

SPECTRAL PARTITIONING, EIGENVALUE BOUNDS, AND CIRCLE PACKINGS FOR GRAPHS OF BOUNDED GENUS*

JONATHAN A. KELNER[†]

Abstract. In this paper, we address two long-standing questions about finding good separators in graphs of bounded genus and degree:

1. It is a classical result of Gilbert, Hutchinson, and Tarjan [*J. Algorithms*, 5 (1984), pp. 391–407] that one can find asymptotically optimal separators on these graphs if given both the graph *and* an embedding of it onto a low genus surface. Does there exist a simple, efficient algorithm to find these separators, given only the graph and not the embedding?
2. In practice, spectral partitioning heuristics work extremely well on these graphs. Is there a theoretical reason why this should be the case?

We resolve these two questions by showing that a simple spectral algorithm finds separators of cut ratio $O(\sqrt{g/n})$ and vertex bisectors of size $O(\sqrt{gn})$ in these graphs, both of which are optimal. As our main technical lemma, we prove an $O(g/n)$ bound on the second smallest eigenvalue of the Laplacian of such graphs and show that this is tight, thereby resolving a conjecture of Spielman and Teng. While this lemma is essentially combinatorial in nature, its proof comes from continuous mathematics, drawing on the theory of circle packings and the geometry of compact Riemann surfaces.

Key words. bounded genus, circle packing, graph separators, Laplacian, partitioning, spectral partitioning

AMS subject classifications. 05C85, 68R10, 68W40, 68Q25, 68W05

DOI. 10.1137/S0097539705447244

1. Introduction. Spectral methods have long been used as a heuristic in graph partitioning. They have had tremendous experimental and practical success in a wide variety of scientific and numerical applications, including mapping finite element calculations on parallel machines [24, 29], solving sparse linear systems [7, 8], partitioning for domain decomposition [8], and VLSI circuit design and simulation [6, 16, 3]. However, it is only recently that people have begun to supply formal justification for the efficacy of these methods [15, 25]. In [25], Spielman and Teng used the results of Mihail [22] to show that the quality of the partition produced by the application of a certain spectral algorithm to a graph can be established by proving an upper bound on the Fiedler value of the graph (i.e., the second smallest eigenvalue of its Laplacian). They then provided an $O(1/n)$ bound on the Fiedler value of a *planar* graph with n vertices and bounded maximum degree. This showed that spectral methods can produce a cut of ratio $O(\sqrt{1/n})$ and a vertex bisector of size $O(\sqrt{n})$ in a bounded degree planar graph.

In this paper, we use the theory of circle packings and conformal mappings of compact Riemann surfaces to generalize these results to graphs of positive genus. We prove that the Fiedler value of a genus g graph of bounded degree is $O(g/n)$ and demonstrate that this is asymptotically tight, thereby resolving a conjecture of Spielman and Teng. We then apply this result to obtain a spectral partitioning algorithm that finds separators whose cut ratios are $O(\sqrt{g/n})$ and vertex bisectors of size $O(\sqrt{gn})$, both of which are optimal. To our knowledge, this provides the only

*Received by the editors September 30, 2004; accepted for publication (in revised form) April 18, 2005; published electronically February 21, 2006.

<http://www.siam.org/journals/sicomp/35-4/44724.html>

[†]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 (kelner@mit.edu).

truly practical algorithm for finding such separators and vertex bisectors for graphs of bounded genus and degree. While there exist other asymptotically fast algorithms for this, they all rely on being given an embedding of the graph in a genus g surface (e.g., [14]). It is not always the case that we are given such an embedding, and computing one is quite difficult. (In particular, computing the genus of a graph is NP-hard [27], and the best known algorithms for constructing such an embedding are either $n^{O(g)}$ [12] or polynomial in n but doubly exponential in g [11]. Mohar has found an algorithm that depends only linearly on n [21], but it has an uncalculated and very large dependence on g .) The excluded minor algorithm of Alon, Seymour, and Thomas [2] does not require an embedding of the graph, but the separators that it produces are not asymptotically optimal.

The question of whether there exists an efficient algorithm for providing asymptotically optimal cuts without such an embedding was first posed twenty years ago by Gilbert, Hutchinson, and Tarjan [14].¹ We resolve this question here, as our algorithm proceeds without any knowledge of an embedding of the graph, instead relying only on simple matrix manipulations of the adjacency matrix of the graph. While the analysis of the algorithm requires some somewhat involved mathematics, the algorithm itself is quite simple, and it can be implemented in just a few lines of Matlab code. In fact, the algorithm is only a slight modification of the spectral heuristics for graph partitioning that are widely deployed in practice without any theoretical guarantees.

We believe that the techniques that we employ to obtain our eigenvalue bounds are of independent interest. To prove these bounds, we make what is perhaps the first real use of the theory of circle packings and conformal mappings of positive genus Riemann surfaces in the computer science literature. This is a powerful theory, and we believe that it will be useful for addressing other questions in spectral and topological graph theory.

The structure of the paper is as follows. In section 2, we provide the necessary background in graph theory and spectral partitioning, and we state our main results. In section 3, we provide a brief outline of our proof techniques. In section 4, we review the basic theory of circle packings on compact Riemann surfaces. We then use this theory in sections 5 and 6 to prove our main results.

