Quantum Algorithms 4/6/2023

Lecture 20: Fourier Sampling

Instructor: Dieter van Melkebeek

Last lecture we discussed the problem of finding a hidden XOR-shift, presented an efficient quantum algorithm and cast the problem as an instantiation of the hidden subgroup problem. This lecture we view the algorithm as a special case of Fourier sampling and show that the approach efficiently solves the hidden subgroup problem over finite Abelian groups modulo the existence of an efficient quantum Fourier transform. Before doing so, we review the classical Fourier transform and generalize it to finite Abelian and other groups. An efficient quantum Fourier transform for finite Abelian groups will be covered next lecture.

1 Exercise #11

We begin with the solution to the exercise from the last lecture, which posed the following question: Given two one-to-one functions $f_0, f_1 : \{0,1\}^n \to \{0,1\}^n$ where $f_1(x) = f_0(x \oplus s)$ for some $s \in \{0,1\}^n$. Find s, with certainty, using $O(n)$ queries to the unitary multiplexer $V : |b\rangle |\psi\rangle \mapsto |b\rangle_{f_b} |\psi\rangle$.

The problem is closely related to finding a hidden XOR-shift and, in fact, reduces to it. Consider the function $f : \{0,1\}^{n+1} \to \{0,1\}^{n+1} : bx \mapsto f_b(x)0$, where b denotes a bit that selects between the two functions f_0 and f_1. Note that $f(bx_1) = f(bx_2)$ iff $x1 = x_2$ because both f_0 and f_1 are one-to-one. We have that

$$f(bx_1) = f(\bar{b}x_2) \Leftrightarrow f_b(x_1) = f_{\bar{b}}(x_2) \Leftrightarrow f_b(x_1) = f_b(x_2 \oplus s) \Leftrightarrow x_1 \oplus x_2 = s.$$

It follows that

$$f(b_1x_1) = f(b_2x_2) \Leftrightarrow b_1x_1 \oplus b_2x_2 \in \{0^{n+1}, 1s\}.$$

Thus, the function f satisfies the promise of the hidden XOR-shift problem for $n+1$ bits with the string $1s$ as the hidden shift. Note that U_f coincides with $V \otimes I$, where the I accounts for the extra output bit that f and that is always 0. Thus, we can run our algorithm for the finding a hidden XOR-shift and find s with certainty using $O(n)$ queries to V.

2 Standard Fourier Transform

Recall the standard Fourier transform for functions over the reals.

Definition 1. *Let $f : \mathbb{R} \to \mathbb{C}$ be such that $\int_x |f(x)|^2\, dx < \infty$. The Fourier transform f is a function $\hat{f} : \mathbb{R} \to \mathbb{C}$ such that $\hat{f}(\omega) = \int_x f(x) e^{2\pi i \omega x}\, dx$ for all $\omega \in \mathbb{R}$.*

The argument x of f often represents time. The argument ω of $\hat{f}$ is often referred to as a frequency.

Properties. The Fourier transform is a linear transformation: $\widehat{af + bg} = a\hat{f} + b\hat{g}$. The Fourier tranform is also unitary; we can see this property using several equivalent definitions of unitarity. A unitary transformation can be viewed as one that preserves inner products. Another definition, which we have used heavily in this class, is that a linear transform is unitary when its inverse is equal to its complex conjugate transpose. Viewing an (invertible) linear map as a basis transformation, yet another definition of unitarity is transformsing an orthonormal basis of its domain into another orthonormal basis. In more detail:

- Consider the a unitary transformation as one which preserves inner products. If we consider the inner product space of functions from $\mathbb{R}$ to $\mathbb{C}$ with inner product $(f, g) = \int_x f(x)\overline{g(x)}\,dx$, then the Fourier transform preserves inner products, i.e., $(\hat{f}, \hat{g}) = (f, g)$, and is hence a unitary transformation.

- The inverse Fourier transform is given by

$$f(x) = \int_\omega \hat{f}(\omega) e^{-2\pi i \omega x}\,d\omega.$$

  As the function $e^{2\pi i\omega x}$ is symmetric in x and ω, and $e^{-2\pi i\omega x}$ is the complex conjugate of $e^{2\pi i\omega x}$, we can see that the inverse Fourier transform is the conjugate conjugate transpose Fourier transform, which is thus unitary.

- The Fourier transform is unitary as it transforms the standard orthonormal basis (consisting of the Dirac delta functions) into an orthonormal basis, referred to as the Fourier basis, which consists of the harmonics

$$e^{2\pi i\omega x} = \cos(2\pi\omega x) + i\sin(2\pi\omega x).$$

Apart from linearity and unitarity, another important property of the Fourier transform is that it turns convolutions into point-wise products. Convolutions are operations that occur naturally in several areas, including signal processing. When a signal is sent through a filter, the modified signal is the convolution of the original signal with a function describing the characteristics of the filter.

Definition 2. *The convolution of $f : \mathbb{R} \to \mathbb{C}$ with $g : \mathbb{R} \to \mathbb{C}$ is $f * g : \mathbb{R} \to \mathbb{C}$ where*

$$(f * g)(x) = \int_y f(x)g(x - y)\,dy.$$

In symbols, the property states that $\widehat{f * g}(\omega) = \hat{f}(\omega)\hat{g}(\omega)$ for all $\omega \in \mathbb{R}$.

