

Cayley Graphs and the Discrete Fourier Transform

Alan Mackey
Advisor: Klaus Lux

May 15, 2008

Abstract

Given a group G , we can construct a graph relating the elements of G to each other, called the Cayley graph. Using Fourier analysis on a group allows us to apply our knowledge of the group to gain insight into problems about various properties of the graph. Ideas from representation theory are powerful tools for analysis of groups and their Cayley graphs, but we start with some simpler theory to motivate the use of group representations.

1 $\mathbb{Z}/n\mathbb{Z}$ and $L^2(\mathbb{Z}/n\mathbb{Z})$

Before we can define what the discrete Fourier transform (or DFT) is, we have to know what exactly it transforms. To introduce the ideas behind the use of the DFT, we will use $G = \mathbb{Z}/n\mathbb{Z}$ as an example. The DFT doesn't actually operate on the elements of $\mathbb{Z}/n\mathbb{Z}$ themselves, but rather on certain *functions* with $\mathbb{Z}/n\mathbb{Z}$ as their domain.

Definition. The set $\{f : G \rightarrow \mathbb{C}\}$ of all functions from a finite group G to the complex numbers is denoted $L^2(G)$.

Let's start with an example. The group $(\mathbb{Z}/n\mathbb{Z}, +_n)$ has n elements, and so any function from $\mathbb{Z}/n\mathbb{Z}$ to \mathbb{C} is determined exactly by the finite set of values $\{f(a) : a \in \mathbb{Z}/n\mathbb{Z}\}$. Thus, we can define a vector space isomorphism $\phi : L^2(\mathbb{Z}/n\mathbb{Z}) \rightarrow \mathbb{C}^n$ by

$$\phi(f) = (f(0), f(1), \dots, f(n-1)).$$

From here, it isn't hard to see that $L^2(\mathbb{Z}/n\mathbb{Z})$ is an n -dimensional vector space over \mathbb{C} . A convenient basis for $L^2(\mathbb{Z}/n\mathbb{Z})$ is the set of delta functions $\delta_1, \dots, \delta_n$, each defined by

$$\delta_i(j) = \begin{cases} 1, & \text{if } i \equiv j \pmod{n}, \\ 0, & \text{otherwise.} \end{cases}$$

Then for any function f ,

$$f = f(1)\delta_1 + f(2)\delta_2 + \dots + f(n)\delta_n$$

Additionally, $L^2(\mathbb{Z}/n\mathbb{Z})$ has the following inner product:

$$\langle f, g \rangle = \sum_{x \in \mathbb{Z}/n\mathbb{Z}} f(x)\overline{g(x)} \quad \text{for } f, g \in L^2(\mathbb{Z}/n\mathbb{Z}).$$

With respect to the inner product above,

$$\langle \delta_i, \delta_j \rangle = \begin{cases} 1, & \text{if } i \equiv j \pmod{n}, \\ 0, & \text{otherwise.} \end{cases}$$

So the delta functions actually form an orthonormal basis for $L^2(\mathbb{Z}/n\mathbb{Z})$.

We can also define a way of "multiplying" functions in $L^2(\mathbb{Z}/n\mathbb{Z})$, called convolution.

Definition. The *convolution* ($f * g$) of two functions f and g in $L^2(\mathbb{Z}/n\mathbb{Z})$ is defined by

$$(f * g)(x) = \sum_{y \in \mathbb{Z}/n\mathbb{Z}} f(y)g(x-y) \quad \text{for } x \in \mathbb{Z}/n\mathbb{Z}.$$

The operation of convolution makes $L^2(\mathbb{Z}/n\mathbb{Z})$ a ring, called the group ring of $\mathbb{Z}/n\mathbb{Z}$ over \mathbb{C} , which we will denote $\mathbb{C}(\mathbb{Z}/n\mathbb{Z})$. If we view $\mathbb{Z}/n\mathbb{Z}$ as a multiplicative cyclic group of the elements $\{1, x, x^2, \dots, x^{n-1}\}$ generated by x , then we may use the group structure to define multiplication of functions.

To make this clear, note that $\delta_a * \delta_b = \delta_{a+nb}$. This is because

$$(\delta_a * \delta_b)(x) = \sum_{y \in \mathbb{Z}/n\mathbb{Z}} \delta_a(y)\delta_b(x-y),$$

and the term

$$\delta_a(y)\delta_b(x-y)$$

is zero whenever $\delta_a(y)$ or $\delta_b(x-y)$ is zero, and so we need only consider the term when $y \equiv a$ and $x-y \equiv b$ for some $y \in \mathbb{Z}/n\mathbb{Z}$. Then we can add the two congruences and see that this only occurs if $x \equiv a+b$. Also, if $x \equiv a+b$, then $y \equiv a$ is the only element of $\mathbb{Z}/n\mathbb{Z}$ which solves the equation.

Thus, the convolution of the two functions is

$$(\delta_i * \delta_j) = \begin{cases} 1, & \text{if } x \equiv a+b \pmod{n}, \\ 0, & \text{otherwise.} \end{cases}$$

However, this function is just δ_{a+nb} . Additionally it is clear that convolution is commutative, associative and linear; that is,

$$\begin{aligned} f * g &= g * f \\ f * (g * h) &= (f * g) * h \\ (f + \lambda g) * h &= (f * h) + \lambda(g * h). \end{aligned}$$

Essentially, $L^2(\mathbb{Z}/n\mathbb{Z})$ is the group ring of $\{\delta_0, \delta_1, \dots, \delta_{n-1}\}$ over \mathbb{C} , where the general element f may be expressed as $f = f(1)\delta_1 + f(2)\delta_2 + \dots + f(n)\delta_n$ and where multiplication in the group $\{\delta_0, \delta_1, \dots, \delta_{n-1}\}$ is convolution of delta functions.

