Fiedler's Theory of Spectral Graph Partitioning

Brian Slininger
Department of Computer Science
University of California, Davis
bslininger@ucdavis.edu

Abstract

Given a connected graph G, it is sometimes desired to divide G into two subgraphs in such a way to minimize the number of edges which have one end vertex in each of the two subgraphs. There are multiple methods of accomplishing this. One method, proposed and proved by Fiedler, is called spectral graph partitioning. It involves using the Laplacian matrix associated with G and dividing the vertices of G into the two subgraphs by using one of the eigenvectors of the Laplacian matrix. In this paper, we restate Fiedler's theorem that the two subgraphs produced by this method are connected, and walk through Demmel's version of the proof in a straightforward way. Afterward, we examine small graphs and apply Fiedler's method to them to observe how his method effectively partitions the graphs.

1 Introduction

Spectral graph partitioning is a method of partitioning a graph into two subgraphs in such a way that the subgraphs have a nearly equal number of vertices (as close to equal as is possible) while also minimizing the number of edges between the two subgraphs. The *spectrum* of a matrix is the set of that matrix's eigenvalues; therefore, spectral partitioning of a graph uses the eigenvalues of a matrix associated with the graph to perform the optimal partition.

We present the theory and method used by Fiedler to perform a spectral graph partition and restate the proof of his theorem that the subgraphs generated by a spectral graph partition are, in fact, connected.

2 Theory

2.1 Notation and Preliminary Terms

To begin, we define terms and notation that will be used in the rest of this document.

A *graph* in this document will always refer to a simple, undirected graph $G = (V, E)$ which contains a set of *vertices* V and a set of *edges* E. Edges connect vertices and are labeled by their endpoints: for $u, v \in V$, if an edge connects vertices u and v, then $(u, v) \in E$. Vertices will each be assigned a unique positive integer n, where $1 \leq n \leq |V|$. That is, the vertices will be labeled with numbers $1, 2, \ldots, |V|$, with no two vertices sharing the same label. The *degree* of a vertex v, written $d(v)$, is equal to the number of edges incident on v, that is, the number of edges in E of the form (v, k), where $k \in V$.

Given a graph $G = (V, E)$, multiple matrices can be defined from it. The *adjacency matrix* of G, $A(G) = (a_{i,j})$ is defined such that

$$a_{i,j} = \begin{cases} 1 & , \quad (i,j) \in E \\ 0 & , \quad (i,j) \notin E \end{cases}$$

(Notice that the labels of vertices are being uses as matrix indices.)

Another matrix that can be defined on G is the *degree matrix*, $D(G) = (d_{i,j})$, defined such that

$$d_{i,j} = \begin{cases} d(i) & , \quad i = j \\ 0 & , \quad i \neq j \end{cases}$$

Finally, the *Laplacian matrix* of the graph G is the matrix $L(G) = (l_{i,j})$, defined such that $L(G) = D(G) - A(G)$ [4].

Given a square matrix $A \in \mathbb{R}^{n \times n}$ and a vector $\mathbf{z} \in \mathbb{R}^n$, the product $A\mathbf{z}$ is another vector in $\mathbb{R}^n$. There are special vectors $\mathbf{x}$ for which multiplying A by $\mathbf{x}$ results in a vector that is in the same direction as $\mathbf{x}$ but has a different magnitude; its length has been *scaled* by a scalar factor λ. In this case, $\mathbf{x}$ is called an *eigenvector* of A and λ is the *eigenvalue* corresponding to the eigenvector $\mathbf{x}$. More formally, $\mathbf{x}$ is an eigenvalue of A with corresponding eigenvector λ if it is true that $A\mathbf{x} = \lambda\mathbf{x}$. An $n \times n$ matrix will have at most n of these eigenvector-eigenvalue pairs; multiple eigenvectors can have the same eigenvalue, but any repeated eigenvector will not have different eigenvalues for each repetition (and they are considered to be the same eigenvector). Note that λ can be complex.

To calculate the eigenvalues of an $n \times n$ matrix A, find all possible solutions λ of the equation $\det(A - \lambda I_n)$, where I_n is the $n \times n$ identity matrix. To calculate the eigenvector corresponding to the eigenvalue λ, solve the equation $(A + \lambda I_n)\mathbf{x} = \mathbf{0}$, where $\mathbf{0}$ is the vector of appropriate size (here of length n) where every entry is 0.

2.2 Method of Spectral Graph Partitioning

Fiedler's theory of spectral graph partitioning is based upon a very simple idea. For a graph $G = (V, E)$ which is to be partitioned, the Laplacian matrix $L(G)$ is formed, and the eigenvector-eigenvalue pairs of $L(G)$ are calculated. Label the eigenvalues so that $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$. The eigenvector $\mathbf{x}_2$, which corresponds to λ_2 (the second-smallest eigenvalue), is known as the *Fiedler vector*, and is used to partition the vertices [6].

Let $n = |V|$. Recall that each of the n vertices of G is assigned a label, which is a unique number in the set $\{1, 2, \ldots, n\}$. Also note that $L(G) \in \mathbb{R}^{n \times n}$ and that $\mathbf{x}_2 \in \mathbb{R}^n$. Therefore each vertex corresponds to a single entry of $\mathbf{x}_2$. More precisely, the vertex with label i corresponds to entry i of $\mathbf{x}_2$, which we will call x_i.