Next we discuss the more general form of the Fourier transform including over finite Abelian groups of size N, which is of particular interest in developing quantum algorithms. This form is also the one most commonly used in practical applications, which rely on the $O(N \log N)$ complexity of the fast Fourier transform, an efficient algorithm to compute the discrete Fourier transform over the group $\mathbb{Z}_N, +$. Due to the convolution property of the Fourier transform, the fast Fourier transform makes it possible to perform convolutions in time $O(N \log N)$ which would take $O(N^2)$ if done straightforwardly in the time domain. For this reason, the Fourier transform is heavily used in fields such as digital signal processing, computer vision, and statistics.

3 General Fourier Transform

In order to apply the Fourier Transform to quantum algorithms, we need to generalize it to a transformation of functions whose domain is a more general group. We define a Fourier transform for a general group as a transformation of complex function on the group with the above three properties.

Definition 3. *Let G be a group. A Fourier Transform on G is a transformation on the space of functions $\{f : G \to \mathbb{C}\}$, mapping f to $\hat{f}$, that is linear and unitary, and turns convolutions into point-wise products, i.e., $\widehat{f * g}(x) = \hat{f}(x)\hat{g}(x)$ for $f, g : G \to \mathbb{C}$ and $x \in G$.*

For finite groups the convolution of f and g is given by $(f * g)(x) = \sum_y f(y)g(x - y)$, where the subtraction refers to the inverse of the group operation; the other operations are in $\mathbb{C}$.

A Fourier transform exists for many important groups (for example, $\mathbb{R}$ under addition as above), though not for all groups. We show in this lecture that it is guaranteed to exist for an important class of groups, finite Abelian groups. The Fourier transform is also unique (up to permutations of the basis elements) for this class of groups.

In constructing the Fourier transform for finite Abelian G, the place of the harmonics is taken by the *characters*, which we discuss next.

4 Characters of a Group

The characters of a group G as the homomorphisms from G to the multiplicative group $\mathbb{C}, \cdot$.

Definition 4. *A character of a group G is a mapping $\chi : G \to \mathbb{C}$ such that*

$$\chi(x + y) = \chi(x) \cdot \chi(y) \tag{1}$$

holds for all $x, y \in G$.

Note that "+" on the left-hand side of (1) denotes the operation of G (written additively), whereas "$\cdot$" on the right-hand side denotes multiplication of complex numbers.

4.1 Properties

Two properties of characters of finite groups are of interest to us: Their values being roots of unity and their orthogonality.

Roots of unity. All elements in the range of a character of G are roots of unity and, in particular, $|G|$-th roots of unity.

Fact 1. *$\chi(x)^{|G|} = 1$ for all $x \in G$.*

Over the reals, the only roots of unity (i.e., of 1) are 1 and -1 (for integers k with even powers). However, given some positive integer n, there are n roots of unity in $\mathbb{C}$: specifically $e^{2k\pi i/n}$ for $0 \leq k < n$. Visualizing them in the complex plane, these values form the vertices of a regular n-gon inscribed in the unit circle, with the point 1 as one of the vertices : see Figure 1 for an example with $n = 6$.

We wish to show that $\chi(x)$ is a $|G|$-th root of unity for every character χ of a finite group G.

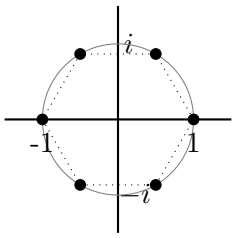

Figure 1: The sixth roots of unity

Proof (of Fact 1). We wish to show that $\chi(x)$ is a $|G|$-th root of unity for every character χ of a finite group G.

First, we show that $\chi(0) = 1$, which follows from the homomorphism property of χ. We have that $\chi(0) = \chi(0+0) = \chi(0)^2$. Since $\chi(0)$, an element of the multaplicative group $\mathbb{C}$, is invertible, we must have $\chi(0) = 1$.

Suppose that $x \in G$. We wish to show that $\chi(x)^{|G|} = 1$, i.e., that x is a $|G|$-th root of unity. Let $\langle x \rangle = \{x, x+x, x+x+x, \dots\} = \{1 \cdot x, 2 \cdot x, 3 \cdot x, \dots\}$ be the subgroup of G generated by x, where we write $n \cdot x$ to represent $x + x + \cdots + x$ when x appears n times.

As $|G|$ is finite, $|\langle x \rangle| \leq |G|$ is finite as well, so we have some positive integer k such that $k \cdot x = 0$. Take the smallest such k, which we call the order of x (and which equals $|\langle x \rangle|$), and consider $\chi(x)^k$.

As χ is a homomorphism, $\chi(x)^k = \chi(k \cdot x) = \chi(0) = 1$. Now, from group theory we have that the order of a subgroup of a finite group divides the size of the group, so k divides $|G|$. Hence, $\chi(x)^{|G|} = 1$. □

Orthogonality. Distinct characters of a a finite group are orthogonal to each other:

Fact 2. *Consider two characters χ, χ' of a finite group G. If $\chi \neq \chi'$, then $(\chi, \chi') = 0$.*

The inner product in Fact 2 is the standard one: For $f, g : G \to \mathbb{C}$ we have

$$(f, g) = \sum_{x \in G} f(x)\overline{g(x)}.$$

Proof (of Fact 2). Suppose that $a \in G$. As a is invertible, we have that $x = y$ if and only if $a + x = a + y$. Now, as G is closed under addition, we have that

$$\sum_{x \in G} \chi(x) = \sum_{a+x \in G} \chi(a+x) = \sum_{x \in G} \chi(a+x) = \sum_{x \in G} \chi(a)\chi(x) = \chi(a) \sum_{x \in G} \chi(x)$$

as χ is a homomorphism.