2. Background in graph theory and spectral partitioning. In this section we provide the basic definitions and results from graph theory and spectral partitioning that we shall require in what follows.

2.1. Graph theory definitions. Throughout the remainder of this paper, let $G = (V, E)$ be a finite, connected, undirected graph with n vertices, m edges, and no loops. In this section, we shall define two objects associated with G : its *Laplacian* and its *genus*.

Let the *adjacency matrix* $A(G)$ be the $n \times n$ matrix whose (i, j) th entry equals 1 if $(i, j) \in E$ and equals 0 otherwise. Let $D(G)$ be the $n \times n$ diagonal matrix whose i th diagonal entry equals the degree of the i th vertex of G .

DEFINITION 2.1. *The Laplacian $L(G)$ is the $n \times n$ matrix given by*

$$L(G) = D(G) - A(G).$$

Since $L(G)$ is symmetric, it is guaranteed to have an orthonormal basis of real eigenvectors and exclusively real eigenvalues. Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenval-

¹Djidjev claimed in a brief note to have such an algorithm [10], but it has never appeared in the literature.

ues of $L(G)$, and let v_1, \dots, v_n be a corresponding orthonormal basis of eigenvectors. For any G , the all-ones vector will be an eigenvector of eigenvalue 0. It is not difficult to see that all of the other eigenvalues will always be nonnegative, so that $v_1 = (1, \dots, 1)^T$ and $\lambda_1 = 0$.

There has been a great deal of work relating the eigenvalues of $L(G)$ to the structure of G . In the present paper, we shall concern ourselves exclusively with λ_2 , also known as the *algebraic connectivity* or *Fiedler value* of G . We call the vector v_2 the *Fiedler vector* of G . As we shall see in section 2.2, the Fiedler value of a graph is closely related to how well connected the graph is.

A different measure of the connectivity of a graph is provided by its *genus*, which measures the complexity of the simplest orientable surface on which the graph can be embedded so that none of its edges cross. Standard elementary topology provides a full classification of the orientable surfaces without boundary. Informally, they are all obtained by attaching finitely many “handles” to the sphere, and they are fully topologically classified (i.e., up to homeomorphism) by the number of such handles. This number is called the *genus* of the surface. The genus 0, 1, 2, and 3 surfaces are shown in Figure 2.1.




FIG. 2.1. *The surfaces of genus 0, 1, 2, and 3.*

DEFINITION 2.2. *The genus g of a graph G is the smallest integer such that G can be embedded on a surface of genus g without any of its edges crossing one another.*

In particular, a planar graph has genus 0. By making a separate handle for each edge, it is easy to see that $g = O(m)$, where m is the number of edges in G .

Using these definitions, we can now state our main technical result, as follows.

THEOREM 2.3. *Let G be a graph of genus g and bounded degree. Its Fiedler value obeys the inequality*

$$\lambda_2 \leq O\left(\frac{g}{n}\right),$$

and this is asymptotically tight.

The constant in this bound depends on the degree of the graph. The proof that we provide yields a polynomial dependence on the degree, but no effort is made to optimize this polynomial. Finding the optimal such dependence is an interesting open question.

2.2. Spectral partitioning. We recall that a *partition* of a graph G is a decomposition $V = A \cup \bar{A}$ of the vertices of G into two disjoint subsets. For such a partition, we let $\delta(A)$ be the set of edges (i, j) such that $i \in A$ and $j \in \bar{A}$, and we call $|\delta(A)|$ the *cut size* of our partition. The *ratio* of our partition is defined to be

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

If our partition splits the graph into two sets that differ in size by at most one, we call it a *bisection*.

Spectral methods aim to use the Fiedler vector to find a partition of the graph with a good ratio. A theorem that begins to address why these work was proven by Mihail and restated in a more applicable form by Spielman and Teng, as follows.

THEOREM 2.4 (see [22, 25]). *Let G have maximum degree Δ . For any vector x that is orthogonal to the all-ones vector, there is a value s so that the partition of G into $\{i : x_i \leq s\}$ and $\{i : x_i > s\}$ has ratio at most*

$$\sqrt{2\Delta \frac{x^T L(G)x}{x^T x}}.$$

If x is an eigenvector of $L(G)$, the fraction $\frac{x^T L(G)x}{x^T x}$ is equal to its eigenvalue. So, if we find the eigenvector with eigenvalue λ_2 , we will thus quickly be able to find a partition of ratio $\sqrt{2\Delta\lambda_2}$. By Theorem 2.3, finding the second eigenvector of the Laplacian thus allows us to find a partition of ratio $O(\sqrt{g/n})$ for a graph of bounded degree. There is no guarantee that this partition has a similar number of vertices in each of the two sets. However, a theorem of Lipton and Tarjan [19] implies that a simple method based on repeated application of this algorithm can be used to give a biselector of size $O(\sqrt{gn})$.

