

The Path Resistance Method for Bounding λ_2 of a Laplacian

Stephen Guattery*
ICASE
smg@icase.edu

Tom Leighton†
Massachusetts Institute of Technology
ftl@math.mit.edu

Gary L. Miller‡
Carnegie Mellon University
glmiller@cs.cmu.edu

September 24, 1996

Abstract

We introduce the path resistance method for lower bounds on the smallest nontrivial eigenvalue of the Laplacian matrix of a graph. The method is based on viewing the graph in terms of electrical circuits; it uses clique embeddings to produce lower bounds on λ_2 and star embeddings to produce lower bounds on the smallest Rayleigh quotient when there is a Dirichlet boundary condition. The method assigns priorities to the paths in the embedding; we show that, for an unweighted tree T , using uniform priorities for a clique embedding produces a lower bound on λ_2 that is off by at most an $O(\log \text{diameter}(T))$ factor. We show that the best bounds that this method can produce for clique embeddings are the same as for a related method that uses clique embeddings and edge lengths to produce bounds.

1 Introduction

In this paper we consider methods based on graph embeddings for estimating the smallest nontrivial eigenvalue of the Laplacian matrix representation of a graph. The Laplacian is one of many ways to view a graph as a matrix; it is defined as follows: Let $G = (V, E)$ be an undirected graph with vertices v_1, \dots, v_n . Then the **Laplacian** of G is an $n \times n$ matrix L such that

$$l_{ij} = \begin{cases} \text{degree}(v_i) & \text{if } i = j \\ -1 & \text{if } (i, j) \in E \\ 0 & \text{otherwise} \end{cases}$$

It is not hard to see that L is positive semidefinite (all eigenvalues are ≥ 0). Since the row sums of L are all zero, the smallest eigenvalue is zero. If G is connected the second smallest eigenvalue λ_2 is positive.

The study of the connection between Laplacian spectra (particularly with respect to λ_2) and properties of the associated graphs dates back to Fiedler's work in

the 1970's (see, e.g., [Fie73] and [Fie75]). These properties have been used in graph algorithms, particularly algorithms for finding small separators [PSL90].

The Laplacian also has an important role in representing physical problems. It often occurs in finite difference, finite element, and control volume representations of problem involving elliptic partial differential equations. These problems often include a Dirichlet boundary condition that specifies that the values at a set of vertices are zero. To represent this condition in the Laplacian, the rows and columns corresponding to the boundary vertices are deleted from the matrix. The resulting matrix is positive definite, and it is the smallest eigenvalue of the matrix that is of interest.

Bounds on the smallest nonzero eigenvalues of both forms of the Laplacian have other important applications. Since the matrices are symmetric, their extreme eigenvalues can be used in computing their condition numbers, which are used in the study of iterative linear system solvers to estimate rates of convergence [HY81], and to analyze the quality of preconditioners [Axe92,

Procedure: The Path Resistance Method

Method:

1. Construct a path embedding of K_n into G , i.e., construct a path between every pair of vertices of G .
2. For each edge e_{ij} compute the congestion $c_{ij} = |\{P | e_{ij} \in P\}|$
3. For each path P and each edge e_{ij} on P allocate a resistor of size c_{ij} to P .
4. For each path P compute its resistance, i.e., $\sum_{e_{ij} \in P} c_{ij}$. Let r be the maximum resistance over all paths.
5. return " $n/r \leq \lambda_2$ ".

*This work supported by NAS contract NAS1-19480.

†This work supported by DARPA contract N00014-95-1246 and Army Contract DAAH04-95-1-0607.

‡This work supported in part by NSF Grant CCR-9505472, DARPA contract N00014-95-1246, and Army Contract DAAH04-95-1-0607.

Figure 1: The Uniform Priority Path Resistance Method

GMZ95]. Bounds on λ_2 are useful in the analysis of spectral partitioning, both because λ_2 occurs in bounds on cut quality [Moh89], and because they can be used in isolating structural properties of the eigenvectors used in making the cuts [GM95, ST96]. The eigenvalue λ_2 has been related to expansion properties of graphs, and can be used in determining if a graph is an expander [AM85a, Alo86].

In many of the preceding applications, it is necessary to show bounds for *classes* of graphs in order to state results in terms of asymptotic algorithm behavior. For example, in the analysis of spectral partitioning, [GM95] uses eigenvalue bounds on a family of bounded-degree graphs to prove facts about the structure of eigenvectors used in partitioning; [ST96] gives an upper bound on planar graph eigenvalues that can be applied in bounds on the cut quotient of the resulting cut. The embedding techniques we present are well-suited to producing such general results, and can be used with known results about embeddings. We have used them to generate lower bounds for families of graphs where the critical path resistance can be parameterized in terms of the size of the family member.

Note that the applications listed above involve both upper and lower bounds. Upper bounds on λ_2 are usually given by exhibiting small Rayleigh quotients. If $\mathbf{x} \neq 0$ is an n -vector a Rayleigh quotient is $\frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$. In fact $\lambda_2 = \min_{\sum \mathbf{x} = 0} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$. Thus a Rayleigh quotient with $\sum \mathbf{x} = 0$ gives an upper bound on λ_2 . As a beautiful example, Spielman and Teng have shown that $\lambda_2 = O(1/n)$ for bounded degree planar graphs by constructing a small Rayleigh quotient for each such graph [ST96].

In the Dirichlet boundary case we are interested in the smallest eigenvalue λ . Consider the Laplacian L with the rows and columns of the boundary vertices deleted. Then $\lambda = \min_{\mathbf{x}} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$. In this case the Rayleigh quotient has a continuous analog: Let Ω be a compact domain in \mathbb{R}^d with boundary Γ , e.g. the unit square in the plane, and let $u \neq 0$ vary over all differentiable functions defined on Ω such that $u = 0$ on Γ . Then the Rayleigh quotient is

$$c = \min_u \frac{\int_{\Omega} (\nabla u)^2}{\int_{\Omega} u^2},$$

where ∇u is the gradient of u . This lower bound on c is known as a Poincare inequality. This analog provides useful insight on the problem and helps with the proofs to follow.