To form the partition, simply create two graphs G_1 and G_2, and for each vertex i of G, if $x_i < 0$, put vertex i in G_1; otherwise, put it in G_2 [2]. This defines a partition of G into two sets, but just by examining this method, it is not evident that this partition attempts to minimize the number of edges between G_1 and G_2, or even that the two subgraphs are connected.

2.3 Reasoning for Correctness of Partition

2.3.1 Why using the second-smallest eigenvalue effectively divides the matrix

To begin, it is not even clear at first why using the second-smallest eigenvalue would produce such a partition. In fact, the use of eigenvectors and eigenvalues to perform a graph partition is not at all an obvious thing to try.

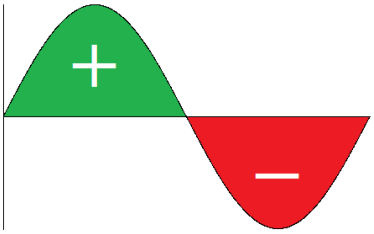

Figure 2.1: A string vibrating at a frequency related to the second-smallest eigenvalue and its corresponding eigenvector of a Laplacian graph will vibrate with a standing wave equal to this one: a sine wave whose wavelength is equal to the length of the string. Half of the string is above the equilibrium line, corresponding to points with positive entries in the eigenvector, and the other half is below it, corresponding to points with negative entries.

To explain this, Demmel uses an example of a vibrating string [2]. When a string is plucked, waves propogate along its length at a certain frequency. If the frequency causes the wave to appear as if it does not move, the wave is called a *standing wave*. If the length of the string is l, this will occur when the wavelength is $\frac{2l}{n}$ for $n \in \mathbb{N}$. It turns out that the frequencies that generate these standing waves are tied directly to the eigenvalues and eigenvectors of the matrix that describes the motion of individual points along the string. (The points are the vertices and the section of string between two points forms an edge, so the string is thought of as a chain graph; the matrix that is used is the Laplacian matrix of this chain graph multiplied by a spring constant k which does not affect the eigenvectors or which eigenvalue is the second-smallest.) A more in-depth analysis of the physics involved in determining the eigenvectors is provided by Demmel at `http://www.cs.berkeley.edu/~demmel/cs267/lecture20/lecture20.html`.

The second-smallest eigenvalue corresponds to a standing wave that is equal to one whole wavelength. As shown in Figure 2.1, this results in a wave with half of the points below the equilibrium point, and half above. The value of the entry for each vertex in the corresponding eigenvector determines the "height" of the vertex above or below the line. Using the sign of this "height" as the tool of determining which subgraph to place the vertices in, it is apparent that about half of the vertices would go into each subgraph. The actual number depends on their position on the string. For general graphs, the same analogy applies, but in a more complicated way. Demmel suggests visualizing a surface for the case of planar graphs; about one half of the surface would be below an equilibrium plane and the other half would be above it.

2.3.2 The connectedness of the subgraphs

The most important result proved by Fiedler is that the two subgraphs resulting from the spectral partition are always connected. Before getting into his proof, we will follow Demmel once again and introduce some preliminary definitions and lemmas [3].

Definition 2.1: A matrix A is *nonnegative*, denoted by $A \geq 0$, if, for every element $a_{i,j}$, we have $a_{i,j} \geq 0$. Similar notions exist for positive, nonpostive, and nonnegative matrices using the appropriate relational symbols. This definition can also be extended to vectors.

Definition 2.2: The *spectral radius* of an $n \times n$ matrix A, denoted by $\rho(A)$, is given by $\rho(A) = \max_i |\lambda_i|$, where $\lambda_1, \lambda_2, \ldots, \lambda_n$ are the eigenvalues of A. That is, $\rho(A)$ is the eigenvalue of A with maximum absolute value.

Definition 2.3: An $n \times n$ matrix A is called *positive-definite* if, for all vectors $\mathbf{x} \in \mathbb{R}^n$ where $\mathbf{x} \neq \mathbf{0}$, $\mathbf{x}^T A\mathbf{x} > 0$. The matrix A is *positive-semidefinite* if $\mathbf{x}^T A\mathbf{x} \geq 0$.

Remark 2.4: The matrix A is symmetric and positive-definite if and only if every eigenvalue of A is positive. Also, A is symmetric and positive-semidefinite if and only if every eigenvalue of A is nonnegative. This can be shown by performing the eigenvalue decomposition of A: $A = Q\Lambda Q^{-1}$, where $Q^T = Q^{-1}$. Letting $\mathbf{z} \in \mathbb{R}^n$ be any nonzero vector, $\mathbf{z}^T A\mathbf{z} > 0$ if and only if $\mathbf{z}^T(Q^{-1}\Lambda Q)\mathbf{z} = (Q\mathbf{z})^T\Lambda(Q\mathbf{z}) > 0$. Since Λ is a diagonal matrix, this will always be true only if each of its diagonal elements is positive. Since Λ contains the eigenvalues of A, each eigenvalue must be positive. A similar argument works for showing positive-semidefiniteness.

