SIAM J. MATRIX ANAL. APPL.
Vol. 21, No. 3, pp. 703–723
© 2000 Society for Industrial and Applied Mathematics

GRAPH EMBEDDINGS AND LAPLACIAN EIGENVALUES*

STEPHEN GUATTERY† AND GARY L. MILLER‡

Abstract. Graph embeddings are useful in bounding the smallest nontrivial eigenvalues of Laplacian matrices from below. For an $n \times n$ Laplacian, these embedding methods can be characterized as follows: The lower bound is based on a clique embedding into the underlying graph of the Laplacian. An embedding can be represented by a matrix Γ; the best possible bound based on this embedding is $n/\lambda_{\max}(\Gamma^T\Gamma)$, where $\lambda_{\max}$ indicates the largest eigenvalue of the specified matrix. However, the best bounds produced by embedding techniques are not tight; they can be off by a factor proportional to $\log^2 n$ for some Laplacians.

We show that this gap is a result of the representation of the embedding: By including edge directions in the embedding matrix representation Γ, it is possible to find an embedding such that $\Gamma^T\Gamma$ has eigenvalues that can be put into a one-to-one correspondence with the eigenvalues of the Laplacian. Specifically, if λ is a nonzero eigenvalue of either matrix, then n/λ is an eigenvalue of the other. Simple transformations map the corresponding eigenvectors to each other. The embedding that produces these correspondences has a simple description in electrical terms if the underlying graph of the Laplacian is viewed as a resistive circuit. We also show that a similar technique works for star embeddings when the Laplacian has a zero Dirichlet boundary condition, though the related eigenvalues in this case are reciprocals of each other. In the zero Dirichlet boundary case, the embedding matrix Γ can be used to construct the inverse of the Laplacian. Finally, we connect our results with previous techniques for producing bounds and provide an example.

Key words. Laplacian matrices, graph eigenvalues and eigenvectors, graph embeddings

AMS subject classifications. 05C50, 15A18, 15A42

PII. S0895479897329825

1. Introduction.

In this paper we present an exact relationship between graph embeddings and the eigenvalues and eigenvectors of Laplacian matrices. The study of the connection between Laplacian spectra (particularly with respect to λ_2, the second smallest (and smallest nonzero) eigenvalue of the Laplacian) and properties of the associated graphs dates back to Fiedler's work in the 1970s (see, e.g., [10] and [11]). The Laplacian also has an important role in representing physical problems. It occurs in finite difference, finite element, and control volume representations of problems 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

*Received by the editors November 10, 1997; accepted for publication (in revised form) by S. A. Vavasis January 27, 1999; published electronically February 10, 2000.
http://www.siam.org/journals/simax/21-3/32982.html

†Department of Computer Science, Bucknell University, Lewisburg, PA 17837 (guattery@bucknell.edu). The research of this author was supported by the National Aeronautics and Space Administration under NASA contract NAS1-19480 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681.

‡School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 (glmiller@cs.cmu.edu). The work of this author was supported in part by NSF grant CCR-9505472, DARPA contract N00014-95-1246, and Army contract DAAH04-95-1-0607.

of iterative linear system solvers to estimate rates of convergence [18] and to analyze the quality of preconditioners [4, 13]. Bounds on λ_2 are useful in the analysis of spectral partitioning, both because λ_2 occurs in bounds on cut quality [24] and because they can be used in isolating structural properties of the eigenvectors used in making the cuts [16, 28]. The eigenvalue λ_2 has been related to expansion properties of graphs and can be used in determining if a graph is an expander [1, 2].

One common class of techniques for computing such lower bounds uses properties of graph embeddings [9, 15, 20, 22, 26, 27]. In such methods, a graph H is embedded into the graph G under study; that is, vertices of H are identified with vertices of G, and paths in G are specified to correspond to edges in H. Specific metrics on the embedding such as congestion and dilation are computed; they are then used to calculate lower bounds. H is usually chosen to have some specific structure; most often it is a clique or a star. The bounds produced in general are not tight, however. Kahale [22] has shown that for some graphs, the gap can be a factor of $\log^2 n$, where $n = |V(G)|$, the number of vertices in G.

In this paper, we prove an exact relationship between graph embeddings and Laplacian eigenvalues. In particular, we present a matrix representation of embeddings that differs from the representation used in the methods mentioned above. This representation allows any embedding of H into G to be expressed as an $|E(H)|\times|E(G)|$ matrix Γ, where $|E(H)|$ and $|E(G)|$ are the number of edges in H and G, respectively. We also introduce the current flow embedding, which is based on routing unit currents between pairs of vertices when edge weights are viewed as conductances. We denote the embedding matrix for the current flow matrix as Γ_{cf}. We show that, for a clique embedding based on routing a unit current between each pair of vertices, if the Laplacian $L(G)$ has a nonzero eigenvalue λ with eigenvector $\mathbf{u}$, then $\Gamma_{cf}^T\Gamma_{cf}$ has eigenvalue n/λ; the corresponding eigenvector is $B\mathbf{u}$, where B is the edge-vertex incidence matrix of the Laplacian.

In the zero Dirichlet boundary case, we show that, for a star embedding based on routing a unit current between every vertex and the boundary, if the Laplacian $L(G)$ has an eigenvalue λ with eigenvector $\mathbf{u}$, then $\Gamma_{cf}^T\Gamma_{cf}$ has eigenvalue $1/\lambda$; the corresponding eigenvector is $B\mathbf{u}$, where B is the edge-vertex incidence matrix of the Laplacian. In this case, the Laplacian is positive definite; we show that $\Gamma_{cf}\Gamma_{cf}^T$ is L^{-1}.

In both cases, we show the extension of these results to weighted graphs.

In addition to their utility with respect to the problems listed above, our results are interesting because they provide further illustration of the utility of looking at Laplacian spectra problems with respect to electrical circuits.

The rest of this paper is organized as follows: Section 2 provides an overview of work in this area. Section 3 covers the notation we use and background information on graphs and matrices. Section 4 covers general lemmas about our embedding representation. Section 5 covers the current flow embedding for the clique case. Section 6 covers the current flow embedding for the star case with zero Dirichlet boundary and shows how to construct inverses using embeddings. Section 7 presents related results, including the connection between the results in this paper and previous lower bound techniques.

2. Previous work. 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 1970s (see, e.g., [10] and [11]). These properties have been used in graph algorithms, particularly algorithms for finding small separators [17, 25, 28].

The relationship between graph embeddings and matrix representations has been

the subject of much interesting research. A large proportion of this work has been aimed at bounding the second largest eigenvalues of time-reversible Markov chains in order to bound the mixing time for random walks. The use of clique embeddings to bound eigenvalues arose in the analysis of mixing times for Markov chains by Jerrum and Sinclair [20, 27]. Further work in this direction was done by Diaconis and Stroock [9] and by Sinclair [26]. Kahale [22] generalized this work in terms of methods assigning lengths to the graph edges and showed 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 (Kahale [22] also cited unpublished work along these lines by Fill and Sokal). He also gave a semidefinite programming formulation for a model allowing fractional paths and showed that the bound is off by at most a factor of $\log^2 n$. He showed that this gap is tight; he also noted that the results can be applied to Laplacians with suitable modifications.

Guattery, Leighton, and Miller [15] presented a lower bound technique for λ_2 of a Laplacian. It assigns priorities to paths in the embedding and uses these to compute congestions of edges in the original graph with respect to the embedding. Summing the congestions along the edges in a path gives the path congestion; the lower bound is a function of the reciprocal of the maximum path congestion taken over all paths. For the clique case, they showed that this method is the dual of the method presented in [22]; the best lower bounds produced by these methods are the same. They also showed how to apply their method in the zero Dirichlet boundary case by using star embeddings. In the clique case, they showed that using uniform priorities for any tree T gives a lower bound that is within a factor proportional to the logarithm of the diameter of the tree.

Gremban and Miller [14] have shown how to use embeddings to generate support numbers, which also are bounds on the largest and smallest generalized eigenvalues (and hence the spectral condition number) of preconditioned linear systems involving a generalized definition of Laplacians. They also defined the support tree preconditioner and used the support number bounds to prove properties about the quality of these preconditioners. Gremban, Miller, and Zagha have evaluated the performance of these techniques [13].

The construction of the embedding we use below is related to the effective resistances between pairs of vertices in the graph when the graph is viewed as a network of unit conductances. Chandra et al. [8] have defined the maximum such value taken over all distinct pairs as the electrical resistance of a graph and have used that quantity to bound the commute and cover times of random walks on the graph.

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 as to which graph is being referred to. When it is clear which graph we are referring to, we use n to denote $|V|$ and m to denote $|E|$.

We use the term *path graph* for a tree that has exactly two vertices of degree 1 (that is, exactly two vertices with exactly one incident edge). 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 nonleaf 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 ith entry. The notation $\mathbf{x} = 0$ indicates that all entries of the vector $\mathbf{x}$ are zero; $\mathbf{1}$ indicates the vector that has 1 for every entry. Unless specifically noted otherwise, we index the eigenvalues of an $n \times n$ matrix in nondecreasing order: λ_1 represents the smallest eigenvalue, and λ_n the largest. We use the notation $\lambda_i(A)$ (respectively, $\lambda_i(G)$) to indicate the ith eigenvalue of matrix A (respectively, of the Laplacian of graph G) if there is any ambiguity about which matrix (respectively, graph) the eigenvalue belongs to. The notation $\mathbf{u}_i$ represents the eigenvector corresponding to λ_i.

