\left\| \sum_{h \in H} \rho_H(h) \right\|^2 = |H|^2 \langle \chi_{\rho_H}, \chi_{\mathbb{1}_H} \rangle_H,$$

because $\langle \chi_{\rho_H}, \chi_{\mathbb{1}_H} \rangle_H$ counts the number of trivial irreducible representations ρ_H contains by Theorem 3(b). Hence, putting this back into Equation 3.1, we have proved that the probability of measuring ρ is $\frac{d_\rho}{|G||H|}|H|^2 \langle \chi_{\rho_H}, \chi_{\mathbb{1}_H} \rangle_H = \frac{|H|}{|G|} d_\rho \langle \chi_{\rho_H}, \chi_{\mathbb{1}_H} \rangle_H$. □

Recall that two subgroups $H_1, H_2 \le G$ are **conjugate** if there is an element $g \in G$ for which $H_2 = \{ghg^{-1} : h \in H_1\} = gHg^{-1}$.

Example 8. Let $\sigma \in S_n$ and $\tau \in S_n$ be of the same cycle type and $H_1 = \langle \sigma \rangle$ and $H_2 = \langle \tau \rangle$ be subgroups generated by those elements respectively. Then H_1 and H_2 are conjugate.

Corollary 2. *Suppose $H_1, H_2 = gH_1g^{-1}$ are conjugate subgroups in G. Then the weak Fourier sampling algorithm cannot distinguish between H_1 and H_2, in the sense that weak Fourier sampling produces the same distribution on representations whether a function f hides H_1 or f hides H_2.*

Proof. If H_1 and $H_2 = gH_1g^{-1}$ are conjugate, there is an isomorphism $\phi : H_1 \to H_2$ given by $\phi(x) = gxg^{-1}$. Hence, $\sum_{h_2 \in H_2} \chi_\rho(h_2) = \sum_{h_1 \in H_1} \chi_\rho(gh_1g^{-1}) = \sum_{h_1 \in H_1} \chi_\rho(h_1)$ since characters are constant on conjugacy classes (Theorem 3(a)). Hence, the probability distribution on representations induced from hiding H_1 is the same as that of hiding H_2 by Theorem 6. □

3.2 Weak Fourier Sampling Fails for Rigid Graph Automorphism

Note that Corollary 2 does not rule out an efficient solution to the rigid graph automorphism problem using weak Fourier sampling, since the trivial group cannot be conjugate to any non-trivial group. Our main goal in this section is to prove that weak Fourier sampling fails to solve the rigid graph automorphism problem.

From Theorem 6, we get that if the hidden subgroup is trivial, then the probability that representation ρ is measured is $\frac{d_\rho}{|G|}\chi_\rho(e) = \frac{d_\rho^2}{|G|}$. If H is some non-trivial subgroup, we say that H is distinguishable if the probability distribution on representations induced from hiding H has sufficiently large L^1 distance from this distribution, specifically when a *polylog*$(|G|)$ number of samples can be used to solve the hidden subgroup problem under the promise that the given function hides either H or the identity.

Definition 7. A subgroup $H \le G$ is **distinguishable from the identity** using weak Fourier sampling if there is a constant c for which

$$\mathcal{D}_H = \frac{1}{|G|} \sum_\rho d_\rho \left| \sum_{h \in H \; h \neq e} \chi_\rho(h) \right| \ge (\log |G|)^{-c}$$

Otherwise, H is **indistinguishable**. Since $|S_n| = n!$, we need $\mathcal{D}_H \ge \frac{1}{poly(n)}$ for some polynomial in n if H is a distinguishable group of S_n.

Theorem 7 ([HRT00, Theorem 6])**.** *Suppose $H = \{e, \sigma\} \le S_{2n}$ where $\sigma = (1n)(2\ n+1)\dots(n\ 2n)$. Then H is indistinguishable from the identity for sufficiently large n.*

The proof uses the character theory of S_n, especially the **Murnaghan-Nakayama rule**. Recall from Section 2 that cycle types and irreducible representations of S_n are labelled by integer partitions. If λ, ρ are partitions of n, we write $\chi_\lambda(\rho)$ as the value of the character χ_λ evaluated at any permutation of cycle type ρ.

