



Spectral partitioning works: Planar graphs and finite element meshes [☆]

Daniel A. Spielman ^{a,1}, Shang-Hua Teng ^{b,*,2}

^a *Department of Computer Science, Yale University, New Haven, CT 06520, United States*

^b *Department of Computer Science, Boston University, Boston, MA 02215, United States*

Received 28 February 2006; accepted 24 July 2006

Available online 2 October 2006

Submitted by L. Hogben

Abstract

Spectral partitioning methods use the Fiedler vector—the eigenvector of the second-smallest eigenvalue of the Laplacian matrix—to find a small separator of a graph. These methods are important components of many scientific numerical algorithms and have been demonstrated by experiment to work extremely well. In this paper, we show that spectral partitioning methods work well on bounded-degree planar graphs and finite element meshes—the classes of graphs to which they are usually applied. While naive spectral bisection does not necessarily work, we prove that spectral partitioning techniques can be used to produce separators whose ratio of vertices removed to edges cut is $O(\sqrt{n})$ for bounded-degree planar graphs and two-dimensional meshes and $O(n^{1/d})$ for well-shaped d -dimensional meshes. The heart of our analysis is an upper bound on the second-smallest eigenvalues of the Laplacian matrices of these graphs: we prove a bound of $O(1/n)$ for bounded-degree planar graphs and $O(1/n^{2/d})$ for well-shaped d -dimensional meshes. © 2006 Elsevier Inc. All rights reserved.

Keywords: Spectral methods; Spectral analysis; Graph partitioning; Eigenvalue problems; Graph embedding

[☆] An extended abstract of this work appeared at [74] and an earlier version of this paper can be found in U.C. Berkeley Technical Report UCB//CSD-96-898.

* Corresponding author.

E-mail addresses: spielman@cs.yale.edu (D.A. Spielman), steng@cs.bu.edu (S.-H. Teng).

¹ The work was done while the author was at U.C. Berkeley.

² This work was done while the author was at the University of Minnesota and Xerox PARC.

1. Introduction

Spectral partitioning has become one of the most successful heuristics for partitioning graphs and matrices. It is used in many scientific numerical applications, such as mapping finite element calculations on parallel machines [70,82], solving sparse linear systems [68], and partitioning for domain decomposition [21,25]. It is also used in VLSI circuit design and simulation [26,46,2]. Substantial experimental work has demonstrated that spectral methods find good partitions of the graphs and matrices that arise in many applications [18,48,49,66,70,82]. However, the quality of the partition that these methods should produce has so far eluded precise analysis. In this paper, we will prove that spectral partitioning methods give good separators for the graphs to which they are usually applied.

The size of the separator produced by spectral methods can be related to the Fiedler value—the second smallest eigenvalue of the Laplacian—of the adjacency structure to which they are applied. By showing that well-shaped meshes in d dimensions have Fiedler value at most $O(1/n^{2/d})$, we show that spectral methods can be used to find bisectors of these graphs of size at most $O(n^{1-1/d})$. While a small Fiedler value does not immediately imply that there is a cut along the Fiedler vector that is a balanced separator, it does mean that there is a cut whose ratio of vertices separated to edges cut is $O(n^{1/d})$. By removing the vertices separated by this cut, computing a Fiedler vector of the new graph, and iterating as necessary, one can find a bisector of $O(n^{1-1/d})$ edges. In particular, we prove that maximum-degree Δ planar graphs have Fiedler value at most $8\Delta/n$, which implies that spectral techniques can be used to find bisectors of size at most $O(\sqrt{n})$ in these graphs. These bounds are the best possible for well-shaped meshes and planar graphs.

1.1. History

The spectral method of graph partitioning was born in the works of Hall [47] who applied quadratic programming to design a placement algorithm, and in the work of Donath and Hoffman [28,29] who first suggested using the eigenvectors of adjacency matrices of graphs to find partitions and point placement. Fiedler [33–35] associated the second-smallest eigenvalue of the Laplacian of a graph with its connectivity and suggested partitioning by splitting vertices according to their value in the corresponding eigenvector. Thus, we call this eigenvalue the *Fiedler value* and a corresponding vector a *Fiedler vector*.

A few years later, Barnes and Hoffman [10,14] used linear programming in combination with an examination of the eigenvectors of the adjacency matrix of a graph. In a similar vein, Boppana [17] analyzed eigenvector techniques in conjunction with convex programming. However, the use of linear and convex programming made these techniques impractical for most applications.

By recognizing a relation between the Fiedler value and the Cheeger constant [20] of continuous manifolds, Alon [3] and Sinclair and Jerrum [72] demonstrated that if the Fiedler value of a graph is small, then directly partitioning the graph according to the values of vertices in the eigenvector will produce a cut with a good ratio of cut edges to separated vertices (see also [4, 36,30,55,57]). Around the same time, improvements in algorithms for approximately computing eigenvectors, such as the Lanczos algorithm, made the computation of eigenvectors practical [67, 70]. In the next few years, a wealth of experimental work demonstrated that spectral partitioning methods work well on graphs that usually arise in practice [18,48,66,70,82]. Spectral partitioning became a standard tool for mesh partitioning in many areas [49]. Still, researchers were unable to prove that spectral partitioning techniques would work well on the graphs encountered in

practice. This failure is partially explained by results of Guattery and Miller [42] demonstrating that naive applications of spectral partitioning, such as spectral bisection, will fail miserably on some graphs that could conceivably arise in practice. By bounding the Fiedler values of the graphs of interest in scientific applications—bounded-degree planar graphs and well-shaped meshes—we are able to show that spectral partitioning methods will successfully find good partitions of these graphs.

In a related line of research, algorithms were developed along with proofs that they always find small separators in various families of graphs. The seminal work in this area was that of Lipton and Tarjan [54], who constructed a linear-time algorithm that produces a $1/3$ -separator of $\sqrt{8n}$ nodes in any n -node planar graph. Their result improved a theorem of Ungar [81] which demonstrated that every planar graph has a separator of size $O(\sqrt{n} \log n)$. Gilbert, Hutchinson, and Tarjan [39] extended these results to show that every graph of genus at most g has a separator of size $O(\sqrt{gn})$. Another generalization was obtained by Alon, Seymour, and Thomas [7], who showed that graphs that do not have an h -clique minor have separators of $O(h^{3/2} \sqrt{n})$ nodes. Plotkin et al. [65] reduced the dependency on h from $h^{3/2}$ to h . Using geometric techniques, Miller, Teng, Thurston, and Vavasis [58–60,62,63,76] extended the planar separator theorem to graphs embedded in higher dimensions and showed that every well-shaped mesh in \mathbb{R}^d has a $1/(d+2)$ -separator of size $O(n^{1-1/d})$. Using multicommodity flow, Leighton and Rao [53] designed a partitioning method guaranteed to return a cut whose ratio of cut size to vertices separated is within logarithmic factors of optimal. Arora et al. [8] recently improved the approximation ratio to $O(\sqrt{\log n})$. For planar graphs, a partition within a constant factor of the optimal can be found in polynomial time [38]. While spectral methods have been favored in practice, they lacked a proof of effectiveness.

1.2. Outline of paper

In Section 2, we introduce the concept of a graph partition, review some facts from linear algebra that we require, and describe the class of spectral partitioning methods.

In Section 3, we prove the *embedding lemma*, which relates the quality of geometric embeddings of a graph with its Fiedler value. We then show (using the main result of Section 4) that every planar graph has a “nice” embedding as a collection of spherical caps on the surface of a unit sphere in three dimensions. By applying the embedding lemma to this embedding, we prove that the Fiedler value of every bounded-degree planar graph is $O(1/n)$.

In Section 4, we show that, for almost every arrangement of spherical caps on the unit sphere in \mathbb{R}^d , there is a sphere-preserving map that transforms the caps so that the center of the sphere is the centroid of their centers. It is this fact that enables us to find nice embeddings of planar graphs.

In Sections 5 and 6, we extend our spectral planar separator theorem to the class of overlap graphs of k -ply neighborhood systems embedded in any fixed dimension. This extension enables us to show that the spectral method finds cuts of ratio $O(1/n^{1/d})$ for k -nearest neighbor graphs and well-shaped finite element meshes.

In Section 7, we discuss some recent extensions to our work.

2. Introduction to spectral partitioning

In this section, we define the spectral partitioning method and introduce the terminology that we will use throughout the paper.

2.1. Graph partitioning

Throughout this paper, $G = (V, E)$ will be a connected, undirected graph on n vertices.

