Spectral Graph Theory Lecture 5

Rings, Paths, and Cayley Graphs

Daniel A. Spielman September 16, 2014

Disclaimer

These notes are not necessarily an accurate representation of what happened in class. The notes written before class say what I think I should say. I sometimes edit the notes after class to make them way what I wish I had said.

There may be small mistakes, so I recommend that you check any mathematically precise statement before using it in your own work.

These notes were last revised on September 16, 2015.

5.1 Overview

This lecture is devoted to an examination of some special graphs and their eigenvalues.

5.2 The Ring Graph

The ring graph on n vertices, R_n, may be viewed as having a vertex set corresponding to the integers modulo n. In this case, we view the vertices as the numbers 0 through $n-1$, with edges $(i, i+1)$, computed modulo n.

Lemma 5.2.1. *The Laplacian of R_n has eigenvectors*

$$\begin{aligned} \boldsymbol{x}_k(u) &= \cos(2\pi k u/n), \text{ and} \\ \boldsymbol{y}_k(u) &= \sin(2\pi k u/n), \end{aligned}$$

for $0 \leq k \leq n/2$, ignoring $\boldsymbol{y}_0$ which is the all-zero vector, and for even n ignoring $\boldsymbol{y}_{n/2}$ for the same reason. Eigenvectors $\boldsymbol{x}_k$ and $\boldsymbol{y}_k$ have eigenvalue $2 - 2\cos(2\pi k/n)$.

Note that $\boldsymbol{x}_0$ is the all-ones vector. When n is even, we only have $\boldsymbol{x}_{n/2}$, which alternates ± 1.

Proof. We will first see that $\boldsymbol{x}_1$ and $\boldsymbol{y}_1$ are eigenvectors by drawing the ring graph on the unit circle in the natural way: plot vertex u at point $(\cos(2\pi u/n), \sin(2\pi u/n))$.

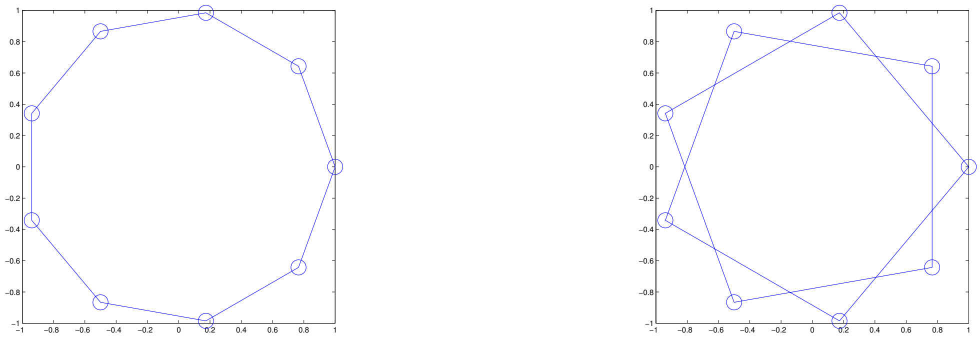

(a) The ring graph on 9 vertices. (b) The eigenvectors for $k = 2$.

Figure 5.1:

You can see that the average of the neighbors of a vertex is a vector pointing in the same direction as the vector associated with that vertex. This should make it obvious that both the x and y coordinates in this figure are eigenvectors of the same eigenvalue. The same holds for all k.

Alternatively, we can verify that these are eigenvectors by a simple computation.

$$\begin{aligned}
(L_{R_n}\boldsymbol{x}_k)(u) &= 2\boldsymbol{x}_k(u) - \boldsymbol{x}_k(u+1) - \boldsymbol{x}_k(u-1) \\
&= 2\cos(2\pi k u/n) - \cos(2\pi k(u+1)/n) - \cos(2\pi k(u-1)/n) \\
&= 2\cos(2\pi k u/n) - \cos(2\pi k u/n)\cos(2\pi k/n) + \sin(2\pi k u/n)\sin(2\pi k/n) \\
&\qquad\qquad - \cos(2\pi k u/n)\cos(2\pi k/n) - \sin(2\pi k u/n)\sin(2\pi k/n) \\
&= 2\cos(2\pi k u/n) - \cos(2\pi k u/n)\cos(2\pi k/n) - \cos(2\pi k u/n)\cos(2\pi k/n) \\
&= (2 - 2\cos(2\pi k/n))\cos(2\pi k u/n) \\
&= (2 - \cos(2\pi k/n))\boldsymbol{x}_k(u).
\end{aligned}$$

The computation for $\boldsymbol{y}_k$ follows similarly. □

5.3 The Path Graph

We will derive the eigenvalues and eigenvectors of the path graph from those of the ring graph. To begin, I will number the vertices of the ring a little differently, as in Figure 5.2.

