Spectral Graph Theory and its Applications September 9, 2004

Lecture 3

Lecturer: Daniel A. Spielman

3.1 Introduction

There are two topics that I want to cover in this lecture. The first is a technique for lower bounding the second-smallest eigenvalue of a Laplacian (one could also apply it to the other small eigenvalues).

The second is a proof that the kth eigenvector of a weighted path graph osciallates at most $k-1$ times. There is a chance that we don't get to this today.

3.2 Lower bounding λ_2

It is simple to prove an upper bound on λ_2: one need merely present any test vector orthogonal to $\mathbf{1}$ and observe that λ_2 is less than the Rayleigh quotient of that vector.

Proving a lower bound is not as easy. In this part of the lecture, I will present one general-purpose technqiue for lower bounding λ_2: by comparing the Laplacian of the given graph to the Laplacian of another graph whose spectra we already understand. Let's begin with some crucial notation used in the Optimization community. For a symmetric matrix A, we write

$$A \succcurlyeq 0$$

if A is positive semidefinite. We similarly write

$$A \succcurlyeq B$$

if

$$A - B \succcurlyeq 0.$$

I find it convenient to overload this notation by defining it for graphs as well. Thus, I'll write

$$G \succcurlyeq H$$

if $L_G \succcurlyeq L_H$. This notation is most powerful when we consider some multiple of a graph. Thus, I could write

$$c \cdot G \succcurlyeq H.$$

Using the Courant-Fischer Theorem, we can prove

Lemma 3.2.1. *If G and H are graphs such that*

$$c \cdot G \succcurlyeq H,$$

then

$$c \cdot \lambda_k(G) \geq \lambda_k(H),$$

for all k.

Proof. By part (2) of the Courant-Fischer Theorem, we have

$$c \cdot \lambda_k(G) = \max_{S \text{ of dim } n-k-1} \min_{x \in S} \frac{c \cdot x^T L_G x}{x^T x} \geq \max_{S \text{ of dim } n-k-1} \min_{x \in S} \frac{\cdot x^T L_H x}{x^T x} = \lambda_k(H).$$

$\square$

Now, the question is: how do we prove that $c \cdot G \succcurlyeq H$ for some graph G and H? Not too many ways are known. We'll do it by proving some identities of this form for some of the simplest graphs, and then extending them to more general graphs. For example, we will prove

$$n \cdot P_n \succcurlyeq G_{1,n}. \tag{3.1}$$

That is, n times the path of length n from vertex 1 to n is greater than the edge from 1 to n. This statement is equivalant to the statement that n serially connected unit resistors have resistance n. To prove (3.1), we will have to consider what happens when we connect resistors of varying resistance. Speaking graph theoretically, this means that we are going to have to consider weighted graphs. Let me define them now.

A weighted graph $G = (V, E, w)$ consists of vertex set V, a set of edges, each of which is an unordered pair of vertices, and a weight function $w : E \to \mathbb{R}$. The Laplacian of G has the form

$$L_G = \sum_{(u,v) \in E} w(u,v) \cdot L_{(u,v)},$$

where I recall that $L_{(u,v)}$ denotes the Laplacian of the graph containing just the edge of unit weight between u and v. Note that if $G = (V, E, w)$ and $H = (G, V, w')$ are two weighted graphs with the same set of edges such that $w(u,v) \geq w'(u,v)$ for all $(u,v) \in E$, then

$$G \succcurlyeq H.$$

Lemma 3.2.2. *Let $c_1, \ldots, c_{n-1} > 0$. Then,*

$$c \cdot \left(\sum_{i=1}^{n-1} c_i L_{(i,i+1)} \right) \succcurlyeq L_{(1,n)},$$

where

$$c = \sum_i (1/c_i).$$

Proof. We will prove the equivalent statement:

$$\sum_{i=1}^{n-1} c_i L_{(i,i+1)} \succcurlyeq \left(\frac{1}{\sum_i (1/c_i)} \right) L_{(1,n)}.$$

We first note that if we can prove this lemma for $n = 3$, then we can prove it for all n. To see this, observe that we could then prove

$$\begin{aligned} \sum_{i=1}^{n-1} c_i L_{(i,i+1)} &\succcurlyeq \left(\frac{1}{\sum_{i=1}^{n-2} (1/c_i)} \right) \cdot L_{(1,n-1)} + c_{n-1} \cdot L_{(n-1,n)} \text{ (assuming proved for } n-1) \\ &\succcurlyeq \left(\frac{1}{\sum_{i=1}^{n-1} (1/c_i)} \right) \cdot L_{(1,n)}, \end{aligned}$$

where in this last inequality, we use the case for $n = 3$.