For every g , Gilbert, Hutchinson, and Tarjan exhibited a class of bounded degree graphs that have no biselectors smaller than $O(\sqrt{gn})$ [14]. This implies that our algorithm gives the best results possible, in general. Furthermore, it establishes the asymptotic tightness of our eigenvalue bound, as a smaller bound would show that every genus g graph has a partition of size $o(\sqrt{gn})$.

Putting all of this together yields our main algorithmic result, the following.

THEOREM 2.5. *Let G be a genus g graph of bounded maximum degree. There is a polynomial time algorithm that produces cuts of ratio $O(\sqrt{g/n})$ and vertex biselectors of size $O(\sqrt{gn})$ in G , and both of these values are optimal.*

All that remains of the proof of Theorem 2.5 is the eigenvalue bound set forth in Theorem 2.3, which is the goal of the remainder of this paper.

3. Outline of the proof of Theorem 2.3. The proof of Theorem 2.3 necessitates the introduction of a good deal of technical machinery. Before launching into several pages of definitions and background theorems, we feel that a brief roadmap of where we're going will be helpful.

The basic motivation for our approach comes from an observation made by Spielman and Teng [25]. They noted that one can obtain bounds on the eigenvalues of a graph G from a nice representation of G on the unit sphere in \mathbb{R}^3 , known as a *circle packing* for G . This is a presentation of the graph on the sphere so that the vertices are the centers of a collection of circles, and the edges between vertices correspond to tangencies of their respective circles, as shown in Figure 4.1. Only planar graphs can be embedded as such if we require that the circles have disjoint interiors. However, if we allow the circles to overlap, as shown in Figure 4.2, we can represent nonplanar graphs as well. This will give rise to a weaker bound in which the eigenvalue bound is multiplied by the maximum number of circles containing a given point (i.e., the number of layers of circles on the sphere).

There is a well developed theory of circle packings, both on the sphere and on higher genus surfaces. The portions of it that we shall use will tell us two main things:

1. We can realize our graph as a circle packing of circles with disjoint interiors on some genus g surface.
2. The theory of discrete circle packings can be thought of as a discrete analogue of classical complex function theory, and many of the results of the latter carry

over to the former.

In classical complex analysis, you can put a complex analytic structure on a genus g surface to obtain a Riemann surface. Any genus g Riemann surface has a map to the sphere that is almost everywhere k -to-one for $k = O(g)$, with only $O(g)$ bad points at which this fails. With this as motivation, we shall try to use the representation of G as a circle packing on a genus g surface to obtain a representation of it as a circle packing on the sphere with $O(g)$ layers.

Unfortunately, the discrete theory is more rigid than the continuous one, and finding such a representation will turn out to be impossible. Instead, we shall actually pass to the continuous theory to prove our result. To do this, we shall provide a subdivision lemma that shows that it suffices to prove Theorem 2.3 for graphs that have circle packings with very small circles. We shall then show that the smooth map that we have from the Riemann surface to the sphere will take *almost* all of the circles of our circle packing to curves on the sphere that are *almost* circles. We will then show that this representation of our graph as an approximate circle packing is enough to provide our desired bounds.

4. Introduction to circle packings. Our proof of Theorem 2.3 operates by obtaining a nice geometric realization of G . We obtain this realization using the theory of circle packings. In this section, we shall review the basics of circle packing theory and quote the main results that our proof will employ. For a more comprehensive treatment of this theory and a historical account of its origins, see [26].

Loosely speaking, a circle packing is a collection of circles on a surface with a given pattern of tangencies. We remark at the outset that the theory that we are discussing *is not the same* as the classical theory of sphere packing. Our theory is concerned with the combinatorics of the tangency patterns, *not* with the maximum number of circles that one can fit in a small region. The coincidence of nomenclature is just an unfortunate historical accident.




FIG. 4.1. A univalent circle packing with its associated graph.

4.1. Planar circle packings. For simplicity, we begin by discussing circle packings in the plane.

DEFINITION 4.1. A planar circle packing \mathcal{P} is a finite collection of (possibly overlapping) circles C_1, \dots, C_n of respective radii r_1, \dots, r_n in the complex plane \mathbb{C} . If all of the C_i have disjoint interiors, we say that \mathcal{P} is univalent.

The associated graph $A(\mathcal{P})$ of \mathcal{P} is the graph obtained by assigning a vertex v_i to each circle C_i and connecting v_i and v_j by an edge if and only if C_i and C_j are mutually tangent.

This is illustrated in Figures 4.1 and 4.2.

We thus associate a graph with every circle packing. It is clear that every graph




FIG. 4.2. A nonunivalent circle packing with its associated graph.

associated with a univalent planar circle packing is planar. A natural question to ask is whether every planar graph can be realized as the associated graph of some planar circle packing. This is answered in the affirmative by the Koebe–Andreev–Thurston theorem [18, 1, 28].

THEOREM 4.2 (Koebe–Andreev–Thurston). *Let G be a planar graph. There exists a planar circle packing \mathcal{P} such that $A(\mathcal{P}) = G$.*

This theorem also contains a uniqueness result, but we have not yet developed the machinery to state it. We shall generalize this theorem in section 4.3, at which point we shall have the proper terminology to state the uniqueness part of the theorem.

