Quantum algorithms	**Monsoon 2007**
Lecture 9	
Lecturer: Pranab Sen	*Date: August 27*

In the last lecture we saw an instance of finding a strict period in the the abelian group $\mathbb{Z}_2^n$ viz. Simon's problem. In this and the next lecture, we generalise Simon's problem and its solution to finding strict periods in any finite abelian group G, also known as the *hidden subgroup problem* in G. For that, we will have to first develop a little bit of the theory of characters of finite abelian groups.

1 The hidden subgroup problem

Definition 1.1 (Hidden subroup problem (HSP)). *Let G be a group, S a set and $f : G \to S$ a function. We are given an oracle for a reversible version of f, $F(f) : |x\rangle|s\rangle \mapsto |x\rangle|s \oplus f(x)\rangle$, where the first and second registers denote elements of G and S respectively, and the operation $\oplus$ is a bitwise XOR of binary strings. The function f satisfies the promise that there exists a subgroup $H \leq G$, called the* hidden subgroup, *such that f is constant on the left cosets of H and distinct on distinct cosets. The aim is to find a generating set for H by making queries to $F(f)$. Ideally, we would require the total running time of the algorithm to be* $\mathrm{polylog}(|G|)$.

The hidden subgroup problem is nothing but finding the subgroup H of periods of a function $f : G \to S$ under the promise that f is strictly periodic, that is, for all $x, y \in G$, $f(x) = f(y)$ iff $y = xh$ for some $h \in H$. The 'if' direction is the periodic part of the definition (it may be clearer to see this for abelian G with the group law written additively), and the 'only if' direction is the strict part of the definition (which in the case of $G = \mathbb{Z}_n$ is equivalent to saying that f is one-one within a periodic interval). Note that the set of periods of a function f defined over a group G always forms a subgroup $H \leq G$.

Many important computational problems reduce to instances of the HSP. For example, integer factoring reduces to HSP in $\mathbb{Z}$, discrete logarithm over $\mathbb{Z}_p$ reduces to HSP in $\mathbb{Z}_{p-1} \times \mathbb{Z}_{p-1}$ and isomorphism of two n-vertex graphs reduces to HSP in S_{2n}, the group of permutations of $2n$ letters. The classical randomised query complexity of finding H is typically $|G|^{\Omega(1)}$, even for simple cases like $G = \mathbb{Z}_n$, $G = \mathbb{Z}_{2^k}$ etc. The lower bound argument for Simon's problem typically generalises to many groups G. On the other hand, the quantum query complexity of this problem is $O(\log |G|)$ for any G, H as shown by Ettinger, Høyer and Knill. Moreover, for the special case of abelian G, we can get a quantum algorithm with total running time $O(\log^3 |G|)$. This fact lies at the heart of most problems where quantum computation seems to offer superpolynomial speedup over classical computation, including Shor's famous quantum algorithms for integer factoring and discrete logarithm. A notable exception where quantum computation seemingly gives superpolynomial speedups over classical computation that does not fit into the hidden subgroup paradigm is a class of quantum algorithms based on knot theory, discovered very recently.

For the rest of the lecture, we restrict ourselves to finite groups. For a subset $X \subseteq G$, we define $|X\rangle := |X|^{-1/2} \sum_{x \in X} |x\rangle$, the uniform superposition over elements of X. The main approach to solving HSP is the so-called *coset state* approach. In this approach, we start off by preparing the coset state σ_H^G

of the hidden subgroup H in the ambient group G defined as $\sigma_H^G = \sum_{x \in G/H} |xH\rangle\langle xH|$, where G/H is a set of representatives of left cosets of H in G. If we can generate the uniform superposition $|G\rangle$ over all elements of G, which can be done efficiently for many groups of interest, then it is easy to prepare the coset state σ_H^G as follows:

$$|G|^{-1/2} \sum_{g \in G} |g\rangle_G |0\rangle_S \stackrel{F(f)}{\longmapsto} |G|^{-1/2} \sum_{g \in G} |g\rangle_G |f(g)\rangle_S \stackrel{\mathrm{Tr_S}}{\longmapsto} \sigma_H^G.$$

We can prepare several independent copies of σ_H^G and then perform a, in general, joint measurement on them in order to discover H. Indeed this was precisely the approach of Ettinger, Høyer and Knill for solving the HSP with $\log |G|$ independent copies of σ_H^G. Unfortunately, it is not known how to implement their algorithm efficiently. However for abelian G, it turns out that we can efficiently perform a certain measurement on a *single* copy of σ_H^G and repeat the procedure $O(\log |G|)$ times in order to discover H with high probability. In the next section, we shall try to understand some properties of this special measurement.