Lemma 2.5: If A is an $n \times n$ symmetric matrix with eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$, then $\lambda_1 = \min_{\mathbf{v}\neq 0} \frac{\mathbf{v}^T A\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$ and $\lambda_n = \max_{\mathbf{v}\neq 0} \frac{\mathbf{v}^T A\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$. That is, λ_1 and λ_n are the minimum and maximum values of the expression $\frac{\mathbf{v}^T A\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$ taken over all nonzero vectors $\mathbf{v} \in \mathbb{R}^{n\times n}$.

Proof: Perform the eigenvalue decomposition of A: $A = Q\Lambda Q^{-1}$, where $Q^{-1} = Q^T$, meaning that this can also be wrtten $A = Q\Lambda Q^T$. Note that $QQ^T = I_n$. Now, $\frac{\mathbf{v}^T A\mathbf{v}}{\mathbf{v}^T\mathbf{v}} = \frac{\mathbf{v}Q\Lambda Q^T\mathbf{v}}{\mathbf{v}^T QQ^T\mathbf{v}}$. Let $\mathbf{y} = Q^T\mathbf{v}$. Then, this expression is equal to $\frac{\mathbf{y}\Lambda\mathbf{y}^T}{\mathbf{y}\mathbf{y}^T}$ which is equal to $\frac{\lambda_1 {y_1}^2+\lambda_2 {y_2}^2+\cdots+\lambda_n {y_n}^2}{{y_1}^2+{y_2}^2+\cdots+{y_n}^2}$. Let this expression be equal to α. Then, since $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$, we have that $\alpha \geq \frac{\lambda_1 {y_1}^2+\lambda_1 {y_2}^2+\cdots+\lambda_1 {y_n}^2}{{y_1}^2+{y_2}^2+\cdots+{y_n}^2}$ (with all instances of λ_i replaced by λ_1) and $\alpha \leq \frac{\lambda_n {y_1}^2+\lambda_n {y_2}^2+\cdots+\lambda_n {y_n}^2}{{y_1}^2+{y_2}^2+\cdots+{y_n}^2}$ (with all instances of λ_i replaced by λ_n). Recall that $\mathbf{y} = Q\mathbf{v}$. If $\mathbf{v} = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \end{bmatrix}^T$, then $\alpha = \lambda_1$. Also, if $\mathbf{v} = \begin{bmatrix} 0 & 0 & \dots & 0 & 1 \end{bmatrix}^T$, then $\alpha = \lambda_n$.

Theorem 2.6 (*Cauchy Interlacing Theorem*)**:** Given an $n \times n$ matrix A, let $A(i\!:\!j, i\!:\!j)$ be the $(j-i+1)\times(j-i+1)$ submatrix of A that consists only of rows $i, i+1, \dots, j$ and columns $i, i+i, \dots, j$ of A. Let A be symmetric and let its eigenvalues be $\lambda_1 \leq \lambda_2 \leq \dots \lambda_n$. Further, let the eigenvalues of $A(i\!:\!j, i\!:\!j)$ be $\chi_1 \leq \chi_2 \leq \cdots \leq \chi_{j-i+1}$. Then, for some $k \leq j-i+1$, the matrix A has at least k eigenvalues less than or equal to χ_k. The proof of this can be found in "Matrix Analysis" by R. Horn and C. Johnson (1988).

Corollary: If A is positive-definite, then $A(i\!:\!j, i\!:\!j)$ is also positive-definite since all eigenvalues of A are positive and, by using the theorem with $k = 1$, $\lambda_1 \leq \chi_1$. Using the facts that $\lambda_1 > 0$ and $\chi_1 \leq \chi_2 \leq \cdots \leq \chi_{j-i+1}$, we have that $0 < \lambda_1 \leq \chi_1 \leq \chi_2 \leq \cdots \leq \chi_{j-i+1}$ and therefore all eigenvalues of $A(i\!:\!j, i\!:\!j)$ are positive, making $A(i\!:\!j, i\!:\!j)$ positive-definite from Remark 2.4.

Lemma 2.7: Let A be an $n \times n$ matrix and let X be any $n \times n$ nonsingular matrix. Then, $X^T AX$ is symmetric positive-definite if and only if A is symmetric positive-definite.

Proof: From Lemma 2.5, the smallest eigenvalue of $X^T AX$ is $\min_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T X^T AX\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$. From here:

$$
\begin{array}{rcl|l}
\min\limits_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T X^T AX\mathbf{v}}{\mathbf{v}^T\mathbf{v}} & = & \min\limits_{\mathbf{v}\neq\mathbf{0}} \left(\frac{\mathbf{v}^T X^T AX\mathbf{v}}{\mathbf{v}^T X^T X\mathbf{v}} \cdot \frac{\mathbf{v}^T X^T X\mathbf{v}}{\mathbf{v}^T\mathbf{v}} \right) & \text{introducing a factor of } \frac{\mathbf{v}^T X^T X\mathbf{v}}{\mathbf{v}^T X^T X\mathbf{v}} = 1 \\
 & \geq & \min\limits_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T X^T AX\mathbf{v}}{\mathbf{v}^T X^T X\mathbf{v}} \cdot \min\limits_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T X^T X\mathbf{v}}{\mathbf{v}^T\mathbf{v}} & \text{the 2 factors could be smaller} \\
 & & & \text{using different vectors } \mathbf{v} \\
 & = & \min\limits_{\mathbf{y}\neq\mathbf{0}} \frac{\mathbf{y}^T A\mathbf{y}}{\mathbf{y}^T\mathbf{y}} \cdot \min\limits_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T X^T X\mathbf{v}}{\mathbf{v}^T\mathbf{v}} & \text{where } \mathbf{y} = X\mathbf{v} \\
 & = & \lambda_1(A) \cdot \lambda_1(X^T X) & \text{by Lemma 2.5}
\end{array}
$$