We use the following notational shorthand: Let D be a diagonal matrix with nonnegative entries on the diagonal. Then $\sqrt{D}$ represents the diagonal matrix with $[\sqrt{D}]_{ii} = \sqrt{d_{ii}}$. If all diagonal entries of D are positive, it is obvious that $\sqrt{D}^{-1} = \sqrt{D^{-1}}$.

I_k denotes the $k \times k$ identity matrix.

For two matrices A and B with the same dimensions, $A \geq B$ if for every entry a_{ij} of A, we have $a_{ij} \geq b_{ij}$.

The following is a standard result (see, e.g., [3] for a version related to Laplacians).

LEMMA 3.1. *Let X be a real matrix. Let $Y = X^T X$ and $Z = XX^T$. If $\mathbf{u}$ is an eigenvector of Y with eigenvalue $\lambda > 0$, then $X\mathbf{u}$ is an eigenvector of Z with eigenvalue λ. Likewise, if $\mathbf{v}$ is an eigenvector of Z with eigenvalue $\mu > 0$, then $X^T\mathbf{v}$ is an eigenvector of Y with eigenvalue μ. Thus Y and Z have the same nonzero eigenvalues.*

Proof. $Z(X\mathbf{u}) = XX^T(X\mathbf{u}) = XY\mathbf{u} = \lambda X\mathbf{u}$, so the first claim holds. Likewise, $Y(X^T\mathbf{v}) = X^TX(X^T\mathbf{v}) = X^TZ\mathbf{v} = \mu X^T\mathbf{v}$, so the second claim holds. Note that Y and Z are positive semidefinite, so any nonzero eigenvalues are positive and the final claim holds. □

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} equal to the number of edges incident to vertex v_i for $v_i \in V(G)$ and all off-diagonal entries zero. Let A be the adjacency matrix for G ($[A]_{ij} = 1$ if and only if $(v_i, v_j) \in E(G)$, 0 otherwise). Then the Laplacian of G is the matrix $L = D - A$.

The following are useful facts about the Laplacian:

- The Laplacian is symmetric positive semidefinite, so all its eigenvalues are greater than or equal to 0 (see, e.g., [3, 10]).
- A graph G is connected if and only if 0 is an eigenvalue of the Laplacian of G with multiplicity 1 (see, e.g., [10]).
- The following characterization of λ_2 holds (see, e.g., [10]):

$$\lambda_2 = \min_{\mathbf{x}^T\mathbf{1}=0} \frac{\mathbf{x}^T L\mathbf{x}}{\mathbf{x}^T\mathbf{x}}. \tag{3.1}$$

- For any vector $\mathbf{x}$ and Laplacian matrix L of the graph G, we have (see, e.g., [19]) the following:

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

An *edge-weighted* graph is a graph for which a real, nonzero weight w_{ij} is associated with each edge (v_i, v_j). Fiedler extended the notion of the Laplacian to graphs with positive edge weights [11]; 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 (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 subspaces with respect to the Laplacian. It is well known (see, e.g., [6, 7]) that the Laplacian of a graph G can be expressed as the product $B^T B$ of G's *edge-vertex incidence matrix* B. B is constructed as follows: Arbitrarily, direct the edges of G and index them from 1 to m. B is an $m \times n$ matrix, where row i has a -1 in the column corresponding to the vertex at its tail, a 1 in the column corresponding to the vertex at its head, and zeros in all other columns.

Generalized Laplacians can also be expressed as a product involving the edge-vertex incidence matrix B. However, we must introduce an edge weight matrix W (also referred to as the *conductance matrix*). W is a diagonal matrix with entry $[W]_{ii}$ equal to the weight of edge i. It is easy to check that the generalized Laplacian $L = B^T W B$ and that L is independent of the directions of the edges in B.

It is common to think of the set of vectors of length n as the *vertex space*; each entry of the vector $\mathbf{x}$ assigns the value x_i to corresponding vertex i (such an assignment of values is often referred to as a *valuation*). Likewise, the set of all vectors of length m is the *edge space*. Since we will use an electrical analogy below, we sometimes refer to the values assigned to vertices as *potentials*. Note that multiplying a vector of potentials on the left by B produces a vector of *potential differences* in the edge space: each entry in the result is the difference in potentials at the head and tail of some edge in G.

The edge space of the graph can be partitioned into two orthogonal subspaces (see, e.g., [5, 7]). The first is the *cycle space*, the subspace spanned by vectors representing cycles in G. A cycle C is represented by a vector as follows: The edges of C have a natural direction, where each edge is directed toward its successor in the cyclic order. For each edge e in C, the vector $\mathbf{v}_C$ has entry 1 if the direction of e with respect to C and the direction of e used in defining the edge-vertex incidence matrix B are the same, and -1 if they are opposite. The entries for edges not in C are 0.

The cycle space is the orthogonal complement of the *cocycle* or *cut space*, the subspace spanned by vectors representing cuts. Let $(S, \overline{S})$ represent a partition of the vertices of a graph G into two sets. The set of edges with one end in S and the other in $\overline{S}$ is called a *cut*. Cuts are represented as vectors as follows: for each edge in the cut, if the edge has its tail in S (i.e., the edge is directed (according to B) from S to $\overline{S}$), the vector entry for that edge is 1. If the edge has its head in S (i.e., the edge is directed from $\overline{S}$ to S), the vector entry for that edge is -1. Entries for edges not in the cut are 0.

It is well known that, for a connected graph G with m edges on n vertices, the dimension of the cycle space is $m - n + 1$ and the dimension of the cut space is $n - 1$ (see, e.g., [7, Theorem 5, Section II.3]). Some simple but useful results about the

relation of the cycle space and cut spaces to the edge-vertex incidence matrix B are stated in the following lemma (the proofs are easy and left to the reader).

LEMMA 3.2. *Let G be a connected graph with edge-vertex incidence matrix B. Then*

- *the span of the columns of B is the cut space of G, and any $n-1$ columns of B form a basis for the cut space;*
- *the cycle space of G is the null space of B^T.*

3.4. Laplacians and electrical circuits. In this paper we consider the graph of a Laplacian as a resistive electrical circuit, with the edge weights representing conductances between the vertices or nodes (conductance is the reciprocal quantity to resistance). We also calculate quantities such as the current flows on the edges of the graph when unit currents are injected at one point and removed at others. These calculations can be defined in terms of three electrical laws.

The first of these is *Ohm's law*, which connects potential differences and currents. We use it mainly in defining *currents*: As noted above, we can consider any valuation of the vertices as a vector of potentials or *voltages*. For any pair of vertices connected by an edge, we define the *potential difference* across the edge as the value at the head of the edge minus the value at the tail. Thus, given a vector $\mathbf{v}$ of voltages, $B\mathbf{v}$ gives the vector of potential differences, which is in the edge space. Ohm's law says that the current on an edge is the potential difference times the conductance. Thus, for $\mathbf{v}$ we can define the vector $\mathbf{i}$ of currents as $\mathbf{i} = WB\mathbf{v}$, where W is the conductance matrix as discussed in section 3.3.

The second is *Kirchoff's voltage law (KVL)*, which states that the potential drops around any cycle in the graph sum to zero. This requires that for any cycle $\mathbf{c}$, $\mathbf{c}^T B\mathbf{v} = 0$. The second point in Lemma 3.2 shows that this holds for our representation, which is thus consistent with KVL.

The third law is *Kirchoff's current law (KCL)*, which states that the net current flow at any vertex is zero; that is, the sum of the currents over all edges incident to a vertex is the negative of the current injected at that vertex. Note that multiplying a vector of currents (in the edge space) by B^T sums the currents at the vertices. Therefore, if $\mathbf{i}_{\text{ext}}$ is a vector of external currents injected into the circuit, KCL says that the resulting currents in the circuit (represented by the vector $\mathbf{i}$) must obey

$$B^T\mathbf{i} = \mathbf{i}_{\text{ext}}. \tag{3.3}$$

We frequently use KCL as a requirement in calculating flows in the circuit given a unit current injected at one vertex and extracted at another. This is equivalent to solving for $\mathbf{i}$ in (3.3). Note that for a solution to exist, the amount of current injected and extracted in $\mathbf{i}_{\text{ext}}$ must be equal; we must have $\mathbf{i}_{\text{ext}}^T\mathbf{1} = 0$. This is equivalent to the condition that $\mathbf{i}_{\text{ext}}$ be in the column space of B^T, or that it is orthogonal to the null space of B, which is spanned by $\mathbf{1}$ when the graph is connected.

3.5. Graph embeddings. 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 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 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.

4. Representing graph embeddings with matrices. Let G be a connected graph, and let Γ be an embedding of H into G. Consider the following observation: Each vertex appears at most once on any path $\gamma \in \Gamma$. Therefore, we can choose one end of the embedded edge $e = (u, v) \in E(H)$ as a source node (say, u) and the other (say, v) as the sink node and order path edges in sequence from u to v. A path (or set of fractional paths) for e in the embedding can thus be viewed as a unit flow from u to v, where the net flow into any intermediate vertex is zero.

These flows can be represented as vectors in the edge space. For the purposes of our representation, we assume that directions have been assigned arbitrarily to the edges of G, and that the vertices and edges of G have been indexed from 1 to n and from 1 to m, respectively. Let $e = (u, v)$ be an edge in H; e is represented in the embedding as a unit flow in G. In the vector for a flow, each edge gets a value with magnitude equal to the amount of the flow it carries. The entry for each edge in the flow also has a sign: positive if the flow direction agrees with the edge direction, and negative if the flow is opposite the edge direction.