Hence, we have either that $\sum_{x \in G} \chi(x) = 0$ or $\chi(a) = 1$ for all $a \in G$.

Noting that the conjugate of a root of unity is its inverse, we have, by the property shown above, that $\overline{\chi} = \chi^{-1}$ for all characters χ of G.

Suppose that χ_1, χ_2 are distinct characters of G. Now, let $\chi = \chi_1 \cdot \overline{\chi_2}$. As the conjugate of a character and the product of two characters both satisfy the homomorphism properties, they are also characters of G, and consequently, χ is a character of G. If χ is identically equal to 1, then we must have that $\overline{\chi_2} = \chi_1^{-1}$, and, by the above, that $\chi_1 = \chi_2$, a contradiction.

Thus, we must instead have

$$0 = \sum_{x \in G} \chi(x) = \sum_{x \in G} \chi_1(x)\overline{\chi_2(x)} = (\chi_1, \chi_2)$$

and we are done. □

Fact 1 implies that $\chi(x)\overline{\chi(x)} = 1$ for all $x \in G$ and characters χ of G. Combining this with Fact 2 we have:

Corollary 3. *The normalized characters* $\frac{1}{\sqrt{|G|}}\chi$ *are orthonormal.*

4.2 Fourier basis

We now try to use the characters of a group as a Fourier basis. If $f : G \to \mathbb{C}$ can be written as

$$f = \frac{1}{\sqrt{|G|}} \sum_{\chi} \hat{f}(\chi)\overline{\chi} \tag{2}$$

for some $\hat{f}$, then $\hat{f}$ is our candidate Fourier transform of f. This is a linear mapping. We now argue that both unitarity and the convolution property (modulo a constant) are satisfied, as well.

Unitarity. Suppose that f and g can be written in the form (2). Then, by the orthogonality property of the characters χ, and the fact that $(\chi, \chi) = |G|$ for all characters χ of G we must have that

$$\begin{aligned}
(f, g) &= \frac{1}{|G|} \sum_{x \in G} \sum_{\chi_1, \chi_2} \hat{f}(\chi_1)\chi_1(x)\overline{\hat{g}(\chi_2)\chi_2(x)} \\
&= \frac{1}{|G|} \sum_{\chi_1, \chi_2} \sum_{x \in G} \hat{f}(\chi_1)\overline{\hat{g}(\chi_2)}\chi_1(x)\overline{\chi_2(x)} \\
&= \frac{1}{|G|} \sum_{\chi_1, \chi_2} \hat{f}(\chi_1)\overline{\hat{g}(\chi_2)} \sum_{x \in G} \chi_1(x)\overline{\chi_2(x)} \\
&= \frac{1}{|G|} \sum_{\chi_1, \chi_2} \hat{f}(\chi_1)\overline{\hat{g}(\chi_2)}(\chi_1, \chi_2) \\
&= \sum_{\chi} \hat{f}(\chi)\overline{\hat{g}(\chi)} \\
&= (\hat{f}, \hat{g})
\end{aligned}$$

and so our candidate Fourier transform preserves inner products (and thus the 2-norm) and is unitary.

We can also show that our candidate Fourier Transform is unitary by showing that its inverse is equal to its complext conjugate transpose. By the orthogonality of the characters χ, we must also havethat our candidate Fourier transform satisfies

$$\hat{f}(\chi) = (f, \overline{\chi}) = \frac{1}{\sqrt{|G|}} \sum_{x \in G} f(x)\chi(x) \tag{3}$$

Recall that the mapping $f \to \hat{f}$ is given by equation (2). From (3), we can see that the inverse mapping $\hat{f} \to f$ is the complex conjugate transpose of the forward mapping, showing again that our candidate Fourier transform is unitary.

Convolution property. If f and g can be written in the form (2), then so can $f * g : G \to \mathbb{C}$, defined by $(f * g)(x) = \sum_{y \in G} f(y)g(x - y)$. Moreover,

$$\widehat{f * g}(\chi) = c(G) \cdot \hat{f}(\chi) \cdot \hat{g}(\chi), \tag{4}$$

where $c(G) = \sqrt{|G|}$. Note that the constant $c(G)$ in (4) is not 1, as required by the convolution property stated before. However, the constant does not affect the usefulness. In general, nontrivial constants appear either in the formula for the Fourier transform, or its inverse, or the convolution property. It is a matter of taste where to allow them.