In the case $n = 3$, it suffices to prove the inequality in the case $1/c_1 + 1/c_2 = 1$. If we now let

$$\begin{aligned} a &= (x_1 - x_2), \text{ and} \\ b &= (x_2 - x_3), \end{aligned}$$

then the inequality reduces to

$$c_1 a^2 + c_2 b^2 \geq (a + b)^2.$$

This is Cauchy's inequality. □

3.2.1 Path Graphs

Now, let's use (3.1) to lower bound λ_2 of the path graph on n vertices. We'll do it by comparing the path graph to the complete graph on n vertices. For each edge $(u, v) \in K_n$, with $u < v$, we will apply the inequality

$$(v - u) \sum_{i=u}^{v-1} L_{(i,i+1)} \succcurlyeq L_{(u,v)}.$$

Summing over all pairs u and v, we obtain

$$\sum_{1 \leq u < v \leq n} (v - u) \sum_{i=u}^{v-1} L_{(i,i+1)} \succcurlyeq \sum_{u<v} L_{(u,v)} = L_{K_n}.$$

On the other hand, we also have

$$\sum_{1 \leq u < v \leq n} (v - u) = \sum_{i=1}^{n-1} i(n - i) \leq n^3/4.$$

So,

$$(n^3/4) \cdot L_{P_n} \succcurlyeq \sum_{1 \leq u < v \leq n} (v - u) \sum_{i=u}^{v-1} L_{(i,i+1)},$$

and

$$(n^3/4) \cdot P_n \succcurlyeq K_n.$$

As $\lambda_2(K_n) = n$, this implies

$$\lambda_2(P_n) \geq 4/n^2.$$

This bound isn't so far off from the easy upper bound we proved last lecture of $12/n^2$.

3.2.2 The Complete Binary Tree

Now, let's apply this to a graph that we haven't yet analyzed: the complete binary tree.

The complete binary tree on $n = 2^d - 1$ nodes, T_n, is the graph with edges of the form $(u, 2u)$ and $(u, 2u+1)$ for $u < n/2$. Pictorially, these graphs look like this:

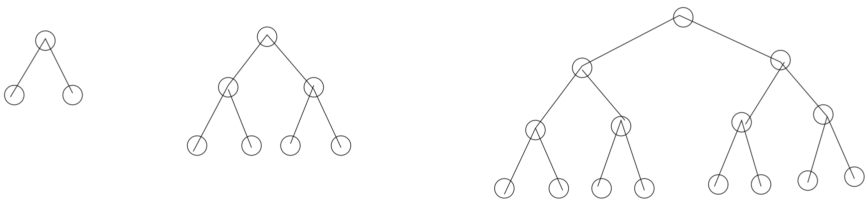

Figure 3.1: T_3, T_7 and T_{15}. Node 1 is at the top, 2 and 3 are its children.

Let's first upper bound $\lambda_2(B_n)$ by constructing a test vector x. Set $x(1) = 0$, $x(2) = 1$, and $x(3) = -1$. Then, for every vertex u that we can reach from node 2 without going through node 1, we set $x(u) = 1$. For all the other nodes, we set $x(u) = -1$. We then have

$$\lambda_2 \leq \frac{\sum_{(u,v)\in B_n} (x_u - x_v)^2}{\sum_u x_u^2} = \frac{(x_1 - x_2)^2 + (x_1 - x_3)^2}{n} = 2/n.$$

We will again prove a lower bound by comparing T_n to the complete graph. For each edge $(u, v) \in K_n$, let $T_n(u, v)$ denote the unique path in T from u to v. This path will have length at most $2\log_2 n$. So, we have

$$L_{K_n} = \sum_{u<v} L_{u,v} \preccurlyeq \sum_{u<v} (2\log_2 n) T_n(u, v) \preccurlyeq \binom{n}{2} (2\log_2 n) T_n.$$

So, we obtain the bound

$$\binom{n}{2} (2\log_2 n) \lambda_2(T_n) \geq n,$$

which implies

$$\lambda_2(T_n) \geq \frac{1}{(n-1)\log_2 n}.$$

In the problem set, I will ask you to improve this lower bound to $1/cn$ for some constant c.

