

Lecture 11: Cayley Graphs of Abelian Groups

In which we show how to find the eigenvalues and eigenvectors of Cayley graphs of Abelian groups, we find tight examples for various results that we proved in earlier lectures, and, along the way, we develop the general theory of harmonic analysis which includes the Fourier transform of periodic functions of a real variable, the discrete Fourier transform of periodic functions of an integer variable, and the Walsh transform of Boolean functions.

Earlier, we proved the Cheeger inequalities

$$\frac{\lambda_2}{2} \leq \phi(G) \leq \sqrt{2\lambda_2}$$

and the fact that Fiedler's algorithm, when given an eigenvector of λ_2 , finds a cut $(S, V - S)$ such that $\phi(S, V - S) \leq 2\sqrt{\phi(G)}$. We will show that all such results are tight, up to constants, by proving that

- The dimension- d hypercube H_d has $\lambda_2 = 1 - \frac{2}{d}$ and $h(H_d) = \frac{1}{d}$, giving an infinite family of graphs for which $\frac{\lambda_2}{2} = \phi(G)$, showing that the first Cheeger inequality is exactly tight.
- The n -cycle C_n has $\lambda_2 = O(n^{-2})$, and $\phi(C_n) = \frac{2}{n}$, giving an infinite family of graphs for which $\phi(G) = \Omega(\sqrt{\lambda_2})$, showing that the second Cheeger inequality is tight up to a constant.
- There is an eigenvector of the 2nd eigenvalue of the hypercube H_d , such that Fiedler's algorithm, given such a vector, outputs a cut $(S, V - S)$ of expansion $\phi(S, V - S) = \Omega(1/\sqrt{d})$, showing that the analysis of the Fiedler's algorithm is tight up to a constant.

In this lecture we will develop some theoretical machinery to find the eigenvalues and eigenvectors of *Cayley graphs of finite Abelian groups*, a class of graphs that includes the cycle and the hypercube, among several other interesting examples.

For readers familiar with the Fourier analysis of Boolean functions, or the discrete Fourier analysis of functions $f : \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$, or the standard Fourier analysis of periodic real functions, this theory will give a more general, and hopefully interesting, way to look at what they already know.

1 Characters

We will use additive notation for groups, so, if Γ is a group, its unit will be denoted by 0, its group operation by $+$, and the inverse of element a by $-a$. Unless, noted otherwise, however, the definitions and results apply to non-abelian groups as well.

Definition 1 (Character) *Let Γ be a group (we will also use Γ to refer to the set of group elements). A function $f : \Gamma \rightarrow \mathbb{C}$ is a character of Γ if*

- f is a group homomorphism of Γ into the multiplicative group $\mathbb{C} - \{0\}$.
- for every $x \in \Gamma$, $|f(x)| = 1$

Though this definition might seem to not bear the slightest connection to our goals, the reader should hang on because we will see that finding the eigenvectors and eigenvalues of the cycle C_n is immediate once we know the characters of the group $\mathbb{Z}/n\mathbb{Z}$, and finding the eigenvectors and eigenvalues of the hypercube H_d is immediate once we know the characters of the group $(\mathbb{Z}/2\mathbb{Z})^d$.

Remark 2 (About the Boundedness Condition) *If Γ is a finite group, and a is any element, then*

$$\underbrace{a + \cdots + a}_{|\Gamma| \text{ times}} = 0$$

and so if $f : \Gamma \rightarrow \mathbb{C}$ is a group homomorphism then

$$1 = f(0) = f(\underbrace{a + \cdots + a}_{|\Gamma| \text{ times}}) = f(a)^{|\Gamma|}$$

and so $f(a)$ is a root of unity and, in particular, $|f(a)| = 1$. This means that, for finite groups, the second condition in the definition of character is redundant.

In certain infinite groups, however, the second condition does not follow from the first, for example $f : \mathbb{Z} \rightarrow \mathbb{C}$ defined as $f(n) = e^n$ is a group homomorphism of $(\mathbb{Z}, +)$ into $(\mathbb{C} - \{0\}, \cdot)$ but it is not a character.

