

Quantum algorithms (CO 781/CS 867/QIC 823, Winter 2013)

Andrew Childs, University of Waterloo

LECTURE 6: Fourier analysis in nonabelian groups

We have seen that hidden subgroup states contain sufficient information to determine the hidden subgroup. Now we would like to know whether this information can be extracted efficiently. In this lecture, we will introduce the theory of Fourier analysis over general groups, an important tool for getting a handle on this problem.

A brief introduction to representation theory

To understand nonabelian Fourier analysis, we first need to introduce some notions from group representation theory. For further information on this subject, a good basic reference is the book *Linear Representations of Finite Groups* by Serre.

A *linear representation* (or simply *representation*) of a group G over the vector space \mathbb{C}^n is a *homomorphism* $\sigma: G \rightarrow \text{GL}(\mathbb{C}^n)$, i.e., a map from group elements to nonsingular $n \times n$ complex matrices satisfying $\sigma(x)\sigma(y) = \sigma(xy)$ for all $x, y \in G$. Clearly, $\sigma(1) = 1$ and $\sigma(x^{-1}) = \sigma(x)^{-1}$. We call \mathbb{C}^n the *representation space* of σ , where n is called its *dimension* (or *degree*), denoted d_σ .

Two representations σ and σ' with representation spaces \mathbb{C}^n are called *isomorphic* (denoted $\sigma \sim \sigma'$) if there is an invertible linear transformation $M \in \mathbb{C}^{n \times n}$ such that $M\sigma(x) = \sigma'(x)M$ for all $x \in G$. Otherwise they are called *non-isomorphic* (denoted $\sigma \not\sim \sigma'$). In particular, representations of different dimensions are non-isomorphic. Every representation of a finite group is isomorphic to a *unitary representation*, i.e., one for which $\sigma(x)^{-1} = \sigma(x)^\dagger$ for all $x \in G$. Thus we can restrict our attention to unitary representations without loss of generality.

The simplest representations are those of dimension one, such that $\sigma(x) \in \mathbb{C}$ with $|\sigma(x)| = 1$ for all $x \in G$. Every group has a one-dimensional representation called the *trivial representation*, defined by $\sigma(x) = 1$ for all $x \in G$.

Two particularly useful representations of a group G are the *left regular representation* and the *right regular representation*. Both of these representations have dimension $|G|$, and their representation space is the *group algebra* $\mathbb{C}G$, the $|G|$ -dimensional complex vector space spanned by basis vectors $|x\rangle$ for $x \in G$. The left regular representation L satisfies $L(x)|y\rangle = |xy\rangle$, and the right regular representation R satisfies $R(x)|y\rangle = |yx^{-1}\rangle$. In particular, both regular representations are *permutation representations*: each of their representation matrices is a permutation matrix.

Given two representations $\sigma: G \rightarrow V$ and $\sigma': G \rightarrow V'$, we can define their *direct sum*, a representation $\sigma \oplus \sigma': G \rightarrow V \oplus V'$ of dimension $d_{\sigma \oplus \sigma'} = d_\sigma + d_{\sigma'}$. The representation matrices of $\sigma \oplus \sigma'$ are block diagonal, of the form

$$(\sigma \oplus \sigma')(x) = \begin{pmatrix} \sigma(x) & 0 \\ 0 & \sigma'(x) \end{pmatrix} \tag{1}$$

for all $x \in G$.

A representation is called *irreducible* if it cannot be decomposed as the direct sum of two other representations. Any representation of a finite group G can be written as a direct sum of irreducible representations (or *irreps*) of G .

Another way to combine two representations is with the *tensor product*. The tensor product of $\sigma: G \rightarrow V$ and $\sigma': G \rightarrow V'$ is $\sigma \otimes \sigma': G \rightarrow V \otimes V'$, a representation of G of dimension $d_{\sigma \otimes \sigma'} = d_\sigma d_{\sigma'}$.

The *character* of a representation σ is the function $\chi_\sigma: G \rightarrow \mathbb{C}$ defined by $\chi_\sigma(x) := \text{tr } \sigma(x)$. We have

- $\chi_\sigma(1) = d_\sigma$ (since $\sigma(1)$ is I_d , the d -dimensional identity matrix)
- $\chi_\sigma(x^{-1}) = \chi_\sigma(x)^*$ (since we can assume that σ is unitary), and
- $\chi_\sigma(yx) = \chi_\sigma(xy)$ for all $x, y \in G$ (since the trace is cyclic).

In particular, $\chi_\sigma(yxy^{-1}) = \chi_\sigma(x)$, so characters are constant on conjugacy classes. For two representations σ, σ' , we have $\chi_{\sigma \oplus \sigma'} = \chi_\sigma + \chi_{\sigma'}$ and $\chi_{\sigma \otimes \sigma'} = \chi_\sigma \cdot \chi_{\sigma'}$.