The notation $\lambda_1(M)$ refers to the smallest eigenvalue of matrix M. The number $\mathbf{v}^T X^T X\mathbf{v}$ is an inner product of a vector with itself and is therefore nonnegative. The same is true of $\mathbf{v}^T\mathbf{v}$. Therefore, $\frac{\mathbf{v}^T X^T X\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$ is nonnegative, and $\lambda_1(X^T X) \geq 0$. Also, since X is nonsingular, $X^T X$ is also nonsingular, and therefore can not have an eigenvalue of 0 (since the determinant of a matrix is the product

of its eigenvalues, and the determinant of a nonsingular matrix is nonzero). So, $\lambda_1(X^TX) > 0$. The chain of equations and inequalities shows that the smallest eigenvalue of X^TAX is equal to $\lambda_1(A) \cdot \lambda_1(X^TX)$. The matrix A is positive-definite, so $\lambda_1(A) > 0$, and it was just stated that $\lambda_1(X^TX) > 0$. So the smallest eigenvalue of X^TAX is positive, and therefore so are all of its eigenvalues, making it positive-definite. To show the other direction, follow the equations and inequalities in the other direction to show that the smallest eigenvalue of X^TAX is positive.

Lemma 2.8: Let A be an $n \times n$ symmetric matrix. If $\rho(A) < 1$, meaning that the value of any eigenvalue of A is on the interval $(-1, 1)$, then $I_n - A$ is nonsingular and $(I_n - A)^{-1} = \sum_{i=0}^{\infty} A^i$.

Proof: The eigenvalues of the matrix $I_n - A$ are on the inverval $(0, 2)$ since every eigenvalue of I_n is 1 and the eigenvalues of A are on the inverval $(-1, 1)$. This is true only because I_n is diagonal and possibly because A is symmetric, and is not true in general cases. Since the eigenvalues of $I_n - A$ are all positive, $I_n - A$ is nonsingular. With the eigenvalue decomposition $A = Q\Lambda Q^T$, we see that $A^i = \underbrace{(Q\Lambda Q^T)(Q\Lambda Q^T)\dots(Q\Lambda Q^T)}_{i \text{ times}} = Q\Lambda^i Q^T$ since $QQ^T = I_n$ and the Q^T at the end of one term cancels with the Q at the beginning of the next term. The matrix Λ is diagonal, meaning that each diagonal element of Λ^i is equal to the corresponding element of Λ multiplied by itself i times. Since each element of Λ is less than 1 in absolute value, Λ^i goes to a matrix where every element is 0 as i goes to ∞, and therefore A^i does as well. Therefore, $\sum_{i=0}^{\infty} A^i$ is convergent. Now notice that $(I_n - A) \cdot \sum_{i=0}^{m} A^i = (I_n - A) + (A - A^2) + (A^2 - A^3) + \dots + (A^m - A^{m+1}) = I_n - A^{m+1}$, since it is a telescoping series. As m goes to ∞, $I_n - A^{m+1} = (I_n - A) \cdot \sum_{i=0}^{m} A^i$ goes to I_n, and therefore $\sum_{i=0}^{m} A^i$ converges to $(I_n - A)^{-1}$.

With all of these facts now shown, we can prove the important theorem that the subgraphs created from a spectral graph partition are connected.

Theorem 2.9 (*Fiedler's theorem of connectivity of spectral graph partitions*)**:** Let G be a connected graph and let $L(G)$ be its Laplacian matrix. Create the subgraphs G_1 and G_2 using the method described in Section 2.2. The two subgraphs G_1 and G_2 are both connected.

Proof: To prove this, we use a proof by contradiction. Assume to a contradiction that G_1 is composed of 2 connected components. Let $\mathbf{x}$ be the eigenvector corresponding to λ_2, second-smallest eigenvalue of $L(G)$. Both $L(G)$ and $\mathbf{x}$ can be written in a block form:

$$L(G) = \begin{bmatrix} L_{1,1} & O & L_{1,3} \\ O & L_{2,2} & L_{2,3} \\ {L_{1,3}}^T & {L_{2,3}}^T & L_{3,3} \end{bmatrix} \quad \text{and} \quad \mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix}$$

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are positive and $\mathbf{x}_3$ is negative. Here, we use simply $\mathbf{x}$ for the Fiedler vector, since the subscripts are reserved for the blocks of that vector. The symbol O represents a matrix of appropriate size which has 0 as every entry. Recall that $L(G)$ is symmetric, so the diagonal blocks are also symmetric, and the off-diagonal block in position (i, j) is the transpose of the block in position (j, i). Also, due to the nature of the Laplacian matrix, all off-diagonal blocks are nonpositive. Due to the way the graph partition is constructed using the values of $\mathbf{x}$, this construction puts the vertices corresponding to the first 2 blocks of $\mathbf{x}$ into G_1, but the matrix $L(G)$

shows that there are no edges from any vertex of $\mathbf{x}_1$ to any vertex of $\mathbf{x}_2$, meaning that G_1 really does consist of 2 connected components.