The flow vectors can be assembled into a matrix representing an embedding. We will use the symbol Γ to represent the embedding matrix as well as the embedding; the use should be clear from context. The matrix Γ includes one row for each edge e in H; that row is the flow vector for the embedding of e into G. Thus Γ is a $|E(H)| \times |E(G)|$ matrix. Note that Γ depends on both G and H; these graphs will be clear from context, so we will not introduce any notation to express this dependency. This matrix representation is similar to the representation presented in [21], though it differs by including negative entries. We discuss the relationship between these two representations in section 7 below.

The following result holds for arbitrary H embedded into connected graph G.

LEMMA 4.1. *For any embedding matrix Γ representing an embedding of H into G, the product $\Gamma B(G) = B(H)$, where the vertex set of $B(H)$ is $V(G)$.*

Proof. Consider the result of multiplying a row of Γ and a column of $B(G)$. The column of $B(G)$ represents a vertex v of G. It is easy to see from the construction of the rows of Γ that this product gives the net flow out of v. We thus have the following:

- If v is the source of the flow, then the product is -1.
- If v is the sink of the flow, then the product is 1.
- Otherwise, the value is 0.

Thus, the row of $\Gamma B(G)$ corresponding to the flow for edge (u, v) has a -1 in the column for u, a 1 in the column for v, and zeros elsewhere. This is the same as the row for (u, v) in $B(H)$, where the edge directions correspond to the directions of the flows used in the embedding. ∎

It is useful to define the matrix $M = \Gamma^T \Gamma$. M is $|E(G)| \times |E(G)|$.

The results in the rest of the paper are based on a special embedding we will refer to as the *current flow embedding*. This embedding is defined for any generalized Laplacian as follows: The graph G serves as a network of conductances, where the conductance of each edge is its (positive) weight. For each edge (u, v) in H, the flow from u to v is the same as the current flow in G when a unit current is injected at u and extracted at v. Since the currents obey Kirchoff's laws, the set of flows clearly forms an embedding of H into G. We note that this embedding can be extended to cases with a zero Dirichlet boundary condition by adding in a ground node; details of this extension are deferred to section 6 below.

We denote this embedding Γ_{cf}. M_{cf} denotes the matrix M for the embedding matrix Γ_{cf}.

In the remainder of this paper, the following results about Γ_{cf} prove useful.

LEMMA 4.2. *Let G be an unweighted connected graph, and let H be a graph on $V(G)$ with a nonempty edge set. Let Γ_{cf} be as defined above for the embedding of H into G. Then the cycle space of G is in the null space of Γ_{cf}.*

Proof. This lemma is a consequence of each flow in the embedding obeying Kirchoff's laws.

Note that each row of Γ_{cf} is the transpose of a current vector $\mathbf{i}$. Because each such current vector is consistent with KVL and KCL, there exists a potential vector $\mathbf{v}$ (which is not unique) such that $\mathbf{i} = WB\mathbf{v}$ by Ohm's law. To see this, note that the external current flow vector $\mathbf{i}_{\text{ext}}$ used to define Γ_{cf} injects a unit current at one vertex and extracts it at another. Thus $\mathbf{i}_{\text{ext}}$ is orthogonal to the null space of the Laplacian (the vector of all ones), and there is a vector of potentials $\mathbf{v}$ such that $\mathbf{i}_{\text{ext}} = L\mathbf{v} = B^T WB\mathbf{v}$. However, by KCL, $\mathbf{i}_{\text{ext}} = B^T\mathbf{i}$, so either $\mathbf{i} = WB\mathbf{v}$ or $\mathbf{i} = WB\mathbf{v} + \mathbf{c}$, where $\mathbf{c}$ is a nonzero vector in the null space of B^T, which is the cycle space of the graph (Lemma 3.2). Assume the latter. The potential drops for the currents in $\mathbf{i}$ can be written as $W^{-1}\mathbf{i}$; since $\mathbf{c}$ is in the cycle space, KVL implies that

$$0 = \mathbf{c}^T W^{-1}\mathbf{i} = \mathbf{c}^T W^{-1}\left(WB\mathbf{v} + \mathbf{c}\right) = \mathbf{c}^T W^{-1}\mathbf{c}.$$

Since W is a diagonal matrix with positive diagonal entries, the last term cannot be 0, so we must have that $\mathbf{i} = WB\mathbf{v}$, which demonstrates the claim.

For an unweighted graph, $W = I_m$, so this simplifies to $\mathbf{i} = B\mathbf{v}$. As already noted, the cycle space is the null space of B^T and, hence, the left null space of B, so for any vector $\mathbf{c}$ in the cycle space, $\mathbf{c}^T\mathbf{i} = \mathbf{c}^T B\mathbf{v} = 0$. This holds for every row in Γ_{cf}, and hence Lemma 4.2 holds. □

A similar result holds if G is a weighted graph.

LEMMA 4.3. *Let G be a weighted connected graph with conductance matrix W, and let H be a graph on $V(G)$ with a nonempty edge set. Let Γ_{cf} be as defined above for the embedding of H into G. Then for any vector $\mathbf{c}$ in the cycle space of G, $W^{-1}\mathbf{c}$ is in the null space of Γ_{cf}.*

Proof. The proof is the same as that for Lemma 4.2, except that W is no longer the identity matrix. However, the vector $\mathbf{c}$ is now scaled by W^{-1}, and the reader can check that the terms W and W^{-1} cancel each other in the argument. □

This result can be stated in slightly altered form; it is no longer consistent with the electrical analogy, but it is useful in the algebraic arguments below.

COROLLARY 4.4. *Let G be a weighted connected graph with conductance matrix W, and let H be a graph on $V(G)$ with a nonempty edge set. Let Γ_{cf} be as defined above for the embedding of H into G. Then for any vector $\mathbf{c}$ in the cycle space of G, $\sqrt{W^{-1}}\mathbf{c}$ is in the null space of $\Gamma_{cf}\sqrt{W^{-1}}$.*

Proof. Since W^{-1} is a diagonal matrix with positive real entries on the diagonal, we can factor it into $\sqrt{W^{-1}} \cdot \sqrt{W^{-1}}$. The result follows from Lemma 4.3. □

We next consider the special properties of the current flow embedding with respect to embedding particular graphs into G.

5. The clique embedding Γ_{cf}. We now show that there is an exact connection between clique embeddings and Laplacian eigenvalues. Let G be a connected graph on n vertices with positive edge weights and Laplacian $L = B^T WB$, where B is the edge-vertex incidence matrix and W is the conductance matrix as defined in section 3.3. Let Γ_{cf} be the embedding matrix for the current flow embedding of the complete graph into G, with M_{cf} defined as in the previous section. For clarity's sake, we start with the result in the case where G is unweighted (that is, W is the identity matrix).

THEOREM 5.1. *Let G be an unweighted connected graph on n vertices with Laplacian L. Let $M_{cf} = \Gamma_{cf}^T \Gamma_{cf}$, where Γ_{cf} is the current flow embedding of the complete graph on n vertices into G. For any $\lambda > 0$, λ is an eigenvalue of L if and only if $\frac{n}{\lambda}$ is an eigenvalue of M_{cf}. Further, if λ has eigenvector $\mathbf{u}$ for L, then $B\mathbf{u}$ is an eigenvector of M_{cf} with eigenvalue $\frac{n}{\lambda}$.*

Proof. By Lemma 4.1, $B^T \Gamma_{cf}^T \Gamma_{cf} B = K$, where K is the Laplacian of the complete graph on $n = |V(G)|$. It is easy to show that, for any vector $\mathbf{x}$ such that $\mathbf{x}^T \mathbf{1} = 0$, $K\mathbf{x} = n\mathbf{x}$. Thus for any such $\mathbf{x}$, we have the following:

$$B^T M_{cf} B\mathbf{x} = B^T \Gamma_{cf}^T \Gamma_{cf} B\mathbf{x} = n\mathbf{x}. \tag{5.1}$$

Since K has rank $n-1$, the preceding result implies that the ranks of M_{cf}, Γ_{cf}, and Γ_{cf}^T are all at least $n-1$.

Lemma 4.2 shows that the null space of Γ_{cf} contains the cycle space of G. The dimension of the null space is thus at least $m - n + 1$; combining this with the result from the previous paragraph shows that the null space is exactly the cycle space and that the rank of Γ_{cf} (and hence of Γ_{cf}^T) is $n-1$. The null space of M_{cf} contains the null space of Γ_{cf}, and the same results hold for its rank and null space. M_{cf} is clearly symmetric positive semidefinite from its definition.

Let $\mathbf{x}$ be any vector in the cut space of G. Since M_{cf} is symmetric with the cycle space of G as its null space, $M_{cf}\mathbf{x}$ is orthogonal to the cycle space.

Let $\mathbf{u}$ be an eigenvector of L with eigenvalue $\lambda > 0$. Then

$$B^T M_{cf} B\mathbf{u} = n\mathbf{u} = \frac{n}{\lambda} L\mathbf{u} = B^T B \frac{n\mathbf{u}}{\lambda}.$$

Thus

$$B^T \left(M_{cf} B\mathbf{u} - \frac{n}{\lambda} B\mathbf{u} \right) = 0,$$

which implies that $M_{cf} B\mathbf{u} - \frac{n}{\lambda} B\mathbf{u}$ is in the null space of B^T. By Lemma 3.2, this means it is in the cycle space of G. But we showed above that $M_{cf} B\mathbf{u}$ is orthogonal to the cycle space. Further, Lemma 3.2 also implies that $B\mathbf{u}$ lies in the cut space of G and is therefore orthogonal to the cycle space. Thus we must have that $M_{cf} B\mathbf{u} = \frac{n}{\lambda} B\mathbf{u}$. Thus $B\mathbf{u}$ is an eigenvector of M_{cf} with eigenvalue $\frac{n}{\lambda}$.