The most useful result in representation theory is probably *Schur's Lemma*, which can be stated as follows:

Theorem (Schur's Lemma). *Let σ and σ' be two irreducible representations of G , and let $M \in \mathbb{C}^{d_\sigma \times d_{\sigma'}}$ be a matrix satisfying $\sigma(x)M = M\sigma'(x)$ for all $x \in G$. Then if $\sigma \not\sim \sigma'$, $M = 0$; and if $\sigma = \sigma'$, M is a scalar multiple of the identity matrix.*

Schur's Lemma can be used to prove the following orthogonality relation for irreducible representations:

Theorem (Orthogonality of irreps). *For two irreps σ and σ' of G , we have*

$$\frac{d_\sigma}{|G|} \sum_{x \in G} \sigma(x)_{i,j}^* \sigma'(x)_{i',j'} = \delta_{\sigma,\sigma'} \delta_{i,i'} \delta_{j,j'}, \quad (2)$$

where we interpret $\delta_{\sigma,\sigma'}$ to mean 1 if $\sigma \sim \sigma'$, and 0 otherwise.

This implies a corresponding orthogonality relation for the *irreducible characters* (i.e., the characters of the irreducible representations):

Theorem (Orthogonality of characters). *For two irreps σ and σ' of G , we have*

$$(\chi_\sigma, \chi_{\sigma'}) := \frac{1}{|G|} \sum_{x \in G} \chi_\sigma(x)^* \chi_{\sigma'}(x) = \delta_{\sigma,\sigma'}. \quad (3)$$

The characters of G supply an orthonormal basis for the space of *class functions*, functions that are constant on conjugacy classes of G . (Recall that the characters themselves are class functions.) This is expressed by the orthonormality of the *character table* of G , the square matrix whose rows are labeled by irreps, whose columns are labeled by conjugacy classes, and whose entries are the corresponding characters. The character orthogonality theorem says that the rows of this matrix are orthonormal, provided each entry is weighted by the square root of the size of the corresponding conjugacy class divided by $|G|$. In fact the columns are orthonormal in the same sense.

Any representation of G can be broken up into its irreducible components. The regular representations of G are useful for understanding such decompositions, since they contain every possible irreducible representation of G , with each irrep occurring a number of times equal to its dimension. Let \hat{G} denote a complete set of irreps of G (which are unique up to isomorphism). Then we have

$$L \cong \bigoplus_{\sigma \in \hat{G}} (\sigma \otimes I_{d_\sigma}), \quad R \cong \bigoplus_{\sigma \in \hat{G}} (I_{d_\sigma} \otimes \sigma^*). \quad (4)$$

In fact, this holds with the same isomorphism for both L and R , since the left and right regular representations commute. This isomorphism is simply the Fourier transform over G , which we discuss further below.

Considering $\chi_L(1) = \chi_R(1) = |G|$ and using this decomposition, we find the well-known identity

$$\sum_{\sigma \in \hat{G}} d_\sigma^2 = |G|. \quad (5)$$

Also, noting that $\chi_L(x) = \chi_R(x) = 0$ for any $x \in G \setminus \{1\}$, we see that

$$\sum_{\sigma \in \hat{G}} d_\sigma \chi_\sigma(x) = 0. \quad (6)$$

In general, the multiplicity of the irrep $\sigma \in \hat{G}$ in an arbitrary representation τ of G is given by $\mu_\sigma^\tau := (\chi_\sigma, \chi_\tau)$. This gives the decomposition

$$\tau \cong \bigoplus_{\sigma \in \hat{G}} \sigma \otimes I_{\mu_\sigma^\tau}. \quad (7)$$

Characters also provide a simple test for irreducibility: for any representation σ , $(\chi_\sigma, \chi_\sigma)$ is a positive integer, and is equal to 1 if and only if σ is irreducible.

Any representation σ of G can also be viewed as a representation of any subgroup $H \leq G$, simply by restricting its domain to elements of H . We denote the resulting *restricted representation* by $\text{Res}_H^G \sigma$. Even if σ is irreducible over G , it may not be irreducible over H .

Fourier analysis for nonabelian groups

The *Fourier transform* is a unitary transformation from the group algebra, $\mathbb{C}G$, to a complex vector space whose basis vectors correspond to matrix elements of the irreps of G , $\bigoplus_{\sigma \in \hat{G}} (\mathbb{C}^{d_\sigma} \otimes \mathbb{C}^{d_\sigma})$. These two spaces have the same dimension by (5).