We note that if we map the plane onto the sphere by stereographic projection, circles in the plane will be sent to circles on the sphere, so this theorem can be interpreted as saying that every genus 0 graph can be represented as a circle packing on the surface of a genus 0 surface. This suggests that we attempt to generalize this theorem to surfaces of higher genus. The theory of circle packings on surfaces of arbitrary genus acts in many ways like a discrete analogue of classical Riemann surface theory. As such, a basic background in Riemann surfaces is necessary to state or motivate many of its results. It is to this that we devote the next section.

4.2. A very brief introduction to Riemann surface theory. In this section, we provide an informal introduction to Riemann surface theory. Our goal is to provide geometric intuition, not mathematical rigor. We assume some familiarity with the basic concept of a manifold, as well as with the basic definitions of complex analysis. For a more complete exposition of the theory, see [13].

We recall that an n -dimensional manifold is a structure that looks locally like \mathbb{R}^n . More formally, we write our manifold M as a topological union of open sets S_i , each endowed with a homeomorphism $\varphi_i : S_i \rightarrow B_n$, where B_n is the ball $\{|x| < 1 \mid x \in \mathbb{R}^n\}$. Furthermore, we require a compatibility among these maps to avoid cusps and such. To this end, we mandate that the compositions $\varphi_j \circ \varphi_i^{-1} : \varphi_i(S_i \cap S_j) \rightarrow \varphi_j(S_i \cap S_j)$ be diffeomorphisms. The orientable 2-dimensional manifolds are precisely the genus g surfaces described above.

An n -dimensional complex manifold is the natural complex analytic generalization of this. We write our manifold M as a union of open sets S_i and endow each such set with a homeomorphism $\varphi_i : S_i \rightarrow B_{\mathbb{C}^n}$, where $B_{\mathbb{C}^n}$ is the complex unit ball $\{|x| < 1 \mid x \in \mathbb{C}^n\}$. Now, instead of requiring that the compositions of these functions obey a smooth compatibility condition, we require that they obey an analytic one: we demand that the compositions $\varphi_i \circ \varphi_j^{-1}$ be biholomorphic maps.

As such, an n -dimensional complex manifold M is a $2n$ -dimensional real manifold with additional complex analytic structure. This structure allows us to transfer over many of the definitions from standard complex analysis. The basic idea is that we

define these notions as before on the S_i , and the compatibility condition allows them to make sense as global definitions. In particular, if $M = (S_i^M, \phi_i^M)$ and $N = (S_j^N, \phi_j^N)$ are complex manifolds of the same dimension, we say that a function $f : M \rightarrow N$ is holomorphic if its restriction to a map $f_{ij} : S_i^M \rightarrow S_j^N$ is holomorphic for all i and j . Since the compositions $\varphi_i^M \circ (\varphi_j^M)^{-1}$ and $\varphi_i^N \circ (\varphi_j^N)^{-1}$ are holomorphic, this notion makes sense where the regions overlap.

DEFINITION 4.3. *A Riemann surface is a 1-dimensional complex manifold.*

In this paper, we shall take all of our Riemann surfaces to be compact. Since there is a natural way to orient the complex plane, we note that the complex structure can be used to define an orientation on the manifold. As such, all complex manifolds, and in particular all Riemann surfaces, are orientable. Compact Riemann surfaces are thus, topologically, 2-dimensional orientable real manifolds. Every compact Riemann surface is therefore topologically one of the genus g surfaces discussed above. The complex structure imposed by the φ_i , however, varies much more widely, and there are many different such structures that have the same underlying topological space.

Nothing in the definition of a Riemann surface supplies a metric on the surface. Indeed, there is no requirement that the different ϕ_i agree in any way about the distance between two points in their intersection. One can assign many different metrics to the surface. However, it turns out that there is way to single out a unique metric on the surface, called the *metric of constant curvature*. This allows us to supply an intrinsic notion of distance on any Riemann surface. In particular, this allows us to define a circle on our Riemann surface to be a simple closed curve that is contractible on the surface and all of whose points lie at a fixed distance from some center.

One particularly important Riemann surface that we shall consider is the Riemann sphere, which we denote $\widehat{\mathbb{C}}$. It is topologically a sphere. It should be thought of as being obtained by taking the complex plane and adjoining a single point called ∞ . One way of visualizing its relation to \mathbb{C} is to consider the stereographic projection away from the North Pole of a sphere onto a plane. The North Pole corresponds to ∞ , and the rest of the sphere corresponds to \mathbb{C} .

We recall from single variable complex analysis that the requirement that a map be analytic is quite a stringent one, and that it imposes a significant amount of local structure on the map. Let $f : \mathbb{C} \rightarrow \mathbb{C}$ be nonconstant and analytic in a neighborhood of the origin, and assume without loss of generality that $f(0) = 0$. There is some neighborhood of the origin in which f can be expressed as a power series $f(z) = a_1z + a_2z^2 + a_3z^3 + \dots$. If $a_1 \neq 0$, $f(z)$ is analytically invertible in some neighborhood of the origin, and thus it is locally an isomorphism. In particular, it is *conformal*—it preserves the angles between intersecting curves, and the image of an infinitesimal circle is another infinitesimal circle.