Since $\mathbf{x}$ is an eigenvector of $L(G)$ and λ_2 is its eigenvalue, we have that $L(G)\mathbf{x} = \lambda_2\mathbf{x}$. Multiplying this out for only the first block $\mathbf{x}_1$ of $\mathbf{x}$, we see that $L_{1,1}\mathbf{x}_1 + L_{1,3}\mathbf{x}_3 = \lambda_2\mathbf{x}_1$.

Recall that every eigenvalue of $L(G)$ is nonnegative. By the corollary to the Cauchy Interlacing Theorem (Theorem 2.6), every eigenvalue of $L_{1,1}$ is also nonnegative. Let $\varepsilon > 0$. Add $\varepsilon\mathbf{x}_1$ to both sides of the above equation, resulting in $\varepsilon\mathbf{x}_1 + L_{1,1}\mathbf{x}_1 + L_{1,3}\mathbf{x}_3 = \varepsilon\mathbf{x}_1 + \lambda_2\mathbf{x}_1$, or, more concisely, $(\varepsilon I + L_{1,1})\mathbf{x}_1 + L_{1,3}\mathbf{x}_3 = (\varepsilon + \lambda_2)\mathbf{x}_1$, where I is of appropriate size.

By the same properties of I and symmetric matrices as was mentioned at the beginning of the proof of Lemma 2.8, the eigenvalues of $\varepsilon I + L_{1,1}$ are all at least ε and therefore positive, so $\varepsilon I + L_{1,1}$ is positive-definite. Using Demmel's method of constructing the proof, let $\varepsilon I + L_{1,1} = D - N$, where the matrix D is diagonal, and each diagonal element is equal to the corresponding element on the diagonal of $\varepsilon I + L_{1,1}$, and where $-L$ contains the other (off-diagonal) elements of $\varepsilon I + L_{1,1}$, and every diagonal element of $-L$ is 0. This can be further decomposed into $D^{\frac{1}{2}}(I - (D^{\frac{1}{2}})^{-1}N(D^{\frac{1}{2}})^{-1})D^{\frac{1}{2}}$, where $D^{\frac{1}{2}}$ is equal to the matrix D after every element of it has had the square root taken of it. It is true that $D^{\frac{1}{2}}D^{\frac{1}{2}} = D$ since D is diagonal and nonnegative. Multiply this expression through to see that it really is equal to $D - N$. Finally, let $M = (D^{\frac{1}{2}})^{-1}N(D^{\frac{1}{2}})^{-1}$, so that $\varepsilon I + L_{1,1} = D - N = D^{\frac{1}{2}}(I - M)D^{\frac{1}{2}}$.

Recall that since $D^{\frac{1}{2}}$ is diagonal, $(D^{\frac{1}{2}})^T = D^{\frac{1}{2}}$. By Lemma 2.7, since $\varepsilon I + L_{1,1} = D^{\frac{1}{2}}(I - M)D^{\frac{1}{2}} = (D^{\frac{1}{2}})^T(I - M)D^{\frac{1}{2}}$ is positive-definite, $I - M$ is also positive-definite. By the same reasoning before involving eigenvalues of I and symmetric matrices, if λ is an eigenvalue of M, then $1 - \lambda$ is an eigenvalue of $I - M$. Each eigenvalue of $I - M$ is also greater than -1 by Lemma 2.5:

$\lambda_1(M)$	$= \min_{\mathbf{v}\neq\mathbf{0}} \frac{\mathbf{v}^T M\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$	λ_1 is the smallest eigenvalue of M
	$\geq \min_{\mathbf{v}\neq\mathbf{0}} \left(-\frac{\lvert\mathbf{v}\rvert^T M\lvert\mathbf{v}\rvert}{\mathbf{v}^T\mathbf{v}}\right)$	since M is positive and nothing is known about the signs of the elements of $\mathbf{v}$
	$= -\max_{\mathbf{v}\neq\mathbf{0}} \frac{\lvert\mathbf{v}\rvert^T M\lvert\mathbf{v}\rvert}{\mathbf{v}^T\mathbf{v}}$	since $\min_{a\in S}(-a) = -\max_{a\in S} a$
	$\geq -\max_{\mathbf{v}\neq\mathbf{0}} \frac{\lvert\mathbf{v}\rvert^T M\lvert\mathbf{v}\rvert}{\mathbf{v}^T\mathbf{v}}$	since M is positive and nothing is known about the signs of the elements of $\mathbf{v}$
	$= -\lambda_{n_1}(M)$	by Lemma 2.5; here, n_1 is the number of elements in $\mathbf{x}_1$
	> -1	all eigenvalues of M are less than 1

So, the absolute value of every eigenvalue of M is less than 1.