A *partition* of a graph G is a division of its vertices into two disjoint subsets, A and \bar{A} . Without loss of generality, we can assume that $|A| \leq |\bar{A}|$. Let $E(A, \bar{A})$ be the set of edges with one endpoint in A and the other in \bar{A} . The *cut size* of the partition (A, \bar{A}) is simply $|E(A, \bar{A})|$. The *cut ratio*, or simply the *ratio* of the cut, denoted $\phi(A, \bar{A})$, is equal to the ratio of the size of the cut to the size of A , namely,

$$\phi(A, \bar{A}) = \frac{|E(A, \bar{A})|}{\min(|A|, |\bar{A}|)}.$$

The isoperimetric number of a graph, which measures how good a ratio cut one can hope to find, is defined to be

$$\phi(G) = \min_{|A| \leq n/2} \frac{|E(A, \bar{A})|}{|A|}.$$

Theorem 2.1 states a relation between the isoperimetric number of a graph and its Fiedler value.

A partition is a *bisection* of G if A and \bar{A} differ in size by at most 1. For δ in the range $0 < \delta \leq 1/2$, a partition is called a δ -*separator* if $\min(|A|, |\bar{A}|) \geq \delta n$. We use the word *cut* to refer to a partition separating any number of vertices and reserve the word *separator* for partitions that are δ -separators for some $\delta > 0$. Given an algorithm that can find cuts of ratio ϕ in G and its subgraphs, we can find a bisector of G of size $O(\phi n)$ (see Lemma A.1).

2.2. Laplacians and Fiedler vectors

The adjacency matrix, $A(G)$, of a graph G is the $n \times n$ matrix whose (i, j) th entry is 1 if $(i, j) \in E$ and 0 otherwise. The diagonal entries are defined to be 0. Let D be the $n \times n$ diagonal matrix with entries $D_{i,i} = d_i$, where d_i is the degree of the i th vertex of G . The *Laplacian*, $L(G)$, of the graph G is defined to be $L(G) = D - A$.

Let M be an $n \times n$ matrix. An n -dimensional vector \vec{x} is an *eigenvector* of M if there is a scalar λ such that $M\vec{x} = \lambda\vec{x}$. λ is the *eigenvalue* of M corresponding to the eigenvector \vec{x} . If M is a real symmetric matrix, then all of its n eigenvalues are real. The only matrices we consider in this paper will be the Laplacians of graphs. Notice that the all-ones vector is an eigenvector of any Laplacian matrix and that its associated eigenvalue is 0. Because Laplacian matrices are positive semidefinite, all the other eigenvalues must be non-negative. We will focus on the second smallest eigenvalue, λ_2 , of the Laplacian and an associated eigenvector \vec{u} . Fiedler called this eigenvalue the “algebraic connectivity of a graph”, so we will call it the *Fiedler value* and an associated eigenvector a *Fiedler vector*.

The following properties of Fiedler values and vectors play an important role in this paper:

- The Fiedler value of a graph is greater than zero if and only if the graph is connected.
- A Fiedler vector $\vec{u} = (u_1, \dots, u_n)$ satisfies

$$\sum_{i=1}^n u_i = 0,$$

because all-ones vector is an eigenvector of the Laplacian and the eigenvectors of a symmetric matrix are orthogonal.

- The Fiedler value, λ_2 , of G satisfies

$$\lambda_2 = \min_{\vec{x} \perp (1,1,\dots,1)} \frac{\vec{x}^T L(G) \vec{x}}{\vec{x}^T \vec{x}},$$

with the minimum occurring only when \vec{x} is a Fiedler vector.

- For any vector $\vec{x} \in R^n$, we have

$$\vec{x}^T L(G) \vec{x} = \sum_{(i,j) \in E} (x_i - x_j)^2.$$

Let M be a symmetric $n \times n$ matrix and \vec{x} be an n -dimensional vector. Then, the *Rayleigh quotient* of \vec{x} with respect to M is

$$\frac{\vec{x}^T M \vec{x}}{\vec{x}^T \vec{x}}.$$

For proofs of these statements and many other fascinating facts about the eigenvalues and eigenvectors of graphs consult one of [75,56,19,16,23].

The Fiedler value, λ_2 , of a graph is closely linked to its isoperimetric number. If G is a graph on more than three nodes, then one can show [4,57,71] (also see Theorem 2.1)

$$\frac{\lambda_2}{2} \leq \phi(G) \leq \sqrt{\lambda_2(2d - \lambda_2)}.$$

2.3. Spectral partitioning methods

Let $\vec{u} = (u_1, \dots, u_n)$ be a Fiedler vector of the Laplacian of a graph G . The idea of spectral partitioning is to find a *splitting value* s and partition the vertices of G into the set of i such that $u_i > s$ and the set such that $u_i \leq s$. We call such a partition a *Fiedler cut*. There are several popular choices for the splitting value s :

- **bisection:** s is the median of $\{u_1, \dots, u_n\}$.
- **ratio cut:** s is the value that gives the best ratio cut.
- **sign cut:** s is equal to 0.
- **gap cut:** s is a value in the largest gap in the sorted list of Fiedler vector components.

Other variations have been proposed.

In this paper, we will analyze the spectral method that uses the splitting value that achieves the best ratio cut. We will show that, for bounded-degree planar graphs and well-shaped meshes, it always finds a good ratio cut. In fact, by the following result of Mihail [55], it is not necessary to use a Fiedler vector; an approximation will suffice.

Theorem 2.1 (Mihail). *Let $G = (V, E)$ be a graph on n nodes of maximum degree d , let Q be its Laplacian matrix, and let ϕ be its isoperimetric number. For any vector $\vec{x} \in R^n$ such that $\sum_{i=1}^n x_i = 0$,*

$$\frac{\vec{x}^T Q \vec{x}}{\vec{x}^T \vec{x}} \geq \frac{\phi^2}{2d}.$$

Moreover, there is an s so that the cut $(\{i : x_i \leq s\}, \{i : x_i > s\})$ had ratio at most $\sqrt{2d \frac{\vec{x}^T Q \vec{x}}{\vec{x}^T \vec{x}}}$.

If a vector \vec{x} that is orthogonal to the all-ones vector has a small Rayleigh quotient with respect to the Laplacian of G , then \vec{x} can be used to find a good ratio cut of G .

Guattery and Miller [42] showed that there exist bounded-degree planar graphs on n vertices with constant-size separators for which spectral bisection and spectral sign cuts give separators that cut $n/3$ edges.

We will show that planar graphs have Fiedler cuts of ratio $O(1/\sqrt{n})$. By Lemma A.1, our result implies that a bisector of size $O(\sqrt{n})$ can be found by repeatedly finding Fiedler cuts. One can extend the results of Guattery and Miller to show that this repeated application of Fiedler cuts is necessary, even for some quite natural graphs. In particular, one can show that, for any constant δ in the range $0 < \delta \leq 1/2$, there are natural families of well-shaped two-dimensional meshes that have no Fiedler cut of small ratio that is also a δ -separator.

3. The eigenvalues of planar graphs

In this section, we will prove that the Fiedler value of every bounded-degree planar graph is $O(1/n)$. Our proof establishes and exploits a connection between the Fiedler value and geometric embeddings of graphs. We obtain the eigenvalue bound by demonstrating that every planar graph has a “nice” embedding in Euclidean space.

A bound of $O(1/\sqrt{n})$ can be placed on the Fiedler value of any planar graph by combining the planar separator theorem of Lipton and Tarjan [54] with the fact that $\lambda_2/2 \leq \phi(G)$. Bounds of $O(1/n)$ on the Fiedler values of planar graphs were previously known for graphs such as regular grids [66], quasi-uniform graphs [41], and bounded-degree trees. Bounds on the Fiedler values of regular grids and quasi-uniform graphs essentially follow from the fact that the diameters of these graphs are large (see [22]). Bounds on trees can be obtained from the fact that every bounded-degree tree has a δ -separator of size 1 for some constant δ in the range $0 < \delta < 1/2$ that depends only on the degree. However, in order to estimate the Fiedler value of general bounded-degree planar graphs and well-shaped meshes, we need different techniques.

We denote the standard l_2 norm of a vector \vec{x} in Euclidean space by $\|\vec{x}\| = \sqrt{x^T x}$. We relate the quality of an embedding of a graph in Euclidean space with its Fiedler value by the following lemma:

Lemma 3.1 (Embedding lemma). *Let $G = (V, E)$ be a graph. Then λ_2 , the Fiedler value of G , is given by*

$$\lambda_2 = \min \frac{\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^n \|\vec{v}_i\|^2},$$

where the minimum is taken over vectors $\{\vec{v}_1, \dots, \vec{v}_n\} \subset R^n$ such that

$$\sum_{i=1}^n \vec{v}_i = \vec{0},$$

where $\vec{0}$ denotes the all-zeroes vector.

Remark. While we state this lemma for vectors in R^n , it applies equally well for vectors in R^m for any $m \geq 1$.

Proof. Because the all-ones vector is the eigenvector of $L(G)$ corresponding to the eigenvalue 0, λ_2 can be characterized by

$$\lambda_2 = \min \frac{\sum_{(i,j) \in E} (x_i - x_j)^2}{\sum_{i=1}^n x_i^2},$$

where the minimum is taken over real x_i 's such that $\sum_{i=1}^n x_i = 0$. The minimum is achieved precisely when (x_1, \dots, x_n) is an eigenvector.