We start by describing a very simple yet powerful method for bounding λ_2 from below, shown in Figure 1. We show that the Path Resistance Method correctly

bounds λ_2 from below. Note that the quality of the lower bound improves with the quality of the path embedding since a better embedding should decrease the congestion. We also show that it is sometimes possible to improve the lower bound by assigning a priority to each path and allocating the edge conductances based on the priorities. Note that the path resistance method as stated in the figure does not mention priorities; we shall refer to that version as the **unit-priority** case.

The clique embedding version of the Path Resistance Method is closely related to another clique embedding technique presented in [Kah96] that assigns edge lengths and computes the sum of the lengths of all paths incident to each edge. We show that these methods are duals of one another, and that the best lower bounds computed by these methods are the same.

We also present a version of the Path Resistance Method that is applicable in the Dirichlet boundary case. Instead of embedding a clique, we embed a star into the graph, with the boundary vertices consolidated into the vertex at the center of the star.

Finally, we show that the clique embedding version of the Path Resistance Method with uniform priorities is robust when applied to trees. We produce upper and lower bounds to show that for a tree T using uniform priorities, our estimate is off by at most a factor of $O(\log \text{diameter}(T))$. It is somewhat surprising that such a simple way of setting the priorities gives such a good bound. It is open as to how well the uniform priority method does for general graphs.

2 Previous Work

Prior work related to the applications of the lower bounds has been cited in the Introduction. The use of clique embeddings to bound eigenvalues arose in the analysis of mixing times for Markov chains by Jerrum and Sinclair [JS89] [SJ89]. Further work in this direction was done by Diaconis and Strook [DS91] and by Sinclair [Sin92]. Kahale [Kah96] generalizes this work in terms of methods assigning lengths to the graph edges, and shows that the best bound over all edge length assignments is the largest eigenvalue of the matrix $\Gamma^T \Gamma$, where Γ is a matrix representing the path embedding ([Kah96] also cites unpublished work by Fill and Sokal in these directions). He also gives a semidefinite programming formulation for a model allowing fractional paths, and shows that the bound is off by at most a factor of $\log^2 n$. He shows this gap is tight; he also notes that the results can be applied to Laplacians with suitable modifications.

3 Terminology, Notation, and Background Results

We assume that the reader is familiar with the basic definitions of graph theory (in particular, for undirected graphs), and with the basic definitions and results of matrix theory. A graph consists of a set of vertices V and a set of edges E ; we denote the vertices (respectively edges) of a particular graph G as $V(G)$ (respectively $E(G)$) if there is any ambiguity about which graph is referred to. When it is clear which graph we are referring to, we use n to denote $|V|$.

We use the term **path graph** for a tree that has exactly two vertices of degree one. That is, a path graph is a graph consisting of exactly its maximal path. A **star** is a tree with exactly one vertex that is not a leaf. We call the non-leaf vertex the **center** of the star.

3.1 Matrices and Matrix Notation We use capital letters to represent matrices and bold lower-case letters for vectors. For a matrix A , a_{ij} or $[A]_{ij}$ represents the element in row i and column j ; for the vector \mathbf{x} , x_i or $[\mathbf{x}]_i$ represents the i^{th} entry in the vector. The notation with square brackets is useful in cases where the matrix or vector name is already subscripted. The notation $\mathbf{x} = 0$ indicates that all entries of the vector \mathbf{x} are zero; $\vec{1}$ indicates the vector that has 1 for every entry. For Laplacians, we index the eigenvalues of an $n \times n$ matrix in non-decreasing order. λ_1 represents the smallest eigenvalue, and λ_n the largest. For $1 < i < n$, we have $\lambda_{i-1} \leq \lambda_i \leq \lambda_{i+1}$. We use the notation $\lambda_i(L)$ (respectively $\lambda_i(G)$) to indicate the i^{th} eigenvalue of matrix L (respectively of the Laplacian of graph G) if there is any ambiguity about which matrix (respectively graph) the eigenvalue belongs to. \mathbf{u}_i represents the eigenvector corresponding to λ_i .

3.2 The Laplacian Matrix Representation of a Graph A common matrix representation of graphs is the **Laplacian**. Let D be the matrix with $d_{ii} = \text{degree}(v_i)$ for $v_i \in V(G)$, and all off-diagonal entries equal to zero. Let A be the adjacency matrix for G ($a_{ij} = 1$ iff $(v_i, v_j) \in E(G)$, 0 otherwise). Then the Laplacian representation of G is the matrix $L = D - A$.

The following are some useful facts about the Laplacian matrix:

- The Laplacian is symmetric positive semidefinite, so all its eigenvalues are greater than or equal to 0 (see e.g. [AM85b]).
- A graph G is connected if and only if 0 is a simple eigenvalue of its Laplacian (see e.g. [AM85b]).
- The following characterization of λ_2 holds (see

e.g. [Fie73]):

$$(3.1) \quad \lambda_2 = \min_{\mathbf{x} \perp \vec{1}} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}.$$

- For any vector \mathbf{x} and Laplacian L of the graph G , we have (see e.g. [Fie75]):

$$(3.2) \quad \mathbf{x}^T L \mathbf{x} = \sum_{(v_i, v_j) \in E(G)} (x_i - x_j)^2$$

An **edge-weighted** graph is a graph for which a real, nonzero weight w_{ij} is associated with each edge (v_i, v_j) (we consider a zero edge weight to indicate the lack of an edge). Fiedler extended the notion of the Laplacian to graphs with positive edge weights [Fie75]; he referred to this representation as the **generalized Laplacian**. Let w_{ij} be the (positive) weight of edge (i, j) in graph G . Then the entries of the generalized Laplacian L of G are defined as follows: l_{ii} is the sum of the weights of the edges incident to vertex v_i ; for $i \neq j$ and $(v_i, v_j) \in E(G)$, $l_{ij} = -w_{ij}$, and $l_{ij} = 0$ otherwise.

With the exception of equation (3.2), the properties listed above also apply to generalized Laplacians. A slightly modified version of (3.2) holds for the generalized Laplacian L :

$$\mathbf{x}^T L \mathbf{x} = \sum_{(v_i, v_j) \in E(G)} w_{ij} (x_i - x_j)^2$$

We usually use L to denote the Laplacian of a graph G , and K to represent the Laplacian of the clique K (whether K refers to the graph or the Laplacian will be clear from context).