Just by looking at the definition, it might look like a finite group might have an infinite number of characters; the above remark, however, shows that a character of a finite group Γ must map into $|\Gamma|$ -th roots of unity, of which there are only $|\Gamma|$, showing a finite $|\Gamma|^{|\Gamma|}$ upper bound to the number of characters. Indeed, a much stronger upper bound holds, as we will prove next, after some preliminaries.

Lemma 3 *If Γ is finite and χ is a character that is not identically equal to 1, then $\sum_{a \in \Gamma} \chi(a) = 0$*

PROOF: Let b be such that $\chi(b) \neq 1$. Note that

$$\chi(b) \cdot \sum_{a \in \Gamma} \chi(a) = \sum_{a \in \Gamma} \chi(b+a) = \sum_{a \in \Gamma} \chi(a)$$

where we used the fact that the mapping $a \rightarrow b+a$ is a permutation. (We emphasize that even though we are using additive notation, the argument applies to non-abelian groups.) So we have

$$(\chi(b) - 1) \cdot \sum_{a \in \Gamma} \chi(a) = 0$$

and since we assumed $\chi(b) \neq 1$, it must be $\sum_{a \in \Gamma} \chi(a) = 0$. \square

If Γ is finite, given two functions $f, g : \Gamma \rightarrow \mathbb{C}$, define the inner product

$$\langle f, g \rangle := \sum_{a \in \Gamma} f(a)g^*(a)$$

Lemma 4 *If $\chi_1, \chi_2 : \Gamma \rightarrow \mathbb{C}$ are two different characters of a finite group Γ , then*

$$\langle \chi_1, \chi_2 \rangle = 0$$

We will prove Lemma 4 shortly, but before doing so we note that, for a finite group Γ , the set of functions $f : \Gamma \rightarrow \mathbb{C}$ is a $|\Gamma|$ -dimensional vector space, and that Lemma 4 implies that characters are orthogonal with respect to an inner product, and so they are linearly independent. In particular, we have established the following fact:

Corollary 5 *If Γ is a finite group, then it has at most $|\Gamma|$ characters.*

It remains to prove Lemma 4, which follows from the next two statements, whose proof is immediate from the definitions.

Fact 6 *If χ_1, χ_2 are characters of a group Γ , then the mapping $x \rightarrow \chi_1(x) \cdot \chi_2(x)$ is also a character.*

Fact 7 *If χ is a character of a group Γ , then the mapping $x \rightarrow \chi^*(x)$ is also a character, and, for every x , we have $\chi(x) \cdot \chi^*(x) = 1$.*

To complete the proof of Lemma 4, observe that:

- the function $\chi(x) := \chi_1(x) \cdot \chi_2^*(x)$ is a character;

- the assumption of the lemma is that there is an a such that $\chi_1(a) \neq \chi_2(a)$, and so, for the same element a , $\chi(a) = \chi_1(a) \cdot \chi_2^*(a) \neq \chi_2(a) \cdot \chi_2^*(a) = 1$
- thus χ is a character that is not identically equal to 1, and so

$$0 = \sum_a \chi(a) = \langle \chi_1, \chi_2 \rangle$$

Notice that, along the way, we have also proved the following fact:

Fact 8 *If Γ is a group, then the set of characters of Γ is also a group, with respect to the group operation of pointwise multiplication. The unit of the group is the character mapping every element to 1, and the inverse of a character is the pointwise conjugate of the character.*

The group of characters is called the Pontryagin dual of Γ , and it is denoted by $\hat{\Gamma}$.

We now come to the punchline of this discussion.

Theorem 9 *If Γ is a finite abelian group, then it has exactly $|\Gamma|$ characters.*

PROOF: We give a constructive proof. We know that every finite abelian group is isomorphic to a product of cyclic groups

$$(\mathbb{Z}/n_1\mathbb{Z}) \times (\mathbb{Z}/n_2\mathbb{Z}) \times \cdots \times (\mathbb{Z}/n_k\mathbb{Z})$$

so it will be enough to prove that

1. the cyclic group $\mathbb{Z}/n\mathbb{Z}$ has n characters;
2. if Γ_1 and Γ_2 are finite abelian groups with $|\Gamma_1|$ and $|\Gamma_2|$ characters, respectively, then their product has $|\Gamma_1| \cdot |\Gamma_2|$ characters.