The embedding lemma now follows from a component-wise application of this fact. Write \vec{v}_i as $(v_{i,1}, \dots, v_{i,n})$. Then, for all $\{\vec{v}_1, \dots, \vec{v}_n\}$ such that $\sum_{i=1}^n \vec{v}_i = \vec{0}$, we have

$$\begin{aligned} \frac{\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^n \|\vec{v}_i\|^2} &= \frac{\sum_{(i,j) \in E} \sum_{k=1}^n (v_{i,k} - v_{j,k})^2}{\sum_{i=1}^n \sum_{k=1}^n v_{i,k}^2} \\ &= \frac{\sum_{k=1}^n \sum_{(i,j) \in E} (v_{i,k} - v_{j,k})^2}{\sum_{k=1}^n \sum_{i=1}^n v_{i,k}^2}. \end{aligned}$$

But, for each k ,

$$\frac{\sum_{(i,j) \in E} (v_{i,k} - v_{j,k})^2}{\sum_{i=1}^n v_{i,k}^2} \geq \lambda_2,$$

so

$$\frac{\sum_{k=1}^n \sum_{(i,j) \in E} (v_{i,k} - v_{j,k})^2}{\sum_{k=1}^n \sum_{i=1}^n v_{i,k}^2} \geq \lambda_2$$

(this follows from the fact that $\sum_i x_i / \sum_i y_i \geq \min_i x_i / y_i$, for $x_i, y_i > 0$). \square

Our method of finding a good geometric embedding of a planar graph is similar to the way in which Miller et al. [59] directly find good separators of planar graphs.

We first find an embedding of the graph on the plane by using the “kissing disk” embedding of Koebe, Andreev, and Thurston [52,5,6,79]:

Theorem 3.2 (Koebe–Andreev–Thurston). *Let G be a planar graph with vertex set $V = \{1, \dots, n\}$ and edge set E . Then, there exists a set of disks $\{D_1, \dots, D_n\}$ in the plane with disjoint interiors such that D_i touches D_j if and only if $(i, j) \in E$.*

Such an embedding is called a *kissing disk* embedding of G .

The analogue of a disk on the sphere is a *cap*. A *cap* is given by the intersection of a half-space with the sphere, and its boundary is a circle. We define *kissing caps* analogously with kissing disks. Following [59], we use stereographic projection to map the kissing disk embedding of the graph on the plane to a kissing cap embedding on the sphere (see Section 4 for more information on stereographic projection). In Theorem 4.2, we will show that we can find a sphere preserving map that sends the *centroid* (also known as the *center of gravity* or *center of mass*) of the centers of the caps to the center of the sphere. Using this theorem, we can bound the eigenvalues of planar graphs:

Theorem 3.3. *Let G be a planar graph on n nodes of degree at most Δ . Then, the Fiedler value of G is at most*

$$\frac{8\Delta}{n}.$$

Accordingly, G has a Fiedler cut of ratio $O(1/\sqrt{n})$, and one can iterate Fiedler cuts to find a bisector of size $O(\sqrt{n})$.

Proof. By Theorems 3.2 and 4.2, there is a representation of G by kissing caps on the unit sphere so that the centroid of the centers of the caps is the center of the sphere. Let $\vec{v}_1, \dots, \vec{v}_n$ be the centers of these caps. Make the center of the sphere the origin, so that $\sum_{i=1}^n \vec{v}_i = \vec{0}$.

Let r_1, \dots, r_n be the radii of the caps. If cap i kisses cap j , then the edge from \vec{v}_i to \vec{v}_j will have length at most $(r_i + r_j)^2$. As this is at most $2(r_i^2 + r_j^2)$, we can divide the contribution of this edge between the two caps. That is, we write

$$\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2 \leq 2\Delta \sum_{i=1}^n r_i^2.$$

But, because the caps do not overlap,

$$\sum_{i=1}^n \pi r_i^2 \leq 4\pi.$$

Moreover, $\|\vec{v}_i\| = 1$ because the vectors are on the unit sphere.

Applying the embedding lemma, we find that the Fiedler value of G is at most

$$\frac{\sum_{(i,j) \in E} \|\vec{v}_i - \vec{v}_j\|^2}{\sum_{i=1}^n \|\vec{v}_i\|^2} \leq \frac{8\Delta}{n}.$$

Given the bound on the Fiedler value, the ratio achievable by a Fiedler cut follows immediately from Theorem 2.1 and the corresponding bisector size follows Lemma A.1. \square

Remark. One can prove a slightly weaker result by using a result of Miller et al. [59] to find a circle-preserving map that makes the center of the sphere a centerpoint [31] of the images of particular points in the caps. If the center of the sphere is a centerpoint, then the centroid must be far away from at least a constant fraction of the centers of the caps. Thus, the numerator of the Rayleigh quotient will be the same as in Theorem 3.3, and the denominator will be $O(n)$.

Remark. The embedding lemma can be viewed as a semi-definite relaxation of an integer program for minimum balanced cut. The integer program would be

$$\begin{aligned} \min \quad & \sum_{(i,j) \in E} (x_i - x_j)^2 \\ \text{s.t.} \quad & \sum_{i=1}^n x_i = 0, \text{ and} \\ & x_i \in \{\pm 1\}. \end{aligned}$$

4. Sphere-preserving maps

Let B^d be the unit ball in d dimensions: $\{(x_1, \dots, x_d) \mid \sum_{i=1}^n x_i^2 \leq 1\}$. Let S^d denote the sphere defining the surface of B^d . This section is concerned with *sphere-preserving* maps from S^d to S^d . A *sphere-preserving* map from S^d to S^d is a continuous function that sends every sphere (of lower

dimension) contained in S^d to a sphere in S^d and such that every sphere in S^d has a pre-image under the map that is also a sphere. Familiar sphere-preserving maps include rotations and the map that sends each point to its antipode.

We will make use of a slightly larger family of sphere-preserving maps. We obtain this family by first considering sphere-preserving maps between the sphere and the plane. Let H^d be the hyperplane tangent to S^d at $(-1, 0, \dots, 0)$. One can map H^d to S^d by *stereographic projection*: $\Pi : H^d \rightarrow S^d$ by

$$\Pi(z) = \text{the intersection of } S^d \text{ with the line connecting } z \text{ to } (1, 0, \dots, 0).$$

Similarly, one defines a map $\Pi^{-1} : S^d \rightarrow H^d$ that sends a point $z \in S^d$ to the intersection of H^d with the line through z and $(1, 0, \dots, 0)$. Note that Π^{-1} is not well-defined at $(1, 0, \dots, 0)$. To fix this, we add the point ∞ to the hyperplane H^d , and define $\Pi^{-1}(1, 0, \dots, 0) = \infty$ as well as $\Pi(\infty) = (1, 0, \dots, 0)$.

For any point $\alpha \in S^d$, we define Π_α to be the stereographic projection from the plane perpendicular to S^d at α , and we let Π_α^{-1} be its inverse (so, $\Pi(\infty) = -\alpha$). One can show that the maps Π_α and Π_α^{-1} are sphere-preserving maps (see [44] or [59] for a proof).

Sphere-preserving maps in the plane include rigid motions of the plane as well as dilations (and other mobius transformations). We will obtain sphere-preserving maps in the sphere by applying a projection onto a plane, then applying a dilation of the plane, and then mapping back by stereographic projection. Thus, for $\alpha \in S^d$ and $a \geq 0$, we define D_α^a to be the map that dilates the hyperplane perpendicular to S^d at α by a factor of a (note that $D_\alpha^a(\infty) = \infty$). For example,

$$D_{(-1,0,\dots,0)}^a : (-1, x_2, \dots, x_d) \mapsto (-1, ax_2, \dots, ax_d).$$

As the composition of sphere-preserving maps is again a sphere-preserving map, we can now define the sphere-preserving maps that we will use. For any α such that $\|\alpha\| < 1$, define $f_\alpha(z)$ by

$$f_\alpha(z) = \Pi_{\alpha/\|\alpha\|} \left(D_{\alpha/\|\alpha\|}^{1-\|\alpha\|} \left(\Pi_{\alpha/\|\alpha\|}^{-1}(z) \right) \right).$$

It is routine to verify that f_α is continuous. We wish to extend the definition of f_α to α on S^2 , even though the resulting maps will not be continuous. For $\|\alpha\| = 1$, we define

$$f_\alpha(z) = \begin{cases} -\alpha & \text{if } z = -\alpha, \text{ and} \\ \alpha & \text{otherwise.} \end{cases}$$

We will now examine the effect of the maps f_α on arrangements of spherical caps on S^d . Recall that a spherical cap on S^d is a connected region of S^d whose boundary is a $(d - 1)$ -dimensional sphere. Thus, the image of a cap under a map f_α is determined by the image of its boundary along with a point in its interior. For a cap C on S^d , let $p(C)$ denote the point on S^d that is the center of C (i.e., the point inside C that is equidistant from its boundary). We want to show that, for any arrangement of caps $\{C_1, \dots, C_n\}$ on S^d , there is an $\alpha \in S^d$ so that the centroid of $\{p(f_\alpha(C_1)), \dots, p(f_\alpha(C_n))\}$ is the origin. But first, we must exclude some degenerate cases:

Definition 4.1. An arrangement of caps $\{C_1, \dots, C_n\}$ in S^d is *well-behaved* if there is no point that belongs to at least half of the caps.

