0.1 Non Abelian Hidden Subgroup

The input to the hidden subgroup problem is a function $f : G \to \mathbb{C}$ that is constant on all cosets of a subgroup $H \subset G$ and takes distinct values on different cosets. Given oracle access to f, the task is to output a set of generators for H. We recall the steps to solve the abelian hidden subgroup problem using the quantum Fourier transform:

- Construct a random coset state $|gH\rangle = \sqrt{\frac{1}{|H|}} \sum_{h \in H} |gH\rangle$.
- Apply the Fourier transform to map the coset state $|gH\rangle \to \sqrt{\frac{1}{|H^{\perp}|}} \sum_{h \in H^{\perp}} \chi_g(h) |h\rangle$. The QFT maps all coset states to superpositions over the dual group $H^{\perp} := \{k \in \widehat{G} \mid \chi_k(h) = 1 \, \forall \, h \in H\}$ differing only in phase.
- Measure to obtain a random element of $H^{\perp}$. The linear constraints corresponding to elements of $H^{\perp}$ determine H. After sufficiently many measurements a spanning set for $H^{\perp}$ is known, and hence H can be reconstructed.

This algorithm works for abelian groups and it is natural to ask is whether quantum algorithms can offer similar speedups when G is non abelian. Solving the non abelian hidden subgroup problem on the symmetric group S_n yields a solution to the graph isomorphism problem. One way to see this is to observe that the automorphism group of graph G is the hidden subgroup for the function $f(P) := P^{-1} A_G P$ where A_G is the adjacency matrix and P is a permutation matrix. Graphs G and H are isomorphic if and only if there is a generator of $Aut(H \cup G)$ that maps vertices of G to vertices of H.

There are strong negative results showing that it is unlikely that the Fourier sampling approach can be used to design an efficient quantum algorithm solving the non abelian hidden subgroup problem. However, we begin with a positive result by Ettinger, Hoyer and Knill showing that the quantum query complexity of the hidden subgroup problem is polynomial.

Theorem 12.1: *There is a quantum algorithm that outputs generators for a subgroup H of G with probability $1 - 2^{-log|G|^2}$ making $O(\log |G|^4)$ queries to a quantum oracle f known to be H-periodic.*

Proof: Suppose H is a subgroup of G generated by $\{1, g_1, g_2, \cdots, g_k\}$. Let H_i be the group generated by $\{1, g_1, g_2, \cdots, g_i\}$. $|H_{i+1}| \geq 2|H_i|$ as H_{i+1} is the union of cosets of the form $g_{i+1}^k H_i$. It follows that $k \leq \log(|G|)$ yielding an upper bound of $\binom{|G|}{\log |G|} = O(2^{\log(|G|)^2})$ on the total number of subgroups of G.

Let E be an enumeration of the subgroups of G in descending order of size. The algorithm tests whether H is a hidden subgroup for f, for subgroups H ordered according to E. It stops at the first H for which the test accepts. The descending order of size in E is required because a function that is H-periodic is also periodic for a subgroup $H^{'}$ of H.

The first step of the algorithm is the preparation of $m = O(\log(|G|)^4)$ random coset states to obtain the state $|\Phi\rangle = \otimes |g_i H\rangle, 1 \leq i \leq m$. Measurements M_H corresponding to subgroups $H \subset G$ are

performed sequentially on $|\Phi\rangle$ with the outcome of M_H revealing if the hidden subgroup is H. Measurements M_H are such that $|||\Phi\rangle - M_H|\Phi\rangle||_2 \leq 2^{-m}$, so that the state is not disturbed much by the measurement.

The measurement M_H on a single register is described by an orthogonal projection to $(A, A^{\perp})$ where $A := \text{span}(|H\rangle, |g_1 H\rangle, \cdots, |g_k H\rangle)$, where $\{g_1, g_2, \cdots, g_k\}$ is a set of distinct coset representatives for H. If the state being measured is a coset state $|gH\rangle$ then the measurement outcome will be 0. If the state being measured is a coset state for a group $H^{'} \neq H$ then we show that the measurement outcome is 1 with probability at least $1/2$.