2 Characters of finite abelian groups

We will use additive notation for abelian groups. Let $x+H$ be a particular coset in the abelian group G of the hidden subgroup H. If we could somehow ensure that x were 0, then just by measuring in the computational basis, we would have sampled a random element of H. By repeating this procedure $O(\log |G|)$ times, with high probability, we would have got a generating set for H. The problem with this is that all we have is a uniformly random element x of G/H, and the probability of x being 0 is typically very small. Thus, we would like to find a way to 'forget' x, and generate a probability distribution that depends only on H. The abelianness of G implies the existence of such a procedure, as explained below.

Let $\mathbb{C}[G]$ denote the vector space of functions from G to $\mathbb{C}$ under pointwise addition and pointwise scalar multiplication. In the literature, $\mathbb{C}[G]$ is sometimes referred to as the *group algebra* of G over $\mathbb{C}$, since there is a natural multiplication of elements of $\mathbb{C}[G]$, sometimes called *convolution*, viewing elements of $\mathbb{C}[G]$ as formal linear combinations of elements of G. In addition, we can define an inner product of any pair of elements $a, b \in \mathbb{C}[G]$ by $\langle a|b\rangle := \sum_{g \in G} \overline{a(g)} b(g)$, which makes $\mathbb{C}[G]$ into an inner product space. The *Dirac point mass* basis, which will also be our computational basis, is defined as $|g\rangle : x \mapsto \delta_{g,x}$ for every $g \in G$. It is easy to see that $\{|g\rangle\}_{g \in G}$ is an orthonormal basis for $\mathbb{C}[G]$. Thus, elements of $\mathbb{C}[G]$ can be naturally thought of as $|G|$-tuples.

For every $x \in G$, define the unitary operator $T_x : |y\rangle \mapsto |x + y\rangle$ on $\mathbb{C}[G]$; T_x translates $\mathbb{C}[G]$ by x. Since G is abelian, the operators $\{T_x\}_{x \in G}$ commute. Thus, they have a common orthonormal eigenbasis for $\mathbb{C}[G]$ that we call $\{\chi_1, \ldots, \chi_{|G|}\}$ for now. For $x \in G$ and $i \in [n]$, let $\lambda_{x,i}$ be the eigenvalue of T_x for eigenvector χ_i; $|\lambda_{x,i}| = 1$ since T_x is unitary. The uniform superposition over H can be expressed in terms of the common eigenbasis as $|H\rangle = \sum_{i=1}^{|G|} \alpha_{H,i} \chi_i$. Then the uniform superposition over a coset can be expressed as

$$|x + H\rangle = T_x |H\rangle = \sum_{i=1}^{|G|} \lambda_{H,i} \alpha_{H,i} \chi_i.$$

If we now measure $|x + H\rangle$ in the common eigenbasis, then

$$\Pr[\text{“}i''] = |\lambda_{x,i}|^2 |\alpha_{H,i}|^2 = |\alpha_{H,i}|^2,$$

which depends only on H and not on the particular coset $x + H$. We will see in the next lecture that $O(\log |G|)$ iterations of producing a fresh independent coset state σ_H^G and measuring in the common eigenbasis gives sufficient information to recover H efficiently.