If $a_1 = 0$ and a_n is the first nonzero coefficient in its power series, f has a *branch point of order n* at the origin. In this case, f operates, up to a scale factor and lower order terms, like the function $f(z) = z^n$. This function is n -to-one on a small neighborhood of the origin, excluding the origin itself. It sends only 0 to 0, however. The preimages of the points in this small neighborhood thus trace out n different “sheets” that all intersect at 0. This confluence of sheets is the only sort of singularity that can appear in an analytic map. We note that the angles between curves intersecting at the branch point are not preserved, but they are instead divided by n .

This local behavior is identical for Riemann surfaces. From this, we can deduce that if $f : M \rightarrow N$ is an analytic map of Riemann surfaces, it has some well-defined

degree k . For all but finitely many points p in N , $\#f^{-1}(p) = k$. The preimage of each of these points looks like a collection of k sheets, and f has nonzero derivative at all of them. There exist some points $q \in M$ at which $f'(q) = 0$. At each such point there is a branch point, so the sheets intersect, and $f(q)$ has fewer than k preimages.

However, the global structure of Riemann surfaces provides further constraints on maps between them, and there are, generally speaking, very few functions $f : M \rightarrow N$ of a given degree. For example, topological arguments, using the local form of analytic maps described above, show that there are no degree 1 maps from the torus to the sphere, and no degree 2 maps from the genus 2 surface to the sphere.

There is a deep theory of maps of Riemann surfaces that describes rather precisely when a map of a given degree exists between two Riemann surfaces, and, if it exists, where and how such a map must branch. Of this theory we shall require only one main result, which is a direct corollary of the celebrated Riemann–Roch theorem.

THEOREM 4.4. *Let M be a Riemann surface of genus g . There exists an analytic map $f : M \rightarrow \widehat{\mathbb{C}}$ of degree $O(g)$ and with $O(g)$ branch points.*

4.3. Circle packings on surfaces of arbitrary genus. We now have the machinery in place to deal with general circle packings. Throughout this section, let G be a graph of genus g , and suppose that it is embedded on a genus g surface S so that none of its edges cross. The graph G divides S into faces. We say that G is a *fully triangulated* graph if all of these faces are triangles, in which case we say that it gives a *triangulation* of S . If G is not fully triangulated, one can clearly add edges to it to make it so. It will follow immediately from (5.2) in section 5 that this will only increase $\lambda_2(G)$, so we shall assume for convenience that G gives a triangulation of S . We are now ready to define our primary objects of study, as follows.

DEFINITION 4.5. *Let S be a compact Riemann surface endowed with its metric of constant curvature. A circle packing \mathcal{P} on S is a finite collection of (possibly overlapping) circles C_1, \dots, C_n of respective radii r_1, \dots, r_n on the surface of S . If all of the C_i have disjoint interiors, we say that \mathcal{P} is univalent.*

The associated graph $A(\mathcal{P})$ of \mathcal{P} is the graph obtained by assigning a vertex v_i to each circle C_i and connecting v_i and v_j by an edge if and only if C_i and C_j are mutually tangent. Alternatively, we say that \mathcal{P} is a circle packing for $A(\mathcal{P})$ on S .

The main result on circle packings that we shall use is the circle packing theorem, which is the natural extension of the Koebe–Andreiev–Thurston theorem to this more general setting. It was originally proven in a restricted form by Beardon and Stephenson [4] and then proven in full generality by He and Schramm [17].

THEOREM 4.6 (circle packing theorem). *Let G be a triangulation of a surface of genus g . There exists a Riemann surface S of genus g and a univalent circle packing \mathcal{P} such that \mathcal{P} is a circle packing for G on S . This packing is unique up to automorphisms of S .*

If G is embedded in a surface of genus g but is not fully triangulated, the Riemann surface and circle packing guaranteed by the theorem still exist, but they need not be unique.

The complex structure on the Riemann surface allows us to define the angle at which two edges of a face meet. If the points u, v , and w are the vertices of a face, we denote the angle between the edges \overline{uv} and \overline{vw} at v by $\langle uvw \rangle$. We can thus define the angle sum at a vertex to be $\sum \langle uvw \rangle$, where the sum is taken over all faces containing v . If \mathcal{P} is a univalent sphere packing, the angle sum at any vertex of $A(\mathcal{P})$ is clearly 2π .

In a nonunivalent circle packing, it is possible for the circles at a point to wrap around the point more than once. In the case of a nonunivalent circle packing, the

edges of its associated graph may intersect, but we can still define an associated triangulation of the surface—there just may be more than one triangle covering a given point. We can therefore compute the angle sum at a point. In this case, it need not be 2π . However, the circles must wrap around the vertex an integral number of times, so it must be some multiple $2\pi k$. (See Figure 4.2.) We then say that the vertex is a *discrete branch point of order k* .

These discrete branch points behave very much like the continuous branch points present on Riemann surfaces. In fact, there is an extensive theory that shows that a large portion of the theory of Riemann surfaces has an analogue in the discrete realm of circle packing. One can define maps of circle packings, just as one can define maps of Riemann surfaces. They consist of a correspondence of the circles on one surface to those on another in a way that commutes with tangency. While analytic maps send infinitesimal circles to infinitesimal circles, maps of circle packings send finite circles to finite circles. The analogue of branched covering maps in Riemannian geometry takes univalent circle packings and places them as nonunivalent circle packings on other surfaces. Unfortunately, these maps are somewhat rarer than their continuous analogues.