The following group theoretic lemma is required which we leave for the reader to prove:

Lemma 12.1: *Let H and $H^{'}$ be two subgroups of G with $H \cap H^{'} = K$. If the intersection of cosets aH and $bH^{'}$ is non empty then $|aH \cap bH^{'}| = |K|$.* The probability of obtaining a vector in A on measuring $|aH^{'}\rangle$ is given by:

$$\sum_{i=1}^{k} |\langle aH^{'} | g_i H\rangle|^2 = \frac{1}{|H||H^{'}|} \sum_{i=1}^{k} |aH^{'} \cap g_i H|^2 = \frac{1}{|H||H^{'}|} \frac{|H^{'}|}{|K|} |K|^2 \leq 1/2 \tag{1}$$

Lemma 0.1 has been used to evaluate the sum. The final inequality holds as K is a proper subgroup of H.

The measurement M_H is performed on all the m registers $|Hg_i\rangle$ with outcome 0 if all registers project to the subspace A and 1 otherwise. If H is the hidden subgroup the outcome is always 0 while if $H^{'}$ is the hidden subgroup the outcome is 0 with probability at most $1/2^m$ by (1). If the answer is correct, the distance between the states $|||M_H\Phi\rangle - |\Phi\rangle||_2 \leq 2^{-m}$. We bound the probability of obtaining a correct answer in the i th measurement made.

$$\Pr[\, M_H(\Phi_i) = 0 \mid H = H_i] \geq \Pr[\, M_H(\Phi) = 0 \mid H = H_i] - |||\Phi_i\rangle - |\Phi\rangle|| \geq 1 - i2^{-m}$$

The success probability for the algorithm is at least $1 - 2^{-log|G|^2}$ if m is taken to be $O(\log |G|^4)$. The error probability is negligible small. □

0.2 Representation Theory

A matrix representation of a group G is a group homomorphism $\rho : G \to GL_d(\mathbb{C})$. A representation is irreducible if there is no invariant subspace $V \subset \mathbb{C}^d$ such that $\rho(V) \subseteq V$. Matrix representations of a finite group can be decomposed into a direct sum of irreducible representations.

The group algebra $\mathbb{C}^G$ is a $|G|$ dimensional vector space over $\mathbb{C}$ generated by $|g\rangle, g \in G$ with multiplication defined by $|g\rangle|h\rangle = |gh\rangle$. The regular representation $\sigma : G \to Aut(\mathbb{C}^G)$ maps group elements to operators corresponding to left multiplication:

$$\sigma_g |h\rangle = |gh\rangle$$

The regular representation decomposes into a direct sum of d_ρ copies of irreducible representation ρ with ρ ranging over the set of non isomorphic irreducible representations of G. It follows that the number of non isomorphic irreps of G is finite and satisfies $\sum_\rho d_\rho^2 = |G|$.

The Fourier transform for abelian groups can be viewed as a change of basis over $\mathbb{C}^G$ that diagonalizes the operators σ_g. This is not possible for non abelian groups because diagonal matrices commute. The Fourier transform for non abelian groups is a change of basis over $\mathbb{C}^G$ such that the operators σ_g are block diagonal with the block sizes given by the dimensions d_ρ of the irreps of G. The Fourier transform is given by:

$$|\widehat{g}\rangle = \sum_{\rho,i,j} \sqrt{\frac{d_\rho}{|G|}} \rho_{ij}(g) |\rho, i, j\rangle$$

Assuming that the irreps ρ are unitary the normalization factor $\sqrt{d_\rho/|G|}$ ensures that $|\widehat{g}\rangle$ is a unit vector. Using orthogonality relations for group characters it can be proved that the Fourier transform is a unitary operator and hence its rows and columns are orthogonal.

Unlike the case for Abelian groups, the Fourier transform on non abelian groups depends on the choice of basis used for the irreps. Basis independent sampling reveals the name of the representation ρ and is referred to as weak sampling. Strong Fourier sampling is the measurement of ρ, i, j in a chosen basis. The next result shows that *HSP* on the symmetric group can not be solved using strong Fourier sampling.

Theorem 12.2: *The hidden subgroup problem on the symmetric group can not be solved efficiently by strong Fourier sampling over any basis by an algorithm that is allowed to perform arbitrary measurements on two coset states.* Strong Fourier sampling makes an arbitrary measurement on a single coset state. It is known that an algorithm that solves the *HSP* on the symmetric group efficiently, would have to make measurements on $\Omega(\log(|G|)$ coset states.

0.2.1 Characters and orthogonality

In class it was mentioned that people who study representation theory love it because the most beautiful statements that one could think of turn out to be theorems. We attempt to substantiate this claim by proving some theorems from representation theory.

A representation ρ over vector space V is unitary with respect to the inner product $\langle v | w \rangle_\rho = \sum_g \langle \rho_g v | \rho_g w \rangle$, where $\langle v | w \rangle$ denotes the standard inner product. Changing basis to construct an equivalent representation, wlog we assume that representations of a finite group are unitary.