Continuing to follow Demmel, let $Y = (\varepsilon I + L_{1,1})^{-1}$. Recall that $\varepsilon I + L_{1,1} = D^{\frac{1}{2}}(I - M)D^{\frac{1}{2}}$. Therefore, $Y = (D^{\frac{1}{2}}(I-M)D^{\frac{1}{2}})^{-1} = (D^{\frac{1}{2}})^{-1}(I-M)^{-1}(D^{\frac{1}{2}})^{-1}$. By Lemma 2.8, $(I-M)^{-1} = \sum_{i=0}^{\infty} M^i$, so $Y = (D^{\frac{1}{2}})^{-1}\left(\sum_{i=0}^{\infty} M^i\right)(D^{\frac{1}{2}})^{-1}$.

The matrix Y is nonnegative since M is nonnegative and $D^{\frac{1}{2}}$ is positive. In fact, since M is symmetric and describes a graph, M is positive and therefore so is Y.

Recall that $(\varepsilon I + L_{1,1})\mathbf{x}_1 + L_{1,3}\mathbf{x}_3 = (\varepsilon + \lambda_2)\mathbf{x}_1$. Multiplying both sides of this by Y gives $\mathbf{x}_1 + YL_{1,3}\mathbf{x}_3 = (\varepsilon + \lambda_2)Y\mathbf{x}_1$, since $Y(\varepsilon I + L_{1,1}) = I$. Further multiplying that equation by ${\mathbf{x}_1}^T$ gives ${\mathbf{x}_1}^T\mathbf{x}_1 + {\mathbf{x}_1}^T YL_{1,3}\mathbf{x}_3 = (\varepsilon + \lambda_2){\mathbf{x}_1}^T Y\mathbf{x}_1$. Once again, let n_1 be the number of elements in

vector $\mathbf{x}_1$. This equation and Lemma 2.5 can be used to show that:

$$\begin{array}{rcl|c} (\varepsilon+\lambda_2)\lambda_{n_1}(Y) & = & \max\limits_{\mathbf{v}\neq\mathbf{0}} \left((\varepsilon+\lambda_2)\cdot\frac{\mathbf{v}^T Y\mathbf{v}}{\mathbf{v}^T\mathbf{v}}\right) & \text{by Lemma 2.5} \\ & \geq & (\varepsilon+\lambda_2)\cdot\frac{{\mathbf{x}_1}^T Y\mathbf{x}_1}{{\mathbf{x}_1}^T\mathbf{x}_1} & \begin{array}{c}\text{using } \mathbf{x}_1 \text{ for the vector } \mathbf{v} \text{ could only} \\ \text{result in a value that is no larger than} \\ \text{the maximum obtained with any } \mathbf{v}\end{array} \\ & = & \frac{{\mathbf{x}_1}^T\mathbf{x}_1+{\mathbf{x}_1}^T Y L_{1,3}\mathbf{x}_3}{{\mathbf{x}_1}^T\mathbf{x}_1} & \begin{array}{c}\text{substituting } {\mathbf{x}_1}^T\mathbf{x}_1+{\mathbf{x}_1}^T Y L_{1,3}\mathbf{x}_3 \\ \text{for } (\varepsilon+\lambda_2){\mathbf{x}_1}^T Y\mathbf{x}_1 \\ \text{as in the above equation}\end{array} \\ & = & 1+\frac{{\mathbf{x}_1}^T Y L_{1,3}\mathbf{x}_3}{{\mathbf{x}_1}^T\mathbf{x}_1} & \frac{{\mathbf{x}_1}^T\mathbf{x}_1}{{\mathbf{x}_1}^T\mathbf{x}_1}=1 \end{array}$$

The vector $L_{1,3}\mathbf{x}_3$ is positive since $L_{1,3}$ is nonpositive and $\mathbf{x}_3$ is negative, and if $L_{1,3}$ were zero, G would not be connected. Recall that the matrix Y and the vector $\mathbf{x}_1$ are positive as well. So, ${\mathbf{x}_1}^T Y L_{1,3}\mathbf{x}_3$ is positive. Also, ${\mathbf{x}_1}^T\mathbf{x}_1$ is also positive. Therefore, by the equation and inequality chain above, $(\varepsilon+\lambda_2)\lambda_{n_1}(Y)>1$. All eigenvalues of Y are positive and since $Y=(\varepsilon I+L_{1,1})$, the eigenvalues of Y are the reciprocals of $\varepsilon I+L_{1,1}$. The largest eigenvalue of Y is equal to the reciprocal of the smallest eigenvalue of $\varepsilon I+L_{1,1}$. Therefore, $\frac{\varepsilon+\lambda_2}{\lambda_1(\varepsilon I+L_{1,1})}>1$. By the properties of sums of eigenvalues involving I and symmetric matrices mentioned before, $\lambda_1(\varepsilon I+L_{1,1})=\varepsilon+\lambda_1(L_{1,1})$. So, $\frac{\varepsilon+\lambda_2}{\lambda_1(\varepsilon I+L_{1,1})}=\frac{\varepsilon+\lambda_2}{\varepsilon+\lambda_1(L_{1,1})}>1$, or, rearranged, $\lambda_1(L_{1,1})<\lambda_2$.

All of these steps can be performed again on the block $L_{2,2}$, using other corresponding vectors and matrices, to see that $\lambda_1(L_{2,2})<\lambda_2$. Therefore, the matrix

$$\begin{bmatrix} L_{1,1} & O \\ O & L_{2,2} \end{bmatrix}$$

which is a leading submatrix of $L(G)$ has two eigenvalues which are less than λ_2. By Theorem 2.6, the Cauchy Interlacing Theorem, $L(G)$ has 2 eigenvalues whose values are less than λ_2. However, this contradicts λ_2 being the second-smallest eigenvalue of $L(G)$. Due to this contradiction, the original assumption must be false, and the two subgraphs G_1 and G_2 must be connected [3].