We now proceed to understand what this common eigenbasis really is; it will turn out that they are the so-called *characters* of G. A character of the finite abelian group G is a homomorphism $\chi : G \to S^1$ from G to the multiplicative group S^1 of unit absolute value complex numbers. Note that the group law for G is written additively whereas the group law for S^1 is written multiplicatively; thus, $\chi(x + y) = \chi(x)\chi(y)$ for all $x, y \in G$ and $\chi(0) = 1$ for any character χ of G. This implies that the image of $\chi(x)$ for any $x \in G$ is a $|G|$th root of unity. The characters of G form an abelian group under pointwise multiplication denoted by $\widehat{G}$. The identity element of $\widehat{G}$ is the *trivial* or *identity character* $\mathbb{1} : x \mapsto 1$ for all $x \in G$. The characters of an abelian group G are in fact the so-called *one-dimensional representations* of G.

We now list several properties of characters of abelian groups.

Lemma 2.1. $\widehat{\mathbb{Z}_n} \cong \mathbb{Z}_n$.

Proof: We write $\mathbb{Z}_n$ additively. Let $x \in \mathbb{Z}_n$. Define $\chi_x \in \widehat{\mathbb{Z}_n}$ by $\chi_x(y) := \exp(2\pi i x y/n)$ for all $y \in \mathbb{Z}_n$. Note that $\chi_x = (\chi_1)^x$. Finally, since the values taken by any character $\chi \in \widehat{\mathbb{Z}_n}$ are nth roots of unity, and because χ is determined by the value of $\chi(1) = \exp(2\pi i x'/n)$, we have $\chi = \chi_{x'}$. This completes the proof of the lemma. □

Proposition 2.1. *For groups G_1, G_2, $\widehat{G_1 \times G_2} \cong \widehat{G_1} \times \widehat{G_2}$.*

Proof: Let $\chi \in \widehat{G_1}$, $\tau \in \widehat{G_2}$. Define the character $(\chi, \tau) \in \widehat{G_1 \times G_2}$ by $(\chi, \tau)(x_1, x_2) := \chi(x_1)\tau(x_2)$ for all $x_1 \in G_1$, $x_2 \in G_2$. It is easy to verify that this gives an homomorphism from $\widehat{G_1} \times \widehat{G_2}$ into $\widehat{G_1 \times G_2}$. The homomorphism is injective since its kernel is trivial viz. $(\chi, \tau) = \mathbb{1}$ iff $\chi = \mathbb{1}$ and $\tau = \mathbb{1}$ as is easy to check. To show that it is also surjective, consider $\omega \in \widehat{G_1 \times G_2}$. Define characters $\chi \in \widehat{G_1}$ and $\tau \in \widehat{G_2}$ by $\chi(x_1) := \omega(x_1, 1)$ and $\tau(x_2) := \omega(1, x_2)$ for all $x_1 \in G_1$ and $x_2 \in G_2$. It is easy to see that $(\chi, \tau) = \omega$. This completes the proof of the lemma. □

Proposition 2.2. *For any finite abelian group G, $\widehat{G} \cong G$.*

Proof: Since every finite abelian group is isomorphic to a direct product of cyclic groups, we can use Lemma 2.1 and Proposition 2.1 to derive the result. □

Remarks:

1. The isomorphism in Proposition 2.2 is not 'natural', in the sense that it depends on the cyclic decomposition of G as well as on the choice of a primitive nth root of unity for a cyclic factor $\mathbb{Z}_n$.
2. Suppose $G \cong \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_k}$ is the chosen cyclic decomposition of G in Proposition 2.2. For any $x \in G$, let χ_x denote its corresponding character under the isomorphism of Proposition 2.2. Let us represent elements of G as k-tuples according to this decomposition. Then for any $x := (x_1, \ldots, x_k), y := (y_1, \ldots, y_k) \in G$, $\chi_x(y) = \exp(2\pi i (x \cdot y)/|G|)$, where $x \cdot y := \sum_{j=1}^{k} \frac{|G|}{n_j} x_j y_j$ is the 'dot product' of x and y according to G.
3. The isomorphism of Proposition 2.2 has the additional property that $\chi_x(y) = \chi_y(x)$ for all $x, y \in G$.