We still need to show that M_{cf} does not have any other nonzero eigenvalues. Let $\mathbf{u}$ be any eigenvector of M_{cf} with eigenvalue μ. Since the null space of M_{cf} is the cycle space of G, $\mathbf{u}$ is in the cut space. By Lemma 3.2, this is in the column space of B, and we can find a vector $\mathbf{x}$ such that $B\mathbf{x} = \mathbf{u}$; further, we can choose $\mathbf{x}$ such that $\mathbf{x}^T \mathbf{1} = 0$. Then

$$\mu\mathbf{u} = M_{cf}\mathbf{u} = M_{cf} B\mathbf{x} = \mu B\mathbf{x}. \tag{5.2}$$

Since $\mathbf{x}$ is orthogonal to $\mathbf{1}$, we can use (5.1) with (5.2) as follows:

$$n\mathbf{x} = B^T M_{cf} B\mathbf{x} = \mu B^T B\mathbf{x} = \mu L\mathbf{x}.$$

Thus $\mathbf{x}$ is an eigenvector of L with eigenvalue $\lambda = \frac{n}{\mu}$. That is, for any eigenvalue μ of M_{cf}, there is an eigenvalue λ of L such that $\mu = \frac{n}{\lambda}$, and the correspondence between eigenvalues is consistent with the theorem statement. ∎

Note that if $\mathbf{v}$ is an eigenvector corresponding to a nonzero eigenvalue of M_{cf}, $B^T\mathbf{v}$ is an eigenvector of L.

5.1. Extending the results to the weighted case. Extending the results about M_{cf} to the weighted case is straightforward. We now need to take into account the matrix W in the calculations; otherwise, the argument is similar.

THEOREM 5.2. *Let G be a connected graph on n vertices with positive edge weights, conductance matrix W, and generalized Laplacian L. Let $M_{cf} = \Gamma_{cf}^T \Gamma_{cf}$, where Γ_{cf} is the current flow embedding of the complete graph on n vertices into G. For any $\lambda > 0$, λ is an eigenvalue of L if and only if $\frac{n}{\lambda}$ is an eigenvalue of $M_{cf}W^{-1}$. Further, if λ has eigenvector $\mathbf{u}$ for L, then $WB\mathbf{u}$ is an eigenvector of $M_{cf}W^{-1}$ with eigenvalue $\frac{n}{\lambda}$.*

Proof. The theorem is stated so as to emphasize the connection with the electrical interpretation; we will actually prove the result for the similar matrix

$$\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} = \sqrt{W^{-1}} \left(M_{cf} W^{-1} \right) \sqrt{W}.$$

This similarity transform corresponds to a change of variable in the theorem statement. We will actually prove that for any eigenvector $\mathbf{u}$ of L with eigenvalue $\lambda > 0$, $\sqrt{W}B\mathbf{u}$ is an eigenvector of $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ with eigenvalue $\frac{n}{\lambda}$. We will also show that $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ has no other eigenvalues. It is easy to verify that these results imply the theorem.

We again use Lemma 4.1 to show that

$$B^T \sqrt{W} \left(\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} \right) \sqrt{W} B = B^T M_{cf} B = B^T \Gamma_{cf}^T \Gamma_{cf} B = K,$$

and hence that for any vector $\mathbf{x}$ such that $\mathbf{x}^T \mathbf{1} = 0$,

$$B^T \sqrt{W} \left(\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} \right) \sqrt{W} B \mathbf{x} = n\mathbf{x}.$$

Since K has rank $n-1$, this result implies that the ranks of $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$, $\Gamma_{cf} \sqrt{W^{-1}}$, and $\sqrt{W^{-1}} \Gamma_{cf}^T$ are all at least $n-1$.

By Corollary 4.4, the null space of $\Gamma_{cf} \sqrt{W^{-1}}$ contains the cycle space of G multiplied by $\sqrt{W^{-1}}$. Because $\sqrt{W^{-1}}$ is an invertible matrix, the dimension of the null space is thus at least $m - n + 1$. Combining this with the result from the previous paragraph shows that this is exactly the null space and that the ranks of $\Gamma_{cf} \sqrt{W^{-1}}$, $\sqrt{W^{-1}} \Gamma_{cf}^T$, and $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ are $n-1$.

Through the rest of this proof, we will use the notation C_0 to denote $\sqrt{W^{-1}}$ times the cycle space of G; that is, the null space of $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ is C_0. Using Lemma 3.2 and the fact that W is invertible, it is easy to show that C_0 is the null space of $B^T \sqrt{W}$ and is orthogonal to the column space of $\sqrt{W} B$.

Because $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ is symmetric with C_0 as its null space, the matrix $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} \mathbf{x}$ is orthogonal to C_0 for any vector $\mathbf{x}$ in the edge space.

Let $\mathbf{u}$ be an eigenvector of L with eigenvalue $\lambda > 0$. Then

$$B^T \sqrt{W} \left(\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} \right) \sqrt{W} B \mathbf{u} = B^T M_{cf} B \mathbf{u} = n\mathbf{u} = \frac{n}{\lambda} L \mathbf{u} = B^T W B \frac{n\mathbf{u}}{\lambda}.$$

Thus

$$B^T \sqrt{W} \left(\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} (\sqrt{W} B \mathbf{u}) - \frac{n}{\lambda} (\sqrt{W} B \mathbf{u}) \right) = 0,$$

which implies that $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}} (\sqrt{W} B \mathbf{u}) - \frac{n}{\lambda} (\sqrt{W} B \mathbf{u})$ is in C_0, the null space of $B^T \sqrt{W}$. But we showed above that the first product is orthogonal to C_0. Further,

the second product is in the column space of $\sqrt{W}B$, which we noted above is orthogonal to C_0. Thus the two products must be equal, and $\sqrt{W}B\mathbf{u}$ is an eigenvector of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ with eigenvalue $\frac{n}{\lambda}$.

We still need to show that $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ does not have any other nonzero eigenvalues. Define C_1 to be the subspace that results from transforming the cut space of G by multiplying it by $\sqrt{W}$. Since $\sqrt{W}$ is invertible, C_1 has dimension $n-1$; we noted above that it is orthogonal to C_0. Because (the symmetric matrix) $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ has null space C_0, any eigenvector of a nonzero eigenvalue must come from C_1.

Let $\mathbf{u} \in C_1$ be any eigenvector of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ with eigenvalue $\mu > 0$. By the definition of C_1, there is a vector $\mathbf{z}$ in the cut space of G such that $\mathbf{z} = \sqrt{W^{-1}}\mathbf{u}$. We showed in the proof of Theorem 5.1 that for any vector $\mathbf{z}$ in the cut space, there is a vector $\mathbf{x}$ in the vertex space such that $B\mathbf{x} = \mathbf{z}$ and $\mathbf{x}^T\mathbf{1} = 0$. Thus we can write $\mathbf{u} = \sqrt{W}B\mathbf{x}$ for the appropriate such $\mathbf{x}$.

Then

$$\mu\mathbf{u} = \sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\mathbf{u} = \sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}(\sqrt{W}B\mathbf{x}) = \mu\sqrt{W}B\mathbf{x}. \tag{5.3}$$

Since $\mathbf{x}$ is orthogonal to $\mathbf{1}$, we can use (5.1) with (5.3) as follows:

$$n\mathbf{x} = B^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\right)\sqrt{W}B\mathbf{x} = \mu B^T W B\mathbf{x} = \mu L\mathbf{x}.$$

Thus $\mathbf{x}$ is an eigenvector of L with eigenvalue $\lambda = \frac{n}{\mu}$. That is, for any eigenvalue μ of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$, there is an eigenvalue λ of L such that $\mu = \frac{n}{\lambda}$, and the correspondence between eigenvalues is consistent with the theorem statement. □

6. The star embedding Γ_{cf}. We now turn to the case of a graph with a zero Dirichlet boundary condition. Such a graph has a set S_b of boundary vertices; the values of these vertices are constrained to be zero. We wish to solve the problem of determining the set of valuations on the nonboundary (interior) vertices that have the property that the values are scaled by a constant when the Laplacian operator is applied to the valuation.

This problem can be described in matrix terms as follows. Let G be a connected graph with boundary set S_b. Without loss of generality, we can assume that G has no edges between pairs of boundary vertices; such edges can be deleted without changing the solution of the problem. In this section, we will use n to denote the number of interior vertices of G, and m to denote the number of edges of G. Let B be the edge-vertex incidence graph subject to the conditions that all edges between boundary and interior vertices are directed away from the boundary and that the boundary vertices are numbered last. Define $\hat{B}$ as the restricted version of B with all columns corresponding to vertices in S_b deleted. The Laplacian $\hat{L}$ is defined as $\hat{B}^T\hat{B}$. The problem stated in the previous paragraph is equivalent to finding the eigenvalues and eigenvectors of $\hat{L}$.

If the edges of G have positive weights assigned to them, we represent the weights in a conductance matrix W. In the weighted case, $\hat{L} = \hat{B}^T W\hat{B}$.

Note that once the columns of B have been deleted to form $\hat{B}$, the problem does not depend on the number of boundary vertices. For edges incident to the boundary, it is no longer possible to distinguish to which boundary vertex they connect. Because of this, it is useful to think of the problem in terms of a graph G_0 that consists of the interior vertices of G plus a single boundary vertex v_0. The edges of G_0 include

all edges between interior vertices of G, plus an edge (v_i, v_0) for every edge between an interior vertex v_i and some boundary vertex. (If the graph is weighted, the edge in G and the corresponding edge in G_0 have the same weight.) G_0 has m edges and $n+1$ vertices. Thus the size of its cycle space is $m-n$ and the size of its cut space is n. It is easy to check that $\hat{L}(G_0) = \hat{L}(G)$ and $\hat{B}(G_0) = \hat{B}(G)$, so we do not need to distinguish these matrices. Thus working with G_0 does not change the solution we are looking for.