Remark. All of the arrangements of caps obtained from graphs contained in the other sections of this paper are well-behaved. By Theorem 3.2, every point on the sphere can belong to at most two caps. For k -nearest neighbor graphs considered in Section 5, each point is covered by at most $\tau_d k$ caps, where τ_d is a parameter depends only on d . Similar bounds hold for the well-shaped meshes

considered in Section 6. Thus, as long as n is sufficiently large, the arrangements associated with k -nearest neighbor graphs and well-shaped meshes are well-behaved.

Theorem 4.2. *For any well-behaved arrangement of caps $\{C_1, \dots, C_n\}$ in S^d , there is an α so that $\|\alpha\| < 1$ and*

$$\frac{\sum_{i=1}^n p(f_\alpha(C_i))}{n} = \vec{0}.$$

Proof. Consider the map from α to the centroid of $\{p(f_\alpha(C_1)), \dots, p(f_\alpha(C_n))\}$. We want to show that $\vec{0}$ has a preimage under this map. This would be easier if the map were continuous, but it is not continuous for $\|\alpha\| = 1$: as $-\alpha$ crosses the boundary of C_i , $p(f_\alpha(C_i))$ jumps from one side of the sphere to the other.

To fix this problem, we construct a slightly modified map that is continuous. Because the set of caps is well-behaved, we can choose an $\epsilon > 0$ so that, for all α such that $\|\alpha\| \geq 1 - \epsilon$, most of the caps $\{f_\alpha(C_1), \dots, f_\alpha(C_n)\}$ are entirely contained within the ball of radius $1/2n$ around $\alpha/\|\alpha\|$. In particular, this implies that f_α does not map the centroid of the centers of the caps to the origin. For $\alpha \in B^d$, we now define the map

$$\phi(\alpha) = \frac{\sum_{i=1}^n w(C_i, \alpha) f_\alpha(p(C_i))}{n},$$

where the weight function w is given by

$$w(C, \alpha) = \begin{cases} (2 - d(\alpha, C))/\epsilon & \text{if } d(\alpha, C) \geq 2 - \epsilon, \text{ and} \\ 1 & \text{otherwise,} \end{cases}$$

where by $d(\alpha, C)$, we mean the greatest distance from α to a point in the cap C (for example, if $-\alpha \in C$, then $d(\alpha, C) = 1 + \|\alpha\|$). We have chosen w to be a continuous function of α that goes to zero as $-\alpha$ approaches the boundary of a cap; so, $\phi(\alpha)$ is also a continuous function.

From the fact that $\{C_1, \dots, C_n\}$ is well-behaved, it is easy to verify that, for $\alpha \in S^d$, $\phi(\alpha)$ lies on the line connecting $\vec{0}$ to α and is closer to α than it is to $-\alpha$. By combining this fact with some elementary algebraic topology, we can use Lemma 4.3 to show that there is an α such that $\phi(\alpha) = \vec{0}$.

By our choice of ϵ , $\|\alpha\| < 1 - \epsilon$, so all of the terms $w(\alpha, C_i)$ are 1, which implies that f_α is the map that we were looking for. \square

Lemma 4.3. *Let $\phi : B^d \rightarrow B^d$ be a continuous function so that, for $\alpha \in S^d$, $\phi(\alpha)$ lies on the line connecting α with $\vec{0}$ and is closer to α than it is to $-\alpha$. Then, there exists an $\alpha \in B^d$ such that $\phi(\alpha) = \vec{0}$.*

Proof. Assume, by way of contradiction, that there is no point $\alpha \in B^d$ such that $\phi(\alpha) = \vec{0}$. Now, consider the map $b(\phi(\alpha))$, where $b : B^d - \{\vec{0}\} \rightarrow S^d$ by

$$b(z) = z/\|z\|.$$

Since b is a continuous map, $b \circ \phi$ is a continuous map of B^d onto S^d that is the identity on S^d . Then $z \mapsto -b(\phi(z))$ is a map from B^d onto S^d that has no fixed point. This contradicts Brouwer’s Fixed Point Theorem, which says that every continuous map from B^d into B^d has a fixed point. \square

We have shown that, for most collections of balls in H^d , there is a sphere preserving map from H^d to S^d such that the centroid of the centers of the caps is the origin. We now show that one can find such a map by performing a rigid motion of H^d followed by a dilation of H^d followed by stereographic projection.

Definition 4.4. An arrangement of balls $\{D_1, \dots, D_n\}$ in H^d is *well-behaved* if there is no point that belongs to at least half of the balls.

Theorem 4.5. Let $\{D_1, \dots, D_n\}$ be a well-behaved collection of balls. Then, there is a point $x \in H^d$ and an $a > 0$ such that the sphere preserving map

$$g_{x,a} : z \mapsto \Pi(a(z - x))$$

sends the balls to a collection of caps, the centroid of whose centers is the origin.

Proof (sketch). For an $\alpha \in S^d$, consider the map $g_{\Pi^{-1}(\alpha), (1-\|\alpha\|)}$ followed by a rotation of the sphere that sends $(-1, 0, \dots, 0)$ to α . As we did in the proof of Theorem 4.2, we can construct a continuous map from α to a weighted centroid of the centers of the caps, which for $\alpha \in S^d$ sends α to a point on the line segment between α and $\bar{0}$. We can then apply Lemma 4.3 to prove that there is some map α such that the map $g_{\Pi^{-1}(\alpha), (1-\|\alpha\|)}$ sends the centroid of the centers of the caps to the origin. \square

5. The spectra of k -nearest neighbor graphs

We extend our spectral planar separator theorem to graphs embedded in three and more dimensions. We show that Fiedler cuts of small ratio can be found in α -overlap graphs of k -ply neighborhood systems. One corollary of this extension is that the spectral method finds small ratio cuts for k -nearest neighbor graphs and well-shaped finite element meshes in any fixed dimension. In this section, we analyze intersection graphs and nearest neighbor graphs. Results on overlap graphs and well-shaped meshes will be given in the next section.

In this section and the next, we will use the following notation: We use capital letters to denote balls in \mathbb{R}^d . If A is a ball in \mathbb{R}^d , then we will use A' to denote its image on the sphere S^{d+1} under stereographic projection. If α is a positive real and A is a ball of radius r , then $\alpha \cdot A$ is the ball with the same center as A and radius αr . Similarly, if A' is a spherical cap of spherical radius r , then $\alpha \cdot A'$ is the spherical cap with the same center as A' and radius αr . Let V_d be the volume of a unit d -dimensional ball. Let A_d be the surface volume of a unit d -dimensional ball.

5.1. Intersection graphs

The graphs that we consider are defined by neighborhood systems. A *neighborhood system* is a set of closed balls in Euclidean space. A k -ply neighborhood system is one in which no point is contained in the interior of more than k of the balls. Given a neighborhood system, $\Gamma = \{B_1, \dots, B_n\}$, we define the *intersection graph* of Γ to be the undirected graph with vertex set $V = \{B_1, \dots, B_n\}$ and edge set

$$E = \{(B_i, B_j) : B_i \cap B_j \neq \emptyset\}.$$

For example, the Koebe–Andreev–Thurston embedding theorem says that every planar graph is isomorphic to the intersection graph of some 1-ply neighborhood system in two dimensions.

Let $P = \{p_1, \dots, p_n\}$ be a point set in \mathbb{R}^d . For each $p_i \in P$, let $N_k(p_i)$ be the set of k points closest to p_i in P (if there are ties, break them arbitrarily). A k -nearest neighbor graph of P is a graph with vertex set $\{p_1, \dots, p_n\}$ and edge set

$$E = \{(p_i, p_j) : p_i \in N_k(p_j) \text{ or } p_j \in N_k(p_i)\}.$$

Miller et al. [60] show that every k -nearest neighbor graph in \mathbb{R}^d is a subgraph of an intersection graph of a $\tau_d k$ -ply neighborhood system, where τ_d is the kissing number in d dimensions—the maximum number of nonoverlapping unit balls in \mathbb{R}^d that can be arranged so that they all touch a central unit ball [24]. Moreover, the maximum degree of a k -nearest neighbor graph is bounded by $\tau_d k$.