Just from the fact that $\delta_a * \delta_b = \delta_{a+nb}$, we may use the distributive property of $L^2(\mathbb{Z}/n\mathbb{Z})$ to convolve two functions in a manner nearly identical to the way we would multiply polynomials. In fact, the function ϕ defined by

$$\phi : L^2(\mathbb{Z}/n\mathbb{Z}) \rightarrow \mathbb{C}[x]/\langle x^n - 1 \rangle$$

$$\phi(f) = f(0) + f(1)x + f(2)x^2 + \dots + f(n-1)x^{n-1}$$

is an isomorphism; that is,

$$L^2(\mathbb{Z}/n\mathbb{Z}) \cong \mathbb{C}[x]/\langle x^n - 1 \rangle.$$

Convolution in the general case is defined analogously, though multiplication in the ring $L^2(G)$ is not as simple.

Definition. Given a group G , the *convolution* $(f * g)$ of two functions f and g in $L^2(G)$ is defined by

$$(f * g)(x) = \sum_{y \in G} f(y)g(xy^{-1}) \quad \text{for } x \in G.$$

2 The Cayley Graph

Consider the following situation: we have an ordinary deck of cards, and are interested in shuffling the cards. We can consider the set of functions which permute the order of the cards as a group under function composition, since there is an identity permutation, the composition of permutations is associative, and each permutation has an inverse. Then, the problem of shuffling is really just the problem of picking a random element from the Group S_{52} .

However, it might be more efficient to just apply elements from a set of generators at random. Then, the question becomes: how effectively can we generate a random element from a group by applying elements at random from a set of generators? This is essentially a random walk through the group, and we can visualize it as a problem in graph theory.

Definition. Given a group G and a symmetric set of generators S (meaning S is closed under taking inverses) of G , we can form the Cayley graph $X(G, S)$ of G as follows: Create a graph whose vertices are the elements of G such that each element of G has a vertex. Then, for two elements x and y in G , connect the two vertices x and y if there exists an element $s \in S$ such that $x = sy$.

Note that the set S being symmetric implies that the graph is not directed; that is, if you can get from x to y by applying an element of S (or equivalently, the vertex x is connected to the vertex y), then you can get from y to x by applying another element from S (or equivalently, the vertex y is connected to the vertex x). Also, S being a generating set implies that the Cayley Graph is connected.

Now, starting with a given element of G and applying elements from a set of generators at random is equivalent to a random walk on $X(G, S)$. By analyzing the Cayley graph, we can gain insight into the structure of the group.

In terms of graph theory, we say that the *degree* of a vertex x of $X(G, S)$ is the number of other vertices connected to x . The graph is called *k-regular* if each vertex has degree k . The *distance* between two vertices x and y is the least number of edges in a path connecting x to y . The maximum distance over all pairs of vertices x and y is called the *diameter* of the graph, and the *girth* of the graph is the length of its smallest circuit. One of the most useful tools we have for analyzing a graph is called its *adjacency matrix*.

Definition. Suppose the vertices of a graph are v_1, \dots, v_n . Then the *adjacency matrix* of the graph is the matrix whose (i, j) entry is 1 if v_i is connected to v_j and 0 otherwise.

For our nondirected Cayley graphs, the adjacency matrix is always symmetric.

The adjacency matrix of a graph has several special properties. Let X be a graph with adjacency matrix A . Then:

1. The (i, j) entry of A^r is the total number of walks of length r connecting the i^{th} and j^{th} vertices of X .
2. The diameter of X is the smallest integer d such that all the entries of the matrix $(I + A)^d$ are nonzero.
3. If the diameter of X is d , then A has at least $d + 1$ distinct eigenvalues.

Proof. 1) It is clear that the (i, j) entry of A^1 is the number of length 1 walks connecting the i^{th} and j^{th} vertices of X . Now assume the statement is true for the case when $r = k$. Now, the (i, j) entry of A^{k+1} is

$$\sum_{l=1}^n b_{il} a_{lj}$$

where b_{il} is the (i, l) entry of A^k and a_{lj} is the (l, j) entry of A .

If either $b_{il} = 0$ or $a_{lj} = 0$, the term is just 0 and does not affect the sum. Now, say $b_{il} = c$ and $a_{lj} = 1$. Then this means that there exist c paths of length r from the vertex i (denoted v_i) to the vertex v_l and one path connecting the v_l and v_j . Thus, there exist c paths of length $r + 1$ going from v_i to v_j which pass through v_l immediately before reaching j . By summing over all l , we obtain every path of length $r + 1$ since each path must pass through some vertex immediately before reaching v_j . Also, no path can pass

through two vertices immediately before reaching v_j , and so each $r + 1$ length path was counted exactly once. Thus, the (i, j) entry of A^{k+1} is the number of paths of length $k + 1$ connecting v_i and v_j . By induction, the result holds for all $r \in \mathbb{Z}^+$.

2) Let d be the diameter of the graph with adjacency matrix A . By the binomial theorem, we have that

$$(I + A)^d = I + \binom{d}{1}A + \binom{d}{2}A^2 + \dots + A^d$$

Note that all entries of all powers of A are nonnegative, since all entries of A are nonnegative, and consider the (i, j) entry of $(I + A)^d$. By the definition of d , there exists some integer $r \leq d$ such that there is a path of length r connecting the vertices v_i and v_j . Thus the (i, j) entry of A^r is nonzero, and so the (i, j) entry of $(I + A)^d$ must be nonzero since that entry is the sum of a nonzero entry and other nonnegative entries. Thus, all entries of $(I + A)^d$ are nonzero.