3.3 Weighted Path Graphs

We will now prove the following theorem of Fiedler:

Theorem 3.3.1. *Let P be a weighted path graph on n vertices, and let v_k be the kth eigenvector of its Laplacian. Then, v_k changes sign $k-1$ times.*

The main ingredient in the proof will be Sylvester's law of intertia, which I will first recall and prove.

Theorem 3.3.2 (Sylvester's Law of Intertia). *Let A be any symmetric matrix and let B be any non-degenerate matrix. Then, the matrix $B^T AB$ has the same number of positive, negative and zero eigenvalues as A.*

Proof. We first recall that every non-degenerate matrix B can be factored into the product of an orthogonal matrix Q and an upper-triangular matrix R with positive diagonals[1] Now, since $Q^T = Q^{-1}$, $Q^T AQ$ has exactly the same eigenvalues as A. Let R_t be the matrix $t * R + (1-t)I$, and consider the family of matrices $M_t = R_t^T Q^T AQR_t$, as t goes from 0 to 1. At $t = 0$, the matrix has the same eigenvalues as A. At $t = 1$, we get $B^T AB$. As the eigenvalues of a matrix are continuous functions, if the number of positive, negative or zero eigenvalues of $B^T AB$ differs from that of A, then there must be some t for which M_t has more zero eigenvalues than does A. But, as none of the matrices R_t are degenerate, this cannot happen. □

Proof of Theorem 3.3.1. We will just consider the case in which v_k has no zero entries. In this case, we wish to show that the number of i for which $v_k(i)v_k(i+1) < 0$ equals $k-1$.

Let L denote the Laplacian of P, and let V_k denote the matrix with v_k on the diagonal that is zero elsewhere. Let λ_k denote the k-th eigenvalue of L. Consider the matrix

$$M = V_k^T (L - \lambda_k I)V_k.$$

By Sylvester's law of intertia, we know that M has $k-1$ negative eigenvalues, one zero eigenvalue, and $n-k$ positive eigenvalues. Note that

$$M\mathbf{1} = \mathbf{0},$$

and that for every i such that $v_k(i)v_k(i+1) < 0$, $M_{i,i+1}$ is positive. Thus, by Lemma 3.3.3, which we will state and prove momentarily, there are exactly $k-1$ such i. □

Lemma 3.3.3. *Let M be a symmetric tri-diagonal matrix with $2p$ positive off-diagonal entries such that*

$$M\mathbf{1} = \mathbf{0}. \tag{3.2}$$

Then, M has p negative eigenvalues.

[1]This is the QR-factorization. It follows from Gram-Schmidt orthonormalization.

Proof. From (3.2), we have

$$x^T M x = \sum_{i=1}^{n-1} -M_{i,i+1}(x_i - x_{i+1})^2.$$

We now apply a change of variables from $x_1, \ldots, x_n$ to $\delta_1, \delta_2, \ldots, \delta_2$, where

$$x_i = \delta_1 + \delta_2 + \cdots + \delta_i.$$

This change of variables is realized by the lower-triangular matrix L which has 1's on and below the diagonal:

$$x = L\delta.$$

By Sylvester's law of intertia, we know that

$$L^T M L$$

has the same number of positive, negative, and zero eigenvalues as M. On the other hand,

$$\delta^T L^T M L \delta = \sum_{i=2}^{n} -M_{i,i+1}\delta_i^2,$$

so this matrix clearly has one zero eigenvalue, and as many negative eigenvalues as there are negative $M_{i,i+1}$. $\square$