3.3 Graph Embeddings We start with the notion of embedding one graph into another. Let G and H be connected graphs such that the vertex set of H is a subset of the vertex set of G . An **embedding** of H into G is a collection Γ of path subgraphs of G such that for each edge $(v_i, v_j) \in E(H)$, the embedding contains a simple path γ_{ij} from v_i to v_j in G . For full generality, we will allow fractional paths in our embeddings: i.e., an edge $(v_i, v_j) \in E(H)$ can be associated with a finite collection of simple paths from v_i to v_j in G ; each such path has a positive fractional weight associated with it such that the weights add up to 1. If a path γ includes edge e , we say that γ is incident to e .

In the unit-priority case, the **congestion** $\text{cong}(e)$ of an edge $e \in E(G)$ is the number of paths in the embedding incident to e (fractional paths using e contribute their fractional weight to this count). The **resistance** $r(e)$ of an edge is its congestion divided by its weight: $r(e) = \frac{\text{cong}(e)}{w_e}$. The resistance of a path

$r(\gamma)$ is computed from the edge resistances in the same way that the resistance of a series of resistors would be computed in an electrical network. That is, $r(\gamma) = \sum_{e \in \gamma} r(e)$. The largest path resistance r_{\max} is defined as $r_{\max} = \max_{\gamma \in \Gamma} r(\gamma)$. The corresponding path is called the **critical path**. Note that r_{\max} is always defined with respect to an embedding; the particular embedding will be clear from the context, and thus we won't specify it in the notation.

The definitions change slightly when we use path priorities. Let p_{ij} be the path priority for γ_{ij} , and let \mathbf{p} be the vector of priorities. The congestion $\text{cong}(e)$ of edge e is the sum of the priorities of e 's incident paths. Each edge has a resistance factor $r(e)$, defined as the congestion of the edge divided by its weight. The resistance of edge e with respect to γ_{ij} is $\frac{r(e)}{p_{ij}}$. The resistance $r(\gamma)$ of a path γ is the sum of the resistance factors of the edges in the path divided by the path priority (i.e., $r(\gamma_{ij}) = \frac{1}{p_{ij}} \sum_{e \in \gamma_{ij}} r(e)$). The maximum path resistance taken over all paths in Γ is denoted r_{\max} . The corresponding path is called the **critical path**.

We note that using these modified definitions, the general method for computing lower bounds using priorities is clear: Compute r_{\max} according to the revised definitions. Then $\frac{n}{r_{\max}} \leq \lambda_2$.

4 Lower Bounds on λ_2 for Generalized Laplacians

4.1 A Bound Based on a Clique Embedding

We now consider the case in which G is a graph for which we would like to bound λ_2 of G 's (generalized) Laplacian from below. We will use a clique embedding to decompose the graph. The proof is based on a view of Laplacians as electrical circuits.

We start with an easy lemma that helps convey the nature of the bounding technique. Let L be the Laplacian of (edge-weighted) graph G , and let K be the Laplacian of the complete graph on $n = |V(G)|$ vertices.

LEMMA 4.1. *Let α be a positive real number. If $\alpha L - K$ is positive semidefinite, then $\lambda_2(L) \geq \frac{n}{\alpha}$.*

Proof. Assume that $\alpha L - K$ is positive semidefinite. Consider $\mathbf{u}_2^T (\alpha L - K) \mathbf{u}_2$, where \mathbf{u}_2 is the eigenvector corresponding to $\lambda_2(L)$. Without loss of generality, assume that \mathbf{u}_2 has been normalized to length 1. Note that every vector orthogonal to $\vec{1}$ is an eigenvector of K with eigenvalue n . Therefore we have

$$\mathbf{u}_2^T (\alpha L - K) \mathbf{u}_2 = \alpha \mathbf{u}_2^T L \mathbf{u}_2 - \mathbf{u}_2^T K \mathbf{u}_2 = \alpha \lambda_2 - n \geq 0,$$

where the last inequality holds by the definition of a positive semidefinite matrix. The lemma follows. \square

Thus, if we can demonstrate an α sufficiently large to make $\alpha L - K$ positive semidefinite, we can bound λ_2

from below. To do this, we form an embedding (Γ, \mathbf{p}) of the complete graph into G , then break $\alpha L - K$ into pieces consisting of edges in K and the corresponding paths in Γ . We use the following additional notation: E_{ij} denotes the Laplacian of the graph on $V(G)$ that has only the single edge (i, j) . $L(\gamma_{ij})$ represents the generalized Laplacian on $V(G)$ whose edges are the edges of γ_{ij} , and whose edge weights are set as follows: each path γ_{ij} using an edge (g, h) gets a share of w_{gh} proportional to $\frac{p_{ij}}{\text{cong}(g, h)}$. This works out to

$$\frac{w_{gh} p_{ij}}{\text{cong}(g, h)} = \frac{p_{ij}}{r_{gh}}.$$

It is easy to see that the weights assigned to all incident paths for (g, h) sum to w_{gh} . (If fractional paths are involved, then the matrix $L(\gamma_{ij})$ is multiplied by the fractional coefficient of path γ_{ij} .)

We now state the theorem that gives us the technique for bounding λ_2 :

THEOREM 4.1. *For any generalized Laplacian L of a connected graph G with positive edge weights, and any clique embedding into G ,*

$$\lambda_2 \geq \frac{n}{r_{\max}}.$$

Proof. Any embedding of the complete graph into G defines a natural decomposition of the Laplacian L in terms of the l_{ij} 's: define the matrix $L_{nonpath}$ as the Laplacian of the graph on $V(G)$ that includes any edges of G not used in any path γ . It is easy to see that

$$L = \sum_{\gamma \in \Gamma} L(\gamma) + L_{nonpath}.$$

Since $L_{nonpath}$ is positive semidefinite, the following inequality holds for any positive α and real vector \mathbf{x} :

$$(4.3) \quad \mathbf{x}^T \left(\sum_{i < j} (\alpha L(\gamma_{ij}) - E_{ij}) \right) \mathbf{x} \leq \mathbf{x}^T (\alpha L - K) \mathbf{x}.$$

Applying properties of linearity and rewriting gives

$$(4.4) \quad \sum_{i < j} (\alpha \mathbf{x}^T L(\gamma_{ij}) \mathbf{x} - \mathbf{x}^T E_{ij} \mathbf{x}) \leq \mathbf{x}^T (\alpha L - K) \mathbf{x}.$$