Additionally, if c is an integer strictly less than d , then there exist some vertices v_i and v_j such that there is no path of length c or less connecting v_i and v_j . Thus the (i, j) entry of each of I, A, A^2, \dots, A^c are all zero. Thus, the (i, j) entry of

$$(I + A)^c = \sum_{i=0}^c \binom{c}{i} A^i$$

is also zero, and so d is the smallest integer such that all the entries of the matrix $(I + A)^d$ are nonzero.

3) If the diameter of X is d , then we know that there are two vertices (call them v_1 and v_{d+1}) such that the shortest path connecting them has length d . Also, say this path passes through the vertices v_2, v_3, \dots, v_d in order. That means the $(1, d + 1)$ entry of A^d is nonzero, but the $(1, d + 1)$ entry of A^{d-1} is zero. Similarly, the $(1, d)$ entry of A^{d-1} is nonzero, but the $(1, d)$ entry of A^{d-2} is zero. Proceeding in this fashion, we see that the $d + 1$ matrices I, A, A^2, \dots, A^d are linearly independent. To see why, suppose

$$\lambda_0 I + \lambda_1 A + \dots + \lambda_d A^d = 0$$

We have shown that for all $n = 2, 3, \dots, d$ that A^n has a nonzero entry where all lower powers of A have a zero entry. This means that we must have $\lambda_d = \lambda_{d-1} = \dots = \lambda_2 = 0$. If λ_0 were not 0, that would imply that A

is a multiple of the identity matrix, contradicting that the Cayley graph is connected. Similarly $\lambda_1 \neq 0$ would imply that A is a multiple of the identity matrix, contradicting that the Cayley graph is connected. Thus the minimal polynomial of A must have degree at least $d + 1$, and so A must have at least $d + 1$ distinct eigenvalues. \square

3 The Discrete Fourier Transform on $L^2(\mathbb{Z}/n\mathbb{Z})$

When $L^2(\mathbb{Z}/n\mathbb{Z})$ is viewed as a vector space over \mathbb{C} , the delta functions form an orthonormal basis (this was shown earlier). However, it may be prudent to change the basis to a set of functions which make understanding some aspects about $\mathbb{Z}/n\mathbb{Z}$ simpler. Particularly, we would like to choose a basis which diagonalizes the adjacency matrix for the Cayley Graph of $\mathbb{Z}/n\mathbb{Z}$. This is possible if we choose as the new basis the set of eigenvectors of the adjacency matrix; that way, the eigenvalues will appear across the diagonal.

Above, we found that the adjacency matrix for a given graph encodes a lot of valuable information about the graph. Ultimately, one of our main goals here is to analyze random walks on the Cayley Graph, and this task and others are made much easier if the adjacency matrix is diagonal. We noted earlier that the adjacency matrix A of a graph is symmetric, and thus self-adjoint. The Spectral Theorem for real, self-adjoint operators then guarantees that such a basis of eigenvectors exists. That is, the matrix A is diagonalizable. The DFT is the transformation which will diagonalize A .

We will use C_n to denote the cyclic group of order n of the elements $\{1, x, x^2, \dots, x^{n-1}\}$ generated by x under multiplication, and refer to this (isomorphic) group rather than $\mathbb{Z}/n\mathbb{Z}$. This is because the operations of $+$ in the group $\mathbb{Z}/n\mathbb{Z}$ and the ring $L^2(\mathbb{Z}/n\mathbb{Z})$ are different, and viewing the cyclic group multiplicatively makes the distinction more clear. All results for C_n hold exactly for $\mathbb{Z}/n\mathbb{Z}$. Additionally, we will use the notation ϕ_k to denote the homomorphism from C_n into the multiplicative group of the complex roots of unity (denoted \mathbb{C}^\times) such that $\phi_k(x) = \omega^k$, where ω is $e^{-2\pi i/n}$.

Now consider the matrix whose columns are indexed by the elements

$$\{1, x, x^2, \dots, x^{n-1}\}$$

of C_n , and whose rows are indexed by the set $\phi_0, \phi_1, \dots, \phi_{n-1}$. The entry in the ϕ_k row and x^j column is just $\phi_k(x^j)$, or ω^{kj} . This is called the character table of C_n . The character table for C_4 is

	1	x	x^2	x^3
ϕ_0	1	1	1	1
ϕ_1	1	$-i$	-1	i
ϕ_2	1	-1	1	-1
ϕ_3	1	i	-1	$-i$

Definition. The *Discrete Fourier Transform*, or DFT, on $\mathbb{Z}/n\mathbb{Z}$ is the linear map from $L^2(\mathbb{Z}/n\mathbb{Z})$ into itself whose matrix with respect to the basis of δ functions defined in section 1 is the matrix whose entries are those of the character table for $\mathbb{Z}/n\mathbb{Z}$. This matrix is denoted \mathcal{F} .

The matrix of the DFT for C_4 (and $\mathbb{Z}/4\mathbb{Z}$) has the same entries as the character table above:

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix}$$

We claim that \mathcal{F} diagonalizes the adjacency matrix of a Cayley graph for $\mathbb{Z}/n\mathbb{Z}$.

To see why, consider the following (currently empty) table for C_4 :

	1	x	x^2	x^3
1				
x				
x^2				
x^3				

We fill in the table. Given an element $y \in C_4$, we place a 1 in the (i, j) entry of the table if $yx^j = x^i$ and 0 otherwise. For example, the matrix in $M_n(\mathbb{C})$ corresponding to the generator x of the cyclic group of order 4 under multiplication is

$$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

We call

$$T : C_n \rightarrow GL_n(\mathbb{C})$$

the map which associates each element of $\mathbb{Z}/n\mathbb{Z}$ with its corresponding element of $GL_n(\mathbb{C})$ by the procedure of the example above. That is, $T(y)$ is the matrix whose i, j entry is 1 if $yx^j = x^i$ for some $x \in C_n$ and 0 otherwise.