To argue the convolution property, we start with the convolution of f and g at x which is equal to the sum over all y's of $f(y)$ times $g(z)$ where $z = x - y$

$$(f * g)(x) \doteq \sum_{y \in G} f(y)g(x - y).$$

By rewriting f and g as a linear combination of the characters, the right-hand side becomes

$$\frac{1}{|G|} \sum_{y \in G} \left(\sum_{\chi_1} \hat{f}(\chi_1) \cdot \overline{\chi_1}(y) \right) \cdot \left(\sum_{\chi_2} \hat{g}(\chi_2) \cdot \overline{\chi_2(x - y)} \right).$$

Note that we introduce a normalization factor of $\frac{1}{\sqrt{G}}$ for both terms and combine them to $\frac{1}{|G|}$. Next, we can rewrite $\chi_2(x - y)$ using the properties of characters. Namely, that the character of a sum equals the product of the character values. Additionally, applying a character to $-y$ is the same as taking the inverse, but since all characters are on the unit sphere, taking the inverse is the same as taking the complex conjugate. So, $\chi_2(x - y) = \chi_2(x) \cdot \overline{\chi_2(y)}$ and taking the complex conjugate of the whole thing leaves us with:

$$\frac{1}{|G|} \sum_{y \in G} \left(\sum_{\chi_1} \hat{f}(\chi_1) \cdot \overline{\chi_1}(y) \right) \cdot \left(\sum_{\chi_2} \hat{g}(\chi_2) \cdot \overline{\chi_2(x)} \cdot \chi_2(y) \right).$$

Then after rearranging terms:

$$\frac{1}{|G|} \sum_{\chi_1, \chi_2} \hat{f}(\chi_1) \cdot \hat{g}(\chi_2) \cdot \overline{\chi_2(x)} \cdot \sum_{y \in G} \chi_2(y) \cdot \overline{\chi_1}(y).$$

We then notice that $\sum_{y \in G} \chi_2(y) \cdot \overline{\chi_1}(y)$ is exactly our definition of the inner product, giving us:

$$\frac{1}{|G|} \sum_{\chi_1, \chi_2} \hat{f}(\chi_1) \cdot \hat{g}(\chi_2) \cdot \overline{\chi_2(x)} \cdot (\chi_2, \chi_1).$$

Finally, we notice that whenever χ_2 differs from χ_1 then $(\chi_2, \chi_1) = 0$, therefore the sum is only nonzero when $\chi_2 = \chi_1$ and we can then simply sum over χ. Furthermore, when $\chi_1 = \chi_2$ then

$(\chi_2, \chi_1) = |G|$ which cancels with $\frac{1}{|G|}$ and all that's left is to rewrite the answer to match the given form.

$$\sum_\chi \hat{f}(\chi) \cdot \hat{g}(\chi) \cdot \overline{\chi(x)} \doteq \frac{1}{\sqrt{|G|}} \sum_\chi \widehat{f * g}(\chi) \cdot \overline{\chi(x)}$$

so $\widehat{f * g}(\chi) = c(G) \cdot \hat{f}(\chi) \cdot \hat{g}(\chi)$ where $c(G) = \sqrt{|G|}$.

4.3 Uniqueness of Fourier basis

We have shown that the characters of a finite group G form a valid Fourier basis for the set of functions $f : G \to \mathbb{C}$ that are in the linear span of the characters. The remaining question is whether every function $f : G \to \mathbb{C}$ can be written in the form (2), i.e., whether the characters of G form a basis for all of $\{f : G \to \mathbb{C}\}$. This is not the case for all finite groups G, but it is for all Abelian finite groups, as we will show in the next section. Here we show that whenever the characters form a basis, they are the unique Fourier basis according to our definition.

Theorem 4. *If the characters of a group G span the space of all functions $f : G \to \mathbb{C}$, then the normalized characters form the unique Fourier basis up to a permutation of the basis elements and global phase.*

Proof. We have already shown above that, if the characters span the space of all functions $f : G \to \mathbb{C}$ that they form a Fourier basis; it remains to show uniqueness.

Suppose that χ_1 and χ_2 are characters of G. From the convolution property,

$$\widehat{\chi_1 * \chi_2} = c(G) \cdot \hat{\chi}_1 \cdot \hat{\chi}_2.$$

By the definition of convolutions of $f : G \to \mathbb{C}$ and the homomorphism properties of χ_2

$$\begin{aligned}(\chi_1 * \chi_2)(x) &= \sum_{y \in G} \chi_1(y)\chi_2(x-y) \\ &= \sum_{y \in G} \chi_1(y)\chi_2(x)\overline{\chi_2(y)} \\ &= (\chi_1, \chi_2) \cdot \chi_2(x)\end{aligned}$$

as $\chi_2(-y) = \chi_2(y)^{-1} = \overline{\chi_2(y)}$.

Hence, if $\chi_1 \neq \chi_2$, then we have

$$c(G) \cdot \hat{\chi}_1 \cdot \hat{\chi}_2 = \widehat{\chi_1 * \chi_2} = (\chi_1, \chi_2) \cdot \hat{\chi}_2 = 0$$

and so $\operatorname{supp}(\hat{\chi}_1) \cap \operatorname{supp}(\hat{\chi}_2) = \emptyset$.

As the vector space of functions $f : G \to \mathbb{C}$ is $|G|$-dimensional, and as the characters span the set of all such functions, we must have at least $|G|$ characters. Furthermore, because the characters are orthogonal by the above, we can have no more than $|G|$ characters and thus there exist exactly $|G|$ distinct characters of G. Since $\hat{\chi}(\chi) = (\chi, \chi) = |G| \neq 0$ for all characters of G, we have $|\operatorname{sup}(\hat{\chi})| \geq 1$ for all χ and hence

$$|G| \leq \sum_\chi |\operatorname{sup}(\hat{\chi})|.$$

But as, by the above, the supports of distinct χ_1 and χ_2 are disjoint, we must also have that

$$\sum_{\chi} |\sup \hat{\chi}| \leq |G| = |\cup_{\chi} \operatorname{supp}(\hat{\chi})|.$$

Hence, $\sum_{\chi} |\sup \hat{\chi}| = |G|$ and we must have $|\sup(\hat{\chi})| = 1$ for all χ.