In particular, if we have a circle packing on a genus g surface S , there is no known analogue of the Riemann–Roch theorem, and thus no analogue of Theorem 4.4. We are therefore not guaranteed that there is a nonunivalent circle packing on the sphere carrying the same associated graph. Intuitively, this comes from the fact that the analytic maps from S to $\widehat{\mathbb{C}}$ are required to be branched over a very restricted locus of points. The discrete maps, however, can only be branched over the centers of circles. If there does not exist an admissible set of branch points among the centers of the circles, we will have difficulty constructing a discrete analytic map. This will lie at the root of many of the technical difficulties that we shall face in the remainder of this paper.

5. An eigenvalue bound. In this section, we prove Theorem 2.3. The proof will assume a technical lemma whose proof we shall postpone until section 6.

We begin by recalling the expression of the Fiedler value of G as a so-called *Rayleigh quotient*:

$$(5.1) \quad \lambda_2 = \min_{x \perp (1, \dots, 1)^T} \frac{x^T L(G)x}{x^T x}.$$

A straightforward calculation shows that for $x = (x_1, \dots, x_n)^T \in \mathbb{R}^n$,

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

so that (5.1) becomes

$$(5.2) \quad \lambda_2 = \min_{x \perp (1, \dots, 1)^T} \frac{\sum_{(i,j) \in E} (x_i - x_j)^2}{x^T x}.$$

As noted by Spielman and Teng [25], it follows easily from (5.2) that we can replace the scalar values x_i with vectors $v_i \in \mathbb{R}^k$, so that

$$(5.3) \quad \lambda_2 = \min \frac{\sum_{(i,j) \in E} \|v_i - v_j\|^2}{\sum_{i=1}^n \|v_i\|^2},$$

where the minimum is taken over all sets of n -vectors such that $\sum v_i = (0, \dots, 0)^T$ and such that at least one of the v_i is nonzero.

The general goal is thus to find a set of v_i that gives a small value for this quotient. The v_i that we use will almost be the centers of a nonunivalent circle packing on the unit sphere $S^2 \subseteq \mathbb{R}^3$. The efficacy of this follows from the following theorem, which follows easily from the work of Spielman and Teng [25].

THEOREM 5.1. *Let \mathcal{P} be a circle packing on the sphere $S^2 = \{x \in \mathbb{R}^3 \mid \|x\|^2 = 1\}$ so that the graph $A(\mathcal{P})$ has no vertex of degree greater than Δ . Suppose further that the packing is of degree k , so that no point on the sphere is contained in the interior of more than k circles, and that the centroid of the centers of the circles is the origin. Then the Fiedler value*

$$\lambda_2(A(\mathcal{P})) \leq O\left(\frac{\Delta k}{n}\right).$$

Proof. The proof follows from (5.3). Let the circles be C_1, \dots, C_n , and let the corresponding radii be r_1, \dots, r_n . Let $v_i \in \mathbb{R}^3$ be the x, y , and z coordinates of the center of the i th circle. The sum $\sum v_i = 0$ by assumption, so λ_2 is less than or equal to the fraction in (5.3). Since all of the v_i are on the unit sphere, we have $\sum \|v_i\|^2 = n$, so it just remains to bound the numerator. If there is an edge (i, j) , the two circles C_i and C_j must be mutually tangent, so that $\|v_i - v_j\|^2 \leq (r_i + r_j)^2 \leq 2(r_i^2 + r_j^2)$. It thus follows that

$$\sum_{(i,j) \in E} \|v_i - v_j\|^2 \leq \sum_{(i,j) \in E} 2(r_i^2 + r_j^2) \leq 2\Delta \sum_{i=1}^n r_i^2.$$

However, the total area of all of the circles is less than or equal to k times the area of the sphere, since the circle packing is of degree k . We thus have that $\sum_{i=1}^n r_i^2 \leq O(k)$, from which the desired result follows. \square

This suggests that we use the circle packing theorem (Theorem 4.6) to embed our graph on a genus g surface and then try to use some analogue of Theorem 4.4 to obtain a branched circle packing on the sphere of degree $O(g)$. Unfortunately, as previously noted, such a circle packing need not exist, due to the restrictiveness of the discrete theory. Thus, we shall instead show that a certain subdivision process on our graph does not significantly decrease $n\lambda_2$. We shall then show that performing this subdivision enough times causes our discrete circle packing to approximate a continuous structure on the Riemann surface, at which point we can use the continuous theory in addition to the discrete one.

The refinement procedure that we shall use is called “hexagonal refinement.” It operates on a triangulation of a surface by replacing each triangle with four smaller triangles, as shown in Figure 5.1. This process produces another triangulation of the same surface, so we can iterate it arbitrarily many times.