The map T can be best understood as the composition of two other maps. First, let S_n denote the symmetric group of order n . Given a group (G, \cdot) of order n , we define a map

$$\begin{aligned}\pi : G &\rightarrow S_n \\ g &\mapsto \pi_g\end{aligned}$$

Where $\pi_g(y) = g \cdot y$ for $y \in G$. Technically, we must actually number the elements of G for π_g to be an element of S_n , but it's easier to think of π_g as permuting the elements of G rather than $1, 2, \dots, n$. As an example, if

$$\begin{aligned}\pi : C_4 &\rightarrow S_4 \\ g &\mapsto \pi_g\end{aligned}$$

then in cycle notation,

$$\begin{aligned}\pi_1 &= (1) \\ \pi_x &= (1, x, x^2, x^3) \\ \pi_{x^2} &= (1, x^2)(x, x^3) \\ \pi_{x^3} &= (1, x^3, x^2, x).\end{aligned}$$

It is clear that for an arbitrary group G , $\ker(\pi) = \{(1)\}$ and that for $g, h \in G$,

$$\pi_{g \cdot h} = \pi_g \pi_h$$

which is because for $y \in G$,

$$\pi_{g \cdot h}(y) = (g \cdot h) \cdot y = g \cdot (h \cdot y) = g \cdot \pi_h(y) = \pi_g \pi_h(y).$$

Thus

$$G \cong \pi(G).$$

Additionally, consider the map

$$\begin{aligned} M : S_n &\rightarrow GL_n(\mathbb{C}) \\ \pi &\mapsto M(\pi) \end{aligned}$$

where $M(\pi)$ is the matrix whose i, j entry [denoted $M(\pi)_{ij}$] is

$$M(\pi)_{ij} = \begin{cases} 1, & \text{if } \pi(j) = i \\ 0, & \text{otherwise.} \end{cases}$$

If we associate the integer i with the column vector

$$\begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} = e_i$$

where the 1 appears in the i^{th} row, then the matrix $M(\pi)$ permutes the vectors e_1, \dots, e_n as π permutes $1, \dots, n$. That is,

$$M(\pi)(e_i) = e_{\pi(i)}.$$

Again, one can check that the kernel of this map is $\{(1)\}$. Also,

$$M(\pi)M(\pi')(e_i) = M(\pi)(e_{\pi'(i)}) = e_{\pi\pi'(i)} = M(\pi\pi')(e_i)$$

and so $S_n \cong M(S_n)$.

Now, we can show that the map T is a homomorphism.

Proof. Consider the map

$$M \circ \pi : G \rightarrow GL_n(\mathbb{C}).$$

First, $M \circ \pi$ maps an element $g \in G$ to the permutation π_g , then to the matrix $M(\pi_g)$. For convenience, we will index the rows and columns of $M(\pi_g)$ by the elements of G . By the definition of M ,

$$M(\pi_g)_{xy} = \begin{cases} 1, & \text{if } \pi_g(y) = x \\ 0, & \text{otherwise.} \end{cases}$$

or equivalently, by the definition of π_g ,

$$M(\pi_g)_{xy} = \begin{cases} 1, & \text{if } gy = x \\ 0, & \text{otherwise.} \end{cases}$$

For the case when $G = C_n$,

$$M(\pi_g)_{x^i x^j} = \begin{cases} 1, & \text{if } gx^j = x^i \\ 0, & \text{otherwise.} \end{cases}$$

But this is just the definition of T . Since the composition of two homomorphisms is a homomorphism, T is a homomorphism. □

We denote the matrix $T(x)$ by T_x . It is clear that the matrix T_x “shifts” elements in $L^2(\mathbb{Z}/n\mathbb{Z})$ by mapping $f = (a_0, a_1, \dots, a_{n-1})$ to the function $(a_{n-1}, a_0, \dots, a_{n-3}, a_{n-2})$. Formally, this process of shifting is actually convolution; that is, $T_x f = f * \delta_1$ and $T_x^k f = f * \delta_k$.

Proof. We may think of the function f in terms of the delta functions:

$$f = a_0 \delta_0 + a_1 \delta_1 + \dots + a_{n-1} \delta_{n-1}$$

Then

$$\begin{aligned}
f * \delta_1 &= (a_0\delta_0 + a_1\delta_1 + \dots + a_{n-1}\delta_{n-1}) * \delta_1 \\
&= (a_0\delta_0 * \delta_1) + (a_1\delta_1 * \delta_1) + \dots + (a_{n-1}\delta_{n-1} * \delta_1) \\
&= a_0(\delta_0 * \delta_1) + \dots + a_{n-1}(\delta_{n-1} * \delta_1) \\
&= a_0\delta_1 + a_1\delta_2 + \dots + a_{n-2}\delta_{n-1} + a_{n-1}\delta_0 \\
&= a_{n-1}\delta_0 + a_0\delta_1 + \dots + a_{n-2}\delta_{n-1} \\
&= (a_{n-1}, a_0, \dots, a_{n-3}, a_{n-2}).
\end{aligned}$$

Thus the shift matrix T_x accomplishes the same thing as convolution with δ_1 , and we can see by induction that

$$T_{x^k} f = (T_x)^k f = \delta_k * f.$$

□

The resulting function $T_{x^k} f = \delta_k * f$ satisfies

$$(f * \delta_k)(y) = f(y - k).$$

As an example, we will continue to use C_4 . If the set of generators given for this group is $S = \{x, x^3\}$, then the adjacency matrix A for $X(\mathbb{Z}/n\mathbb{Z}, S)$ is

$$\begin{bmatrix}
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0 \\
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0
\end{bmatrix}$$