For the first claim, consider, for every $r \in \{0, \dots, n-1\}$, the function

$$\chi_r(x) := e^{2\pi i r x / n}$$

Each such function is clearly a character (0 maps to 1, $\chi_r(-x)$ is the multiplicative inverse of $\chi_r(x)$), and, recalling that $e^{2\pi i k} = 1$ for every integer k , we also have $\chi_r(a + b \bmod n) = e^{2\pi i r a / n} \cdot e^{2\pi i r b / n}$, and the values of $\chi_r(1)$ are different for different values of r , so we get n distinct characters. This shows that $\mathbb{Z}/n\mathbb{Z}$ has at least n characters, and we already established that it can have at most n characters.

For the second claim, note that if χ_1 is a character of Γ_1 and χ_2 is a character of Γ_2 , then it is easy to verify that the mapping $(x, y) \rightarrow \chi_1(x) \cdot \chi_2(y)$ is a character of $\Gamma_1 \times \Gamma_2$. Furthermore, if (χ_1, χ_2) and (χ'_1, χ'_2) are two distinct pairs of characters, then the mappings $\chi(x, y) := \chi_1(x) \cdot \chi_2(y)$ and $\chi'(x, y) := \chi'_1(x) \cdot \chi'_2(y)$ are two distinct characters of $\Gamma_1 \times \Gamma_2$, because we either have an a such that $\chi_1(a) \neq \chi'_1(a)$, in which case $\chi(a, 0) \neq \chi'(a, 0)$, or we have a b such that $\chi_2(b) \neq \chi'_2(b)$, in which case $\chi(0, b) \neq \chi'(0, b)$. This shows that $\Gamma_1 \times \Gamma_2$ has at least $|\Gamma_1| \cdot |\Gamma_2|$ characters, and we have already established that it can have at most that many \square

This means that the characters of a finite abelian group Γ form an orthogonal basis for the set of all functions $f : \Gamma \rightarrow \mathbb{C}$, so that any such function can be written as a linear combination

$$f(x) = \sum_{\chi} \hat{f}(\chi) \cdot \chi(x)$$

For every character χ , $\langle \chi, \chi \rangle = |\Gamma|$, and so the characters are actually a scaled-up orthonormal basis, and the coefficients can be computed as

$$\hat{f}(\chi) = \frac{1}{|\Gamma|} \sum_x f(x) \chi^*(x)$$

Example 10 (The Boolean Cube) Consider the case $\Gamma = (\mathbb{Z}/2\mathbb{Z})^n$, that is the group elements are $\{0, 1\}^n$, and the operation is bitwise xor. Then there is a character for every bit-vector (r_1, \dots, r_n) , which is the function

$$\chi_{r_1, \dots, r_n}(x_1, \dots, x_n) := (-1)^{r_1 x_1 + \dots + r_n x_n}$$

Every boolean function $f : \{0, 1\}^n \rightarrow \mathbb{C}$ can thus be written as

$$f(x) = \sum_{r \in \{0, 1\}^n} \hat{f}(r) \cdot (-1)^{\sum_i r_i x_i}$$

where

$$\hat{f}(r) = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x) \cdot (-1)^{\sum_i r_i x_i}$$

which is the boolean Fourier transform.

Example 11 (The Cyclic Group) To work out another example, consider the case $\Gamma = \mathbb{Z}/N\mathbb{Z}$. Then every function $f : \{0, \dots, N-1\} \rightarrow \mathbb{C}$ can be written as

$$f(x) = \sum_{r \in \{0, \dots, N-1\}} \hat{f}(r) e^{2\pi i r x / n}$$

where

$$\hat{f}(x) = \frac{1}{N} \sum_x f(x) e^{-2\pi i r x / n}$$

which is the discrete Fourier transform.

2 A Look Beyond

Why is the term "Fourier transform" used in this context? We will sketch an answer to this question, although what we say from this point on is not needed for our goal of finding the eigenvalues and eigenvectors of the cycle and the hypercube.