LEMMA 5.2. *Let G be a graph with n vertices, m edges, and maximum degree Δ that triangulates some surface without boundary, and let G' be the graph with n' vertices and m' edges obtained by performing k successive hexagonal refinements on G . Then*

$$n\lambda_2(G) \leq C(\Delta)n'\lambda_2(G').$$

For the sake of continuity, we defer this proof to section 6.




FIG. 5.1. The hexagonal subdivision procedure applied to a triangulation with two triangles.

The refinement process replaces each triangle in our graph with four smaller triangles. If all of the original triangles remained the same size and shape, this would imply that performing enough hexagonal refinements would give rise to a circle packing whose circles have arbitrarily small radii. However, it is possible for the original triangles to change size and shape as we refine, so this is no longer obvious. Nevertheless, it remains true, as shown by the following lemma.

LEMMA 5.3. *Let G be a graph that triangulates a genus g Riemann surface without boundary, and let $G^{(k)}$ be the graph obtained by performing k hexagonal refinements on G . For every $\epsilon > 0$, there exists some k_ϵ so that for all $\ell \geq k_\epsilon$ every circle in $G^{(\ell)}$ has radius less than ϵ .*

Proof. This was essentially proven by Rodin and Sullivan [23]; their proof, however, was stated for only the genus 0 case. The precise statement above was proven by Bowers and Stephenson [5]. \square

We get a new Riemann surface for each iteration of the refinement procedure. It is intuitive that, as the number of iterations grows and the circles in the refined graph get arbitrarily small, the Riemann surfaces will somehow converge, and the embedding of the graph on these Riemann surfaces will somehow stabilize. This can be made formal by the following lemma.

LEMMA 5.4. *Let G be a graph that triangulates a genus g compact Riemann surface without boundary, let $G^{(k)}$ be the result of performing k hexagonal refinements on G , and let $S^{(k)}$ be the Riemann surface on which $G^{(k)}$ is realized as a circle packing. Further, let $h_k : S^{(k)} \rightarrow S^{(k+1)}$ be the map that takes a triangle to its image under the subdivision procedure by the obvious piecewise-linear map. The sequence of surfaces $\{S^{(k)}\}$ converges in the moduli space of genus g surfaces, and the sequence of maps $\{h_k\}$ converges to the identity.*

Proof. This has been shown by Bowers and Stephenson [5]. \square

We shall also require one last definition, as follows.

DEFINITION 5.5. *Let $f : X \rightarrow Y$ be a map between two locally Euclidean metric spaces. The quantity*

$$H_f(x, r) = \frac{\max_{|x-y|=r} |f(x) - f(y)|}{\min_{|x-y|=r} |f(x) - f(y)|} - 1$$

is called the radius r distortion of f at x .

We are now finally ready to prove Theorem 2.3.

Proof of Theorem 2.3. Using the circle packing theorem (Theorem 4.6), realize the graph $G = G^{(0)}$ as a circle packing on some Riemann surface S of genus g . Let $G^{(k)}$ be the result of performing k hexagonal refinements on G , and let $S^{(k)}$ be the Riemann surface on which it can be realized as a circle packing. By Theorem 4.4, there exists an analytic map $f^{(k)}$ from $S^{(k)}$ to the Riemann sphere of degree $O(g)$

and with $O(g)$ branch points. Embed the Riemann sphere as the unit sphere in \mathbb{R}^3 using the conformal map given by inverse stereographic projection. By the work of Spielman and Teng (Theorem 9 of [25]), postcomposing with a Möbius transformation allows us to assume, without loss of generality, that the centroid of the images of the vertices of each $G^{(k)}$ under $f^{(k)}$ is the origin. By Lemma 5.4, the $S^{(k)}$ converge to some surface $S^{(\infty)}$, and the $f^{(k)}$ can be chosen so as to converge to some continuous limit map $f^{(\infty)}$.

By Lemma 5.2, it suffices to prove the theorem for an arbitrarily fine hexagonal refinement of the original graph. Away from its branch points, a map of Riemann surfaces is conformal, meaning it sends infinitesimal circles to infinitesimal circles. In particular, given a map $f : S \rightarrow \widehat{\mathbb{C}}$, the compactness of S guarantees that for every $\epsilon, \kappa > 0$ there exists a $\delta > 0$ so that the radius δ' distortion $H_f(x, \delta')$ is less than ϵ for every x that is at least distance κ from any branch point and any $\delta' \leq \delta$. In fact, by the convergence results of the last paragraph, there exist some N and δ such that this holds for every $f^{(k)}$ with $k > N$. Fix ϵ and κ , and let δ and N be chosen so that this is true. By possibly increasing N if necessary, we can assume by Lemma 5.3 that all of the circles on $S^{(k)}$ have radius at most δ for all $k > N$.

Let k be at least N . We shall break $S^{(k)}$ into two parts, $S^{(k)} = S_1^{(k)} \cup S_2^{(k)}$, as follows. Construct a ball of radius κ around each branch point of $f^{(k)}$, and let $S_2^{(k)}$ be the union of these balls. Let $S_1^{(k)}$ be the complement $S^{(k)} \setminus S_2^{(k)}$.