2.3.3 Notes on algebraic connectivity

The second-smallest eigenvalue of $L(G)$, $\lambda_2(L(G))$, is often called the *algebraic connectivity* of the graph G [4]. A basic intuition behind the use of this term is that a graph with a higher algebraic connectivity typically has more edges, and can therefore be thought of as being "more connected". As an example, consider the two graphs K_{100}, the complete graph on 100 vertices, and C_{100}, the cycle graph on 100 vertices. By generating the Laplacian matrix of both of these graphs, we can find that the second-smallest eigenvalue of K_{100} is 100 and the second-smallest eigenvalue of C_{100} is about 0.003947. The algebraic connectivity of K_{100} is many orders of magnitude greater than that of C_{100}, which makes sense, since K_{100} has 4950 edges and C_{100} has only 100 edges on the same number of vertices—clearly, K_{100} can be seen as being much more "connected" than C_{100}. In general, the smaller the algebraic connectivity of a graph, the more likely it is that the Fiedler method will find a "good" partition.

3 Examples

At this point, with the general theory fully discussed and proven, we will show some examples to illustrate how using the Fiedler vector (the eigenvector corresponding to the second-smallest

eigenvalue) of the Laplacian matrix of various graphs effectively partitions them in an ideal way in order to minimize the number of edges between the 2 generated subgraphs.

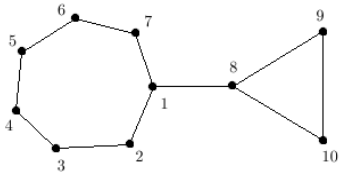

(a) The graph A

$$L(A) = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10 \end{array} \begin{array}{c} \begin{array}{cccccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \end{array} \\ \begin{bmatrix} 3 & -1 & 0 & 0 & 0 & 0 & -1 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 2 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 2 & -1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & -1 & 2 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 3 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & -1 & 2 \end{bmatrix} \end{array}$$

(b) The Laplacian matrix $L(A)$

$$\mathbf{x}_A = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \\ 9 \\ 10 \end{array} \begin{array}{c} \\ \begin{bmatrix} 0.0521 \\ -0.1147 \\ -0.2543 \\ -0.3335 \\ -0.3335 \\ -0.2543 \\ -0.1147 \\ 0.3734 \\ 0.4897 \\ 0.4897 \end{bmatrix} \end{array}$$

(c) The Fiedler vector $\mathbf{x}_A$

Figure 3.1: The graph A and its associated Laplacian matrix and Fiedler vector are depicted here. The algebraic connectivity of A is about 0.2375. Note that the sum of all elements of the Fiedler vector is 0. This is true of the Fiedler vector of any graph.

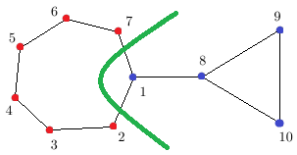

Figure 3.2: This is the spectral partition of the graph A according to Fiedler's method, where each subgraph is shown consisting of vertices of all the same color, divided by a curved green line that forms the cut.

To begin, consider graph A, which is drawn in Figure 3.1(a). By forming the Laplacian matrix $L(A)$ (Figure 3.1(b)) and calculating the Fiedler vector, which we will call $\mathbf{x}_A$ here (Figure 3.1(c)), we can perform a spectral graph partition. As always, we form this partition by placing vertices corresponding to negative elements of $\mathbf{x}_A$ into one subgraph (in this case, the subgraph of red vertices) and placing the vertices corresponding to positive elements of $\mathbf{x}_A$ into the other subgraph (blue vertices). This partition is shown in Figure 3.2. Both $L(A)$ and $\mathbf{x}_A$ are labeled with the labels of the vertices that each row or column of $L(A)$ or entry of $\mathbf{x}_A$ represents. There are some remarks to be made about this partition.

First, no matter in what way we look at this partition, it is not perfect. If we were looking for a partition that simply minimized the number of edges that cross the partition, we could get one by having only edge $(1, 8)$ be in the cut, thereby changing which subgraph the vertex 1 is in. (The *cut* is the set of edges in a graph partition which each have their end vertices in different subgraphs.) On the other hand, if we were looking to form a partition where the number of vertices in each subgraph is as near to equal as possible, then by removing edge $(7, 1)$ from the cut and adding edge $(6, 7)$ to the cut, both subgraphs would have 5 vertices. In fact, if this were done, the number of edges in the cut would not increase at all. The Fiedler spectral graph partition appears to be attempting to balance having a minimum number edges in the cut with having an equal number of vertices in both subgraphs, resulting in this partition. This example shows that this method of spectral graph partition is not flawless, even on small graphs.

Now we consider graph B, which is depicted in Figure 3.3(a). Once again, we form the Laplacian matrix $L(B)$ (Figure 3.3(b)) and compute its Fiedler vector (Figure 3.3(c)). We again use this Fiedler vector to construct a spectral graph partition of B in the same way as for A. This partition is shown in Figure 3.4. In this case, the graph B was formed more nicely, allowing Fiedler's spectral graph partitioning method to easily choose the best partition. We can observe that each subgraph has the same number of vertices, and the cut is of smallest possible size of any cut that could be made on B (disregarding the sizes of the subgraphs).