The Fourier transform of the basis vector $|x\rangle \in \mathbb{C}G$ corresponding to the group element $x \in G$ is a weighted superposition over all irreducible representations $\sigma \in \hat{G}$, namely

$$|\hat{x}\rangle := \sum_{\sigma \in \hat{G}} \frac{d_\sigma}{\sqrt{|G|}} |\sigma, \sigma(x)\rangle, \quad (8)$$

where $|\sigma\rangle$ is a state that labels the irreducible representation, and $|\sigma(x)\rangle$ is a normalized, d_σ^2 -dimensional state whose amplitudes correspond to the entries of the matrix $\sigma(x)/\sqrt{d_\sigma}$:

$$|\sigma(x)\rangle := \sum_{j,k=1}^{d_\sigma} \frac{\sigma(x)_{j,k}}{\sqrt{d_\sigma}} |j, k\rangle. \quad (9)$$

(If σ is one-dimensional, then $|\sigma(x)\rangle$ is simply a phase factor $\sigma(x) = \chi_\sigma(x) \in \mathbb{C}$ with $|\sigma(x)| = 1$.) The Fourier transform over G is the unitary matrix

$$F_G := \sum_{x \in G} |\hat{x}\rangle \langle x| \quad (10)$$

$$= \sum_{x \in G} \sum_{\sigma \in \hat{G}} \sqrt{\frac{d_\sigma}{|G|}} \sum_{j,k=1}^{d_\sigma} \sigma(x)_{j,k} |\sigma, j, k\rangle \langle x|. \quad (11)$$

Note that the Fourier transform over G is not uniquely defined, but rather, depends on a choice of basis for each irreducible representation.

It is straightforward to check that F_G is indeed a unitary transformation. Using the identity

$$\langle \sigma(y) | \sigma(x) \rangle = \text{tr } \sigma^\dagger(y) \sigma(x) / d_\sigma \quad (12)$$

$$= \text{tr } \sigma(y^{-1}x) / d_\sigma \quad (13)$$

$$= \chi_\sigma(y^{-1}x) / d_\sigma, \quad (14)$$

we have

$$\langle \hat{y} | \hat{x} \rangle = \sum_{\sigma \in \hat{G}} \frac{d_\sigma^2}{|G|} \langle \sigma(y) | \sigma(x) \rangle \quad (15)$$

$$= \sum_{\sigma \in \hat{G}} \frac{d_\sigma}{|G|} \chi_\sigma(y^{-1}x). \quad (16)$$

Hence by (5–6) above, we see that $\langle \hat{y} | \hat{x} \rangle = \delta_{x,y}$.

F_G is precisely the transformation that decomposes both the left and right regular representations of G into their irreducible components. Let us check this explicitly for the left regular representation L . Recall that this representation satisfies $L(x)|y\rangle = |xy\rangle$, so we have

$$\hat{L}(x) := F_G L(x) F_G^\dagger \quad (17)$$

$$= \sum_{y \in G} |xy\rangle \langle \hat{y}| \quad (18)$$

$$= \sum_{y \in G} \sum_{\sigma, \sigma' \in \hat{G}} \sum_{j,k=1}^{d_\sigma} \sum_{j',k'=1}^{d_{\sigma'}} \frac{\sqrt{d_\sigma d_{\sigma'}}}{|G|} \sigma(xy)_{j,k} \sigma'(y)_{j',k'}^* |\sigma, j, k\rangle \langle \sigma', j', k'| \quad (19)$$

$$= \sum_{y \in G} \sum_{\sigma, \sigma' \in \hat{G}} \sum_{j,k,\ell=1}^{d_\sigma} \sum_{j',k'=1}^{d_{\sigma'}} \frac{\sqrt{d_\sigma d_{\sigma'}}}{|G|} \sigma(x)_{j,\ell} \sigma(y)_{\ell,k} \sigma'(y)_{j',k'}^* |\sigma, j, k\rangle \langle \sigma', j', k'| \quad (20)$$

$$= \sum_{\sigma \in \hat{G}} \sum_{j,k,\ell=1}^{d_\sigma} \sigma(x)_{j,\ell} |\sigma, j, k\rangle \langle \sigma, \ell, k| \quad (21)$$

$$= \bigoplus_{\sigma \in \hat{G}} (\sigma(x) \otimes I_{d_\sigma}), \quad (22)$$

where in the fourth line we have used the orthogonality relation for irreducible representations.

A similar calculation can be done for the right regular representation defined by $R(x)|y\rangle = |yx^{-1}\rangle$, giving

$$\hat{R}(x) := F_G R(x) F_G^\dagger \quad (23)$$

$$= \bigoplus_{\sigma \in \hat{G}} (I_{d_\sigma} \otimes \sigma(x)^*). \quad (24)$$

This identity will be useful when analyzing the application of the quantum Fourier transform to the hidden subgroup problem.

To use the Fourier transform as part of a quantum computation, we must be able to implement it efficiently by some quantum circuit. Efficient quantum circuits for the quantum Fourier transform are known for many, but not all, nonabelian groups. Groups for which an efficient QFT is known include metacyclic groups (i.e., semidirect products of cyclic groups), such as the dihedral group; the symmetric group; and many families of groups that have suitably well-behaved towers of subgroups. There are a few notable groups for which efficient QFTs are *not* known, such as the general linear group $\text{GL}_n(q)$ of $n \times n$ invertible matrices over \mathbb{F}_q , the finite field with q elements.