The point is that it is possible to set up a definitional framework that unifies both what we did in the previous section with finite Abelian groups, and the Fourier series and Fourier transforms of real and complex functions.

In the discussion of the previous section, we started to restrict ourselves to finite groups Γ when we defined an inner product among functions $f : \Gamma \rightarrow \mathbb{C}$.

If Γ is an infinite abelian group, we can still define an inner product among functions $f : \Gamma \rightarrow \mathbb{C}$, but we will need to define a measure over Γ and restrict ourselves in the choice of functions. A measure μ over (a sigma-algebra of subsets of) Γ is a Haar measure if, for every measurable subset A and element a we have $\mu(a + A) = \mu(A)$, where $a + A = \{a + b : b \in A\}$. For example, if Γ is finite, $\mu(A) = |A|$ is a Haar measure. If $\Gamma = (\mathbb{Z}, +)$, then $\mu(A) = |A|$ is also a Haar measure (it is ok for a measure to be infinite for some sets), and if $\Gamma = (\mathbb{R}, +)$ then the Lebesgue measure is a Haar measure. When a Haar measure exists, it is more or less unique up to multiplicative scaling. All *locally compact topological* abelian groups have a Haar measure, a very large class of abelian groups, that include all finite ones, $(\mathbb{Z}, +)$, $(\mathbb{R}, +)$, and so on.

Once we have a Haar measure μ over Γ , and we have defined an integral for functions $f : \Gamma \rightarrow \mathbb{C}$, we say that a function is an element of $L^2(\Gamma)$ if

$$\int_{\Gamma} |f(x)|^2 d\mu(x) < \infty$$

For example, if Γ is finite, then all functions $f : \Gamma \rightarrow \mathbb{C}$ are in $L^2(\Gamma)$, and a function $f : \mathbb{Z} \rightarrow \mathbb{C}$ is in $L^2(\mathbb{Z})$ if the series $\sum_{n \in \mathbb{Z}} |f(n)|^2$ converges.

If $f, g \in L^2(\Gamma)$, we can define their inner product

$$\langle f, g \rangle := \int_{\Gamma} f(x)g^*(x)d\mu(x)$$

and use Cauchy-Schwarz to see that $|\langle f, g \rangle| < \infty$.

Now we can repeat the proof of Lemma 4 that $\langle \chi_1, \chi_2 \rangle = 0$ for two different characters, and the only step of the proof that we need to verify for infinite groups is an analog of Lemma 3, that is we need to prove that if χ is a character that is not always equal to 1, then

$$\int_{\Gamma} \chi(x)d\mu(x) = 0$$

and the same proof as in Lemma 3 works, with the key step being that, for every group element a ,

$$\int_{\Gamma} \chi(x+a)d\mu(x) = \int_{\Gamma} \chi(x)d\mu(x)$$

because of the property of μ being a Haar measure.

We don't have an analogous result to Theorem 9 showing that Γ and $\hat{\Gamma}$ are isomorphic, however it is possible to show that $\hat{\Gamma}$ itself has a Haar measure $\hat{\mu}$, that the dual of $\hat{\Gamma}$ is isomorphic to Γ , and that if $f : \Gamma \rightarrow \mathbb{C}$ is continuous, then it can be written as the "linear combination"

$$f(x) = \int_{\hat{\Gamma}} \hat{f}(\chi)\chi(x)d\hat{\mu}(\chi)$$

where

$$\hat{f}(\chi) = \int_{\Gamma} f(x)\chi^*(x)d\mu(x)$$

In the finite case, the examples that we developed before correspond to setting $\mu(A) := |A|/|\Gamma|$ and $\hat{\mu}(A) = |A|$.