We can rewrite this matrix as

$$\begin{bmatrix}
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0 \\
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0
\end{bmatrix}
=
\begin{bmatrix}
0 & 0 & 0 & 1 \\
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0
\end{bmatrix}
+
\begin{bmatrix}
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1 \\
1 & 0 & 0 & 0
\end{bmatrix}$$

or equivalently

$$\begin{bmatrix}
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0 \\
0 & 1 & 0 & 1 \\
1 & 0 & 1 & 0
\end{bmatrix}
= T_x + T_{x^3} = T_x + (T_x)^3$$

From here, it seems that if we could just show that the DFT diagonalizes the matrix T_x , it diagonalizes the adjacency matrix, since the adjacency operator is just the sum of the matrices of elements in S , and in the case of a cyclic group the elements of S are just powers of the generating element x .

For this particular example, we may just multiply the matrix of the DFT and the matrix T_x :

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix}^{-1} = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

The proof for C_n (or $\mathbb{Z}/n\mathbb{Z}$) is similar in many ways, but there are ideas which must be generalized.

First, the adjacency matrix A is just the sum of shift matrices.

Proof. Let $f \in L^2(\mathbb{Z}/n\mathbb{Z})$. Then if A_{ij} is the i, j entry of the adjacency matrix and $y \in \mathbb{Z}/n\mathbb{Z}$,

$$(Af)(y) = \sum_{k=1}^n A_{yk} f(k)$$

which is just matrix multiplication with the column vector in \mathbb{C}^n corresponding to $f \in L^2(\mathbb{Z}/n\mathbb{Z})$. Now, the k^{th} term

$$A_{yk} f(k) = \begin{cases} f(k), & \text{if } A_{yk} = 1 \\ 0, & \text{otherwise} \end{cases}$$

However, $A_{yk} = 1$ is 1 iff v_y is connected to v_k , which happens iff $y-k \in S$. Thus

$$(Af)(y) = \sum_{y-k \in S} f(k) = \sum_{k \in S} f(y-k)$$

After a change of variables $(y-k) \rightarrow k$.

However, we saw above that $f(y-k) = (f * \delta_k)(y) = T_{x^k} f$. Thus

$$(Af)(y) = \sum_{k \in S} (T_{x^k} f)(y) = \sum_{k \in S} ([T_x]^k f)(y)$$

And so the adjacency matrix A is just the sum of the matrices $(T_x)^k$ for $k \in S$. \square

4 Diagonalizing A

These are all nice results, but you're probably asking yourself "Okay, so we went to the trouble to define the DFT in a very particular way, but *why?*"

The reason is because the DFT behaves particularly well with respect to convolution. Consider two functions δ_i and δ_j in the basis for $L^2(\mathbb{Z}/n\mathbb{Z})$, and let $y \in \mathbb{Z}/n\mathbb{Z}$. Then

$$[\mathcal{F}(\delta_a * \delta_b)](y) = [\mathcal{F}(\delta_{a+nb})](y) = [\mathcal{F}(\delta_a)](y) \cdot [\mathcal{F}(\delta_b)](y).$$

Proof. Since the matrix of the DFT was defined with respect to the basis of δ functions, $\mathcal{F}(\delta_{a+nb})$ is just the $(a+nb)^{\text{th}}$ column of the DFT matrix, and $[\mathcal{F}(\delta_{a+nb})](y)$ is just the entry in the y^{th} row of this column. However, recall that we defined the rows of the DFT matrix as homomorphic images of $\mathbb{Z}/n\mathbb{Z}$ in \mathbb{C}^\times and that the i, j entry of the DFT matrix is just $\phi_i(x^j)$. Thus

$$[\mathcal{F}(\delta_{a+nb})](y) = \phi_y(x^{a+nb}) = \phi_y(x^a)\phi_y(x^b).$$

But $\phi_y(x^a)$ is the y, a element in the matrix, or $[\mathcal{F}(\delta_a)](y)$, so

$$\begin{aligned}\phi_y(x^a) &= [\mathcal{F}(\delta_a)](y) \text{ and} \\ \phi_y(x^b) &= [\mathcal{F}(\delta_b)](y).\end{aligned}$$

Substituting into the equation above, we have that

$$[\mathcal{F}(\delta_a * \delta_b)](y) = [\mathcal{F}(\delta_a)](y) \cdot [\mathcal{F}(\delta_b)](y).$$

□

There's at least one more handy property of \mathcal{F} : its inverse is simply $n^{-1}\overline{\mathcal{F}}$.

Proof. We begin by noting that \mathcal{F} is symmetric, and that its rows are orthogonal with respect to the inner product defined in section 1. Thus since \mathcal{F} is symmetric, $\mathcal{F}\overline{\mathcal{F}}$ is diagonal.

Formally, consider the i, j entry of $\mathcal{F}\overline{\mathcal{F}}$. This entry is just

$$\sum_{k=0}^{n-1} \phi_i(k)\overline{\phi_j(k)}.$$

Letting ω denote $e^{-2\pi i/n}$, $\phi_m(n)$ is just ω^{mn} , which follows from the definition of ϕ_m . Thus if $i \neq j$,

$$\sum_{k=0}^{n-1} \phi_i(k) \overline{\phi_j(k)} = \sum_{k=0}^{n-1} \omega^{ik} \omega^{-jk} = \sum_{k=0}^{n-1} (\omega^{i-j})^k.$$

$i \neq j$ and $0 \leq i, j \leq n-1 \implies (i-j) \nmid n$, and since ω generates \mathbb{C}^\times , $\omega^{i-j} \neq 1$. Thus the sum is just a geometric series, and

$$\sum_{k=0}^{n-1} (\omega^{i-j})^k = \frac{(\omega^{i-j})^n - 1}{\omega^{i-j} - 1}.$$

An application of Lagrange's Theorem to \mathbb{C}^\times shows that $(\omega^{i-j})^n$ is 1, and thus the sum is 0.