5.2. A spectral bound

Theorem 5.1. *Let G be a subgraph of an intersection graph of a k -ply neighborhood system in \mathbb{R}^d such that the maximum degree of G is Δ . Then, the Fiedler value of $L(G)$ is bounded by $c_d \Delta (k/n)^{2/d}$, where $c_d = 2(A_{d+1}/V_d)^{2/d}$.*

Proof. Let $\Gamma = \{B_1, \dots, B_n\}$ be the k -ply neighborhood system of which G is the intersection graph. By Theorem 4.2, there exists a sphere-preserving map $\phi : \mathbb{R}^d \rightarrow S^d$ such that the centroid of the centers of the images of the B_i 's is the center of the sphere.

Let $\phi(\Gamma) = \{B'_1, \dots, B'_n\}$ be the images of the balls in Γ under ϕ . Then, the balls in $\phi(\Gamma)$ also form a k -ply system. Let r_i be the radius of B'_i . Because $V_d r_i^d \leq \text{volume}(B'_i)$,

$$\sum_{i=1}^n V_d r_i^d \leq \sum_{i=1}^n \text{volume}(B'_i) \leq k A_{d+1}. \tag{1}$$

By Lemma 3.1,

$$\begin{aligned} \lambda_2(L(G)) &\leq \frac{\sum_{i=1}^n 2\Delta r_i^2}{n} \\ &\leq (2\Delta) \frac{(k A_{d+1}/V_d)^{2/d} n^{1-2/d}}{n} \\ &\leq (2\Delta) \left(\frac{A_{d+1}}{V_d}\right)^{2/d} \left(\frac{k}{n}\right)^{2/d}. \end{aligned}$$

Note that the second inequality follows from (1). \square

The next two corollaries follow from Theorems 5.1, 2.1, and Lemma A.1.

Corollary 5.2. *The Fiedler value of a k -nearest neighbor graph of n points in \mathbb{R}^d is bounded by $O(k^{1+2/d}/n^{2/d})$. Therefore, G has a Fiedler cut of ratio $O(k^{1+1/d}/n^{1/d})$, and one can repeatedly take Fiedler cuts to find a bisector of size $O(k^{1+1/d}n^{1-1/d})$.*

Corollary 5.3. *Let G be a subgraph of an intersection graph of a k -ply neighborhood system in \mathbb{R}^d whose maximum degree is Δ . Then, G has a Fiedler cut of ratio $O(\Delta^{1+1/d}/n^{1/d})$, and one can iterate Fiedler cuts to obtain a bisector of size $O(\Delta^{1+1/d}n^{1-1/d})$.*

6. The spectra of well-shaped meshes

One of the main applications of the spectral method is the partitioning of meshes for parallel numerical simulations. Many experiments demonstrate the effectiveness of this method [18,48,49,66,70,82]. In this section, we explain why the spectral method finds such good partitions of well-shaped meshes.

6.1. Well-shaped meshes

Most numerical methods work by approximating continuous problems with discrete problems on finite structures whose solutions can be efficiently computed. The finite structure used is often called a *mesh*. Many such methods have been developed and applied to important problems in mechanics and physics.

Most of these numerical methods can be classified as *equation based methods* (e.g., the finite element, finite difference, and finite volume methods) or *particle methods* (e.g., the N -body simulation method). However different the particular methods may be, a basic principle is common to all—accuracy of approximation is ensured by using meshes that satisfy certain numerical and geometric constraints. Meshes that satisfy these constraints are said to be *well-shaped*.

To motivate our spectral analysis of well-shaped meshes, we review the conditions required of finite element and finite difference meshes. More detailed discussions can be found in several books and papers (for example, see [69,50,13,12,37]). Background material on the particle method can be found in [15,43,45,83].

The *finite element method* approximates a continuous problem by subdividing the domain (a subset of \mathbb{R}^d) of the problem into a *mesh* of polyhedral *elements* and then approximates the continuous function by piecewise polynomial functions on the elements. A common choice for an element is a d -dimensional *simplex*. Accordingly, a d -dimensional finite element mesh is a d -dimensional *simplicial complex*, a collection of d -dimensional simplices that meet only at shared faces [13,12,58].

The computation graph associated with each simplicial complex is often its 1-skeleton or the 1-skeleton of its geometric dual (as used in the finite volume method). In the finite element method, a linear system is defined over a mesh, with variables representing physical quantities at the nodes. The nonzero structure of the coefficient matrix of such a linear system is exactly the adjacency structure of the 1-skeleton of the simplicial complex.

To ensure accuracy, in addition to the conditions that a mesh must conform to the boundaries of the region and be fine enough, each individual element of the mesh must be *well-shaped*. A common shape criterion for the finite element method is that the angles of each element are not too small, or the *aspect ratio* of each element is bounded [9,13,37]. Other numerical formulations require slightly different conditions. For example, the controlled volume formulation [64,61] using a Voronoi diagram requires that the *radius aspect ratio* (the ratio of the circumscribed radius to the shortest edge length of an element in the dual Delaunay diagram) is bounded.

The finite difference method also uses a discrete structure, a *finite difference mesh*, to approximate a continuous problem. Finite difference meshes are often produced by inserting a uniform grid from \mathbb{R}^2 or \mathbb{R}^3 into the domain via a boundary-matching conformal mapping. Notice that, unlike a finite element mesh, a finite difference mesh need not be a collection of simplices or elements, so we cannot analyze it as we do a triangulation. In general, the derivative of the conformal transformation must vary gradually with respect to the mesh size in order to produce

good results (see, for example [80]). This means that the mesh will probably satisfy a *density condition* [11,63].

Let G be an undirected graph and let π be an embedding of its nodes in \mathbb{R}^d . We say π is an embedding of *density* α if the following inequality holds for all vertices v in G : Let u be the node closest to v . Let w be the node farthest from v that is connected to v by an edge. Then

$$\frac{\|\pi(w) - \pi(v)\|}{\|\pi(u) - \pi(v)\|} \leq \alpha.$$

In general, G is an α -density graph in \mathbb{R}^d if there exists an embedding of G in \mathbb{R}^d with density α .

6.2. Modeling well-shaped meshes

We will use the overlap graph to model well-shaped meshes (Miller et al. [59]). An overlap graph is based on a k -ply neighborhood system. The neighborhood system and a parameter, $\alpha \geq 1$, define an overlap graph: Let $\alpha \geq 1$, and let $\Gamma = \{B_1, \dots, B_n\}$ be a k -ply neighborhood system in \mathbb{R}^d . The α -overlap graph of Γ is the graph with vertex set $\{B_1, \dots, B_n\}$ and edge set

$$\{(B_i, B_j) : (B_i \cap (\alpha \cdot B_j) \neq \emptyset) \text{ and } ((\alpha \cdot B_i) \cap B_j \neq \emptyset)\},$$

where by $\alpha \cdot B$, we mean the ball whose center is the same as the center of B and whose radius is larger by a multiplicative factor of α .

Overlap graphs are good models of well-shaped meshes because each well-shaped mesh in two, three, or higher dimensions is a *bounded-degree subgraph* of some overlap graph (for suitable choices of the parameters α and k). For example,

- Let M be a finite element mesh embedded in \mathbb{R}^d in which every element has aspect ratio bounded by a . Then, there is a constant α depending only on d and a so that the 1-skeleton of M is a subgraph of an α -overlap graph of a 1-ply neighborhood system. Moreover its maximum degree is bounded by a constant that also depends only on d and a [59].
- Let M be a Voronoi diagram (from a finite volume method) in \mathbb{R}^d in which the radius aspect ratio of its dual Delaunay diagram is bounded by a . Then there is a constant α depending only on d and a so that the dual Delaunay diagram is an α -density graph [61].
- If G is an α -density graph in \mathbb{R}^d , then the maximum degree of G is bounded by a constant depending only on α and d ; and, G is a subgraph of an α -overlap graph of a 1-ply neighborhood system [63,59].
- The computation/communication graph used in hierarchical N -body simulation methods (such as the Barnes-Hut's treecode method [15] and the fast-multipole method [43]) is a subgraph of an α -overlap graph of an $O(\log n)$ -ply neighborhood system [77].

6.3. Spherical embeddings of overlap graphs

In this section, we show that an α -overlap graph is a subgraph of the intersection graph obtained by projecting its neighborhoods onto the sphere and then dilating each by an $O(\alpha)$ factor. By choosing the proper projection, we are able to use this fact to bound the eigenvalues of these graphs.




Fig. 1. Symmetric stereographic images.

Theorem 6.1. *Let $\alpha \geq 1$. Let A and B be balls in \mathbb{R}^d such that*

$$(A \cap (\alpha \cdot B) \neq \emptyset) \text{ and } ((\alpha \cdot A) \cap B \neq \emptyset).$$

Then, $(\pi\alpha + \alpha + \pi) \cdot A'$ touches $(\pi\alpha + \alpha + \pi) \cdot B'$.

Our proof uses two lemmas that handle orthogonal special cases.