Lemma 5.3.1. *Let $P_n = (V, E)$ where $V = \{1, \ldots, n\}$ and $E = \{(i, i+1) : 1 \leq i < n\}$. The Laplacian of P_n has the same eigenvalues as R_{2n}, excluding 2. That is, P_n has eigenvalues namely $2(1 - \cos(\pi k/n))$, and eigenvectors*

$$\boldsymbol{v}_k(u) = \cos(\pi k u/n - \pi k/2n).$$

for $0 \leq k < n$

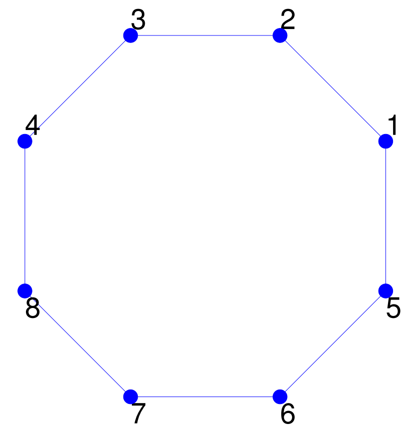

Figure 5.2: The ring on 8 vertices, numbered differently

Proof. We derive the eigenvectors and eigenvalues by treating P_n as a quotient of R_{2n}: we will identify vertex u of P_n with vertices u and $u+n$ of R_{2n} (under the new numbering of R_{2n}). These are pairs of vertices that are above each other in the figure that I drew.

Let $\boldsymbol{I}_n$ be the n-dimensional identity matrix. You should check that

$$\begin{pmatrix} \boldsymbol{I}_n & \boldsymbol{I}_n \end{pmatrix} \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} = 2\boldsymbol{L}_{P_n}.$$

If there is an eigenvector $\boldsymbol{\psi}$ of R_{2n} with eigenvalue λ for which $\boldsymbol{\psi}(u) = \boldsymbol{\psi}(u+n)$ for $1 \leq u \leq n$, then the above equation gives us a way to turn this into an eigenvector of P_n: Let $\boldsymbol{\phi} \in \mathbb{R}^n$ be the vector for which

$$\boldsymbol{\phi}(u) = \boldsymbol{\psi}(u), \text{for } 1 \leq u \leq n.$$

Then,

$$\begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\phi} = \boldsymbol{\psi}, \quad \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\phi} = \lambda \boldsymbol{\psi}, \quad \text{and} \quad \begin{pmatrix} \boldsymbol{I}_n & \boldsymbol{I}_n \end{pmatrix} \boldsymbol{L}_{R_{2n}} \begin{pmatrix} \boldsymbol{I}_n \\ \boldsymbol{I}_n \end{pmatrix} \boldsymbol{\psi} = 2\lambda \boldsymbol{\phi}.$$

So, if we can find such a vector $\boldsymbol{\psi}$, then the corresponding $\boldsymbol{\phi}$ is an eigenvector of P_n of eigenvalue λ.

As you've probably guessed, we can find such vectors $\boldsymbol{\psi}$. I've drawn one in Figure 5.2. For each of the two-dimensional eigenspaces of R_{2n}, we get one such a vector. These provide eigenvectors of eigenvalue

$$2(1 - \cos(\pi k/n)),$$

for $1 \leq k < n$. Thus, we now know $n-1$ distinct eigenvalues. The last, of course, is zero. $\square$

5.4 Cayley Graphs

The ring graph is a type of Cayley graph. In general, the vertices of a Cayley graph are the elements of some group Γ. In the case of the ring, the group is the set of integers modulo n. The edges of a Cayley graph are specified by a set $S \subset \Gamma$, which are called the *generators* of the Cayley graph. The set of generators must be closed under inverse. That is, if $s \in S$, then $s^{-1} \in S$. Vertices $u, v \in \Gamma$ are connected by an edge if there is an $s \in S$ such that

$$u \circ s = v,$$

where $\circ$ is the group operation. In the case of Abelian groups, like the integers modulo n, this would usually be written $u + s = v$. To guarantee that the graph is undirected, we must insist that the inverse of every $s \in S$ also appear in S. In the case of the ring graph, the generators are $\{1, -1\}$.

Many of the most interesting graphs are Cayley graphs. We have seen at least one other: the hybercube.

Cayley graphs over Abeliean groups are particularly convenient because we can find an orthonormal basis of eigenvectors without knowing the set of generators. They just depend on the group[1]. Knowing the eigenvectors makes it much easier to compute the eigenvalues.