If $i = j$, then $\omega^{i-j} = 1$, and the sum is just n . Thus

$$\mathcal{F} \overline{\mathcal{F}} = \begin{bmatrix} n & 0 & \cdots & 0 \\ 0 & n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & n \end{bmatrix}$$

and so $(\mathcal{F})(n^{-1} \overline{\mathcal{F}}) = (n^{-1} \overline{\mathcal{F}})(\mathcal{F}) = I_n$ and $\mathcal{F}^{-1} = n^{-1} \overline{\mathcal{F}}$. □

We noted earlier that for $f \in L^2(\mathbb{Z}/n\mathbb{Z})$,

$$T_x f = \delta_1 * f,$$

so if we take the Fourier Transform,

$$(\mathcal{F} T_x f)(a) = [\mathcal{F}(\delta_1 * f)](a) = \mathcal{F} \delta_1(a) \cdot \mathcal{F} f(a).$$

Now, if we set $h = \mathcal{F} f$, we see that

$$[(\mathcal{F} T_x \mathcal{F}^{-1})h](a) = \mathcal{F} \delta_1(a) \cdot h(a).$$

This means that when we apply the matrix $(\mathcal{F} T_x \mathcal{F}^{-1})$ to a function h , the entry in the a^{th} row of the resulting function is a scalar multiple of the a^{th} row, or $h(a)$, of h . This means that the matrix $\mathcal{F} T_x \mathcal{F}^{-1}$ is diagonal. Additionally, the entries on the diagonal must be $\mathcal{F} \delta_1(a)$, which means that these are the eigenvalues of A .

5 The Discrete Fourier Transform on Finite Abelian Groups

The DFT for finite Abelian groups is closely related to the DFT for the cyclic group. This isn't surprising, because the Fundamental Theorem of Abelian Groups tells us that every finite abelian group is isomorphic to a direct product of cyclic groups. We begin with notions analogous to those presented in Section 1.

Recalling the definition of $L^2(G)$ from Section 1, we see that it is a vector space with inner product

$$\langle f, g \rangle = \sum_{x \in G} f(x) \overline{g(x)} \quad \text{for } f, g \in L^2(G).$$

The DFT for $L^2(G)$ is built in a similar way to the DFT for $L^2(\mathbb{Z}/n\mathbb{Z})$.

Definition. A *character* χ of a finite, Abelian group G is a homomorphism

$$\chi : G \rightarrow \mathbb{C}^\times.$$

Formally, if G is a finite Abelian group,

$$G \cong (\mathbb{Z}/n_1\mathbb{Z}) \oplus \cdots \oplus (\mathbb{Z}/n_k\mathbb{Z}).$$

This means we can think of the group G as

$$G = \{(a_1, \dots, a_n) : a_i \in \mathbb{Z}/n_i\mathbb{Z}\}$$

or, equivalently, $G = \{(x_1, \dots, x_n) : x_i \in C_{n_i}\}$ where the operation of G is the componentwise multiplication of each C_{n_i} , so

$$(x_1, \dots, x_k)(y_1, \dots, y_k) = (x_1y_1, \dots, x_ny_n).$$

Definition. The *dual group* \hat{G} of a group G is the set of all characters of G under the operation of pointwise multiplication of functions. That is, for ϕ and χ in \hat{G} and x in G ,

$$(\chi\phi)(x) = \chi(x)\phi(x).$$

Given $\chi \in \hat{G}$ and $x \in G$, we may write $x = (x_1, \dots, x_k)$ and so

$$\chi(x) = \chi(x_1, \dots, x_k) = \chi_1(x_1)\chi_2(x_2) \cdots \chi_k(x_k)$$

where

$$\chi_i : C_{n_i} \rightarrow \mathbb{C}^\times$$

is just $\chi(1, 1, \dots, x_i, 1, \dots, 1)$.

It is easy to see that the identity element in the group \hat{G} is just the trivial homomorphism (denoted 1) which carries all elements of G to 1, and for $\chi \in \hat{G}$, if $\chi(x) = y = e^{i\alpha}$, then

$$\chi^{-1}(x) = y^{-1} = e^{-i\alpha} = \bar{y} = \overline{\chi(x)}.$$

Now, consider the inner product $\langle \chi, \phi \rangle$. To determine its value, we note that

$$\langle \chi, \phi \rangle = \sum_{x \in G} \chi(x) \overline{\phi(x)} = \sum_{x \in G} \chi(x) \phi^{-1}(x) = \sum_{x \in G} (\chi \phi^{-1})(x) = \langle \chi \phi^{-1}, 1 \rangle$$

and so the problem reduces to the case when $\phi = 1$. Assume that $\chi \neq 1$, so that $\chi(y) \neq 0$ for some fixed $y \in G$. Then $\langle \chi, 1 \rangle = \sum_{x \in G} \chi(x)$ implies that

$$\chi(y) \langle \chi, 1 \rangle = \chi(y) \sum_{x \in G} \chi(x) = \sum_{x \in G} \chi(y) \chi(x) = \sum_{x \in G} \chi(yx) = \sum_{z \in G} \chi(z) = \langle \chi, 1 \rangle.$$

But this means that $\chi(y) \langle \chi, 1 \rangle = \langle \chi, 1 \rangle$ with $\chi(y) \neq 0$, and so it must be that $\langle \chi, 1 \rangle = 0$. In the case when $\chi = 1$, it is clear that $\langle \chi, 1 \rangle = \langle 1, 1 \rangle = |G|$. Thus

$$\langle \chi, \phi \rangle = \begin{cases} |G| & \text{if } \chi = \phi \\ 0 & \text{otherwise} \end{cases}$$

and so the characters of G are all orthogonal.