For any function $f : G\mathbb{C}$, $|\operatorname{supp}(\hat{f})|$ equals the number of the Fourier basis that are needed to express f as a linear combination of them. Thus, $|\operatorname{supp}(\hat{\chi})| = 1$ means that χ is itself an element of the Fourier basis, up to a scalar. As the Fourier basis is orthonormal, $\chi/\sqrt{|G|}$ must, in particular, be a member of the basis up to global phase. Consequently, the Fourier basis consisting of the normalized characters is unique up to a permutation of the basis elements and global phase. □

5 Fourier Transform over Finite Abelian Groups

We show that for finite Abelian groups G, there are as many characters as the dimension of the space $\{f : G \to \mathbb{C}\}$, namely $|G|$. By the previous section, this means that the characters form the unique basis that satisfies our requirements for a Fourier basis. As such, a Fourier transform over finite Abelian groups exists and is unique.

We make use the following result from group theory:

Theorem 5 (Structure Theorem). *Every finite Abelian group is isomorphic to*

$$\mathbb{Z}_{N_1} \times \mathbb{Z}_{N_2} \times \mathbb{Z}_{N_3} \times \cdots \times \mathbb{Z}_{N_k}$$

under component-wise addition for some $N_1, N_2, \ldots, N_k \in \mathbb{N}$.

We first find N characters for $\mathbb{Z}_N$ for $N \in \mathbb{N}$ and then find $|G_1| \cdot |G_2|$ characters of $G_1 \times G_2$ where G_1 and G_2 have $|G_1|$ and $|G_2|$ characters, respectively.

Characters of modular addition. We simply construct the following N distinct characters. Recall that the range of the characters of G is the set of N-th roots of unity (see Figure 1 for an example for $N = 6$). For each element $y \in \mathbb{Z}_N$, we construct a unique character that maps 1 to $\exp(2\pi iy/N)$.

Explicitly, for $y \in \mathbb{Z}_N$, let $\chi_y : \mathbb{Z}_{\mathbb{N}} \to \mathbb{C}$ such that $\chi_y(1) = (e^{2\pi i/N})^y = e^{2\pi iy/N}$ and $\chi_y(x) = \chi_y(1)^x = e^{2\pi ixy/N}$ for $x \in \mathbb{Z}_N$. As χ_y is a homomorphism and distinct for each $y \in \mathbb{Z}_N$, we are done.

For the special case of $N = 2$, the simple group with only two elements, we have that $\chi_y(x) = (-1)^{xy}$. In particular, $\chi_0(x) \equiv 1$ and $\chi_1(x) = (-1)^x$

Characters of direct product. We construct the following $|G_1| \cdot |G_2|$ characters. For $y_1 \in G_1$ and $y_2 \in G_2$ let

$$\chi_{y_1,y_2}(x_1, x_2) = \chi_{y_1}^{(G_1)}(x_1) \cdot \chi_{y_2}^{(G_2)}(x_2).$$

As we have given a distinct $\chi_{y_1}^{(G_1)}$ for each $y_1 \in G_1$ and similarly for G_2, we have $|G_1| \cdot |G_2| = |G_1 \times G_2|$ of these, which are distinct because the $\chi_{y_1}^{(G_1)}$ and $\chi_{y_2}^{(G_2)}$ are.

To show this, note that, where 0_1 and 0_2 are the identities of G_1 and G_2, respectively, we have that $\chi_{y_1,y_2}(0_1, x_2) = \chi_{y_2}^{(G_2)}(x_2)$ and $\chi_{y_1,y_2}(x_1, 0_2) = \chi_{y_1}^{(G_1)}(x_1)$ since homomorphisms map identities

to identities (in this case, to 1). If $(y_1, y_2) \neq (y_1', y_2')$ it follows that χ_{y_1,y_2} and $\chi_{y_1',y_2'}$ will disagree on some point. It remains to show that they are characters, i.e., that they are homomorphisms.

By the definition of χ_{y_1,y_2} and as $\chi_{y_1}^{(G_1)}$ and $\chi_{y_2}^{(G_2)}$ are homomorphisms,

$$\begin{aligned}
\chi_{y_1,y_2}(x_1 + z_1, x_2 + z_2) &= \chi_{y_1}^{(G_1)}(x_1 + z_1) \cdot \chi_{y_2}^{(G_2)}(x_2 + z_2) \\
&= (\chi_{y_1}^{(G_1)}(x_1)\chi_{y_1}^{(G_1)}(z_1)) \cdot (\chi_{y_2}^{(G_2)}(x_2)\chi_{y_2}^{(G_2)}(z_2)) \\
&= (\chi_{y_1}^{(G_1)}(x_1)\chi_{y_2}^{(G_2)}(x_2)) \cdot (\chi_{y_1}^{(G_1)}(z_1)\chi_{y_2}^{(G_2)}(z_2)) \\
&= \chi_{y_1,y_2}(x_1, x_2) \cdot \chi_{y_1,y_2}(z_1, z_2).
\end{aligned}$$