5.5 The Hypercube

The d-dimensional hypercube, H_d, is a Cayley graph over the additive group $(\mathbf{Z}/2\mathbf{Z})^d$: that is the set of vectors in $\{0,1\}^d$ under addition modulo 2. The generators are given by the vectors in $\{0,1\}^d$ that have a 1 in exactly one position. This set is closed under inverse, because every element of this group is its own inverse.

5.6 Generalizing Hypercubes

To generalize the hypercube, we will consider this same group, but with a general set of generators. We will call then $\boldsymbol{g}_1, \dots, \boldsymbol{g}_k$, and remember that each is a vector in $\{0,1\}^d$, modulo 2.

Let G be the Cayley graph with these generators. To be concrete, I set $V = \{0,1\}^d$, and note that G has edge set

$$\left\{ (\boldsymbol{x}, \boldsymbol{x} + \boldsymbol{g}_j) : \boldsymbol{x} \in V, 1 \leq j \leq k \right\}.$$

We will now check that the eigenvectors of H_d that I described in Lecture 3 are eigenvectors of G as well. Knowing these will make it easy to describe the eigenvalues.

[1] More precisely, the characters always form an orthonromal set of eigenvectors, and the characters just depend upon the group. When two different characters have the same eigenvalue, we obtain an eigenspace of dimension greater than 1. These eigenspaces do depend upon the choice of generators.

5.7 Analyzing the Eigenvectors and Eigenvalues

For each $\boldsymbol{b} \in \{0,1\}^d$, define the function $\boldsymbol{\psi_b}$ from V to the reals given by

$$\boldsymbol{\psi_b}(\boldsymbol{x}) = (-1)^{\boldsymbol{b}^T\boldsymbol{x}}.$$

When I write $\boldsymbol{b}^T\boldsymbol{x}$, you might wonder if I mean to take the sum over the reals or modulo 2. As both $\boldsymbol{b}$ and $\boldsymbol{x}$ are $\{0,1\}$-vectors, you get the same answer either way you do it.

While it is natural to think of $\boldsymbol{b}$ as being a vertex, that is the wrong perspective. Instead, you should think of $\boldsymbol{b}$ as indexing a Fourier coefficient (if you don't know what a Fourier coefficient is, just don't think of it as a vertex).

The eigenvectors and eigenvalues of the graph are determined by the following theorem. As this graph is k-regular, the eigenvectors of the adjacency and Laplacian matrices will be the same.

Lemma 5.7.1. *For each $\boldsymbol{b} \in \{0,1\}^d$ the vector $\boldsymbol{\psi_b}$ is a Laplacian matrix eigenvector with eigenvalue*

$$k - \sum_{i=1}^{k} (-1)^{\boldsymbol{b}^T\boldsymbol{g}_i}.$$

Proof of Theorem **??**. We begin by observing that

$$\boldsymbol{\psi_b}(\boldsymbol{x}+\boldsymbol{y}) = (-1)^{\boldsymbol{b}^T(\boldsymbol{x}+\boldsymbol{y})} = (-1)^{\boldsymbol{b}^T\boldsymbol{x}}(-1)^{\boldsymbol{b}^T\boldsymbol{y}} = \boldsymbol{\psi_b}(\boldsymbol{x})\boldsymbol{\psi_b}(\boldsymbol{y}).$$

Let $\boldsymbol{L}$ be the Laplacian matrix of the graph. For any vector $\boldsymbol{\psi_b}$ for $\boldsymbol{b} \in \{0,1\}^d$ and any vertex $\boldsymbol{x} \in V$, we compute

$$\begin{aligned}
(\boldsymbol{L}\boldsymbol{\psi_b})(\boldsymbol{x}) &= k\boldsymbol{\psi_b}(\boldsymbol{x}) - \sum_{i=1}^{k} \boldsymbol{\psi_b}(\boldsymbol{x}+\boldsymbol{g}_i) \\
&= k\boldsymbol{\psi_b}(\boldsymbol{x}) - \sum_{i=1}^{k} \boldsymbol{\psi_b}(\boldsymbol{x})\boldsymbol{\psi_b}(\boldsymbol{g}_i) \\
&= \boldsymbol{\psi_b}(\boldsymbol{x})\left(k - \sum_{i=1}^{k} \boldsymbol{\psi_b}(\boldsymbol{g}_i)\right).
\end{aligned}$$

So, $\boldsymbol{\psi_b}$ is an eigenvector of eigenvalue

$$k - \sum_{i=1}^{k} \boldsymbol{\psi_b}(\boldsymbol{g}_i) = k - \sum_{i=1}^{k} (-1)^{\boldsymbol{b}^T\boldsymbol{g}_i}.$$

□

5.8 A random set of generators

We will now show that if we choose the set of generators uniformly at random, for k some constant multiple of the dimension, then we obtain a graph that is a good approximation of the complete graph. That is, all the eigenvalues of the Laplacian will be close to k. I will set $k = cd$, for some $c > 1$. Think of $c = 2$ or $c = 10$.