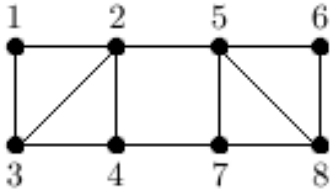

(a) The graph B

$$\begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{array} \begin{array}{c} \begin{array}{cccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{array} \\ \begin{bmatrix} 2 & -1 & -1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 4 & -1 & -1 & -1 & 0 & 0 & 0 \\ -1 & -1 & 3 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & -1 & 3 & 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 & 4 & -1 & -1 & -1 \\ 0 & 0 & 0 & 0 & -1 & 2 & 0 & -1 \\ 0 & 0 & 0 & -1 & -1 & 0 & 3 & -1 \\ 0 & 0 & 0 & 0 & -1 & -1 & -1 & 3 \end{bmatrix} \end{array}$$

(b) The Laplacian matrix $L(B)$

$$\begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{array} \begin{array}{c} \\ \begin{bmatrix} 0.4886 \\ 0.2503 \\ 0.4006 \\ 0.1953 \\ -0.2503 \\ -0.4886 \\ -0.1953 \\ -0.4006 \end{bmatrix} \end{array}$$

(c) The Fiedler vector $\mathbf{x}_B$

Figure 3.3: The graph B and its associated Laplacian matrix and Fiedler vector are depicted here. The algebraic connectivity of A is about 0.6678.

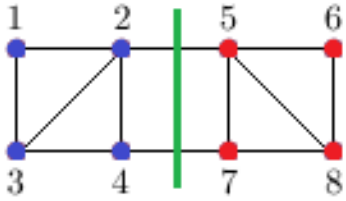

Figure 3.4: This is the spectral partition of the graph B according to Fiedler's method, where each subgraph is shown consisting of vertices of all the same color, divided by a green line that forms the cut.

Finally, Gleich illustrates how spectral graph partitioning using Fiedler's method can work on a randomly-generated matrix, where the edges between vertices are generated probabilistically [5]. In particular, he starts with a set of vertices and then forms the two subgraphs C_1 and C_2. Then, he generates edges between pairs of vertices u and v, using the following probabilities:

- If u and v are both in C_1, the probability of edge (u, v) being generated is 0.5.
- If u and v are both in C_2, the probability of edge (u, v) being generated is 0.4.
- If u and v are in different subgraphs, the probability of edge (u, v) being generated is 0.1.

He then forms the Laplacian matrix and calculates the Fiedler vector. The method is successful in splitting the generated graph into the subgraphs C_1 and C_2. The details can be found at `http://www.cs.purdue.edu/homes/dgleich/demos/matlab/spectral/spectral.html`.

4 Conclusion

We have discussed Fiedler's theory of spectral graph partitioning, including a detailed proof based off that of Demmel. Following that, we noted that, like virtually all other graph partitioning methods, Fiedler's method is not without fault. In fact, a simple example of a graph with 10 vertices showed a sub-optimal partition when the partition was done by Fiedler's method (Figure 3.2).

It is likely that future research will continue optimizing the method of finding a graph partition where the number of edges in the cut is minimized and the number of vertices in each subgraph is as near to equal as possible. Other graph partition methods currently in use include Inertial Partitioning and Random Circle [1]. Some of these methods require coordinates to be assigned to the vertices of the graph, and others, like the Fiedler method, do not. In fact, one of the main reasons the Fiedler method is so acclaimed is due to the fact that the vertices do not need to have coordinates assigned. As all of these methods are refined and new ones are discovered, it is possible that a much better algorithm will be developed. What is more likely, however, is that

in the end, multiple methods will continue to exist, with each of them being the best method in certain situations but without any one method being the best one for every situation.

5 Acknowledgements

The work by Dr. James Demmel in explaining Dr. Miroslav Fiedler's theory was invaluable, and it was that explanation which I reproduced here in a somewhat more explanatory form here. Dr. David Gleich also provided a very good description which helped me explain certain aspects of the spectral graph partition theory that I did not quite understand at first. Finally, my principal investigator Dr. Zhaojun Bai assisted in the structure of this paper.

References

[1] James Demmel. Graph partitioning, part 1. available at `http://www.cs.berkeley.edu/~demmel/cs267/lecture18/lecture18.html`, 1996.

[2] James Demmel. Graph partitioning, part 2. available at `http://www.cs.berkeley.edu/~demmel/cs267/lecture20/lecture20.html`, 1999.

[3] James Demmel. Spectral partitioning - statement and proof of theorem 2. available at `http://www.cs.berkeley.edu/~demmel/cs267/lecture20/lecture20b.html`, 1999.

[4] Miroslav Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25(4):619–633, 1975.

[5] David Gleich. Spectral graph partitioning and the Laplacian with Matlab. available at `http://www.cs.purdue.edu/homes/dgleich/demos/matlab/spectral/spectral.html`, 2006.

[6] Daniel A. Spielman and Shang-Hua Teng. Spectral partitioning works: Planar graphs and finite element meshes. Preliminary draft, 1996.