Observe that G_0 might have multiple edges between some vertex v_i and the boundary vertex. There is nothing in our definitions that makes this a problem. Such edges form cycles between v_i and v_0. (We note in passing that a graph similar in structure to G_0 can be defined to eliminate such multiple edges by combining them into a single edge with a weight equal to the sum of the weights of the combined edges. This leads to the same result, but the change in the size of the graph complicates the analysis.)

In this section we will consider embeddings into G based on star embeddings into G_0. In mapping a star to G_0, G_0's single boundary vertex is mapped to the center of the star, and every other vertex of G_0 is mapped to a leaf. The embedding itself consists of a path (or set of fractional paths with weights summing to 1) from each interior vertex to the boundary vertex.

Any star embedding into G_0 has a corresponding embedding into G. The only difference is that the flow on edges between the interior and the boundary terminates at different boundary nodes if G has vertices with multiple edges to the boundary. The flow on each edge is the same, however. We call such an embedding a *star embedding into* G.

The path (or set of paths) for each interior vertex can be represented in a matrix Γ as described in section 4. Since the flow on each edge is the same, Γ is the same for both G and G_0 for any embedding. Note that each row of Γ corresponds to a vertex in the interior; we require that the ordering of the rows in Γ corresponds to the ordering of the vertices in $\hat{B}$.

For the rest of the section, we will assume that G is a connected graph with zero boundary S_b, n interior vertices, and m edges. We also assume that $\hat{B}$ and $\hat{L}$ have been defined consistent with the restrictions specified above and that any embedding matrix Γ has its row order consistent with the vertex order.

The following lemma is analogous to Lemma 4.1.

LEMMA 6.1. *For any star embedding into* G, $\Gamma\hat{B} = I_n$.

Proof. We work with the graph G_0, since that is sufficient to prove the result, and since G_0 maps the center of the star to a single vertex. By Lemma 4.1, $\Gamma B = B_*$ is the edge-vertex incidence matrix for the star. The conditions on ordering ensure that edge i of the star is an edge from the boundary to vertex i. Thus, $[B_*]_{ii} = 1$. The only other entries in row i are in columns associated with boundary vertices. Dropping the boundary columns to produce $\hat{B}$ results in the dropping of the corresponding columns from the product. □

The following corollary is useful in the proofs below.

COROLLARY 6.2. *For any star embedding into* G, $\hat{B}^T\Gamma^T\Gamma\hat{B} = I_n$.

Proof. The result is obvious since

$$\hat{B}^T\Gamma^T = \left(\Gamma\hat{B}\right)^T = I_n^T. \qquad \square$$

The current flow embedding Γ_{cf} of the star is defined as per section 4. In electrical terms, the boundary vertices serve as grounds. Therefore, these vertices are combined

into a single ground node in the electrical circuit. Each current flow is computed with respect to a unit current injected at an interior vertex and extracted at the ground node. Note that the current flows for star embeddings into G_0 and G are the same.

The following theorem gives the connection between Γ_{cf} and the eigenvalues and eigenvectors of $\hat{L}$.

THEOREM 6.3. *Let $M_{cf} = \Gamma_{cf}^T \Gamma_{cf}$. For any $\lambda > 0$, λ is an eigenvalue of $\hat{L}$ if and only if $\frac{1}{\lambda}$ is an eigenvalue of $M_{cf}W^{-1}$. Further, if λ has eigenvector $\mathbf{u}$ for $\hat{L}$, then $W\hat{B}\mathbf{u}$ is an eigenvector of $M_{cf}W^{-1}$ with eigenvalue $\frac{1}{\lambda}$.*

Proof. Because $\hat{L} = (\sqrt{W}B)^T(\sqrt{W}B)$ and W is diagonal with positive entries on the diagonal, all nonzero eigenvalues of $\hat{L}$ are positive, so the theorem applies to all nonzero eigenvalues of $\hat{L}$.

The argument is essentially the same as that for Theorem 5.2. As for that theorem, the theorem statement emphasizes the connection with the electrical interpretation; we will actually prove the result for the similar matrix

$$\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}} = \sqrt{W^{-1}}\left(M_{cf}W^{-1}\right)\sqrt{W}.$$

This similarity transform corresponds to a change of variable. We will actually prove that for any eigenvector $\mathbf{u}$ of $\hat{L}$ with eigenvalue $\lambda > 0$, $\sqrt{W}\hat{B}\mathbf{u}$ is an eigenvector of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ with eigenvalue $\frac{1}{\lambda}$. We will also show that $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ has no other nonzero eigenvalues. It is easy to verify that these results imply the theorem.

By Corollary 6.2,

$$\hat{B}^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\right)\sqrt{W}\hat{B} = \hat{B}^T M_{cf}\hat{B} = \hat{B}^T\Gamma_{cf}^T\Gamma_{cf}\hat{B} = I_n,$$

and hence that for any vector $\mathbf{x}$ in the vertex space,

$$\hat{B}^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\right)\sqrt{W}\hat{B}\mathbf{x} = \mathbf{x}. \tag{6.1}$$

Since I_n has rank n, this result implies that the ranks of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$, $\Gamma_{cf}\sqrt{W^{-1}}$, and $\sqrt{W^{-1}}\Gamma_{cf}^T$ are all at least n. Since n is one of the dimensions of Γ_{cf}, it is exactly the rank of $\Gamma_{cf}\sqrt{W^{-1}}$ and $\sqrt{W^{-1}}\Gamma_{cf}^T$. $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ is the product of these matrices and thus has rank less than or equal to n; this implies that it too has rank exactly n.

A similar argument shows that $\sqrt{W}\hat{B}$ has rank n.

We can apply Lemma 4.3 and Corollary 4.4 to G_0: For any vector $\mathbf{x}$ in the cycle space of G_0, $W^{-1}\mathbf{x}$ is in the null space of Γ_{cf}. Further, for any vector $\mathbf{x}$ in the cycle space of G_0, $\sqrt{W^{-1}}\mathbf{x}$ is in the null space of $\Gamma_{cf}\sqrt{W^{-1}}$.

Because $\sqrt{W^{-1}}$ is an invertible matrix, the dimension of the null space is thus at least $m - n$. Combining this with the result from the previous paragraph shows that this is exactly the null space. Through the rest of this proof, we will use the notation C_0 to denote the the null space of $\Gamma_{cf}\sqrt{W^{-1}}$ and of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$.

Using Lemma 3.2 and the fact that W is invertible, it is easy to show that C_0 is the null space of $B^T\sqrt{W}$, and is orthogonal to the column space of $\sqrt{W}B$. Since $\hat{B}$ is derived from B by dropping a column (with respect to G_0), and since $\sqrt{W}\hat{B}$ is an $m \times n$ matrix with rank n, we immediately have that C_0 is the null space of $\hat{B}^T\sqrt{W}$ and is orthogonal to the column space of $\sqrt{W}\hat{B}$.

Because $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ is a symmetric matrix with C_0 as its null space, $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\mathbf{x}$ is orthogonal to C_0 for any vector $\mathbf{x}$ in the edge space.

Let $\mathbf{u}$ be an eigenvector of $\hat{L}$ with eigenvalue $\lambda > 0$. Then

$$\hat{B}^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\right)\sqrt{W}\hat{B}\mathbf{u} = \hat{B}^T M_{cf}\hat{B}\mathbf{u} = I_n\mathbf{u} = \mathbf{u} = \frac{1}{\lambda}\hat{L}\mathbf{u} = \hat{B}^T W\hat{B}\frac{\mathbf{u}}{\lambda}.$$

Thus

$$\hat{B}^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}(\sqrt{W}\hat{B}\mathbf{u}) - \frac{1}{\lambda}(\sqrt{W}\hat{B}\mathbf{u})\right) = 0,$$

which implies that $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}(\sqrt{W}\hat{B}\mathbf{u}) - \frac{1}{\lambda}(\sqrt{W}\hat{B}\mathbf{u})$ is in C_0, the null space of $\hat{B}^T\sqrt{W}$. But we showed above that the first product is orthogonal to C_0. Further, the second product is in the column space of $\sqrt{W}\hat{B}$, which we noted above is orthogonal to C_0. Thus the two products must be equal, and $\sqrt{W}\hat{B}\mathbf{u}$ is an eigenvector of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ with eigenvalue $\frac{1}{\lambda}$.

We still need to show that $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ does not have any other nonzero eigenvalues. Define C_1 to be the subspace that results from transforming the cut space of G_0 by multiplying it by $\sqrt{W}$. By Lemma 3.2, the columns of $\hat{B}$ are a basis for the cut space of G_0, so the columns of $\sqrt{W}\hat{B}$ form a basis for C_1. We noted above that $\sqrt{W}\hat{B}$ is orthogonal to C_0. Since $\sqrt{W}$ is invertible, C_1 has dimension n. Because (the symmetric matrix) $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ has null space C_0, any eigenvector of a nonzero eigenvalue must come from C_1.