It's worth noting that the method above constitutes an alternate proof that the homomorphisms ϕ_i defined in Section 3 are orthogonal, a fact critical to the proof of the invertibility of \mathcal{F} on $L^2(\mathbb{Z}/n\mathbb{Z})$.

Nearly all of the other properties of the Fourier Transform which we encountered in the section on cyclic groups have analogues here. Many of the proofs are very similar to those given in the first section, and so in the interest of time are omitted here. Instead, we go on to focus on the DFT for general, non-abelian groups. The theory from the general case applies to the Abelian case as well.

6 The General Case: Non-Abelian Groups, Representations, and Characters

For an Abelian group G , it's not hard to see that the dual group consists of products of complex exponentials corresponding to the cyclic subgroups of G . However, for a non-Abelian group, the homomorphisms $\phi : G \rightarrow \mathbb{C}^\times$ are too simple. Let $GL(n, \mathbb{C})$ denote the set of $n \times n$ matrices with entries in \mathbb{C} .

Definition. A *representation* of a finite group G is a group homomorphism $\pi : G \rightarrow GL(n, \mathbb{C})$. A representation π is *unitary* if $\pi(g)$ is a unitary matrix for all $g \in G$. (A matrix U is unitary if its inverse is its complex conjugate transpose.)

Definition. Given a representation π , another representation ρ is a *subrepresentation* of π if there is some subspace W of \mathbb{C}^n such that $\rho(g)W \subseteq W \forall g \in G$ and $\rho(g)$ is just $\pi(g)$ restricted to W . A representation is *irreducible* if its only subrepresentations are itself and the zero representation.

The map T which we defined on page 9 is a representation. In fact, it was a special case of the *left regular representation*.

Definition. The *left regular representation* L of a finite group G is a map

$$L : G \rightarrow GL(|G|, \mathbb{C})$$

where $g \mapsto L(g)$ and for $g, h \in G$ and $a \in L^2(G)$,

$$[L(g)a](h) = a(g^{-1}h).$$

The left regular representation maps a group into permutation matrices (matrices with zero entries except for one 1 in each row and column). Clearly there's more than one way to represent a group. We may say, in some sense, that all representations are created equal; however, some representations are more equal than others.

Definition. Two representations π and ρ are *equivalent* if they differ by a change of basis in $GL(n, \mathbb{C})$. That is, π and ρ satisfy

$$M\pi(g)M^{-1} = \rho(g)$$

for some $M \in GL(n, \mathbb{C})$ and for all $g \in G$.

In the cases of cyclic groups, we found that we were able to come up with a base change which simultaneously diagonalized all the elements of G . For non-Abelian groups, we can't do so well; if all representations were equivalent to scalar matrices, then all group elements would commute. The best we can hope for is block diagonal matrices. To prove this, we first need some theorems. Proofs are available in any book on representation theory.

Theorem 1. Every representation of a finite group is equivalent to some unitary representation.

Theorem 2 (Maschke's Theorem). Let π be a unitary representation with irreducible subrepresentations $\pi_1, \pi_2, \dots, \pi_m$. Then π is equivalent to the following representation in block diagonal form:

$$(\pi_1 \oplus \dots \oplus \pi_m)(g) = \begin{bmatrix} \pi_1(g) & 0 & \cdots & 0 \\ 0 & \pi_2(g) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \pi_m(g) \end{bmatrix}$$

Theorem 3 (Schur's Lemma). Suppose that T is the matrix of some linear operator on \mathbb{C}^n , and that $\pi : G \rightarrow GL(n, \mathbb{C})$ is an irreducible representation. If $T\pi(g) = \pi T(g) \forall g \in G$, then $\exists \lambda \in \mathbb{C}$ such that $T = \lambda I$, where I is the $n \times n$ identity matrix.

However, the point of all the theory is that we might analyze the Cayley graph of a given group. The thing we're really interested are the eigenvalues of the adjacency matrix A of the Cayley graph, which give us an idea of how quickly random walks on the graph converge to a uniform distribution. Rather than actually computing the Fourier transform to find the eigenvalues of A , we can use some tricks to find them. Additionally, we can choose the generating set for the Cayley graph in such a way that finding the eigenvalues of A becomes easier.

We found out in the case of cyclic groups that the adjacency matrix A is actually just the sum of matrices of group elements in the generating set S . If we manage to choose S such that A commutes with $\pi(g)$ for all $g \in G$, then Schur's Lemma tells us that after a change of basis to put π in the form of Theorem 2, A will be a diagonal matrix. It won't be a multiple of the identity matrix unless π is irreducible, but it will consist of blocks along the diagonal which are multiples of identity matrices. This will tell us the

eigenvalues of A , and the size of each block will tell us the multiplicity of the corresponding eigenvalue.

For A to commute with each group element, we want that $A = xAx^{-1}$ for all $x \in G$. From here, it seems logical that we pick A to be the sum of elements in a conjugacy class. If we think of inclusion in the sum A as membership in a set K , then we would just like that $K = xKx^{-1}$ for all $x \in G$.