Putting things together. We have constructed N distinct characters for each $\mathbb{Z}_N$ for all $N \in \mathbb{N}$. We have also shown how to construct $|G_1| \cdot |G_2|$ characters for the direct product of groups G_1 and G_2, given $|G_1|$ and $|G_2|$ for G_1 and G_2, respectively. The combination of those two constructions allows us to construct for all groups of the form

$$\mathbb{Z}_{N_1} \times \mathbb{Z}_{N_2} \times \mathbb{Z}_{N_3} \times \cdots \times \mathbb{Z}_{N_k},$$

where $N_1, N_2, \ldots, N_k \in \mathbb{N}$, a number of distinct characters equal to

$$N_1 \cdot N_2 \cdot \cdots \cdot N_k = |\mathbb{Z}_{N_1} \times \mathbb{Z}_{N_2} \times \mathbb{Z}_{N_3} \times \cdots \times \mathbb{Z}_{N_k}|.$$

The characters can be indexed by $y = (y_1, y_2, \ldots, y_k) \in \mathbb{Z}_{N_1} \times \mathbb{Z}_{N_2} \cdots \times \mathbb{Z}_{N_k}$ and are given by

$$\chi_y(x) = \prod_{j \in [k]} \exp(2\pi i x_j y_j / N_j) = \exp\left(2\pi i \sum_{j \in [k]} \frac{x_j y_j}{N_j}\right).$$

For the special case of $G = (\mathbb{Z}_2)^n$, the characters for $\mathbb{Z}_2$ are the constant function 1 and the ± 1-valued parity function. The n-fold product of this group is obtained by taking n independent copies:

$$\chi_y(x) = \prod_j (-1)^{x_j y_j} = (-1)^{x \cdot y},$$

which corresponds to the set of all ± 1-valued parity functions.

By the Structure Theorem, all finite Abelian groups are isomorphic to such a group. It follows from our analysis in the previous section that such groups have a unique Fourier basis (up to permutation and global phase), given by the normalized characters $\chi_y / \sqrt{|G|}$. The Fourier transform $\hat{f}$ of a function $f : G \to \mathbb{C}$ is given by

$$\hat{f}(y) = \frac{1}{\sqrt{|G|}} \sum_x f(x)\chi_y(x) = \frac{1}{\sqrt{|G|}} \sum_x f(x) \exp\left(2\pi i \sum_{j \in [k]} \frac{x_j y_j}{N_j}\right).$$

Quantum Fourier transform. The classical Fourier transform is applied to a vector with one component for every element of G and outputs the same number of components. However, the quantum Fourier transform is instead applied to a superposition on $\log |G|$ qubits. The quantum subroutine transforms input $\sum_{x \in G} \alpha(x) |x\rangle$ into output $\sum_{x \in G} \hat{\alpha}(x) |x\rangle$, where $\hat{\alpha}$ is the Fourier transform of α. These transformations can be realized by unitary circuits of size poly $\log |G|$ for every finite Abelian G (and some other groups). For the special case of $\mathbb{Z}_2^n$ under addition, we already know the way to realize the quantum Fourier transform, namely to apply $H^{\otimes n}$. For the case of $\mathbb{Z}_N$, we will see it next lecture.

6 Hidden subgroup problem over finite Abelian groups

Recall the hidden subgroup problem (HSP) for a group G. The input is blackbox $f : G \to R$ for some group G and set R such that for some subgroup H of G:

$$f(x_1) = f(x_2) \Leftrightarrow Hx_1 = Hx_2$$

where $Hx \doteq \{h \cdot x : h \in H\}$ is the right coset of x modulo H, and $\cdot$ denotes the group operation in multiplicative notation. The goal is to find a set S of generators for H, i.e., $S \subseteq G$ such that $H = \langle S \rangle \doteq \{s_1 \cdot s_2 \ldots s_k : s_1, s_2 \ldots, s_k \in S \cup S^{-1}, k \in \mathbb{N}\}$.

We now develop an efficient quantum algorithm for the hidden subgroup problem over finite Abelian groups. More precisely, we establish the result for groups that are the direct product of cyclic groups.

Theorem 6. *Consider the group*

$$G = \mathbb{Z}_{N_1} \times \mathbb{Z}_{N_2} \times \cdots \times \mathbb{Z}_{N_k} \tag{5}$$

under component-wise addition, where $N_1, N_2, \ldots, N_k \in \mathbb{N}$, and suppose that the prime factorizations of the numbers N_j for $j \in [k]$ are given. There exists a quantum algorithm that solves the hidden subgroup problem over G with error bounded by ϵ, runs in time $O(\operatorname{poly}\log(|G|/\epsilon))$, and makes $O(\log(|G|) + \log(1/\epsilon))$ queries to the black-box. If the size of the hidden subgroup H is also given, then the algorithm is exact, runs in time $O(\operatorname{poly}\log|G|)$, and makes $O(\log(|G|/|H|))$ queries to the black-box.

Theorem 6 applies to all of the instantiations of the HSP over the additive group $\mathbb{Z}_2^n$ that we discussed: distinguishing constant and balanced functions for $n = 1$, learning linear functions, and finding a hidden XOR-shift. There is one more instantiation of HSP over finite Abelian groups that we discussed, namely the discrete log problem over $\mathbb{Z}_p$ for prime p. As the underlying group is $G = \mathbb{Z}_{p-1} \times \mathbb{Z}_{p-1}$ under addition, Theorem 6 yields a quantum algorithm that runs in time $O(\operatorname{poly}\log p)$ once we apply the polynomial-time quantum algorithm for factoring integers to the integer $p - 1$. A similar combination of Theorem 6 and the algorithm for factoring integers yields an efficient algorithm for the HSP over more complicated finite Abelian groups provided we can efficiently compute an isomorphism with a product of cyclic groups (5). The existence of an isomorphism is guaranteed by the Structure Theorem for finite Abelian groups; the isomorphism may or may not be efficiently computable.