Let $\mathbf{u} \in C_1$ be any eigenvector of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ with eigenvalue $\mu > 0$. By the definition of C_1, there is a vector $\mathbf{z}$ in the cut space of G_0 such that $\mathbf{z} = \sqrt{W^{-1}}\mathbf{u}$. Since the columns of $\hat{B}$ are a basis for the cut space of G_0, there is a vector $\mathbf{x}$ in the vertex space restricted to the interior vertices such that $\hat{B}\mathbf{x} = \mathbf{z}$. Thus we can write $\mathbf{u} = \sqrt{W}\hat{B}\mathbf{x}$ for the appropriate such $\mathbf{x}$.

Then

$$\mu\mathbf{u} = \sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\mathbf{u} = \sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}(\sqrt{W}\hat{B}\mathbf{x}) = \mu\sqrt{W}\hat{B}\mathbf{x}. \tag{6.2}$$

We can use (6.1) with (6.2) as follows:

$$\mathbf{x} = \hat{B}^T\sqrt{W}\left(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}\right)\sqrt{W}\hat{B}\mathbf{x} = \mu\hat{B}^T W\hat{B}\mathbf{x} = \mu\hat{L}\mathbf{x}.$$

Thus $\mathbf{x}$ is an eigenvector of $\hat{L}$ with eigenvalue $\lambda = \frac{1}{\mu}$. That is, for any eigenvalue μ of $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$, there is an eigenvalue λ of $\hat{L}$ such that $\mu = \frac{1}{\lambda}$, and the correspondence between eigenvalues is consistent with the theorem statement. □

The details of the proof easily provide the following result, which is not new and is easy to show using other techniques.

COROLLARY 6.4. *$\hat{L}$ is positive definite.*

Proof. We can form the matrix $\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}$ for $\hat{L}$ as defined above. As shown in the preceding proof, it has rank n. Because $M_{cf} = \Gamma_{cf}^T\Gamma_{cf}$ and W is diagonal with only positive diagonal entries, this matrix is positive semidefinite. It therefore has n positive eigenvalues, the reciprocal of each of which is an eigenvalue of $\hat{L}$ by the preceding theorem. Thus $\hat{L}$ has n positive eigenvalues; it is also an $n \times n$ symmetric matrix, which proves the corollary. □

6.1. Current flow embeddings and inverses. We can use Theorem 6.3 to prove the following theorem on the inverse of $\hat{L}$.

THEOREM 6.5. *$\Gamma_{cf}W^{-1}\Gamma_{cf}^T$ is the inverse of $\hat{L}$.*

Proof. $\hat{L}$ is an $n \times n$ real symmetric matrix. Therefore we can write any n-vector $\mathbf{x}$ as a weighted sum of eigenvectors of $\hat{L}$. Showing that $\Gamma_{cf} W^{-1} \Gamma_{cf}^T \hat{L}\mathbf{u} = \mathbf{u}$ for every eigenvector $\mathbf{u}$ of $\hat{L}$ is sufficient to prove the theorem.

$\Gamma_{cf} W^{-1} \Gamma_{cf}^T$ can be rewritten as $(\Gamma_{cf}\sqrt{W^{-1}})(\sqrt{W^{-1}}\Gamma_{cf}^T)$. Applying Lemma 3.1 shows that if $\mathbf{v}$ is an eigenvector of $(\sqrt{W^{-1}}\Gamma_{cf}^T)(\Gamma_{cf}\sqrt{W^{-1}}) = \sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ with eigenvalue $\mu > 0$, then $\Gamma_{cf}\sqrt{W^{-1}}\mathbf{v}$ is an eigenvector of $\Gamma_{cf} W^{-1} \Gamma_{cf}^T$ with eigenvalue μ.

Let $\mathbf{u}$ be any eigenvector of $\hat{L}$ with eigenvalue λ. By Corollary 6.4, $\lambda > 0$. By Theorem 6.3, $\sqrt{W}\hat{B}\mathbf{u}$ is an eigenvector of $\sqrt{W^{-1}} M_{cf} \sqrt{W^{-1}}$ with eigenvalue $\frac{1}{\lambda}$. By the result from the preceding paragraph, $\Gamma_{cf}\sqrt{W^{-1}}\sqrt{W}\hat{B}\mathbf{u}$ is an eigenvector of $\Gamma_{cf} W^{-1} \Gamma_{cf}^T$ with eigenvalue $\frac{1}{\lambda}$. Applying Lemma 6.1 gives the following:

$$\Gamma_{cf}\sqrt{W^{-1}}\sqrt{W}\hat{B}\mathbf{u} = \Gamma_{cf}\hat{B}\mathbf{u} = \mathbf{u}.$$

Hence $\mathbf{u}$ is an eigenvector of $\Gamma_{cf} W^{-1} \Gamma_{cf}^T$ with eigenvalue $\frac{1}{\lambda}$.

Combining these results gives the following:

$$\Gamma_{cf} W^{-1} \Gamma_{cf}^T \hat{L}\mathbf{u} = \lambda \Gamma_{cf} W^{-1} \Gamma_{cf}^T \mathbf{u} = \lambda \frac{\mathbf{u}}{\lambda} = \mathbf{u}.$$

This holds for every eigenvalue of $\hat{L}$, which proves the theorem. □

7. Current flow embeddings and lower bound techniques.

7.1. Connections with previous embedding techniques. We now discuss our results in light of previous embedding techniques used to bound Laplacian eigenvalues. For the sake of simplicity, we will focus on unweighted graphs in this section. Details for weighted graphs can be found in the references cited.

We will discuss two general methods. The first, which we call the edge length method, was presented by Kahale in [22]. It works as follows:

- Specify a clique embedding for the graph G and assign each edge in G 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 $\rho_{\max}$ be the maximum such value taken over all the edges. Then $\frac{n}{\rho_{\max}}$ is a lower bound for λ_2.

The original statement is in terms of an upper bound on the second largest eigenvalue of a reversible Markov chain, but, as Kahale notes, can be reformulated so as to apply to Laplacians.

The second is the path resistance method, presented by Guattery, Leighton, and Miller in [15]. It works as follows:

- Construct a clique embedding into G.
- For each edge e_{ij} in G, compute its congestion c_{ij}. In the unweighted case, the congestion of an edge is the number of paths that are incident to it.
- For each path P and each edge e_{ij} on P, allocate a resistor of size c_{ij} to P.
- 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. Then $\frac{n}{r}$ is a lower bound on λ_2.

These methods are duals of each other in the sense that, given an embedding, the best possible bounds produced by the two methods are the same [15]. This is most easily understood in terms of a matrix representation introduced by Kahale [22], who showed that the edge length method could be viewed as an eigenvalue problem. The

representation uses an embedding matrix similar to the one used throughout this paper, except that no edge directions are used, so the embedding matrix is nonnegative. Thus, if Γ is an embedding matrix as defined above, then $|\Gamma|$ is Kahale's form of the embedding matrix for the same embedding. The statement of the bound problem in terms of an eigenvalue problem is as follows: Let Γ be the matrix for a clique embedding into connected graph G. Then

$$\lambda_2(G) \geq \frac{n}{\lambda_{\max}\left(|\Gamma|^T|\Gamma|\right)}. \tag{7.1}$$

The assignment of edge lengths in this method corresponds to multiplying a vector of lengths by the matrix $|\Gamma|^T|\Gamma|$. The eigenvector for $\lambda_{\max}\left(|\Gamma|^T|\Gamma|\right)$ is the best length assignment.

In [15], Guattery, Leighton, and Miller show that the path resistance method can be viewed in terms of the following problem: For Γ as in the paragraph above,

$$\lambda_2(G) \geq \frac{n}{\lambda_{\max}\left(|\Gamma||\Gamma|^T\right)}. \tag{7.2}$$

Note that, by Lemma 3.1, this is the same bound as in (7.1). Note also that these expressions can be applied to weighted graphs (generalized Laplacians) by including the conductance matrix: The matrices in inequalities (7.1) and (7.2) become $|\Gamma|^T|\Gamma|W^{-1}$ and $|\Gamma|W^{-1}|\Gamma|^T$, respectively.

Reference [15] also extends the technique to the zero Dirichlet boundary case through star embeddings: For connected graph G with zero Dirichlet boundary S, let Γ be the matrix for a star embedding. Then

$$\lambda_1(G_0) \geq \frac{1}{\lambda_{\max}\left(|\Gamma||\Gamma|^T\right)}.$$

These results present the problem of bounding an eigenvalue of one matrix in terms of finding an eigenvalue of another; in practice one does not do this, but instead finds a good approximation to the largest eigenvector of $|\Gamma|^T|\Gamma|$. In many cases where the goal is to bound λ_2 asymptotically for a family of graphs with regular structure, a reasonable approximate vector is relatively easy to find. In other cases, the following technique from [21] can be used: Note that the matrices $|\Gamma|^T|\Gamma|$ and $|\Gamma||\Gamma|^T$ are nonnegative. In the clique embedding case, the matrix is irreducible as a result of our assumption that the graph is connected. In the star embedding case, we can consider the connected components that result when the boundary is deleted; the resulting pieces are irreducible. For irreducible nonnegative matrices, the eigenvalue with the largest magnitude is positive; one can start with a positive vector and repeatedly multiply it by the matrix $|\Gamma|^T|\Gamma|$ to get improved estimates.

In the following sections, we present lemmas that connect results from this paper to previous techniques described above. This allows us to show that the earlier techniques can produce exact bounds for certain classes of graphs. Finally, we give an example to demonstrate some of these results. In particular, it shows that including edge directions can give improved bounds when the repeated multiplication method is used.

In the next two sections, we state results for weighted graphs in the interest of generality.