Now consider the terms $\alpha \mathbf{x}^T L(\gamma_{ij}) \mathbf{x} - \mathbf{x}^T E_{ij} \mathbf{x}$. By (3.2), $\mathbf{x}^T E_{ij} \mathbf{x} = (x_i - x_j)^2$. Likewise, the term for the path can be written in terms of its edges and edge weights:

$$\alpha \mathbf{x}^T L(\gamma_{ij}) \mathbf{x} = \alpha \sum_{(g, h) \in \gamma_{ij}} \frac{p_{ij}}{r_{gh}} (x_g - x_h)^2$$

Since all the priorities and resistance factors are positive, the sum in the expression above can be interpreted

as the energy dissipation of a series of resistances (the reciprocals of the $\frac{p_{ij}}{r_{gh}}$'s) given a set of potentials (the x_i 's) at the nodes between the resistances. It is well known (see e.g. [DS84]) that this quantity is minimized when the potentials at the internal nodes are consistent with Kirchoff's law (this is easily seen through an application of the Cauchy-Schwartz inequality). At that minimum the sum is the square of the potential differences of the endpoints divided by the path resistance:

$$\sum_{(g,h) \in \gamma_{ij}} \frac{p_{ij}}{r_{gh}} (x_g - x_h)^2 \geq \frac{(x_i - x_j)^2}{r(\gamma_{ij})}.$$

Thus, if $\alpha = r_{\max}$, we have for every path γ_{ij} that

$$\alpha \mathbf{x}^T L(\gamma_{ij}) \mathbf{x} - \mathbf{x}^T E_{ij} \mathbf{x} \geq r_{\max} \frac{(x_i - x_j)^2}{r(\gamma_{ij})} - (x_i - x_j)^2 \geq 0.$$

For this value of α , the left-hand sides of inequalities (4.3) and (4.4) are nonnegative. Thus $r_{\max} L - K$ is positive semidefinite, and the theorem holds by Lemma 4.1.

If fractional paths are involved, then each $L(\gamma_{ij})$ is multiplied by the fractional coefficient, as is the corresponding E_{ij} . It is easy to see that the theorem also holds in this case. \square

We now compare the path resistance method with the edge length method described in [Kah96]. The method as applied to Laplacians is described as follows:

- Specify a clique embedding for the graph, and assign each edge a positive length.
- Compute the length of each path with respect to the edge lengths.
- For each edge, compute the sum of the lengths of all incident paths divided by the length of that edge. Let ρ_{\max} be the maximum such value taken over all the edges. Then $\frac{n}{\rho_{\max}}$ is a lower bound for λ_2 .

The proof that this produces a lower bound can be done in a fashion similar to the proof given above for the path resistance method; however, the argument manipulates the $\mathbf{x}^T K \mathbf{x}$ term.

The two methods can be thought of in terms of the following electrical analogies: the path resistance method partitions the conductance (i.e., the reciprocal of the resistance) of each edge with respect to the incident paths. The resulting paths have conductances; we find a multiplier sufficient to increase the conductance of every path so that it is at least as big as the conductance of the corresponding clique edge. The edge length method also deals with path conductances, but instead increases the conductance of each edge in the graph such that the conductance of each path in the embedding is equal to the conductance of the associated clique edge.

The two methods are duals of each other in the following sense:

THEOREM 4.2. *Given a path embedding, the best lower bound for the path resistance method taken over all allowed priority assignments and the best lower bound for the edge length method taken over all allowed length assignments is the same.*

Proof. To show this, we will use a representation presented by Kahale in [Kah96]. He defines an embedding matrix Γ as follows: Each row of Γ represents a path between a pair of distinct vertices. A row entry is 1 if the corresponding edge is in the path and zero otherwise (this is easily generalized to fractional paths). It is easy to see that multiplying a vector of edge lengths by Γ computes the length of each path, and that multiplying the result by Γ^T sums the lengths of the incident paths for each edge. Thus multiplying an edge vector by $\Gamma^T \Gamma$ and dividing the result for each edge by the original length and taking the maximum is a way to compute the value ρ_{\max} used in the edge length lower bound. $\Gamma^T \Gamma$ is clearly a nonnegative matrix, and Kahale shows, using the properties of such matrices (see, e.g., [Min88]), that the best lower bound for this method is the largest eigenvalue of this matrix.

The path resistance lower bound can be computed by multiplying a vector of priorities by $\Gamma \Gamma^T$, then dividing the result termwise by the original priorities and setting r_{\max} to the maximum. The same arguments about nonnegative matrices apply. However, $\Gamma \Gamma^T$ and $\Gamma^T \Gamma$ have the same set of nonzero eigenvalues; this is easily seen by showing that for every eigenvector \mathbf{u} of a nonzero eigenvalue of $\Gamma \Gamma^T$, $\Gamma^T \mathbf{u}$ is an eigenvector of $\Gamma^T \Gamma$ with nonzero eigenvalue; likewise, if \mathbf{u} is an eigenvector of a nonzero eigenvalue of $\Gamma^T \Gamma$, $\Gamma \mathbf{u}$ is an eigenvector of $\Gamma \Gamma^T$ with nonzero eigenvalue. Thus the best lower bounds produced are the same. \square

4.2 Star Embeddings Embeddings other than clique embeddings are useful. In particular, we embed a star into graphs with Dirichlet boundary conditions. Such boundary conditions specify a set of vertices whose values are zero. Let G be a graph with such a zero boundary, and let L be its Laplacian. Let \mathcal{X} be the set of non-zero vectors consistent with the boundary restriction. We are interested in that value of the following quantity:

$$\lambda = \min_{\mathbf{x} \in \mathcal{X}} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$$

To find a lower bound, we embed a star into G such that every boundary vertex is mapped to the center of the star, and every other vertex is mapped to a

distinct leaf. Thus we specify a path in G from each non-boundary vertex to a boundary vertex. Priorities, congestions, resistances, and r_{\max} are all defined as before. The following theorem applies:

THEOREM 4.3.

$$\frac{1}{r_{\max}} < \lambda.$$

For brevity, we omit the proof, which is almost the same as the proof in the clique case. For readers wishing to construct the proof themselves, we note that for the star, the Rayleigh quotient of any nonzero vector with zero at the center is exactly 1.