Example 12 (Fourier Series) *The set of characters of the group $[0, 1)$ with the operation of addition modulo 1 is isomorphic to \mathbb{Z} , because for every integer n we can define the function $\chi_n : [0, 1) \rightarrow \mathbb{C}$*

$$\chi_n(x) := e^{2\pi i x n}$$

and it can be shown that there are no other characters. We thus have the Fourier series for continuous functions $f : [0, 1) \rightarrow \mathbb{C}$,

$$f(x) = \sum_{n \in \mathbb{Z}} \hat{f}(n) e^{2\pi i x n}$$

where

$$\hat{f}(n) = \int_0^1 f(x) e^{-2\pi i x n} dx$$

3 Cayley Graphs and Their Spectrum

Let Γ be a finite group. We will use additive notation, although the following definition applies to non-commutative groups as well. A subset $S \subseteq \Gamma$ is *symmetric* if $a \in S \Leftrightarrow -a \in S$.

Definition 13 For a group Γ and a symmetric subset $S \subseteq \Gamma$, the Cayley graph $\text{Cay}(\Gamma, S)$ is the graph whose vertex set is Γ , and such that (a, b) is an edge if and only if $b - a \in S$. Note that the graph is undirected and $|S|$ -regular.

We can also define Cayley *weighted* graphs: if $w : \Gamma \rightarrow \mathbb{R}$ is a function such that $w(a) = w(-a)$ for every $a \in \Gamma$, then we can define the weighted graph $\text{Cay}(G, w)$ in which the edge (a, b) has weight $w(b - a)$. We will usually work with unweighted graphs, although we will sometimes allow parallel edges (corresponding to positive integer weights).

Example 14 (Cycle) The n -vertex cycle can be constructed as the Cayley graph $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{-1, 1\})$.

Example 15 (Hypercube) The d -dimensional hypercube can be constructed as the Cayley graph

$$\text{Cay}((\mathbb{Z}/2\mathbb{Z})^d, \{(1, 0, \dots, 0), (0, 1, \dots, 0), \dots, (0, 0, \dots, 1)\})$$

where the group is the set $\{0, 1\}^d$ with the operation of bit-wise xor, and the set S is the set of bit-vectors with exactly one 1.

If we construct a Cayley graph from a finite abelian group, then the eigenvectors are the characters of the groups, and the eigenvalues have a very simple description.

Lemma 16 Let Γ be a finite abelian group, $\chi : \Gamma \rightarrow \mathbb{C}$ be a character of Γ , $S \subseteq \Gamma$ be a symmetric set. Let A be the adjacency matrix of the Cayley graph $G = \text{Cay}(\Gamma, S)$. Consider the vector $\mathbf{x} \in \mathbb{C}^\Gamma$ such that $x_a = \chi(a)$.

Then \mathbf{x} is an eigenvector of G , with eigenvalue

$$\sum_{s \in S} \chi(s)$$

PROOF: Consider the a -th entry of $M\mathbf{x}$:

$$\begin{aligned} (A\mathbf{x})_a &= \sum_b A_{a,b} x_b \\ &= \sum_{b: b-a \in S} \chi(b) \\ &= \sum_{s \in S} \chi(a+s) \\ &= x_a \cdot \sum_{s \in S} \chi(s) \end{aligned}$$

And so

$$A\mathbf{x} = \left(\sum_{s \in S} \chi(s) \right) \cdot \mathbf{x}$$

□

The eigenvalues of the form $\sum_{s \in S} \chi(s)$, where χ is a character, enumerate all the eigenvalues of the graph, as can be deduced from the following observations:

1. Every character is an eigenvector;
2. The characters are linearly independent (as functions $\chi : \Gamma \rightarrow \mathbb{C}$ and, equivalently, as vectors in \mathbb{C}^Γ);
3. There are as many characters as group elements, and so as many characters as nodes in the corresponding Cayley graphs.

It is remarkable that, for a Cayley graph, a system of eigenvectors can be determined based solely on the underlying group, independently of the set S .

4 The Cycle

The n -cycle is the Cayley graph $\text{Cay}(\mathbb{Z}/n\mathbb{Z}, \{-1, +1\})$. Recall that, for every $n \in \{0, \dots, n-1\}$, the group $\mathbb{Z}/n\mathbb{Z}$ has a character $\chi_r(x) = e^{2\pi i r x/n}$.