Definition 8. Given any partition $\lambda = (\lambda_1, \dots, \lambda_n)$ of n, the **diagram** $D(\lambda)$ associated to λ is an array of left justified boxes with λ_i boxes in row i from the top. A **skew shape** ψ in a diagram is the subset of a diagram whose interior is connected and does not contain any 2×2 square. The **height** of a skew shape ψ,

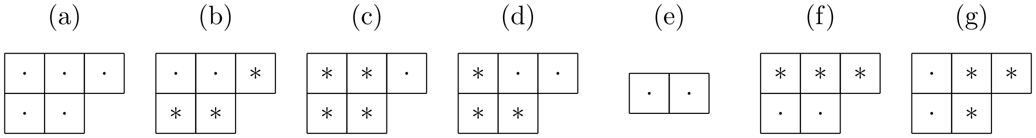

Figure 2: (a) is a diagram of shape $(3, 2)$. The cells indicated by * in (b) and (c) are subsets of the diagram in (a) which are not skew shapes. The cells indicated by * in (d), (f), and (g) are subsets of (a) which is a skew shape in (a) of heights 1, 0, 1 respectively. (e) is the diagram (a) with the skew shape (d) removed and it is a domino.

written as $\mathrm{ht}(\psi)$ is the one less the number of rows it intersects in $D(\lambda)$. If ψ is a skew shape in the diagram of λ whose removal leaves a diagram of partition shape, we write $\lambda \setminus \psi$ for the resulting partition. Finally, we call a skew shape with two cells a **domino**. These concepts are illustrated in Figure 2.

Theorem 8 (Murnaghan-Nakayama Rule, [Sag13, Theorem 4.10.2])**.** *Let $\lambda = (\lambda_1, \ldots, \lambda_n)$ and $\rho = (\rho_1, \ldots, \rho_n)$ be partitions of n and $D(\lambda)$ be the diagram of λ. Then, $\chi_\lambda(\rho) = \sum_{\psi \subseteq D(\lambda)} (-1)^{ht(\psi)} \chi_{\lambda \setminus \psi}(\rho')$, summed over all skew shapes ψ with ρ_1 boxes where $\lambda \setminus \psi$ is a partition, $\rho' = (\rho_2, \ldots, \rho_n)$, and with base case $\chi_\lambda(\rho) = 1$, with both λ and ρ the empty partition.*

Example 9. We have $\chi_{(1,1)}(2) = -1$ and $\chi_{(2)}(2) = 1$ since both $(1, 1)$ and (2) are skew shapes where removing two cells leaves the empty partition.

Now to compute $\chi_{(3,2)}(3, 2)$, note that there are 3 possible skew hooks with 3 boxes that can be removed from $(3, 2)$, namely Figure 2(d), (f) and (g), to leave a partition shape. Since they have heights 1,0, and 1 respectively, then $\chi_{(3,2)}(3, 2) = -\chi_{(2)}(2) + \chi_{(2)}(2) - \chi_{(1,1)}(2) = 1$.

Before proving the main theorem, we make the following observation about dominos and diagrams.

Lemma 1. *For any diagram λ with n cells, there are at most $4\sqrt{n}$ dominos λ_D whose removal from λ yields a diagram of partition shape.*

Proof. We proceed by induction on n and the claim follows for $n \leq 2$. Now note that since λ has n cells, its topmost row has at least $\sqrt{n}$ cells and if not, the leftmost column has at least $\sqrt{n}$ cells for there to be n cells in total. In the first case, there are most two dominos in λ intersecting the topmost row whose removal leaves a partition shape (namely its leftmost or rightmost cell and one to the left or below). Otherwise, in the second case, there is at most two dominos in λ intersecting the leftmost column whose removal leaves a partition shape (namely its bottommost or topmost cell and the one above or to the right). Any other domino with this property must lie entirely in the remaining shape with $\leq n - \sqrt{n}$ cells with either the first row or first column removed from λ. Hence, by the inductive hypothesis, we have that $|\lambda_D| \leq 4\sqrt{n - \sqrt{n}} + 2 \leq 4\sqrt{n}$ since $\sqrt{n} - \sqrt{n - \sqrt{n}} = \frac{\sqrt{n}}{\sqrt{n} + \sqrt{n - \sqrt{n}}} \geq \frac{1}{2}$. □