Lemma 6.2. *Let A and C be balls in \mathbb{R}^d equidistant from the origin and having the same radius. Let A' and C' be their images under stereographic projection onto S^{d+1} . If $\alpha \cdot A$ touches $\alpha \cdot C$, then $(\alpha\pi/2) \cdot A'$ touches $(\alpha\pi/2) \cdot C'$.*

Proof. Let r be the radius of A and C . Because $\alpha \cdot A$ touches $\alpha \cdot C$, the centers of A and C are at distance at most $2\alpha r$ from each other.

Let S be the sphere centered at the origin that passes through the centers of A and C . The geodesic arc between the centers of A and C (on S) has length at most $2\alpha r\pi/2$. The portion of this arc that lies in the interior of A has length at least r (see Fig. 1 for a two dimensional example). Since stereographic projection preserves the relations between the intersections of A and C with S , $(\alpha\pi/2) \cdot A'$ will touch $(\alpha\pi/2) \cdot C'$. \square

Lemma 6.3. *Let A and B be balls in \mathbb{R}^d so that the center of A , the center of B , and the origin are colinear and the origin does not lie on the line segment between the center of A and the center of B . If A is closer to the origin than B and $\alpha \cdot A$ touches B , then $\alpha \cdot A'$ touches B' .*

Proof. We will restrict our attention to the plane through the top of the sphere, the origin, and the centers of A and B (see Fig. 2). Let a denote the interval that is the intersection of A with the plane. Observe that an interval of the same size as a but located further to the right on the line will have a smaller projection on the circle. The lemma follows. \square

Proof of Theorem 6.1. Let A and B be any two balls in \mathbb{R}^d and let A' and B' be their images under stereographic projection on S^{d+1} . Assume that $\alpha \cdot A$ touches B and $\alpha \cdot B$ touches A . We will show that $(\pi\alpha + \alpha + \pi) \cdot A'$ touches $(\pi\alpha + \alpha + \pi) \cdot B'$.




Fig. 2. Restriction to the plane through the top of the sphere, the origin, and the centers of A and B .




Fig. 3. C is the circle congruent to and equidistant from the center to A that is closest to B .

Assume, without loss of generality, that A is at least as large as B . Let C be the disk of the same distance to the origin as A and congruent to A that is closest to B . Then, the centers of C and B are colinear with the origin (see Fig. 3). Let C' be the image of C . Since C is closer to B than A is, $\alpha \cdot C$ touches B and $\alpha \cdot B$ touches A . By Lemma 6.3, $\alpha \cdot C'$ touches $\alpha \cdot B'$.

The distance between the centers of A and B is less than $(\alpha + 1)$ times the radius of A (because we assume that A is at least as large as B). The same holds for the distance between the center of C and the center of B . Therefore, $(\alpha + 1) \cdot A$ touches $(\alpha + 1) \cdot C$, so Lemma 6.2 implies that $\pi(\alpha + 1)/2 \cdot A'$ touches $\pi(\alpha + 1)/2 \cdot C'$. Since A' and C' have the same spherical radius, $\alpha \cdot C' \subset (\pi(\alpha + 1) + \alpha)A'$. Thus, $(\pi\alpha + \alpha + \pi) \cdot A'$ must touch $(\pi\alpha + \alpha + \pi) \cdot B'$. \square

6.4. A spectral bound

We now show that the Fiedler value of a bounded degree subgraph of an α -overlap graph is small.

Theorem 6.4. *If G is a subgraph of an α -overlap graph of a k -ply neighborhood system in \mathbb{R}^d and the maximum degree of G is Δ , then the Fiedler value of $L(G)$ is bounded by $\gamma_d \Delta \alpha^2 (k/n)^{2/d}$, where $\gamma_d = 2(\pi + 1 + \pi/\alpha)^2 (A_{d+1}/V_d)^{2/d}$. Accordingly, G has a Fiedler cut of ratio $O(\Delta \alpha (k/n)^{1/d})$, and one can iterate Fiedler cuts to obtain a bisector of size $O(\Delta \alpha k^{1/d} n^{1-1/d})$.*

Proof. Let $\Gamma = \{B_1, \dots, B_n\}$ be the k -ply neighborhood system whose intersection graph contains G . By Theorem 4.5, there is a stereographic projection Π from R^d onto a particular sphere S^{d+1} so that the centroid of the centers of the images of the neighborhoods is the center of the sphere.

Let $\Pi(\Gamma) = \{B'_1, \dots, B'_n\}$ be the images of the balls in Γ under Π . Let r_i be the radius of B'_i . Because $V_d r^d \leq \text{volume}(B'_i)$, We know that

$$\sum_{i=1}^n V_d r_i^d \leq \sum_{i=1}^n \text{volume}(B'_i) \leq k A_{d+1}.$$

By Theorem 6.1, G is a subgraph graph of the intersection graph of $\{(\pi\alpha + \alpha + \pi) \cdot B'_i : 1 \leq i \leq n\}$. Thus, by Lemma 3.1,

$$\begin{aligned} \lambda_2(L(G)) &\leq \frac{\sum_{i=1}^n 2\Delta(\pi\alpha + \alpha + \pi)^2 r_i^2}{n} \\ &\leq (2\Delta)(\pi\alpha + \alpha + \pi)^2 \left(\frac{A_{d+1}}{V_d}\right)^{2/d} \left(\frac{k}{n}\right)^{2/d}. \end{aligned}$$

Given the bound on the Fiedler value, the ratio achievable by a Fiedler cut follows immediately from Theorem 2.1 and the corresponding bisector size follows Lemma A.1. \square

Remark. Recently, Agarwal and Pach [1] and, independently, Spielman and Teng [73] gave an elementary proof of the sphere separator theorem of Miller et al. [60] on planar graphs and intersection graphs. However, these proofs do not directly extend to overlap graphs. The relation between overlap graphs and intersection graphs established by Theorem 6.1 enables us to prove the overlap graph separator theorem using the intersection graph separator theorem. The same reduction also extends the deterministic linear time algorithm for finding a good sphere separator from intersection graphs to overlap graphs [32].

7. Final remarks

The *genus* of a graph is the minimum number of handles that must be added to the plane to embed the graph without any crossings. Graph with a bounded genus is a natural generalization of planar graphs, which have genus zero. Let G_1, \dots, G_h be h disjoint connected subgraphs of G . The *minor* defined by G_1, \dots, G_h is a graph H whose vertex set is $\{1, 2, \dots, h\}$ and whose edge set contains exactly all those pairs $(i, j) : 1 \leq i, j \leq h$ such that there is an edge in G that connects a vertex of G_i to a vertex of G_j . Clearly, H can be obtained from a subgraph of G by contraction.

The following are two natural extensions of Theorem 3.3.

Conjecture 1 (Graphs with bounded genus). *Let G be a graph of n nodes and genus $g \geq 1$. If the maximum vertex degree of G is Δ , then, the Fiedler value of G is at most $O\left(\frac{\Delta g}{n}\right)$.*

Conjecture 2 (Graphs with bounded forbidden minors). *Let G be a graph of n nodes. If the maximum vertex degree of G is Δ and G does not have K_h , the clique of h vertices, as a minor, then the Fiedler value of G is at most $O\left(\frac{\Delta h^k}{n}\right)$, for some small constant k .*

Kelner [51] recently proved both of these conjectures.

The following conjecture remains open. Suppose G is an undirected graph. Let G_1, \dots, G_h be h disjoint connected subgraphs of G . Let H be a graph minor defined by G_1, \dots, G_h . The *depth* of H is the maximum diameter of G_1, \dots, G_h .

Conjecture 3 (Graphs with bounded shallow excluded minors). *Let G be a graph of n nodes. If the maximum vertex degree of G is Δ and there exists $d \geq 1$ such that for every L , G does not have L^d -clique minor of depth L , then the Fiedler value of G is at most $O(\Delta/n^{2/d})$.*

Plotkin et al. [65] proved that any graph that excludes K_h as a depth L minor has a separator of size $O(lh^2 \log n + n/L)$. In addition, they showed that for any well-shaped meshes G in \mathbb{R}^d , G excludes K_h minors of depth L for $h = \Omega(L^{d-1})$. Teng [78] extended the latter result to k -nearest neighborhood graphs in \mathbb{R}^d . The above conjecture is a natural extension of Theorem 6.4 on well-shaped meshes and Corollary 5.2 on k -nearest neighbor graphs.

To prove Conjecture 3, we may need to develop new techniques. The proofs of this paper as well as of Kelner [51] critically used the fact that every graph considered has a geometric realization (as the intersection of a disk-packings). So far, there is no similar embedding results for graphs with bounded shallow excluded minors.

One approach is to develop a combinatorial characterization of graph eigenvalues. There are several nice results that establish bounds on eigenvalues of graphs based on combinatorial information. For example, Chung [22] showed that graphs with a large diameter have small eigenvalues. However, the existing combinatorial techniques are still too limited to prove the conjectures mentioned above. For example, the hard cases for graphs with bounded shallow excluded minors are graphs with small diameters such as $O(\log n)$.