The trace function satisfies the commutativity property as $Tr(AB) = Tr(BA)$ for matrices A and B. This property makes the trace of a representation a very useful quantity. Combining two representations by taking the direct sum and the tensor product corresponds to addition and multiplication of the trace.

$$Tr(\rho_g \oplus \sigma_g) = Tr(\rho_g) + Tr(\sigma_g), \quad Tr(\rho_g \otimes \sigma_h) = Tr(\rho_g).Tr(\sigma_h)$$

The character of a representation ρ is a function $\chi_\rho : G \to \mathbb{C}$ defined as $\chi_\rho(g) := Tr(\rho_g)$. A character is constant on conjugacy classes as $\chi_\rho(h^{-1}gh) = Tr(\rho_h^{-1}\rho_g\rho_h) = Tr(\rho_h^{-1}\rho_h\rho_g) = Tr(\rho_g) = \chi_\rho(g)$ by the commutativity of the trace. A function constant on conjugacy classes of G is called a class function.

Theorem 12.3: *The characters χ_ρ of distinct irreducible representations of group G form an orthonormal basis for the space of class functions for G.*

Proof: Suppose ρ and σ are distinct irreps acting on spaces A and B such that the characters χ_ρ and χ_σ are not orthogonal. The representation μ on $A \otimes B$ is defined as $\mu_g := \rho_g \otimes \sigma_{g^{-1}}$.

$$\langle \chi_\rho | \chi_\sigma \rangle = \frac{1}{|G|} \sum_{g \in G} \chi_\rho(g) \overline{\chi_\sigma(g)} = \frac{1}{|G|} \sum_{g \in G} Tr(\rho_g).Tr(\sigma_{g^{-1}}) = \frac{1}{|G|} \sum_{g \in G} Tr(\rho_g \otimes \sigma_{g^{-1}}) = \frac{1}{|G|} \sum_{g \in G} Tr(\mu_g) \tag{2}$$

It follows that the linear operator $R := \frac{1}{|G|} \sum_{g \in G} \mu_g$ has non zero trace. Operator R is a projection as $R^2 = \frac{1}{|G|^2} \sum_{h \in G} \sum_{g \in G} \mu_{hg} = R$. A projection operator with non zero trace must have an eigenvector V with eigenvalue 1.

$$||V|| = ||RV|| \leq \frac{1}{|G|} \sum_g \left|\left| \mu_g V \right|\right| = \frac{1}{|G|} \sum_g ||V||$$

The eigenvector $V \in A \otimes B$ is such that $\mu_g V = V$ for all $g \in G$. Decomposing V in the basis $a_i \otimes b_j$ where a_i and b_j are basis vectors for A and B we have

$$\mu_g V = \rho_g \otimes \sigma_{g^{-1}} \sum_{i,j} v_{ij} (a_i \otimes b_j) = |\rho_g\rangle \left(\sum_{i,j} v_{ij} |a_i\rangle \langle b_j|\right) \langle \sigma_{g^{-1}}| = |\rho_g\rangle V \langle \sigma_{g^{-1}}| = V$$

The matrix $V : B \rightarrow A$ commutes with the action of irreps. ρ and σ as $\rho_g V = V \sigma_g$ for all $g \in G$. This implies that the subspaces $\text{Ker}(V) := \{b \in B \mid Vb = 0\}$ and $\text{Im}(V) := \{a \in A \mid a = Vb\}$ are invariant under the action of the representations ρ and σ respectively. The representations do not have non trivial invariant subspaces as they are irreducible and V is assumed to be non zero, so $\text{Ker}(V) = 0$ and $\text{Im}(V) = B$. The matrix V is an isomorphism between ρ and σ contradicting the assumption that they are distinct irreps. □

The argument in the last part of the proof of the above theorem is known as Schur's lemma in the literature.

Lemma 12.2: *Suppose V is a matrix that commutes with irreps. ρ and σ of group G so that $\rho_g V = V \sigma_g$ for all $g \in G$. i) If $\rho = \sigma$ then $V = cI$. ii) If ρ and σ are non isomorphic then $V = 0$.*

Part i) of the lemma follows as an eigenspace for V is invariant under the action of ρ. We use the Schur lemma to prove that the Fourier transform is a unitary operator. This result is known as the great orthogonality theorem.