Lemma 2. *For any partition λ of $2n$, we have $|\chi_\lambda((2^n))| \leq 4^n(\sqrt{2n})^n$.*

Proof. We proceed by induction on n and we already have already verified the case for $n = 1$ by Example 9. Now, let $n > 1$. By the Murnaghan-Nakayama rule, we have that $|\chi_\lambda((2^n))| \leq \sum_{\lambda'} |\chi_{\lambda'}((2^{n-1}))|$ where λ' is a partition of $2n - 2$ created by removing a domino from λ. By Lemma 1 and the inductive hypothesis, $|\chi_\lambda(2^n)| \leq \left(4\sqrt{2n}\right)\left(4^{n-1}\right)\left(\sqrt{2n-2}\right)^{n-1} \leq 4^n(\sqrt{2n})^n$. □

Proof of Theorem 7. We have from Theorem 5 that $d_\rho \leq \sqrt{(2n)!}$ for any representation of S_{2n}, and also that the number of partitions of $2n$ is $2^{c\sqrt{2n}}$ for some constant c ([HRT00, Appendix]). Hence, from these results and Lemma 2, for sufficiently large n,

$$\mathcal{D}_H \leq \frac{2^{c\sqrt{2n}}}{\sqrt{(2n)!}}(4^n(\sqrt{2n})^n) \leq 2^{c'n}\sqrt{\frac{n^n}{(2n)!}} \leq 2^{c'n}\sqrt{\frac{1}{n!}} \leq 2^{-c'n} \leq \frac{1}{poly(n)},$$

letting $c' = c\sqrt{2} + 3$ and since $n! \geq 16^{c'n}$ for sufficiently large n. □

3.3 Weak Fourier Sampling fails to Find Automorphism Groups of Graphs

In the previous section, we proved that if $H \leq S_n$ is a subgroup of constant size, weak Fourier sampling may not be able to determine H efficiently. We may then wonder if weak Fourier sampling may still fail to determine H if H is a subgroup of larger size. Indeed, if we have that $|H| = \frac{n!}{2}$, then H is normal in S_n and hence we have an efficient solution to the hidden subgroup problem by the normal subgroup reconstruction algorithm in [HRT00, Theorem 5]. However, we will prove in this section that there are subgroups of $poly(n)$ size in S_n that may not be distinguishable by the weak Fourier sampling algorithm.

Definition 9. Let $n \geq 3$. The **dihedral group** D_n is set of symmetries of a regular n-gon. As embdedded into S_n, it consists of n rotations (cycles of length n) and n reflections (elements of order 2). If n is odd, every reflection has some fixed point, and if n is even, there are $n/2$ reflections fixing two points and $n/2$ reflections without fixed points.

Note the if C_n is the n-vertex cycle group with vertices $\{1, \dots, n\}$ and edges $(i, i+1)$ for $i = 1, \dots, n-1$ and $(n, 1)$, then the automorphism group of the graph C_n is dihedral group D_n. Figure 3 represents the generators of $Aut(C_4) \simeq D_4$.

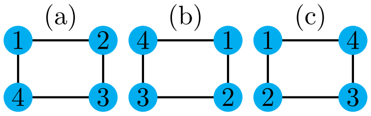

Figure 3: Graph (a) is the cycle C_4, graph (b) represents a rotation automorphism and graph (c) represents a reflection automorphism.

Theorem 9 ([SL17], Theorem 6)**.** *The subgroup $D_n \leq S_n$ is indistinguishable from the identity using weak Fourier sampling for sufficiently large n.*

The original proof of Theorem 9 in [SL17] uses induced representations of D_n to S_n. We provide a shorter proof of Theorem 9 using the following result due to Kempe et al.

Theorem 10 ([KS05], Proposition 1)**.** *Suppose $H \leq G$, G has non-identity conjugacy classes $C_1, \dots, C_k$ and $\mathcal{D}_H$ is as given in Definition 7, then $\mathcal{D}_H \leq \sum_{i=1}^{k} |C_i \cap H||C_i|^{-\frac{1}{2}}$.*

Proof. Using the triangle inequality, we have

$$\mathcal{D}_H \leq \frac{1}{|G|} \sum_{\rho} \sum_{h \in H, h \neq e} d_\rho |\chi_\rho(h)| = \frac{1}{|G|} \sum_{h \in H, h \neq e} \sum_{\rho} d_\rho |\chi_\rho(h)|. \tag{2}$$