One possible direction is to build on our proof and intuition in the spectral analysis of planar graphs to develop an alternative combinatorial proof of Theorem 3.3. It could be an interesting study of how the conformal mapping scheme for planar graphs can be converted into a combinatorial argument. Intuitively, the conformal mapping introduces a distance metric among vertices so that neighboring vertices are close. There might be a combinatorial approximation or characterization of this metric that does not use Koebe embedding. Such a study might provide combinatorial insights about spectral decomposition of graph matrices and might potentially provide a combinatorial framework for establishing tight upper and lower bounds on graph eigenvalues.

Acknowledgments

We thank Michael Mandell for advice on algebraic topology, John Gilbert, Nabil Kahale, Gary Miller, and Horst Simon for helpful discussions on the spectra of graphs, and Edmond Chow, Alan Edelman, Steve Guattery, Bruce Hendrickson, Satish Rao, Ed Rothberg, Dafna Talmor, and Steve Vavasis for comments on an early draft of this paper.

Appendix A. From ratio cuts to bisection

We now explain how to use good ratio cuts to produce a bisection of a graph.

Lemma A.1. Assume that we are given an algorithm that will find a cut of ratio at most $\phi(k)$ in every k -node subgraph of G , for some monotonically decreasing function ϕ . Then repeated application of this algorithm can be used to find a bisection of G of size at most

$$\int_{x=1}^n \phi(x) \, dx.$$

Proof. The following algorithm (see [54,40]) will find the bisection.

1. Initially, let $D^{(0)} = G$, let A and B be empty sets, and let $i = 0$.
2. If $D^{(i)}$ is empty, then return A and B ; otherwise repeat
 - (a) Find a cut of ratio at most $\phi(|D^{(i)}|)$ that divides $D^{(i)}$ into $F^{(i)}$ and $\bar{F}^{(i)}$. We assume that $|F^{(i)}| \leq |\bar{F}^{(i)}|$.
 - (b) If $|A| \leq |B|$, let $A = A \cup F^{(i)}$; otherwise, let $B = B \cup F^{(i)}$.
 - (c) Let $D^{(i+1)} = \bar{F}^{(i)}$, let $i = i + 1$, and return to step (a).

We assume that the algorithm terminates after t iterations. To show that this algorithm produces a bisection, we need to prove that, for all i in the range $0 \leq i < t$, $\min(|A|, |B|) + |F^{(i)}| \leq n/2$. Because $|F^{(i)}| \leq |\bar{F}^{(i)}|$,

$$\min(|A|, |B|) + |F^{(i)}| \leq (|A| + |B| + |F^{(i)}| + |\bar{F}^{(i)}|)/2 = n/2.$$

We now analyze the total cut size. Because the algorithm finds cuts of ratio at most $\phi(|D^{(i)}|)$ at the i th iteration, the number of edges we cut to separate $F^{(i)}$ is at most

$$\begin{aligned} \phi(|D^{(i)}|)|F^{(i)}| &= \sum_{j=1}^{|F^{(i)}|} \phi(|D^{(i)}|) \\ &= \sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(|D^{(i)}|) \\ &\leq \sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(j) \end{aligned}$$

The inequality follows from the fact that ϕ is monotonically decreasing. The total number of edges cut by this algorithm is at most

$$\begin{aligned} \sum_{i=0}^{t-1} \phi(|D^{(i)}|)|F^{(i)}| &\leq \sum_{i=0}^{t-1} \left(\sum_{j=|D^{(i)}|}^{|D^{(i)}|-|F^{(i)}|+1} \phi(j) \right) \\ &= \sum_{j=1}^n \phi(j) \\ &\leq \int_1^n \phi(x) \, dx \end{aligned}$$

The last inequality follows from the assumption that ϕ is monotonically decreasing. \square

Remark. If $\phi(x) = x^{-1/d}$ then

$$\int_1^n \phi(x) dx = \frac{d}{d-1} (n^{1-1/d} - 1).$$

Lipton and Tarjan [54] showed that by repeatedly applying an α -separator of size $\beta\sqrt{n}$, one can obtain a bisection of size $\beta/(1 - \sqrt{1 - \alpha})\sqrt{n}$. Gilbert [40] extended this result to graphs with positive vertex weights at the expense of a $1/(1 - \sqrt{2})$ factor in the bisection bound. Djidjev and Gilbert [27] further generalized this result to graphs with arbitrary weights. Leighton and Rao [53] showed that one can obtain an $O(\alpha)$ -approximation to a 1/3-separator from an α -approximation to a ratio cut.

References

- [1] P. Agarwal, J. Pach, *Combinatorial Geometry*, Wiley-Interscience, 1995.
- [2] C.J. Alpert, A.B. Kahng, Recent directions in netlist partitioning: a survey, *Integration: The VLSI J.* 19 (1995) 1–81.
- [3] N. Alon, Eigenvalues and expanders, *Combinatorica* 6 (2) (1986) 83–96.
- [4] N. Alon, V.D. Milman, λ_1 , isoperimetric inequalities for graphs, and superconcentrators, *J. Combin. Theory Series B* 38 (1985) 73–88.
- [5] E.M. Andreev, On convex polyhedra in lobacevskii space, *Math. USSR Sbornik* 10 (3) (1970) 413–440.
- [6] E.M. Andreev, On convex polyhedra of finite volume in lobacevskii space, *Math. USSR Sbornik* 12 (2) (1970) 259–270.
- [7] Noga Alon, Paul Seymour, Robin Thomas, A separator theorem for graphs with an excluded minor and its applications, *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, ACM, 1990, pp. 293–299.
- [8] S. Arora, S. Rao, U. Vazirani, Expander flows, geometric embeddings and graph partitioning, in: *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, 2004, pp. 222–231.
- [9] I. Babuška, A.K. Aziz, On the angle condition in the finite element method, *SIAM J. Numer. Anal.* 13 (2) (1976) 214–226.
- [10] E.R. Barnes, An algorithm for partitioning of nodes of a graph, *SIAM J. Algebraic Discrete Methods* 4 (3) (1982) 541–550.
- [11] M.J. Berger, S. Bokhari, A partitioning strategy for nonuniform problems on multiprocessors, *IEEE Trans. Comp.* C-36 (1987) 570–580.
- [12] M. Bern, D. Eppstein, Mesh generation and optimal triangulation, in: F.K. Hwang, D.-Z. Du (Eds.), *Computing in Euclidean Geometry*, World Scientific, 1992, pp. 23–90.
- [13] M. Bern, D. Eppstein, J.R. Gilbert, Provably good mesh generation, *J. Comput. Syst. Sci.* 48 (1994) 384–409.
- [14] E.R. Barnes, A.J. Hoffman, Partitioning, spectra and linear programming, *Progr. Combin. Optim.* (1984) 13–25.
- [15] J. Barnes, P. Hut, A hierarchical $O(n \log n)$ force calculation algorithm, *Nature* 324 (1986) 446–449.
- [16] Norman Biggs, *Algebraic Graph Theory*, second ed., Cambridge University Press, New York, NY, 1993.
- [17] R. Boppana, Eigenvalues and graph bisection: an average-case analysis, 28th Annual Symposium on Foundations of Computer Science, Los Angeles, October 1987, IEEE, 1987, pp. 280–285.
- [18] S.T. Barnard, H.D. Simon, A fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems, Technical Report RNR-92-033, NASA Ames Research Center, 1992.
- [19] Dragos M. Cvetkovic, Michael Doob, Horst Sachs, *Spectra of Graphs: Theory and Application*, Academic Press, New York, 1990.
- [20] J. Cheeger, A lower bound for the smallest eigenvalue of the laplacian, in: R.C. Gunning (Ed.), *Problems in Analysis*, Princeton University Press, 1970, pp. 195–199.
- [21] T.F. Chan, D.C. Resasco, A framework for the analysis and construction of domain decomposition preconditioners, Technical Report CAM-87-09, UCLA, 1987.
- [22] F.R.K. Chung, Diameters and eigenvalues, *J. Amer. Math. Soc.* 2 (2) (1989) 187–196.
- [23] F.R.K. Chung, *Spectral Graph Theory*, Amer. Math. Soc., Providence, RI, 1997.
- [24] J.H. Conway, N.J.A. Sloane, *Sphere Packings, Lattices and Groups*, Springer-Verlag, 1988.
- [25] T.F. Chan, B. Smith, Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes, *Contemp. Math.* (1993) 1–14.