7.2. Technical lemmas. Let G be a connected graph with conductance matrix W and edge-vertex incidence matrix B and Laplacian $L = B^T W B$. In the zero Dirichlet boundary case, we assume that the subgraph induced by the interior vertices is connected, and we use the Laplacian $\hat{L}$ as defined in section 6. As in the previous sections, Γ is an embedding matrix and $M = \Gamma^T \Gamma$. We will use Γ_{cf} and M_{cf} when the embedding is the current flow embedding. In all cases, we assume that all embeddings and B for a particular graph use the same indexing. Whether the embedding is a clique or a star will be clear from context.

The following two lemmas and their corollaries show that we can drop the absolute values from the previous statement of the lower bound problem and state it in terms of $\Gamma^T \Gamma W^{-1}$. Further, for the embedding Γ_{cf}, the result is exact.

LEMMA 7.1. *For any clique embedding* Γ, $\lambda_{\max}(MW^{-1}) \geq \lambda_{\max}(M_{cf}W^{-1})$, *and* $n/\lambda_{\max}(MW^{-1}) \leq \lambda_2(G)$.

Proof. As in Theorem 5.2, we will use the symmetric matrix $\sqrt{W^{-1}}M\sqrt{W^{-1}}$, which is similar to MW^{-1}, to prove the result.

The Courant–Fischer minimax theorem (see, e.g., [12] or [29]) implies the following:

$$\lambda_{\max}(\sqrt{W^{-1}}M\sqrt{W^{-1}}) \geq \max_{\mathbf{x}\neq 0} \frac{\mathbf{x}^T(\sqrt{W^{-1}}M\sqrt{W^{-1}})\mathbf{x}}{\mathbf{x}^T\mathbf{x}}.$$

This inequality holds for $\mathbf{x} = \sqrt{W}B\mathbf{u}_2$, where $\mathbf{u}_2$ is the eigenvector corresponding to $\lambda_2(G)$, and B is the edge-vertex incidence matrix for G. Then we have

$$\begin{aligned}\lambda_{\max}\left(\sqrt{W^{-1}}M\sqrt{W^{-1}}\right) &\geq \frac{\left(\sqrt{W}B\mathbf{u}_2\right)^T\left(\sqrt{W^{-1}}M\sqrt{W^{-1}}\right)\left(\sqrt{W}B\mathbf{u}_2\right)}{\left(\sqrt{W}B\mathbf{u}_2\right)^T\left(\sqrt{W}B\mathbf{u}_2\right)} \\ &= \frac{\mathbf{u}_2^T(B^TMB)\mathbf{u}_2}{\mathbf{u}_2^TL\mathbf{u}_2} = \frac{n\,\mathbf{u}_2^T\mathbf{u}_2}{\lambda_2\,\mathbf{u}_2^T\mathbf{u}_2} = \frac{n}{\lambda_2},\end{aligned}$$

where the next-to-last equality follows from an application of Lemma 4.1 to the numerator as it is applied in the proof of Theorem 5.2. This proves the second claim in the theorem statement. Theorem 5.2 implies

$$\lambda_{\max}(\sqrt{W^{-1}}M_{cf}\sqrt{W^{-1}}) = \frac{n}{\lambda_2},$$

which, when combined with the previous result, proves the first claim. □

COROLLARY 7.2. *Let* L *be the generalized Laplacian of connected graph* G *with conductance matrix* W. *For any clique embedding* Γ, $\lambda_2(L) \geq n/\lambda_{\max}(\Gamma^T\Gamma W^{-1})$. *Further, if the current flow embedding* Γ_{cf} *is used, then equality holds in the expression above.*

Proof. This follows directly from Lemma 7.1, with M rewritten as $\Gamma^T\Gamma$. □

Similar results hold in the zero Dirichlet boundary case below.

LEMMA 7.3. *In the zero Dirichlet boundary case, for any star embedding* Γ, $\lambda_{\max}(MW^{-1}) \geq \lambda_{\max}(M_{cf}W^{-1})$ *and* $1/\lambda_{\max}(MW^{-1}) \leq \lambda_2(G_0)$.

Proof. The proof is essentially the same as for Lemma 7.1, though it uses Theorem 6.3 rather than Theorem 5.2, and the eigenvector of the smallest eigenvalue of $\hat{L}$ rather than $\mathbf{u}_2$. Details are left to the reader. □

COROLLARY 7.4. *Let* L *be the generalized Laplacian of connected graph* G *with conductance matrix* W *and zero boundary* S. *For any star embedding* Γ *into* G_0,

$\lambda_1(\hat{L}) \geq 1/\lambda_{\max}(\Gamma^T \Gamma W^{-1})$. *Further, if the current flow embedding* Γ_{cf} *is used, equality holds in the expression above.*

Proof. This follows directly from Lemma 7.3, with M rewritten as $\Gamma^T\Gamma$. □

The following technical lemma will simplify the proof of Lemma 7.6.

LEMMA 7.5. *Let* X *be an* $m \times n$ *matrix. Then* $\lambda_{\max}\left(\left|X^T X\right|\right) \leq \lambda_{\max}\left(|X|^T |X|\right)$.

Proof. It is well known that, for nonnegative matrices $A \geq B$, $\lambda_{\max}(A) \geq \lambda_{\max}(B)$ (see, e.g., [23, p. 38]). Thus the lemma holds if we show that, for all $1 \leq i \leq n$ and $1 \leq j \leq n$, $\left[\left|X^T X\right|\right]_{ij} \leq \left[|X|^T |X|\right]_{ij}$. Let X_i represent column i of X. For any i and j, it is clear that $\left|X_i^T X_j\right| \leq |X_i|^T |X_j|$. Since the terms in the inequality are the values of the desired matrix entries, the lemma holds. □

We now show that we can use $|\Gamma^T\Gamma|W^{-1}$ to compute lower bounds that are at least as good as those computed using $|\Gamma|^T|\Gamma|W^{-1}$.

LEMMA 7.6. *For any embedding* Γ,

$$\frac{n}{\lambda_{\max}(|\Gamma|^T|\Gamma|W^{-1})} \leq \frac{n}{\lambda_{\max}(|\Gamma^T\Gamma|W^{-1})} \leq \lambda_2.$$

Proof. By Lemma 7.5, $\lambda_{\max}(|\Gamma|^T|\Gamma|W^{-1}) \geq \lambda_{\max}(|\Gamma^T\Gamma|W^{-1})$, which proves the first inequality.

To show the second inequality, note that for real symmetric matrix A, $\lambda_{\max}(A) \leq \lambda_{\max}(|A|)$. (To see this, let $\mathbf{u}$ be a unit-length eigenvector for $\lambda_{\max}(A)$, and compare the products $\mathbf{u}^T A \mathbf{u}$ and $|\mathbf{u}|^T|A||\mathbf{u}|$.) Then Lemmas 7.1 and 7.3 imply the desired result. □

Since $|\Gamma^T\Gamma|W^{-1}$ is a nonnegative matrix, it can be used with the repeated multiplication algorithm.

7.3. Classes for which previous techniques give tight bounds. Note that for trees there is a unique clique embedding. Note also that for trees with a single zero-value boundary point, there is a unique star embedding. In certain cases, it is possible to set the edge orientations in the edge-vertex incidence matrix so that the entries in the embedding matrix Γ are all positive. In that case, $\Gamma^T\Gamma$, $|\Gamma^T\Gamma|$, and $|\Gamma|^T|\Gamma|$ are all the same matrix, and the desired eigenvalue can be computed exactly in terms of $\lambda_{\max}(|\Gamma|^T|\Gamma|)$. Thus the results from the previous section imply that the best lower bounds derived using techniques in [21] and [15] are exact. We will show this for two cases: one for the clique embedding and one for the star.

The first case we cover is the Laplacian form of a result shown in the birth-death chain example from [21]: If L is the generalized Laplacian of any weighted path graph with conductance matrix W, then there is an embedding such that $n/\lambda_{\max}(|\Gamma|^T|\Gamma|W^{-1}) = \lambda_2(L)$. To see this, note that we can index the vertices in order along the path and that we can direct every edge from the lower-index vertex to the higher-index one. Likewise, direct every path in the clique embedding from the lower-index vertex to the higher-index one. There is only one possible embedding, and the choices of path and edge directions ensures that all entries in the embedding matrix Γ are positive. Thus

$$\frac{n}{\lambda_2} = \lambda_{\max}(\Gamma_{cf}^T \Gamma_{cf} W^{-1}) = \lambda_{\max}(|\Gamma|^T|\Gamma|W^{-1}).$$

In the star case, a similar result holds for trees with a single zero-value boundary point. In this case, every edge lies on a path from some vertex to the boundary; the vertex at one end of the edge is closer to the boundary than the vertex at the other. Direct each edge toward its endpoint that is closer to the boundary vertex. Likewise,

direct each path in the embedding toward the center of the star. The embedding is unique, and

$$\frac{1}{\lambda_2} = \lambda_{\max}(\Gamma_{cf}^T \Gamma_{cf} W^{-1}) = \lambda_{\max}(|\Gamma|^T |\Gamma| W^{-1}).$$

7.4. A simple example. We now use a cycle on four vertices to demonstrate the case where previous lower bound techniques are not exact. We number the vertices in clockwise order and also direct the edges in clockwise order. The edge-vertex incidence matrix B and the Laplacian L are shown below:

$$B = \begin{bmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{bmatrix}, \qquad L = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ -1 & 0 & -1 & 2 \end{bmatrix}.$$

The eigenvalues of L are $\lambda = 0$, $\lambda = 2$ with multiplicity 2, and $\lambda = 4$.

The previous embedding techniques cannot produce an exact bound for this problem; as shown in the corollary to the following lemma, the best lower bound they can produce is 4/3.

We use d_{ij} to denote the distance (in terms of number of edges) between vertices v_i and v_j.