The algorithm of Theorem 6 generalizes the one we developed for finding a hidden XOR-shift. It consists of two parts:

- An exact quantum subroutine A that outputs a uniform element of $H^\perp$, where

  $$H^\perp \doteq \{g \in G : (\forall\, h \in H)\, \chi_h(g) = 1\}.$$

  The subroutine consists of Fourier sampling. It hinges on an efficient algorithm for the quantum Fourier transform over the finite Abelian group G, which we already know for the case $G = \mathbb{Z}_2^n$ (the n-fold Hadamard transform), and which we will develop in full generality for groups of the form (5) in the next lectures.

- A classical part which uses the quantum subroutine A to construct a set S of generators for H. In the case of finding a hidden XOR-shift we ran A a number of times to find a set of generators of $H^\perp$, and then solved the system of linear equations they define in the components of the hidden shift s. In the general case, the process will similarly first find a set of generators for $H^\perp$, and then solve several systems of modular equations, each one yielding an element of a generating set S for H. More precisely, the process runs the quantum subroutine $O(\log|G| + \log(1/\epsilon))$ times to obtain a generating set for $H^\perp$ with probability at least $1-\epsilon$. Like in the special case of finding a hidden XOR-shift, an exact algorithm when $|H|$ is known can be obtained by amplitude amplification.

6.1 Quantum subroutine - Fourier Sampling

Let F denote the quantum Fourier transform over G. We use the Fourier transform over G because it interacts nicely with the symmetries captured by the group G.

Consider a (right) coset state $|Hg\rangle$, which is the uniform superposition of all elements of the coset Hg,

$$|Hg\rangle \doteq \frac{1}{\sqrt{|H|}} \sum_{h\in H} |hg\rangle .$$

The (quantum) Fourier transform of $|Hg\rangle$ is given by the following formula:

$$\begin{aligned} F|Hg\rangle &= \frac{1}{\sqrt{|H|}} \sum_{h\in H} \frac{1}{\sqrt{|G|}} \sum_{y\in G} \chi_y(hg)\,|y\rangle \\ &= \frac{1}{\sqrt{|H||G|}} \sum_{y\in G} \chi_y(g) \left(\sum_{h\in H} \chi_y(h) \right) |y\rangle . \end{aligned} \tag{6}$$

Exercise 1. *Show that*

$$\sum_{h\in H} \chi_y(h) = \begin{cases} |H| & \text{if } y \in H^\perp \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

Plugging in Equation (7) into Equation (6) gives us:

$$F|Hg\rangle = \sqrt{\frac{|H|}{|G|}} \sum_{y\in H^\perp} \chi_y(g)\,|y\rangle . \tag{8}$$

Thus, the Fourier transform of the coset state $|Hg\rangle$ yields an equally weighted superposition over $H^\perp$. In particular, if $g = 0$ we get a uniform superposition over $H^\perp$.

The quantum subroutine acts on a system with two registers, where the first register contains elements of the domain G of the black-box function $f : G \to R$, and the second register contains elements of the range R. We start with the first register in a uniform superposition over G and the second one in the basis state $|0\rangle$, i.e. the initial state

$$\frac{1}{\sqrt{|G|}} \sum_{g\in G} |g\rangle\,|0\rangle .$$

By applying our blackbox f via U_f, we obtain the transformed state

$$\frac{1}{\sqrt{|G|}} \sum_{g \in G} |g\rangle |f(g)\rangle = \sqrt{\frac{|H|}{|G|}} \sum_{\text{cosets}} |Hg\rangle |f(Hg)\rangle ,$$

where we used the fact that $f(g)$ only depends on the coset Hg. Next, we apply the Fourier transform F to the first register, which by Equation (8) gives us the resulting quantum state

$$\sqrt{\frac{|H|}{|G|}} \sum_{\text{cosets}} F |Hg\rangle |f(Hg)\rangle = \frac{|H|}{|G|} \sum_{\text{cosets}} \sum_{y \in H^\perp} \chi_y(g) |y\rangle |f(Hg)\rangle .$$

As distinct cosets have distinct values under f and $|\chi_y(g)| = 1$ for every $g \in G$, measuring the first register yields a y uniformly at random from $H^\perp$.

6.2 Use of the quantum subroutine

Running the quantum subroutine $O(\log(|H^\perp|) + \log(1/\epsilon)) = O(\log(|G|) + \log(1/\epsilon))$ times and collecting all elements s' yields a generating set S' for $H^\perp$ with error bounded by ϵ. In order to construct a generating set S for H out of S', we make use of the fact that $(H^\perp)^\perp = H$. The fact can be argued as follows.