4.3 Paths with Masses To simplify the computation of quantities such as congestions for embeddings into trees in the unit priority case, we introduce the idea of **mass**. Consider a tree T and a unit priority embedding of a graph H into T (note that this embedding Γ is unique). Using the definitions from Section 3.3 above, we can compute the critical path. We distinguish two types of vertices that do not lie on the critical path: those not in any path containing a critical path edge (noncontributing vertices), and those that are in some such path (contributing vertices). For each contributing vertex v , there is a first critical path vertex v_c on any path that includes both v and a critical path edge. Note that since we are working with a tree, v_c is the same for every such path. We will call v_c the **connection point** for v . Assume that the critical path has k vertices indexed from 1 to k in order. We assign masses to the critical path vertices as follows: Vertex v_i has a mass m_i equal to 1 plus the number of contributing vertices that have v_i as a connection point.

We refer to the graph P consisting of a path on k vertices plus the masses for the critical path as the path-with-masses model. It has an embedding of a subgraph of H consisting of a subset of Γ that includes paths with both endpoints on the critical path. Congestions are computed as before, except that now each path contributes an amount equal to the product of the masses of its endpoints to the edge congestions. It is easy to see that the path congestion of P in this model is equal to the congestion on the critical path of the original embedding.

We can also use the path-with-masses model in computing upper bounds. Let y be a vector of length k . Let $L(P)$ be the Laplacian of P , and let M be the diagonal matrix with entry $[M]_{ii} = m_i$. Let L be the Laplacian of T , the tree we are working with. Construct a vector \mathbf{x} on n vertices as follows: For the i^{th} vertex on the critical path, set the corresponding entry of \mathbf{x} to y_i . For each noncontributing vertex, set the corresponding entry of \mathbf{x} to 0. For each contributing

vertex, set the corresponding entry of \mathbf{x} to the value of y for its connection point. It is easy to show that

$$\frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\mathbf{y}^T L(P) \mathbf{y}}{\mathbf{y}^T M \mathbf{y}}.$$

In the star case with Dirichlet boundary, it is easy to see that if we enforce the boundary condition in \mathbf{y} , then it is enforced in \mathbf{x} , which implies that $\frac{\mathbf{y}^T L(P) \mathbf{y}}{\mathbf{y}^T M \mathbf{y}}$ is an upper bound on λ as defined above. For clique embeddings, note that enforcing the condition $\mathbf{y}^T M \vec{1} = 0$ is equivalent to enforcing the condition $\mathbf{x}^T \vec{1} = 0$; in this case, $\frac{\mathbf{y}^T L(P) \mathbf{y}}{\mathbf{y}^T M \mathbf{y}}$ is an upper bound on λ_2 .

5 Lower Bounds for Trees

In this section we show that, for an arbitrary unweighted tree T with diameter $\text{diam}(T)$, assigning all paths priority 1 in the clique embedding gives a lower bound within a $\log \text{diam}(T)$ factor of λ_2 . The argument works by first showing the bound holds for a tree with a single zero boundary vertex, then applying this fact along the critical path to handle the clique embedding case.

We start with a Dirichlet boundary case, where λ is defined as in Section 4.2. Let k be the length of the critical path.

THEOREM 5.1. *For any unweighted tree with a single zero boundary vertex, let r_{\max} be the maximum path resistance for the star embedding in the unit priority case. Then*

$$\frac{\max(\log k, 1)}{r_{\max}} = \Omega(\lambda).$$

Proof. Note that if the boundary vertex separates the tree into multiple components, the components can be considered separately. We need only consider the component with the critical path.

Let L be the Laplacian of the tree. For any such tree and boundary, we construct a nonzero vector \mathbf{x} consistent with the boundary condition such that

$$(5.5) \quad \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \frac{c \max(\log k, 1)}{r_{\max}},$$

where c is a constant independent of the choice of tree and boundary. By the definition of λ , the theorem then holds. The construction of \mathbf{x} is done with respect to the path-with-masses model of the critical path as described in Section 4.3.

Let P be the critical path. By assumption it has k edges. We index its vertices as follows: vertex 0 is the boundary vertex. Vertex i is the vertex at distance i from the boundary. The maximum index is k . Mass m_i is assigned to v_i as specified in Section 4.3 to account for paths that start off the critical path, but intersect with it.

Because we are dealing with the critical path plus a zero boundary, we can express the value x_i as follows: Let $\delta_i = x_i - x_{i-1}$. then $x_i = \sum_{j=1}^i \delta_j$. This allows us to rewrite some of the quantities from Inequality (5.5) as follows:

$$(5.6) \quad \frac{1}{r_{\max}} = \frac{1}{\sum_{i=1}^k \sum_{j=i}^k m_j} = \frac{1}{\sum_{i=1}^k i m_i},$$

and

$$(5.7) \quad \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\sum_{i=1}^k \delta_i^2}{\sum_{i=1}^k m_i \left(\sum_{j=1}^i \delta_j \right)^2}.$$

We specify \mathbf{x} by setting $\delta_i = i^{-\frac{1}{2}}$. For $k = 1$ it is clear that the quantities on the right-hand sides of (5.6) and (5.7) are equal and the theorem holds. For $k > 1$, we have

$$\sum_{i=1}^k \delta_i^2 = \sum_{i=1}^k i^{-1} = O(\log k).$$

Thus if we show that the quantities in the denominators of the right-hand sides of (5.6) and (5.7) are within a constant factor, then the theorem holds. Note that the two denominators can be thought of as expressions in the m_i 's. If we show for each i that the corresponding coefficients for each m_i are within a constant of each other, that will suffice. Thus, we want to show that $i m_i$ is within a constant factor of $\left(\sum_{j=1}^i j^{-\frac{1}{2}} \right)^2 m_i$. For $i = 1$ these both reduce to m_i . For $i > 1$, the summation can be bounded using integration techniques; in this case the sum lies between $i^{\frac{1}{2}}$ and $2i^{\frac{1}{2}}$. Thus the square of the sum lies between i and $4i$, which gives the desired result. \square

It is possible to find examples for which this gap is tight. Consider a tree consisting of a path graph connecting the center vertices of a series of stars. In particular, let the path have $k + 1$ vertices numbered from 0 to k . Vertex i ($0 < i < k$) on the path is the center of a star with $\lceil k^2/i^2 \rceil - 1$ leaves. For simplicity of reference, we refer to this tree as the **bad tree** for parameter k .