Now, let $h \in H$ be fixed and using the Cauchy-Schwarz inequality,

$$\sum_{\rho} d_\rho |\chi_\rho(h)| \leq \left(\sum_{\rho} d_\rho^2 \right)^{1/2} \left(\sum_{\rho} |\chi_\rho(h)|^2 \right)^{1/2}. \tag{3}$$

We have by Theorem 5 that $\sum_\rho d_\rho^2 = |G|$ and that $\sum_\rho |\chi_\rho(h)|^2 = \frac{|G|}{|C(h)|}$ by Theorem 4(b) where $C(h)$ is the conjugacy class of h. Hence, combining these observations with Equation 2 and Equation 3 yields

$$\mathcal{D}_H \leq \frac{1}{|G|} \sum_{h \in H, h \neq e} |G|^{1/2} \left(\frac{|G|}{|C(h)|} \right)^{1/2} = \frac{1}{|G|} \sum_{i=1}^{k} |H \cap C_i||G|^{1/2} \left(\frac{|G|}{|C_i|} \right)^{1/2} = \sum_{i=1}^{k} |H \cap C_i||C_i|^{-1/2} \tag{4}$$

since every $h \in H$ occurs in exactly one conjugacy class. □

Our main tool in the proof is the following result.

Theorem 11 ([Sta12, Proposition 1.3.2]). *The number of permutations $\sigma \in S_n$ with α_i cycles of length i is* $n! \prod_{i=1}^{n} \frac{1}{i^{\alpha_i} \alpha_i!}$.

Proof of Theorem 9. Let $H = D_n$ and firstly assume that $n \geq 3$ is even. Then since S_n has $(n-1)!$ cycles of length n and $\frac{n!}{2^i i!}$ involutions fixing $n - 2i$ elements, then we have that

$$\mathcal{D}_H \leq \frac{n}{\sqrt{(n-1)!}} + \frac{n}{2}\left(\frac{2^{n/2}(n/2)!}{n!}\right)^{1/2} + \frac{n}{2}\left(\frac{2^{(n-2)/2}((n-2)/2)!}{n!}\right)^{1/2} \leq \frac{n}{\sqrt{(n-1)!}} + n\left(\frac{2^{n/2}(n/2)!}{n!}\right)^{1/2}.$$

We have for $n \geq 6$ that $\frac{(n/2)!}{n!} \leq 2^{-n}$ and also for such n, we have $(n-1)! \geq 2^{n-2}$. Hence, using these bounds,

$$\mathcal{D}_H \leq n2^{-(n-2)/2} + n2^{n/4}2^{-n/2} = 2n(2^{-n/2}) + n2^{-n/4} \leq 3n2^{-n/4} \leq 2^{-n/8}.$$

for sufficiently large n. Otherwise, in the case where n is odd, we have that for $n \geq 7$ using similar estimates that

$$\begin{aligned}\mathcal{D}_H &\leq \frac{n}{\sqrt{(n-1)!}} + n\left(\frac{2^{(n-1)/2}((n-1)/2)!}{n!}\right)^{1/2} \leq 2n2^{-n/2} + n2^{(n-1)/4}2^{-(n+1)/2} \\ &= 2n2^{-n/2} + n2^{-(n+3)/4} \leq 3n2^{-n/4} \leq 2^{-n/8}.\end{aligned}$$

Hence in either case, $\mathcal{D}_H < \frac{1}{poly(n)}$ for sufficiently large n. □

Note that Theorem 10 can also be used to give a proof of Theorem 7 in a similar way.

4 Conclusion and Open Problems

In this paper, we have analyzed the weak Fourier sampling algorithm and proved that it may not give an efficient algorithm for the hidden subgroup problem for subgroups $H \leq S_n$ of polynomial size. A natural question is to exactly characterize the subgroups of S_n that are distinguishable. Towards this, [KS05] provides a characterization of families of all polynomial sized subgroups of S_n that are distinguishable and shows that families of subgroups of size $|H| \geq (n!)^{\epsilon}$ with $\epsilon > 0$ are distinguishable using certain results derived from the classification of finite simple groups. It would be interesting to find more elementary proofs and find exponential size groups which are distinguishable or indistinguishable. We believe that the group $(\mathbb{Z}/2\mathbb{Z})^l$ realized as a subgroup of S_{2l} is a candidate for such an indistinguishable group but proving this may require additional analysis of the characters of S_n along the lines of Section 3.2.