For $\boldsymbol{b} \in \{0,1\}^d$ but not all zero, and for $\boldsymbol{g}$ chosen uniformly at random from $\{0,1\}^d$, $\boldsymbol{b}^T\boldsymbol{g}$ modulo 2 is uniformly distributed in $\{0,1\}$, and so

$$(-1)^{\boldsymbol{b}^T\boldsymbol{g}}$$

is uniformly distributed in ± 1. So, if we pick $\boldsymbol{g}_1, \dots, \boldsymbol{g}_k$ independently and uniformly from $\{0,1\}^d$, the eigenvalue corresponding to the eigenvector $\boldsymbol{\psi}_{\boldsymbol{b}}$ is

$$\lambda_{\boldsymbol{b}} \stackrel{\text{def}}{=} k - \sum_{i=1}^{k} (-1)^{\boldsymbol{b}^T\boldsymbol{g}_i}.$$

The right-hand part is a sum of independent, uniformly chosen ± 1 random variables. So, we know it is concentrated around 0, and thus $\lambda_{\boldsymbol{b}}$ will be concentrated around k. To determine how concentrated the sum actually is, we use a Chernoff bound. There are many forms of Chernoff bounds. I will not use the strongest, but which will give us results that are qualitatively correct.

Theorem 5.8.1. *Let $x_1, \dots, x_k$ be independent ± 1 random variables. Then, for all $h > 0$,*

$$Pr\left[\left|\sum_i x_i\right| \geq h\right] \leq 2e^{-h^2/2k}.$$

This becomes very small when h is a constant fraction of k. In fact, it becomes so small that it is unlikely that any eigenvalue deviates from k by more than h.

Theorem 5.8.2. *With high probability, all of the nonzero eigenvalues of the generalized hypercube differ from k by at most*

$$k\sqrt{\frac{2}{c}},$$

where $k = cd$.

Proof. Let $h = k\sqrt{2/c}$. Then, for every nonzero $\boldsymbol{b}$,

$$\Pr\left[|k - \lambda_{\boldsymbol{b}}| \geq h\right] \leq 2e^{-h^2/2k} \leq 2e^{-k/c} = 2e^{-d}.$$

Now, the probability that there is some $\boldsymbol{b}$ for which $\lambda_{\boldsymbol{b}}$ violates these bounds is at most the sum of these terms:

$$\Pr\left[\exists \boldsymbol{b} : |k - \lambda_{\boldsymbol{b}}| \geq h\right] \leq \sum_{\boldsymbol{b} \in \{0,1\}^d, \boldsymbol{b} \neq 0^d} \Pr\left[|k - \lambda_{\boldsymbol{b}}| \geq h\right] \leq (2^d - 1)2e^{-d},$$

which is always less than 1 and goes to zero exponentially quickly as d grows. $\square$

I initially suggested thinking of $c = 2$ or $c = 10$. The above bound works for $c = 10$. To get a useful bound for $c = 2$, we need to sharpen the analysis. A naive sharpening will work down to $c = 2 \ln 2$. To go lower than that, you need a stronger Chernoff bound.

5.9 Conclusion

We have now seen that a random generalize hypercube probably has all non-zero Laplacian eigenvalues between

$$k(1 - \sqrt{2/c}) \quad \text{and} \quad k(1 + \sqrt{2/c}).$$

If we let n be the number of vertices, and we now multiply the weight of every edge by n/k, we obtain a graph with all nonzero Laplacian eigenvalues between

$$n(1 - \sqrt{2/c}) \quad \text{and} \quad n(1 + \sqrt{2/c}).$$

Thus, this is a $\sqrt{2/c}$ approximation of the complete graph on n vertices. But, the degree of every vertex is only $c \log_2 n$. Expanders are infinite families of graphs like this, but with no dependence on n in their degrees.

One other useful property of this generalized hypercube is that it has a very compact description. The number of bits needed to describe its generators is cd^2, despite the fact that it has 2^d vertices. This allows us to use it for many applications that require us to implicitly deal with a very large graph. A few weeks from now, we will see how to use such graphs to construct pseudo-random generators.

5.10 Acknowledgment

I thank Zeyuan Allen Zhu for pointing out a mistake in my 2009 lecture notes on the path graph.