We show that the gap for the bad tree is tight by showing a set of priorities that give a lower bound within a constant of the upper bound constructed as per the proof of Theorem 4.3. Let T be a bad tree with $k > 1$; let L be the Laplacian of T . The path connecting the stars is obviously the critical path. We assume that the path vertices are numbered as in the definition of the bad tree. Path edges are numbered by the larger of the endpoint indices: i.e., edge i is from v_{i-1} to v_i .

THEOREM 5.2. *Let v_0 be the zero boundary vertex for bad tree T with parameter $k > 1$. Let $\frac{1}{r_{\max}}$ be the*

unit priority lower bound estimate on λ of L given that boundary. Then $\lambda = \Theta\left(\frac{\log k}{r_{\max}}\right)$.

Proof. We again use the ‘‘path with masses’’ model. It is clear that vertex i along the path has mass $\left\lceil \frac{k^2}{i^2} \right\rceil$.

We note that the quantity $\frac{k^2}{i^2}$ is always greater than or equal to 1 in the allowed range, so the ceilings of these quantities are bigger by no more than a factor of 2. Because we are interested in bounds, and because of the forms of the particular functions we are working with, it is easy to verify that dropping the ceilings does not change the results by more than a constant factor.

As in the proof of Theorem 4.3, we can construct a vector \mathbf{x} with $x_0 = 0$ and, for $i > 0$, $x_i = i^{-\frac{1}{2}}$. Values at vertices off the critical path are set as per Section 4.3. Consider the Rayleigh quotient $\frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}}$. As in the previous proof, the numerator is $\Theta(\log k)$. The denominator is

$$\mathbf{x}^T \mathbf{x} = \sum_{i=1}^k m_i \left(\sum_{j=1}^i \delta_j \right)^2 = \sum_{i=1}^k m_i \left(\sum_{j=1}^i j^{-\frac{1}{2}} \right)^2.$$

As noted in the proof to Theorem 4.3, $\sum_{j=1}^i j^{-\frac{1}{2}}$ is \sqrt{i} to within a constant factor, so, to within a constant factor the denominator is

$$\sum_{i=1}^k i m_i = \sum_{i=1}^k i \frac{k^2}{i^2} = \sum_{i=1}^k \frac{k^2}{i} = \Theta(k^2 \log k)$$

(note that $k > 1$ by assumption). Thus the upper bound is $\Theta\left(\frac{1}{k^2}\right)$.

For the unit priority lower bound, it is obvious that the critical path runs from the boundary to vertex v_k . As noted previously, we can write the formula for r_{\max} as follows (once again we drop the ceilings without changing the result by more than a constant factor):

$$r_{\max} = \sum_{i=1}^k i m_i = \sum_{i=1}^k \frac{k^2}{i} = \Theta(k^2 \log k).$$

Thus the upper and lower bounds differ by a $\log k$ factor, the value of the Rayleigh quotient numerator.

We can get a better lower bound and close the gap by assigning priorities to paths. The relative sizes of the priorities determine the relative shares of the conductance that the paths get. In the current problem, the resistance of the critical path is high because it gets very small shares of the conductances of the edges close to the boundary. By increasing the shares of longer paths, we can offset this problem; because the number of shorter paths using edge i grows fast as the boundary is approached, the number of paths getting an increased

share of conductance is relatively small and doesn't increase the resistance of these shorter paths too much.

We now specify a set of priorities that gives us a lower bound that matches the upper bound. Each path that either starts at v_i on the critical path, or that has v_i as its connection point, gets priority $p_i = \sqrt{i}$. With this change, the congestion of edge i is equal to

$$\sum_{j=i}^k p_j m_j = \sum_{j=i}^k j^{\frac{1}{2}} \frac{k^2}{j^2} = k^2 \sum_{j=i}^k j^{-\frac{3}{2}}.$$

For all i greater than 1, we use integration techniques to get the bounds

$$(5.8) \quad k^2 \sum_{j=i}^k j^{-\frac{3}{2}} \geq 2k^2 \left(\frac{1}{\sqrt{i}} - \frac{1}{\sqrt{k+1}} \right)$$

and

$$(5.9) \quad k^2 \sum_{j=i}^k j^{-\frac{3}{2}} \leq 2k^2 \left(\frac{1}{\sqrt{i-1}} - \frac{1}{\sqrt{k}} \right).$$

For $i = 1$ the upper bound is replaced by $2k^2 \left(2 - \frac{1}{\sqrt{k}} \right)$; the lower bound remains unchanged.

Recall that in the case where path γ has priority p , we compute the path resistance as follows:

$$r(\gamma) = \frac{1}{p} \sum_{e \in \gamma} \text{cong}(e).$$

Plugging in the upper and lower bounds on edge congestion in (5.8) and (5.9) above gives upper and lower bounds on the path resistance for a path γ_j starting at v_j . We start with the lower bound; the third line below follows from an application of standard integral techniques for bounding sums:

$$\begin{aligned} \frac{1}{p} \sum_{e \in \gamma} \text{cong}(e) &= \frac{k^2}{\sqrt{j}} \sum_{i=1}^j \sum_{l=i}^k l^{-\frac{3}{2}} \\ &\geq \frac{2k^2}{\sqrt{j}} \sum_{i=1}^j \left(\frac{1}{\sqrt{i}} - \frac{1}{\sqrt{k+1}} \right) \\ &\geq 2k^2 \left(\frac{2\sqrt{j+1} - 2}{\sqrt{j}} - \sqrt{\frac{j}{k+1}} \right). \end{aligned}$$

With some straightforward calculations, the reader can verify that the term in parentheses in the last line is greater than $\frac{1}{12}$ for all j and k such that $k \geq j$.

The calculation for the upper bound is similar:

$$\frac{1}{p} \sum_{e \in \gamma} \text{cong}(e) = \frac{k^2}{\sqrt{j}} \sum_{i=1}^j \sum_{l=i}^k l^{-\frac{3}{2}}$$

$$\begin{aligned} &\leq \frac{2k^2}{\sqrt{j}} \left(2 - \frac{1}{\sqrt{k}} + \sum_{i=2}^j \left(\frac{1}{\sqrt{i-1}} - \frac{1}{\sqrt{k}} \right) \right) \\ &= \frac{2k^2}{\sqrt{j}} \left(2 - \frac{j}{\sqrt{k}} + \sum_{i=1}^{j-1} \frac{1}{\sqrt{i}} \right) \\ &\leq 2k^2 \left(\frac{2 + 2\sqrt{j-1}}{\sqrt{j}} - \sqrt{\frac{j}{k}} \right). \end{aligned}$$

The sum in parentheses is easily shown to be less than 3 for all allowed values of j and k (i.e., $k \geq j$).