Exercise 2. *Show that*

1. *The quotient group G/H is isomorphic to $H^\perp$.*
2. *$|G| = |H| \cdot |H^\perp|$.*
3. *$(H^\perp)^\perp = H$.*

We can efficiently construct a generating set S for $(H^\perp)^\perp$ out of a generating set S' for $H^\perp$ in the following way, with bounded error. The elements $s \in (H^\perp)^\perp$ are exactly those that satisfy $\chi_{s'}(s) = 1$ for all $s' \in S'$. Recall that G is of the form (5), so we can write $s = (s_1, s_2, \ldots, s_k)$ and $s' = (s'_1, s'_2, \ldots, s'_k)$ where $s_j, s'_j \in \mathbb{Z}_{N_j}$ for each $j \in [k]$. Using the formula we derived for the characters of additive groups of the form (5), we have that

$$\chi_{s'}(s) = \prod_{j=1}^{k} \exp\left(2\pi i s'_j s_j / N_j\right) = \exp\left(2\pi i \sum_{j=1}^{k} s'_j s_j / N_j\right) .$$

Thus, $\chi_{s'}(s) = 1$ iff $\sum_{j=1}^{k} s'_j s_j / N_j \in \mathbb{Z}$, which is equivalent to the integral modular equation

$$\sum_{j=1}^{k} \frac{M}{N_j} s'_j \cdot s_j = 0 \bmod M, \tag{9}$$

where $M \doteq \operatorname{lcm}(N_1, N_2, \ldots, N_k)$.

Theorem 7. *Given the prime decomposition of M, we can classically do both of the following in time* $\operatorname{poly}(n, \log M)$ *for a system of at most n linear equations in at most n variables over $\mathbb{Z}_M$:*

(a) Deterministically compute the number of solutions and, in particular, decide whether the system is solvable.

(b) If a solution exists, deterministically compute one as well as generate a solution chosen uniformly at random among all solutions.

Proof. We use the Chinese remainder theorem to independently solve the system modulo $p_j^{e_j}$ where $M = \prod_j p_j^{e_j}$ is the prime factorization of M, and combine those solutions into solutions to the original system. The total number of solutions equals the product of the solutions modulo each $p_j^{e_j}$, and a uniform solution is obtained by combining independent uniform solutions modulo each $p_j^{e_j}$.

To achieve (a) and (b) for a system of linear equations in n variables modulo p^e, we employ the following reduction:

- If there is a coefficient that is not divisible by p, say the coefficient of variable x_k in equation $\sum_{j=1}^n c_j x_j = b \bmod p^e$, use it to express x_k as a linear combination of the other variables:

$$x_k = c_k^{-1} \left(b - \sum_{k \neq j = 1}^{n} c_j x_j \right) \bmod p^e, \tag{10}$$

  where $(c_k)^{-1}$ denotes the inverse of c_k modulo p^e, which exists because $\gcd(c_k, p^e) = 1$, and can be computed efficiently using the extended Euclidean algorithm. Then use (10) to eliminate x_k from the system. The reduced system has one variable less, the number of solutions remains the same, and a uniform solution to the full system is obtained from a uniform solution of the reduced system by extending it via (10).

- If every coefficient and every right-hand side is divisible by p, then replace every equation $\sum_{j=1}^n c_j x_j = b \bmod p^e$ by the equation $\sum_{j=1}^n c'_j x'_j = b' \bmod p^{e-1}$, where $c_j = p \cdot c'_j$ and $b = p \cdot b'$. There is a bijective relationship between solutions x to the original system on the one hand, and solutions x' to the reduced system combined with choices $l_i \in \mathbb{Z}_p$ for each $i \in [n]$ on the other; the connection is given by $x_i = x'_i + l_i \cdot p^{e-1}$. It follows that the number of solutions to the original system equals the number of solutions to the reduced system times p^n, and a uniform solution to the original system is obtained by picking a uniform solution of the reduced system combined with independent uniform choices for $l_i \in \mathbb{Z}_p$.

- If every coefficient is divisible by p but not every right-hand side is, then the system has no solution.[1]

We apply part (b) of Theorem 7 to generate $O(\log |H| + \log(1/\epsilon))$ independent uniformly distributed samples of the solutions to the system of equations (9). With probability at least $1 - \epsilon$, the resulting set S generates $(H^\perp)^\perp = H$.

In case $|H|$ is known, we also know $|H^\perp| = |G|/|H^\perp|$ by part (b) of Exercise 2. In that case, we can make use of amplitude amplifications with known success probability to obtain, with certainty, a generating set S' for $H^\perp$ in time $\operatorname{poly} \log |G|$ using a number of black-box queries bounded by $O(\log |H^\perp|) = O(\log(|G|/|H|))$. We construct the set S' element by element. In each step we obtain, with certainty, an element $s' \in H^\perp$ that is not in the set generated by the current S'.[2]

[1] This case cannot happen for the homogenous system consisting of the equations (9), but the exact algorithm uses another application of the claim, in which the right-hand sides are not all zero.

[2] This makes use of Theorem 7 with right-hand side s'_j modulo N_j.

Once we have the generating set S' for $H^\perp$, we can similarly find a generating set S for H with certainty in the stated time and query complexity. □

Note that the proof of Theorem 6 uses the fact that the underlying group G is finite Abelian in two ways:

- A small number of Fourier samples contains enough information to determine generators for the hidden subgroup H of G (and H can be retrieved efficiently from the samples).
- The quantum Fourier transform over G can be computed efficiently.

The first item hinges on the homomorphic properties of the Fourier basis, and breaks down for more general groups. In particular, for the symmetric group S_n, even though the quantum Fourier transform can be computed efficiently, one needs an exponential number of queries in n in order to obtain a significant statistical distance between the distributions of positive and negative instances of graph isomorphism.