We mentioned the strong Fourier sampling algorithm briefly in Section 2.3. It was shown in [MRRS07] that strong Fourier sampling can solve hidden subgroup problems for groups of type $\mathbb{Z}/p\mathbb{Z} \ltimes \mathbb{Z}/q\mathbb{Z}$ with $p|q-1$, where $\ltimes$ is the semidirect product of two groups, whereas weak Fourier sampling cannot. (Note that the dihedral groups D_n are isomorphic to $\mathbb{Z}_2 \ltimes \mathbb{Z}_n$ so these groups are in general not Abelian). However, in [MRS08], it is shown that strong Fourier sampling fails to solve the rigid graph automorphism problem. Furthermore, if we allow the Fourier sampling algorithm on k independent registers and make entangled measurements on them, it is shown in [HMR$^+$10] that an exponential number of measurements are needed so as long as $k = o(n \log n)$. Surveying these results could be useful to prove further lower bounds on the usefulness of Fourier sampling for different hidden subgroup problems. Indeed Dinh et al. in [DMR10] have derived some criteria for when a subgroup of S_n or $GL_2(\mathbb{Z}/p\mathbb{Z})$ for p a prime is distinguishable by strong Fourier sampling due to the relevance of these groups in the constructing McEliece-type cryptosystems.

Finally, although Fourier sampling has had great successes in developing quantum algorithms, it has also been shown to be a dead end for these instances of the hidden subgroup problem over S_n. However, the celebrated result that these problems have polynomial query complexity in [EHK04] should give us some hope that a polynomial time quantum algorithm is possible. Innovative ideas in quantum algorithms will ultimately be required for these tasks.

References

[Bea97] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 48–53. ACM, 1997.

[DF04] D.S. Dummit and R.M. Foote. *Abstract Algebra.* Wiley, 2004.

[DMR10] Hang Dinh, Cristopher Moore, and Alexander Russell. The McEliece cryptosystem resists Quantum Fourier Sampling attacks. *CoRR*, abs/1008.2390, 2010.

[EHK04] Mark Ettinger, Peter Høyer, and Emanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Inf. Process. Lett.*, 91(1):43–48, 2004.

[ER63] Paul Erdős and Alfréd Rényi. Asymmetric graphs. *Acta Mathematica Academiae Scientiarum Hungarica*, 14(3-4):295–315, 1963.

[HMR$^+$10] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. *J. ACM*, 57(6):34:1–34:33, 2010.

[HRT00] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. Normal subgroup reconstruction and quantum computation using group representations. In F. Frances Yao and Eugene M. Luks, editors, *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 627–635. ACM, 2000.

[Joz01] Richard Jozsa. Quantum factoring, discrete logarithms, and the hidden subgroup problem. *Computing in Science and Engineering*, 3(2):34–43, 2001.

[KS05] Julia Kempe and Aner Shalev. The hidden subgroup problem and permutation group theory. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*, pages 1118–1125. SIAM, 2005.

[MRRS07] Cristopher Moore, Daniel N. Rockmore, Alexander Russell, and Leonard J. Schulman. The power of strong Fourier sampling: Quantum algorithms for affine groups and hidden shifts. *SIAM J. Comput.*, 37(3):938–958, 2007.

[MRS08] Cristopher Moore, Alexander Russell, and Leonard J. Schulman. The symmetric group defies strong Fourier sampling. *SIAM J. Comput.*, 37(6):1842–1864, 2008.

[Sag13] B.E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions.* Graduate Texts in Mathematics. Springer New York, 2013.

[SL17] O. Shehab and S. J. Lomonaco, Jr. Quantum Fourier Sampling is Guaranteed to Fail to Compute Automorphism Groups of Easy Graphs. *ArXiv e-prints*, May 2017.

[SS96] L.L. Scott and J.P. Serre. *Linear Representations of Finite Groups.* Graduate Texts in Mathematics. Springer New York, 1996.

[Sta12] R.P. Stanley. *Enumerative Combinatorics:.* Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2012.

[Wan10] Frédéric Wang. The hidden subgroup problem. *ArXiv e-prints*, 2010.