The combination of these bounds implies that, for every path in the embedding that starts at a vertex v_j , the path resistance is $\Theta(k^2)$. Note that for any path starting at one of the leaves this value is increased by at most 1; thus any path in the embedding has resistance proportional to k^2 . Therefore the lower bound on λ is $\Theta(\frac{1}{k^2})$, which is within a constant of the upper bound. \square

We now show how to combine the star embedding lower bounds for paths with masses to get a lower bound on the uniform priority path resistance lower bound for trees. Consider what happens when we pick any point along critical path of the clique embedding and set it to zero. The zero point splits the path in two, and we get two lower bounds, one for each component.

To minimize subscript conflicts, we introduce the notation r^* to stand for r_{\max} . We have three r^* 's to consider: one from the clique embedding in the original graph, which we denote as r_K^* , and the two zero-boundary bounds with respect to our split point. Since we are working with a path, we can think of the path laid out with vertices in increasing order from left to right; we therefore denote the two path resistances used in the bounds as r_l^* for the left side and r_r^* for the right.

Recall that the lower bound for the clique embedding is $\frac{n}{r_K^*}$; we can relate it to the boundary case as follows: For tree T with a clique embedding and uniform priorities, let P be the critical path with vertices indexed from 1 to $k+1$ and let r_K^* be the path resistance of P . Let $1 \leq s \leq k+1$ be the index of a split point. Removing v_s separates T into subgraphs; let T_l and T_r be the subtrees containing s plus the vertices of P with indices less than s and greater than s respectively. Let r_l^* and r_r^* be the maximum path resistances for the star embeddings of T_l and T_r respectively when v_s is a zero boundary point and uniform priorities are used (if either of these trees is empty, its maximum path resistance is zero).

LEMMA 5.1. *For the situation described above,*

$$\frac{1}{r_l^* + r_r^*} < \frac{n}{r_K^*}$$

Proof. The critical path has k edges and $k + 1$ vertices, with edge i between vertices i and $i + 1$. Since all of the terms involved are positive, we work with the reciprocals of the quantities in the lemma statement. The reciprocal of the left term can be written as follows:

$$r_l^* + r_r^* = \sum_{i=1}^{s-1} \sum_{j=1}^i m_j + \sum_{i=s}^k \sum_{j=i+1}^{k+1} m_j.$$

We can write the reciprocal of the right expression as

$$\begin{aligned} \frac{r_K^*}{n} &= \frac{1}{n} \sum_{i=1}^k \left(\sum_{j=1}^i m_j \cdot \sum_{j=i+1}^{k+1} m_j \right) \\ &= \sum_{i=1}^{s-1} \left(\sum_{j=1}^i m_j \cdot \frac{\sum_{j=i+1}^{k+1} m_j}{n} \right) \\ &\quad + \sum_{i=s}^k \left(\frac{\sum_{j=1}^i m_j}{n} \cdot \sum_{j=i+1}^{k+1} m_j \right) \\ &< \sum_{i=1}^{s-1} \sum_{j=1}^i m_j + \sum_{i=s}^k \sum_{j=i+1}^{k+1} m_j. \end{aligned}$$

To see the final inequality, recall that $n = \sum_{i=1}^{k+1} m_i$. This proves the lemma. \square

We can use this to show how to assemble a vector that gives a Rayleigh quotient for the original Laplacian that is within a $\log \text{diam}(T)$ factor of λ_2 . This is shown in the next theorem:

THEOREM 5.3. *The uniform priority path resistance method produces a lower bound $\frac{n}{r_K^*} \leq \lambda_2$ for the Laplacian of any unweighted tree T that is off by a factor that is at most $O(\log \text{diam}(T))$.*

Proof. Let L be the Laplacian of T . We consider splitting the critical path at various points. Since r_l^* and r_r^* depend on the split point s , we make that clear by writing them as $r_l^*(s)$ and $r_r^*(s)$ respectively in this proof. Note that as s increases, $r_l^*(s)$ increases and $r_r^*(s)$ decreases. We also have that $r_l^*(1) = r_r^*(k+1) = 0$.

Let $\beta(s) = \frac{r_r^*(s)}{r_l^*(s)}$. This ratio is unbounded when $s = 1$ and decreases as s increases, reaching 0 when $s = k + 1$. We consider two cases: when there exists an s such that $2 \geq \beta(s) \geq \frac{1}{2}$, and when there is no such s .

In the first case, Lemma 5.1 gives us the following:

$$(5.10) \quad \frac{1}{r_l^*(s) + r_r^*(s)} < \frac{n}{r_K^*} \leq \lambda_2.$$

Theorem 5.1 gives a way to construct the vectors \mathbf{x}_l and \mathbf{x}_r so that the respective Rayleigh quotients for the left and right sides are $O\left(\frac{\max(\log(s-1), 1)}{r_l^*(s)}\right)$ and

$O\left(\frac{\max(\log(k+1-s), 1)}{r_r^*(s)}\right)$ respectively. Recall that these vectors are positive. The left Rayleigh quotient has value

$$RQ_l = \frac{\sum_{i=1}^{s-1} (x_i - x_{i+1})^2}{\sum_{i=1}^{s-1} m_i x_i^2},$$

the right has value

$$RQ_r = \frac{\sum_{i=s}^k (x_i - x_{i+1})^2}{\sum_{i=s+1}^{k+1} m_i x_i^2}.$$

We construct a vector \mathbf{x} for the whole tree as follows: assign the values from \mathbf{x}_l and the negatives of the values of \mathbf{x}_r to the corresponding vertices on the critical path, with $x_s = 0$. Extend the vector to the rest of the tree by assigning vertices off the critical path the values of their connection points as described in Section 4.3. To use the Rayleigh quotient for \mathbf{x} as an upper bound on λ_2 , its entries must sum to 0. However, this can be done by scaling one of the vectors \mathbf{x}_l or \mathbf{x}_r ; this scaling does not affect the respective Rayleigh quotients RQ_l and RQ_r . By construction, we have the following (remember that $x_s = 0$):

$$\frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\sum_{i=1}^{s-1} (x_i - x_{i+1})^2 + \sum_{i=s}^k (x_i - x_{i+1})^2}{\sum_{i=1}^{s-1} m_i x_i^2 + \sum_{i=s+1}^{k+1} m_i x_i^2}.$$

This is the sum of the numerators of RQ_l and RQ_r divided by the sum of their denominators. It is well known that, for positive values a, b, c , and d ,

$$\frac{a+c}{b+d} \leq \max\left(\frac{a}{b}, \frac{c}{d}\right).$$

Thus we have that

$$(5.11) \quad \lambda_2 \leq \max(RQ_l, RQ_r).$$

By (5.10), (5.11), and the condition on $\beta(s)$, the desired result holds for this case.