Theorem 12.4: *The following orthogonality relations hold for unitary irreps. ρ and σ for group G.*

$$\sum_g \rho_g(ij).\overline{\sigma_g(lk)} = \frac{|G|}{d_\rho} \delta_{(\rho,i,j),(\sigma,l,k)}$$

Proof: The matrix $M = \sum_g \rho_g X \sigma_{g^{-1}}$ where X is an arbitrary $d_\rho \times d_\sigma$ matrix commutes with the ρ and σ. It can be easily verified that $\rho_g M = M \sigma_g$ for all $g \in G$. Applying the two parts of Schur's lemma we have:
i) $\rho = \sigma$: If $V \neq 0$ then $Tr(M) = cd_\rho = |G| Tr(X)$, so c must be equal to $\frac{|G| Tr(X)}{d_\rho}$. Choosing $X = |x_j\rangle \langle x_k|$ we note that

$$M_{il} = \sum_g \rho_g(ij).\rho_{g^{-1}}(kl) = \sum_g \rho_g(ij).\overline{\rho_g(lk)} = \frac{|G|}{d_\rho} \delta_{il} Tr(X) = \frac{|G|}{d_\rho} \delta_{il} \delta_{jk}$$

ii) $\rho \neq \sigma$: Choosing $X = |x_j\rangle\langle x_k|$ we have $M_{il} = \sum_g \rho_g(ij).\sigma_{g^{-1}}(kl) = \sum_g \rho_g(ij).\overline{\sigma_g(lk)} = 0$. □

Using the orthogonality theorem for characters, it is a good exercise to prove that the regular representation decomposes as $\mathbb{C}^G = \oplus \rho^{d_\rho}$.

0.3 HSP on the dihedral group

We have seen that it is unlikely that the hidden subgroup problem for highly non abelian groups like S_n is in *BQP*. This leaves the case of simpler non abelian groups on which it might be possible to solve the hidden subgroup problem. The dihedral group D_n is the group of symmetries of a regular n-gon. It is generated by the rotation element x and the reflection element y satisfying the relations:

$$D_n := < \ x,y \mid \ x^n = 1, \ \ y^2 = 1, \ \ yxyx = 1 \ >$$

The dihedral group D_{2n} has 4 one dimensional irreps. and $(n-1)$ two dimensional irreps. The two dimensional irreducible representations are given by:

$$\rho_j(x) = \begin{bmatrix} \omega^{jl} & 0 \\ 0 & \omega^{-jl} \end{bmatrix}, \quad \rho_j(y) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \ \ 1 \leq j \leq (n-1) \tag{3}$$

The hidden subgroup problem on the dihedral group reduces to the problem of finding a hidden reflection $H = \{1, x^h y\}, 0 \leq h < n$. The hidden subgroup on the cyclic part of D_{2n} can be found classically once the factors of n are known. We discuss Kuperberg's algorithm to solve the HSP on D_{2n} in time $O(2^{\sqrt{n}})$.

We assume that n is a power of 2. The algorithm works by finding the parity of h, thus reducing to the problem to a *HSP* on $D_{n/2}$, which is solved recursively.

A coset state $\frac{1}{\sqrt{2}}(|x^k\rangle + |x^{h+k}y\rangle)$ is prepared and the Fourier transform on D_n is applied to it followed by a measurement revealing the irrep. ρ to obtain the state:

$$\rho_j(x^k) + \rho_j(x^{h+k}y) = \begin{bmatrix} \omega^{jk} & \omega^{j(h+k)} \\ \omega^{-j(h+k)} & \omega^{-jk} \end{bmatrix}$$

The row of the matrix ρ_j (the qubit in register i) is measured to obtain the quantum state $|\Psi_j\rangle = |0\rangle + \omega^{jh}|1\rangle$, where j is distributed uniformly in the range $[1, n-1]$.

The parity of h can be determined by measuring the state $|0\rangle + (-1)^h|1\rangle$ in the Hadamard basis. The sieve algorithm attempts to construct the state $|0\rangle + (-1)^h|1\rangle$ starting from a list L_0 of $2^{O(\sqrt{n})}$ copies of states $|0\rangle + \omega^{jh}|1\rangle$ prepared through Fourier sampling. The algorithm has a parameter m and performs the following steps to determine the parity of h.