- [26] P.K. Chan, M. Schlag, J. Zien, Spectral k -way ratio cut partitioning and clustering, in: *Symp. on Integrated Systems*, 1993.
- [27] H.N. Djidjev, J.R. Gilbert, Separators in graphs with negative and multiple vertex weights, Technical Report CS-92-07, Xerox PARC, Palo Alto, 1992.
- [28] W.E. Donath, A.J. Hoffman, Algorithms for partitioning of graphs and computer logic based on eigenvectors of connection matrices, *IBM Techn. Disclosure Bull.* 15 (1972) 938–944.
- [29] W.E. Donath, A.J. Hoffman, Lower bounds for the partitioning of graphs, *J. Res. Develop.* 17 (1973) 420–425.
- [30] P. Diaconis, D. Strook, Geometric bounds for eigenvalues of Markov chains, *Ann. Appl. Probab.* 1 (1) (1991) 36–61.
- [31] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, NY, 1987.
- [32] D. Eppstein, G.L. Miller, S.-H. Teng, A deterministic linear time algorithm for geometric separators and its applications, *Fund. Inform.* 22 (4) (1995) 309–330.
- [33] M. Fiedler, Algebraic connectivity of graphs, *Czechoslovak Math. J.* 23 (98) (1973) 298–305.
- [34] M. Fiedler, Eigenvectors of acyclic matrices, *Czechoslovak Math. J.* 25 (100) (1975) 607–618.
- [35] M. Fiedler, A property of eigenvectors of nonnegative symmetric matrices and its applications to graph theory, *Czechoslovak Math. J.* 25 (100) (1975) 619–633.
- [36] James Allen Fill, Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process, *Ann. Appl. Probab.* 1 (1) (1991) 62–87.
- [37] I. Fried, Condition of finite element matrices generated from nonuniform meshes, *AIAA J.* 10 (1972) 219–221.
- [38] N. Garg, H. Saran, V.V. Vazirani, Finding separator cuts in planar graphs within twice the optimal, *SIAM J. Comput.* 29 (1) (1999) 159–179.
- [39] J.R. Gilbert, J.P. Hutchinson, R.E. Tarjan, A separation theorem for graphs of bounded genus, *J. Algorith.* 5 (1984) 391–407.
- [40] John Russell Gilbert, Graph Separator theorems and sparse gaussian elimination, Ph.D. thesis, Computer Science Department, Stanford University, 1980.
- [41] J.R. Gilbert, N. Kahale, personal communication, 1995.
- [42] Stephen Guattery, Gary L. Miller, On the performance of the spectral graph partitioning methods, in: *Proceedings of The Second Annual ACM–SIAM Symposium on Discrete Algorithms*, ACM–SIAM, 1995, pp. 233–242.
- [43] L. Greengard, V. Roklin, A fast algorithm for particles simulations, *J. Comput. Phys.* 73 (1987) 325–348.
- [44] D. Hilbert, S. Cohn-Vossen, *Geometry and the Imagination*, Chelsea Publishing Company, New York, 1952.
- [45] R.W. Hackney, J.W. Eastwood, *Computer Simulation Using Particles*, McGraw-Hill, 1981.
- [46] L. Hagen, A.B. Kahng, New spectral methods for ratio cut partitioning and clustering, *IEEE Trans. Computer-Aided Des.* 11 (9) (1992) 1074–1085.
- [47] K.M. Hall, An r -dimensional quadratic placement algorithm, *Manage. Sci.* 17 (1970) 219–229.
- [48] Bruce Hendrickson, Robert Leland, An improved spectral graph partitioning algorithm for mapping parallel computation, Technical Report SAND92-1460, Sandia National Lab., December 1992.
- [49] Bruce Hendrickson, Robert Leland, The Chaco user’s guide, Version 1.0, Technical Report SAND93-2339, Sandia National Lab., 1993.
- [50] C. Johnson, *Numerical Solution of Partial Differential Equations by the Finite Element Method*, Cambridge University Press, 1992.
- [51] J.A. Kelner, Spectral partitioning, eigenvalue Bounds, and circle packings for graphs of bounded genus, in: *Proceedings of the 36th Annual ACM Symposium on the Theory of Computing*, 2004.
- [52] Paul Koebe, Kontaktprobleme der konformen Abbildung, *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse* 88 (1936) 141–164.
- [53] F.T. Leighton, S. Rao, An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms, in: *29th Annual Symposium on Foundations of Computer Science*, 1988, pp. 422–431.
- [54] R.J. Lipton, R.E. Tarjan, A separator theorem for planar graphs, *SIAM J. Appl. Math.* 36 (April) (1979) 177–189.
- [55] M. Mihail, Conductance and convergence of Markov chains—a combinatorial treatment of expanders, in: *30th Annual Symposium on Foundations of Computer Science*, IEEE, 1989, pp. 526–531.
- [56] B. Mohar, The Laplacian spectrum of graphs, in: *Sixth International Conference on the Theory and Applications of Graphs*, 1988, pp. 871–898.
- [57] B. Mohar, Isoperimetric numbers of graphs, *J. Combin. Theory, Series B* 47 (1989) 274–291.
- [58] Gary L. Miller, William Thurston, Separators in two and three dimensions, *Proceedings of the 22th Annual ACM Symposium on Theory of Computing*, Baltimore, May 1990, ACM, 1990, pp. 300–309.
- [59] G.L. Miller, S.-H. Teng, W. Thurston, S.A. Vavasis, Finite element meshes and geometric separators, *SIAM J. Scient. Comput.* 19 (2) (1998).

- [60] G.L. Miller, S.-H. Teng, W. Thurston, S.A. Vavasis, Separators for sphere-packings and nearest neighborhood graphs, *J. ACM* 44 (1) (1997) 1–29.
- [61] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, Noel Walkington, A Delaunay based numerical method for three dimensions: generation, formulation, and partition, *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, Las Vegas, May 1995, ACM, 1995, pp. 683–692.
- [62] G.L. Miller, S.-H. Teng, S.A. Vavasis, An unified geometric approach to graph separators, *32nd Annual Symposium on Foundations of Computer Science*, IEEE, 1991, pp. 538–547.
- [63] Gary L. Miller, Stephen A. Vavasis, Density graphs and separators, in: *Proceedings of the ACM–SIAM Symposium on Discrete Algorithms*, 1991, pp. 331–336.
- [64] R.A. Nicolaides, Direct discretization of planar div–curl problems, *SIAM J. Numer. Anal.* 29 (1) (1992).
- [65] S. Plotkin, S. Rao, W.D. Smith, Shallow excluded minors and improved graph decomposition, in: *5th Symp. Discrete Algorithms*, SIAM, 1994, pp. 462–470.
- [66] A. Pothen, H.D. Simon, K.-P. Liou, Partitioning sparse matrices with eigenvectors of graphs, *SIAM J. Matrix Anal. Appl.* 11 (1990) 430–452.
- [67] B. Parlett, H. Simon, L. Stringer, Estimating the largest eigenvalues with the lanczos algorithm, *Math. Comput.* 38 (1982) 153–165.
- [68] A. Pothen, H.D. Simon, L. Wang, Spectral nested dissection, Technical Report CS-92-01, Pennsylvania State University, Department of Computer Science, 1992.
- [69] G. Strang, G.J. Fix, *An Analysis of the Finite Element Method*, Prentice-Hall, Englewood Cliffs, New Jersey, 1973.
- [70] H.D. Simon, Partitioning of unstructured problems for parallel processing, *Comput. Syst. Eng.* 2 (2/3) (1991) 135–148.
- [71] A.J. Sinclair, M.R. Jerrum, Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, ACM, 1988, pp. 235–244.
- [72] Alistair Sinclair, Mark Jerrum, Approximative counting, uniform generation and rapidly mixing Markov chains, *Inform. Comput.* 82 (1) (1989) 93–133.
- [73] D.A. Spielman, S.-H. Teng, Disk packings and planar separators, in: *Proceedings of 12th ACM Symposium on Computational Geometry*, May 1996, pp. 358–394.
- [74] D.A. Spielman, S.-H. Teng, Spectral partitioning works: planar graphs and finite element meshes, in: *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, 1996, pp. 96–106.
- [75] G. Strang, *Linear Algebra and Its Applications*, Harcourt Brace Jovanovich, Publishers, San Diego, CA, 1988.
- [76] S.-H. Teng, Points, spheres, and separators: a unified geometric approach to graph partitioning, Ph.D. thesis, Carnegie-Mellon University, School of Computer Science, 1991, CMU-CS-91-184.
- [77] Shang-Hua Teng, Provably good partitioning and load balancing algorithms for parallel adaptive n -body simulation, *SIAM J. Scient. Comput.* 19 (2) (1998) 635–656.
- [78] Shang-Hua Teng, Combinatorial aspects of geometric graphs, *Comput. Geom.* 9 (1998) 277–287.
- [79] W.P. Thurston, *The geometry and topology of 3-manifolds*, Princeton University Notes, 1988.
- [80] J.F. Thompson, Z.U.A. Warsi, C.W. Mastin, *Numerical Grid Generation: Foundations and Applications*, North Holland, New York, 1985.
- [81] P. Ungar, A theorem on planar graphs, *J. London Math. Soc.* 26 (1951) 256–262.
- [82] R.D. Williams, Performance of dynamic load balancing algorithms for unstructured mesh calculations, Technical Report, California Institute of Technology, 1990.
- [83] F. Zhao, An $O(n)$ algorithm for three-dimensional n -body simulation, Technical Report TR AI Memo 995, MIT, AI Lab., October 1987.