Now assume that there is no s such that $2 \geq \beta(s) \geq \frac{1}{2}$. Then there is some s such that $\beta(s) > 2$ and $\beta(s+1) < \frac{1}{2}$. By Lemma 5.1, both $\frac{1}{r_l^*(s) + r_r^*(s)}$ and $\frac{1}{r_l^*(s+1) + r_r^*(s+1)}$ are lower bounds on $\frac{n}{r_K^*}$, and hence on λ_2 . The second of these plus the condition on $\beta(s)$ imply that $\frac{2}{3r_l^*(s+1)} < \frac{n}{r_K^*}$; this follows since the condition on $\beta(s+1)$ implies that $r_l^*(s+1) > 2r_r^*(s+1)$. A similar argument shows that $\frac{2}{3r_r^*(s)}$ is also a lower bound.

We construct a vector \mathbf{x} to demonstrate an upper bound. The construction is slightly different this time; we construct \mathbf{x}_l with respect to the boundary being at vertex $s+1$, and \mathbf{x}_r with respect to the boundary being at vertex s . Corresponding values are mapped into \mathbf{x} as in the previous case. Once again, we can scale one side

as necessary to insure that the values in \mathbf{x} sum to zero without affecting the left or right Rayleigh quotients.

Note that if it were the case that

$$\frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \frac{\text{num}(RQ_l) + \text{num}(RQ_r)}{\text{denom}(RQ_l) + \text{denom}(RQ_r)},$$

the argument as in the previous case would show that $\max(RQ_l, RQ_r)$ would be an upper bound on λ_2 . This, combined with the lower bounds for this case demonstrated above, would be sufficient to prove that the theorem statement holds for this case. However, there is no zero point in this case, and the equality does not hold. To solve this problem, we will show that in fact the Rayleigh quotient is no more than a factor of two larger than the right hand side of the inequality.

To see this, note that the denominator of the actual Rayleigh quotient is the same as in the preceding expression. The numerator has the following changes: we lose two edges from the ends of \mathbf{x}_l and \mathbf{x}_r to presumed zero points; this decreases the numerator by x_s^2 and x_{s+1}^2 . We replace these edges by an edge between v_s and v_{s+1} that contributes $(x_s - x_{s+1})^2 < 2x_s^2 + 2x_{s+1}^2$. Hence

$$\begin{aligned} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} &< \frac{2(\text{num}(RQ_l) + \text{num}(RQ_r))}{\text{denom}(RQ_l) + \text{denom}(RQ_r)} \\ &\leq 2 \max(RQ_l, RQ_r), \end{aligned}$$

and there are upper and lower bounds within a log factor of the diameter.

This proves that there is a constant such that the theorem holds for the second case, which implies that the theorem holds for both cases using the larger of the two constants from the cases. \square

6 Open Questions

A number of interesting questions remain open. How good in general is the bound based on uniform priorities? Is there an easy way to set path priorities that works well in general? There is also the question of whether there is an easy way to calculate priorities for trees in both the Dirichlet case and the case without boundary conditions that improves the lower bounds.

References

- [Alo86] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [AM85a] N. Alon and V. D. Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38:73–88, 1985.
- [AM85b] W. N. Anderson, Jr. and T. D. Morley. Eigenvalues of the Laplacian of a graph. *Linear and Multilinear Algebra*, 18:141–145, 1985.
- [Axe92] O. Axelsson. Bounds of eigenvalues of preconditioned matrices. *SIAM Journal on Matrix Analysis and Applications*, 13(3):847–862, July 1992.
- [DS84] P. G. Doyle and J. L. Snell. *Random Walks and Electrical Networks*. The Mathematical Association of America, Washington, D.C., 1984.
- [DS91] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of Markov chains. *Annals of Applied Probability*, 1:36–61, 1991.
- [Fie73] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.
- [Fie75] M. Fiedler. A property of eigenvectors of non-negative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25(100):619–633, 1975.
- [GM95] Stephen Guattery and G.L. Miller. On the performance of spectral graph partitioning methods. In *6th ACM-SIAM Symposium on Discrete Algorithms*, pages 233–242, San Francisco, January 1995. ACM-SIAM.
- [GMZ95] K.D. Gremban, G.L. Miller, and M. Zaghera. Performance evaluation of a parallel preconditioner. In *9th International Parallel Processing Symposium*, pages 65–69, Santa Barbara, April 1995. IEEE.
- [HY81] Lous A. Hageman and Davis M. Young. *Applied Iterative Methods*. Computer Science and Applied Mathematics. Academic Press, Inc, San Diego and London, 1981.
- [JS89] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, December 1989.
- [Kah96] Nabil Kahale. A semidefinite bound for mixing rates of Markov chains. In *Proc. Fifth Integer Programming and Combinatorial Optimization Conference*, pages 190–203, 1996. Lecture Notes in Computer Science 1084.
- [Min88] H. Minc. *Nonnegative Matrices*. Wiley, New York, 1988.
- [Moh89] Bojan Mohar. Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B*, 47:274–291, 1989.
- [PSL90] Alex Pothen, Horst D. Simon, and Kang-Pu Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM Journal on Matrix Analysis and Applications*, 11(3):430–452, July 1990.
- [Sin92] A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1:351–370, 1992.
- [SJ89] A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, 82:93–133, 1989.
- [ST96] Daniel A. Spielman and Shang-Hua Teng. Spectral partitioning works: Planar graphs and finite element meshes. Technical Report UCB CSD-96-898, U.C. Berkeley, 1996. Will also appear in the 37th Annual Symposium on Foundations of Computer Science.