- The list L_i is split into pairs $|\Psi_j\rangle, |\Psi_k\rangle$ such that j and k agree in the m non zero least significant bits (apart from mi trailing zeroes). In the last iteration, agreement is required on all the $n-1$ bits except the most significant bit.
- A measurement projecting onto the space $(|00\rangle, |11\rangle)$ is performed on the joint state $|\Psi_j\rangle \otimes |\Psi_k\rangle$.

$$(|0\rangle + \omega^{jh}|1\rangle) \otimes (|0\rangle + \omega^{kh}|1\rangle) = (|00\rangle + \omega^{(j+k)h}|11\rangle) + \omega^{kh}(|01\rangle + \omega^{(j-k)h}|10\rangle)$$

The measurement yields one of the states $|\Psi_{j+k}\rangle, |\Psi_{j-k}\rangle$ with probability $1/2$.

- The list L_i is obtained from L_{i-1} by retaining states of the form $|\Psi_{j-k}\rangle$ obtained in step 2. If the state $|\Psi_k\rangle \in L_i$ then k ends with a string of at least mi zeroes. The algorithm repeats for $t = (n-1)/m$ iterations.
- The states in the list L_t are either $|\Psi_0\rangle$ or $|\Psi_{2^{N-1}}\rangle$. Measuring the state $|\Psi_{2^{N-1}}\rangle$ in the Hadamard basis gives the parity of h.

The algorithm correctly finds the parity if the final list L_t is non empty. At most 2^m elements remain unmatched in step 1 and the list size for matched elements decreases by a factor of 4.

$$L_{i+1} \geq \frac{L_i - 2^m}{4}$$

This gives a bound of $O(2^{m+2t})$ on the size of L_0 for the algorithm to correctly compute the parity of s. The minimum value of $m + 2(n-1)/m$ equals $3\sqrt{n}$ for $m = \sqrt{n}$, showing that the sieve algorithm runs in time $2^{3\sqrt{n}}$.

0.3.1 Reductions

There are two reductions due to Regev relating the hidden subgroup problem on the dihedral group to the subset sum problem and the unique SVP problem on the lattices. We sketch out these reductions making a number of simplifying assumptions.

- Reconstruction of h can be done if the subset sum problem can be solved classically.
- The unique SVP problem on lattices can be solved if reconstruction of h is possible.

There is an efficient quantum algorithm to reconstruct h that uses an efficient classical algorithm for solving the subset sum problem as a subroutine. Fourier sampling on $r = O(\log |n|)$ copies of the coset state, yields the state $\otimes^r |0\rangle + \omega^{j_i h} |1\rangle$. The elements of $J := \{j_1, j_2, \cdots, j_r\}$ are distributed uniformly in $[1, n-1]$, so the distribution of the subset sums will also be close to uniform. We make the simplifying assumption that every element of $[n]$ has a unique representation as a subset sum of J.

$$\begin{aligned} \otimes^r \left(|0\rangle + \omega^{j_i h} |1\rangle \right) &= \sum_{S \subseteq J} \omega^{\mathrm{Sum}(S)h} |S\rangle \rightarrow \sum_{S \subseteq J} \omega^{\mathrm{Sum}(S)h} |S\rangle |\mathrm{Sum}(S)\rangle \\ &\rightarrow \sum_{S \subseteq J} \omega^{\mathrm{Sum}(S)h} |0\rangle |\mathrm{Sum}(S)\rangle = \sum_k \omega^{kh} |k\rangle \rightarrow^{FT} |h\rangle \end{aligned} \tag{4}$$

Since subsets are identified uniquely by their sums and we have a classical algorithm to solve the subset sum problem, the state $|S\rangle |\mathrm{Sum}(S)\rangle$ can be uncomputed to obtain $|0\rangle |\mathrm{Sum}(S)\rangle$. The actual analysis must handle many more details in the absence of a simplifying assumption.

We next sketch a reduction of the unique SVP on lattices to the problem of reconstructing h. The input to the unique SVP problem is a basis $B = \{b_1, b_2, \cdots, b_n\}$ for a lattice Λ which is known to have a unique shortest vector v such that $|w| \geq \sqrt{n}|v|$ for all $w \in \Lambda$ not parallel to v. The output must be the vector v.

The algorithm guesses a length for the unique shortest vector and generates a superposition over states $|\sum_i a_i b_i\rangle$ where the integers a_i are selected from a sufficiently large range. The space is partitioned into cubes of side length equal to the guess for the length of the shortest vector. These cubes can not have more than one lattice point in a direction other than $\vec{v}$ because of the uniqueness constraint. The lattice is scaled so that not more than two points in the direction of $\vec{v}$ can occur in a cube. A measurement revealing the cube generates with high probability a state of the form $|x\rangle + |x + cv\rangle$ for a suitable constant c and a random vector x. Making several such measurements the direction $\vec{v}$ can be recovered using the *HSP* solver on the dihedral group.