Proof. Let K be the conjugacy class of an element $g \in G$. An element b of xKx^{-1} is of the form xhx^{-1} , where $h \in cl(g) = K$, so we may rewrite

$$b = xhx^{-1} = x(aga^{-1})x^{-1} = (xa)g(xa)^{-1}$$

for some $a \in G$. The above equation clearly implies that $b \in cl(g) = K$. Since $|K| = |xKx^{-1}|$, we have that $K = xKx^{-1}$. \square

So we know that if we choose A to be the sum of matrices of elements in a collection of conjugacy classes, then A will commute with all matrices of group elements and so by Schur's Lemma be diagonal after some change of basis. Using information about G and some more ideas from representation theory, we can determine each eigenvalue of A along with its multiplicity. First, we must update a few definitions and discover a few more theorems.

Definition. Given a finite group G , the *dual group* \hat{G} is the set of all nonequivalent irreducible unitary representations of G .

Definition. For a representation $\pi : G \rightarrow GL(n, \mathbb{C})$, the degree of π (denoted d_π) is n .

Theorem 4. Let $\pi : G \rightarrow GL(n, \mathbb{C})$ be a representation. Let g, h be two elements of G in the same conjugacy class. Then $\text{Trace}(\pi(g)) = \text{Trace}(\pi(h))$.

Proof. If g and h are conjugate, that means $h = aga^{-1}$ for some element a . Since π is a homomorphism, this means that $\pi(h) = \pi(a)\pi(g)\pi(a)^{-1}$. From linear algebra, we know that for any two matrices A and B , $\text{Trace}(AB) = \text{Trace}(BA)$. Thus $\text{Tr}(\pi(h)) = \text{Tr}(\pi(a)\pi(g)\pi(a)^{-1}) = \text{Tr}(\pi(a)\pi(a)^{-1}\pi(g)) = \text{Tr}(\pi(g))$. \square

When working with a group's Cayley graph, we'll always be using the left regular representation. This is because if $|G| = n$, then A is an $n \times n$ matrix and thus the sum of $n \times n$ matrices of group elements.

Theorem 5. For every irreducible representation $\pi \in \hat{G}$, π appears d_π times in L (the left regular representation). So if $\hat{G} = \{\pi_1, \dots, \pi_m\}$, then

$$L \cong d_{\pi_1} \pi_1 \oplus \dots \oplus d_{\pi_m} \pi_m.$$

By Maschke's Theorem, we see that this means there exists some change of basis for \mathbb{C}^n such that we can write

$$L(g) = \begin{bmatrix} \pi_1(g) & 0 & \cdots & 0 \\ 0 & \ddots & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \pi_m(g) \end{bmatrix}$$

Where π_k appears d_{π_k} times. Since $d_L = |G|$, we have as a quick corollary that

$$\sum_{\pi \in \hat{G}} d_\pi^2 = |G|.$$

Now recall the example of section 2. In general, we want to shuffle a deck of n cards by repeatedly applying certain permutations from a subset K of S_n at random. If we choose K carefully, we can use the ideas we've established above to determine how quickly the cards converge to a random shuffle by seeing how quickly random walks on the Cayley graph converge to uniform. The rate of this convergence depends on the eigenvalues of the adjacency matrix A , so these are what we want to find.

It may seem like the scenario when K is made up of conjugacy classes is too specific to be useful, but for S_n there are many choices of conjugacy class which make for an interesting experiment. For example, all 2-cycles in S_n are conjugate and generate S_n , and so we can easily see how effective 2-card transpositions are at shuffling a deck.

In this case, A is a sum of representations of group elements in the same conjugacy class (and generating set) $K = \text{cl}(g)$ for some $g \in G$. From this, it is easy to see that the blocks of A are of the form

$$T_\pi = \sum_{h \in K} \pi(h)$$

where each block T_π appears d_π times. That is,

$$A \cong d_{\pi_1} T_{\pi_1} \oplus \dots \oplus d_{\pi_m} T_{\pi_m}$$

where $\hat{G} = \{\pi_1, \dots, \pi_m\}$.

Additionally, we know that A commutes with each group element and so by Schur's Lemma, each block T_{π_k} is a scalar multiple of an identity matrix. If we say that $T_{\pi_k} = \lambda_k I$, then we want to find λ_k . If we take the trace, we see that $\text{Tr}(T_{\pi_k}) = d_{\pi_k} \lambda_k$, and so

$$\lambda_k = \frac{\text{Tr}(T_{\pi_k})}{d_{\pi_k}}.$$

We can do even better by noting that

$$\text{Tr}(T_{\pi_k}) = \text{Tr} \sum_{h \in K} \pi_k(h) = \sum_{h \in K} \text{Tr}(\pi_k(h)) = \sum_{h \in K} \text{Tr}(\pi_k(g)) = |\text{cl}(g)| \cdot \text{Tr}(\pi_k(g)).$$

Thus we have that

$$\lambda_k = \frac{|\text{cl}(g)| \cdot \text{Tr}(\pi_k(g))}{d_{\pi_k}}$$

where the element g remains a constant representative of the conjugacy class K we're working in.

This may not seem like a great accomplishment, but all the information on the right hand side of the above equation is readily available from a group's character table. With a little more work, change of basis whose existence is asserted by Maschke's Theorem can actually be computed, and this is the Fourier transform. Using the Fourier transform, it is possible to determine the eigenvalues of adjacency matrices for Cayley graphs with generating sets which are more complicated than conjugacy classes.

References

Gallian, Joseph. Contemporary Abstract Algebra. 6th ed. Boston: Houghton Mifflin Company, 2004.

Hungerford, Thomas W. Algebra. Springer, 2003.

Terras, Audrey. Fourier Analysis on Finite Groups and Applications. New York: Cambridge University Press, 1999.

Acknowledgements

I would like to thank Klaus Lux for his help advising me on this project, and Robert Indik and the University of Arizona for choosing to fund it through the VIGRE grant.