LEMMA 7.7. *For L the Laplacian of an unweighted graph, the best lower bound on $\lambda_2(L)$ produced by the edge length or path resistance methods is less than or equal to $mn/\sum_{i,j} d_{ij}^2$.*

Proof. The lower bound for any embedding Γ is $n/\lambda_{\max}(|\Gamma|^T|\Gamma|)$. For any $\mathbf{x}$,

$$\frac{\mathbf{x}^T |\Gamma|^T |\Gamma| \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq \lambda_{\max}(|\Gamma|^T |\Gamma|). \tag{7.3}$$

Set $\mathbf{x} = \mathbf{1}$, the vector of length m of all ones. The entries of $|\Gamma|\mathbf{1}$ are the row sums of $|\Gamma|$. If a row represents a single path, there is an entry of 1 for each edge in the path. The sum in this case is the length of the path, which is at least the distance between the endpoints. If a row represents a set of fractional paths, the row sum is the sum over each fractional path of the product of the fractional weight times the path length. Since the fractional weights must add to 1, the row sum is at least the distance between the endpoints of the path. We get one entry in $|\Gamma|\mathbf{x}$ for each edge in K_n.

Note that the numerator of the left-hand side of (7.3) is just the dot product of $|\Gamma|\mathbf{x}$ with itself, which by the previous paragraph is at least $\sum_{i,j} d_{ij}^2$. (Since $d_{ii} = 0$ for all i, we can add these terms into the sum without penalty.)

The denominator $\mathbf{x}^T\mathbf{x}$ is m, the length of the vector. Thus

$$\frac{\sum_{i,j} d_{ij}^2}{m} \leq \lambda_{\max}(|\Gamma|^T |\Gamma|).$$

Applying (7.1) and the duality of the edge length and path resistance methods gives the result in the lemma statement. ∎

COROLLARY 7.8. *For the cycle on four vertices, the best lower bounds from the path resistance or edge length methods are no bigger than 4/3.*

Proof. Apply Lemma 7.7. The cycle has four pairs of vertices at distance 1 and two pairs at distance 2, so the sum of the squares of the distances is 12. Dividing $mn = 16$ by this number gives the result. ∎

Thus, these methods do not produce exact results for our example cycle. The following (undirected) embedding matrix will produce the lower bound of 4/3:

$$|\Gamma| = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \qquad |\Gamma|^T|\Gamma| = \frac{1}{2} \cdot \begin{bmatrix} 3 & 1 & 1 & 1 \\ 1 & 3 & 1 & 1 \\ 1 & 1 & 3 & 1 \\ 1 & 1 & 1 & 3 \end{bmatrix}.$$

By adding edge directions, however, we can produce an exact lower bound. Consider the following embedding matrix Γ and the absolute value of the product $\Gamma^T\Gamma$:

$$\Gamma = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \end{bmatrix}, \qquad |\Gamma^T\Gamma| = \frac{1}{2} \cdot \begin{bmatrix} 3 & 0 & 1 & 0 \\ 0 & 3 & 0 & 1 \\ 1 & 0 & 3 & 0 \\ 0 & 1 & 0 & 3 \end{bmatrix}.$$

The eigenvalues of the $|\Gamma^T\Gamma|$ are $\lambda = 1$ and $\lambda = 2$, each with multiplicity 2. By the results in section 7.2, we get a lower bound of 2 on λ_2 of the cycle, which is exactly equal to the eigenvalue.

As a final example, we show that forming a nonnegative matrix based on the current flow embedding does not necessarily produce the best bounds. This is true in spite of the exact relationship between the current flow embedding eigenvalues and the Laplacian eigenvalues. We show below the current flow embedding matrix and two nonnegative matrices produced from it:

$$\Gamma = \frac{1}{4} \cdot \begin{bmatrix} 3 & -1 & -1 & -1 \\ -1 & 3 & -1 & -1 \\ -1 & -1 & 3 & -1 \\ -1 & -1 & -1 & 3 \\ 2 & 2 & -2 & -2 \\ -2 & 2 & 2 & -2 \end{bmatrix}, \qquad |\Gamma^T\Gamma| = \frac{1}{4} \cdot \begin{bmatrix} 5 & 1 & 3 & 1 \\ 1 & 5 & 1 & 3 \\ 3 & 1 & 5 & 1 \\ 1 & 3 & 1 & 5 \end{bmatrix},$$

and

$$|\Gamma|^T|\Gamma| = \frac{1}{4} \cdot \begin{bmatrix} 5 & 4 & 4 & 4 \\ 4 & 5 & 4 & 4 \\ 4 & 4 & 5 & 4 \\ 4 & 4 & 4 & 5 \end{bmatrix}.$$

The eigenvalues of $|\Gamma^T\Gamma|$ are $\lambda = 1/2$ with multiplicity 2, $\lambda = 3/2$, and $\lambda = 5/2$, yielding a lower bound of 8/5 on λ_2; the previous example gives a better lower bound. The eigenvalues of $|\Gamma^T\Gamma|$ are $\lambda = 1/4$ with multiplicity 3 and $\lambda = 4\,1/4$, yielding a lower bound of 16/17, which is less than 1 and lower than the bounds for our previous examples.

Acknowledgments. The authors wish to thank John Lafferty of Carnegie Mellon University, who provided many valuable comments and suggestions on a previous version of this paper.

REFERENCES

[1] N. Alon, *Eigenvalues and expanders*, Combinatorica, 6 (1986), pp. 83–96.

[2] N. Alon and V. D. Milman, λ_1*, isoperimetric inequalities for graphs, and superconcentrators*, J. Combin. Theory Ser. B, 38 (1985), pp. 73–88.

[3] W. N. Anderson, Jr., and T. D. Morley, *Eigenvalues of the Laplacian of a graph*, Linear and Multilinear Algebra, 18 (1985), pp. 141–145.

[4] O. Axelsson, *Bounds of eigenvalues of preconditioned matrices*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 847–862.

[5] C. Berge, *Graphs and Hypergraphs*, North–Holland, Amsterdam, 1973.

[6] N. Biggs, *Algebraic Graph Theory*, Cambridge University Press, Cambridge, UK, 1974.

[7] B. Bollobás, *Graph Theory: An Introductory Course*, Springer-Verlag, New York, 1979.

[8] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari, *The electrical resistance of a graph captures its commute and cover times*, in Proceedings 21st Annual Symposium on Theory of Computation, Seattle, 1989, pp. 574–586.

[9] P. Diaconis and D. Stroock, *Geometric bounds for eigenvalues of Markov chains*, Ann. Appl. Probab., 1 (1991), pp. 36–61.

[10] M. Fiedler, *Algebraic connectivity of graphs*, Czechoslovak Math. J., 23 (1973), pp. 298–305.

[11] M. Fiedler, *A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory*, Czechoslovak Math. J., 25 (1975), pp. 619–633.

[12] G. H. Golub and C. F. V. Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 1989.

[13] K. Gremban, G. Miller, and M. Zagha, *Performance evaluation of a parallel preconditioner*, in Proceedings 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995, pp. 65–69.

[14] K. D. Gremban, *Combinatorial Preconditioners for Sparse, Symmetric, Diagonally Dominant Linear Systems*, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, 1996.

[15] S. Guattery, F. T. Leighton, and G. Miller, *The path resistance method for bounding* λ_2 *of a Laplacian*, in Proceedings 8th ACM-SIAM Symposium on Discrete Algorithms, New Orleans, LA, 1997, pp. 201–210.

[16] S. Guattery and G. Miller, *On the performance of spectral graph partitioning methods*, in Proceedings 6th ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, 1995, pp. 233–242.

[17] S. Guattery and G. L. Miller, *On the quality of spectral separators*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 701–719.

[18] L. A. Hageman and D. M. Young, *Applied Iterative Methods*, Computer Science and Applied Mathematics, Academic Press, New York, 1981.

[19] L. Hagen and A. B. Kahng, *New spectral methods for ratio cut partitioning and clustering*, IEEE Trans. Computer-Aided Design, 11 (1992), pp. 1074–1085.

[20] M. Jerrum and A. Sinclair, *Approximating the permanent*, SIAM J. Comput., 18 (1989), pp. 1149–1178.

[21] N. Kahale, *A Semidefinite Bound for Mixing Rates of Markov Chains*, TR 95-41, DIMACS, Piscataway, NJ, 1995.

[22] N. Kahale, *A semidefinite bound for mixing rates of Markov chains*, in Proceedings 5th Integer Programming and Combinatorial Optimization Conference, Lecture Notes in Comput. Sci. 1084, Springer-Verlag, New York, 1996, pp. 190–203.

[23] H. Minc, *Nonnegative Matrices*, John Wiley, New York, 1988.

[24] B. Mohar, *Isoperimetric numbers of graphs*, J. Combin. Theory Ser. B, 47 (1989), pp. 274–291.

[25] A. Pothen, H. D. Simon, and K.-P. Liou, *Partitioning sparse matrices with eigenvectors of graphs*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 430–452.

[26] A. Sinclair, *Improved bounds for mixing rates of Markov chains and multicommodity flow*, Combin. Probab. Comput., 1 (1992), pp. 351–370.

[27] A. Sinclair and M. Jerrum, *Approximate counting, uniform generation and rapidly mixing Markov chains*, Inform. and Comput., 82 (1989), pp. 93–133.

[28] D. A. Spielman and S. H. Teng, *Spectral partitioning works: Planar graphs and finite element meshes*, Techical report UCB CSD-96-898, University of California at Berkeley, Berkeley, CA, 1996.

[29] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, 1965.