Lemma 2.2. *For groups G_1, G_2, $\mathbb{C}[G_1 \times G_2] \cong \mathbb{C}[G_1] \otimes \mathbb{C}[G_2]$.*

Proof: The map $(x_1, x_2) \mapsto |x_1\rangle \otimes |x_2\rangle$, $x_1 \in G_1, x_2 \in G_2$ extended by $\mathbb{C}$-linearity defines an isomorphism from $\mathbb{C}[G_1 \times G_2]$ to $\mathbb{C}[G_1] \otimes \mathbb{C}[G_2]$. Moreover, it is inner product preserving. □

For a character χ of a finite abelian group G, define the unit length vector $|\chi\rangle := |G|^{-1/2} \sum_{y \in G} \chi(y)|y\rangle$.

Proposition 2.3. *For any finite abelian group G, $\{|\chi\rangle\}_{\chi \in \widehat{G}}$ forms an orthonormal basis for $\mathbb{C}[G]$.*

Proof: It is easy to verify for $G = \mathbb{Z}_n$ that $\langle \chi_x | \chi_{x'} \rangle = \delta_{x,x'}$ for $x, x' \in \mathbb{Z}_n$. This shows that $\{|\chi_x\rangle\}_{x \in \mathbb{Z}_n}$ forms an orthonormal basis of $\mathbb{C}[\mathbb{Z}_n]$. Since every finite abelian group is isomorphic to a direct product of cyclic groups, we can use Lemmas 2.1 and 2.2, and Proposition 2.2 to derive the result. □

By Proposition 2.3, for a finite abelian group G, the change of basis from Dirac point mass to the scaled characters is a unitary transformation of $\mathbb{C}[G]$ called the *quantum Fourier transform* over G, defined by $\mathrm{QFT}_G : |x\rangle \mapsto |\chi_x\rangle$ for all $x \in G$. The vectors $|\chi_x\rangle$ are also called the Fourier basis vectors of $\mathbb{C}[G]$.

Proposition 2.4. *For finite abelian G_1, G_2, $\mathrm{QFT}_{G_1 \times G_2} = \mathrm{QFT}_{G_1} \otimes \mathrm{QFT}_{G_2}$.*

Proof: Follows from Proposition 2.1 and the definition of QFT. □

In view of Proposition 2.4, in order to design circuits for QFT_G for finite abelian G, it suffices to build circuits for $\mathrm{QFT}_{\mathbb{Z}_n}$ for any integer $n \geq 2$. We shall see how to do this in subsequent lectures.

The next proposition combined with the arguments at the beginning of this section indicates why the Fourier basis is useful for solving the HSP in abelian groups.

Proposition 2.5. *For a finite abelian group G, $\{|\chi\rangle\}_{\chi \in \widehat{G}}$ form a common eigenbasis for the unitary operators $\{T_x\}_{x \in G}$, the translations by elements of G.*

Proof: It is easy to see that $T_x|\chi_{x'}\rangle = \overline{\chi_x(x')}|\chi_{x'}\rangle$ for all $x, x' \in G$. The result now follows from Proposition 2.3. □

For a subgroup $H \leq G$, define its *orthogonal subgroup* $H^\perp := \{y \in G : \chi_y(h) = 1 \; \forall h \in H\}$.

Proposition 2.6. *For any subgroup H of a finite abelian group G, $|H^\perp| = |G|/|H|$.*

Proof: Let $y \in G$. Then,

$$\sum_{h \in H} \chi_y(h) = \begin{cases} |H| & \text{if } y \in H^\perp \\ 0 & \text{otherwise} \end{cases},$$

where the second equality follows because in that case χ_y restricted to H is a non-trivial character of H which must be orthogonal to the trivial character of H. Thus,

$$|H| \cdot |H^\perp| = \sum_{y \in G} \sum_{h \in H} \chi_y(h) = \sum_{h \in H} \sum_{y \in G} \chi_h(y) = |G|,$$

where the third equality follows because χ_h is a non-trivial character of G if $h \neq 0$ by Proposition 2.2, and so χ_h must be orthogonal to the trivial character of G. This completes the proof of the proposition. □

Proposition 2.7. *For any subgroup H of a finite abelian group G, $(H^\perp)^\perp = H$.*

Proof: It is easy to see that $H \leq (H^\perp)^\perp$. The equality follows because $|(H^\perp)^\perp| = |H|$ by Proposition 2.6. □