This means that for every $r \in \{0, \dots, n-1\}$ we have the eigenvalue

$$\lambda_r = e^{2\pi i r/n} + e^{-2\pi i r/n} = 2 \cos(2\pi r/n)$$

where we used the facts that $e^{ix} = \cos(x) + i \sin(x)$, that $\cos(x) = \cos(-x)$, and $\sin(x) = -\sin(-x)$.

For $r = 0$ we have the eigenvalue 2. For $r = 1$ we have the second largest eigenvalue $2 \cos(2\pi/n) = 2 - \Theta(1/n^2)$. If λ is an eigenvalue of the adjacency matrix, then $1 - \lambda/2$ is an eigenvalue of the normalized Laplacian. From the above calculations, we have that the second smallest Laplacian eigenvalue is $\Theta(n^{-2})$.

The expansion of the cycle is $\phi(C_n) \geq 2/n$, and so the cycle is an example in which the second Cheeger inequality is tight.

5 The Hypercube

The group $\{0, 1\}^d$ with bitwise xor has 2^d characters; for every $r \in \{0, 1\}^d$ there is a character $\chi_r : \{0, 1\}^d \rightarrow \{-1, 1\}$ defined as

$$\chi_r(x) = (-1)^{\sum_i r_i x_i}$$

Let us denote the set S by $\{e^1, \dots, e^d\}$, where we let $e^j \in \{0, 1\}^d$ denote the bit-vector that has a 1 in the j -th position, and zeroes everywhere else. This means that, for every bit-vector $r \in \{0, 1\}^d$, the hypercube has the eigenvalue

$$\sum_j \chi_r(e^j) = \sum_j (-1)^{r_j} = -|r| + d - |r| = d - 2|r|$$

where we denote by $|r|$ the *weight* of r , that is, the number of ones in r .

Corresponding to $r = (0, \dots, 0)$, we have the eigenvalue d .

For each of the d vectors r with exactly one 1, we have the second largest eigenvalue $d - 2$. The second smallest Laplacian eigenvalue is $1 - (d - 2)/d = 2/d$.

Let us compute the expansion of the hypercube. Consider “dimension cuts” of the form $S_i := \{x \in \{0, 1\}^d : x_i = 0\}$. The set S_i contains half of the vertices, and the number of edges that cross the cut $(S_i, V - S_i)$ is also equal to half the number of

vertices (because the edges form a perfect matching), so we have $\phi(S_i, V - S_i) = \frac{1}{d}$ and so $\phi(G) \leq \frac{1}{d}$.

These calculations show that the first Cheeger inequality $\lambda_2/2 \leq \phi(G)$ is tight for the hypercube.

Finally, we consider the tightness of the approximation analysis of Fiedler's algorithm.

We have seen that, in the d -dimensional hypercube, the second eigenvalue has multiplicity d , and that its eigenvectors are vectors $\mathbf{x}^j \in \mathbb{R}^{2^d}$ such that $x_a^j = (-1)^{a_j}$. Consider now the vector $\mathbf{x} := \sum_j \mathbf{x}^j$; this is still clearly an eigenvector of the second eigenvalue. The entries of the vector \mathbf{x} are

$$x_a = \sum_j (-1)^{a_j} = d - 2|a|$$

Suppose now that we apply Fiedler's algorithm using \mathbf{x} as our vector. This is equivalent to considering all the cuts $(S_t, V - S_t)$ in the hypercube in which we pick a threshold t and define $S_t := \{a \in \{0, 1\}^n : |a| \geq t\}$.

Some calculations with binomial coefficients show that the best such "threshold cut" is the "majority cut" in which we pick $t = n/2$, and that the expansion of $S_{n/2}$ is

$$\phi(S_{n/2}, V - S_{n/2}) = \Omega\left(\frac{1}{\sqrt{d}}\right)$$

This gives an example of a graph and of a choice of eigenvector for the second eigenvalue that, given as input to Fiedler's algorithm, result in the output of a cut $(S, V - S)$ such that $\phi(S, V - S) \geq \Omega(\sqrt{\phi(G)})$. Recall that we proved $\phi(S, V - S) \leq 2\sqrt{\phi(G)}$, which is thus tight, up to constants.