We can now use (5.3) to bound λ_2 , just as in the proof of Theorem 5.1. Let $G^{(k)}$ have n_k vertices. The denominator of (5.3) is equal to n_k , so it suffices to bound the numerator. We shall separately consider the circles contained entirely in $S_1^{(k)}$ and those that intersect $S_2^{(k)}$.

We begin with the circles contained in $S_1^{(k)}$. Every circle of the packing gets mapped by f to some connected region on $\widehat{\mathbb{C}}$, and there are at most $O(g)$ such regions covering any point of the sphere. Let C be a circle in $S_1^{(k)}$, let D be the diameter function, which takes a region to the length of the longest geodesic it contains, and let A be the area function. Since the radius δ distortion of f inside of $S_1^{(k)}$ is at most ϵ , and the radius of C is at most δ , the ratio $D^2(f(C))/A(f(C))$ is at most $O(1 + \epsilon)$. Using the same argument as in the proof of Theorem 5.1, the vertex at the center of a circle C cannot contribute more than $O(dD^2(f(C)))$ to the sum, and the total area of the regions from $S_1^{(k)}$ cannot exceed $O(g)$, so the total contribution to the numerator of the vertices in $S_1^{(k)}$ cannot be more than $O(dg(1 + \epsilon))$.

If this were the only term in the numerator, we could complete the proof by setting ϵ to be a constant. It thus remains to show that the contribution from the circles intersecting $S_2^{(k)}$ can be made small. To do this, we need only show that the contribution $\theta^{(k)}(x)$ to the numerator per unit area at a point x from these circles remains bounded as we subdivide, since we can make the area of $S_2^{(k)}$ arbitrarily small by sending κ to zero, and thus the area of the circles intersecting $S_2^{(k)}$ will go to zero as k goes to infinity and the circles get arbitrarily small. Our argument is similar to one used by McCaughan to analyze the recurrence of random walks on circle packings [20].

Let $x_i, i = 1, \dots, 3$, be the coordinate functions on \mathbb{R}^3 , and let $f^{(k)*}x_i$ be their pullbacks along $f^{(k)}$ to $S^{(k)}$. (That is, if y is a point on $S^{(k)}$, $f^{(k)*}x_i(y) = x_i(f^{(k)}(y))$.) In addition, let $C_1^{(k)}$ and $C_2^{(k)}$ be a pair of adjacent circles in $S_2^{(k)}$ with respective radii $r_1^{(k)}$ and $r_2^{(k)}$ and respective centers $c_1^{(k)}$ and $c_2^{(k)}$. The contribution of the correspond-

ing edge in $G^{(k)}$ to the numerator of (5.3) will be

$$(5.4) \quad \left\| \left(f^{(k)*} x_i(c_1^{(k)}) \right)_{i=1}^3 - \left(f^{(k)*} x_i(c_2^{(k)}) \right)_{i=1}^3 \right\|^2 = \sum_{i=1}^3 \left(f^{(k)*} x_i(c_1^{(k)}) - f^{(k)*} x_i(c_2^{(k)}) \right)^2.$$

The distance between $c_1^{(k)}$ and $c_2^{(k)}$ equals $r_1^{(k)} + r_2^{(k)}$. As k goes to infinity, the radii $r_1^{(k)}$ and $r_2^{(k)}$ both go to zero, by Lemma 5.3. By the smoothness of the $f^{(k)}$, their convergence to $f^{(\infty)}$, and the compactness of their domains, we can approximate each term on the right-hand side of (5.4) arbitrarily well by its first order approximation, so that

$$(5.5) \quad \left(f^{(k)*} x_i(c_1^{(k)}) - f^{(k)*} x_i(c_2^{(k)}) \right)^2 \leq (1 + o(1))(r_1^{(k)} + r_2^{(k)})^2 \|\nabla f^{(k)*} x_i(c_1^{(k)})\|^2$$

as k goes to infinity and the distance between $c_1^{(k)}$ and $c_2^{(k)}$ shrinks to zero.

The right-hand side of (5.5) is bounded above by

$$(5.6) \quad (2 + o(1))[(r_1^{(k)})^2 + (r_2^{(k)})^2] \|\nabla f^{(k)*} x_i(c_1^{(k)})\|^2 = O(1)[(r_1^{(k)})^2 \|\nabla f^{(k)*} x_i(c_1^{(k)})\|^2 + (r_2^{(k)})^2 \|\nabla f^{(k)*} x_i(c_2^{(k)})\|^2].$$

The degree of our graph is bounded, so every vertex appears in at most a constant number of edges. If we sum the right-hand side of (5.6) over all of the edges in our graph, the total contribution of terms involving a fixed circle of radius r centered at c is thus bounded above by

$$O(1)r^2 \|\nabla f^{(k)*} x_i(c)\|^2,$$

so the contribution per unit area is bounded above by

$$O(1) \|\nabla f^{(k)*} x_i(c)\|^2.$$

This clearly remains bounded as k goes to infinity and $f^{(k)}$ approaches $f^{(\infty)}$. It thus follows that the contribution to the numerator of (5.3) of the vertices in $S_2^{(k)}$ tends to zero as k goes to infinity and κ is made arbitrarily small. By setting ϵ to be a constant and sending κ to zero, Theorem 2.3 follows. \square

6. The proof of Lemma 5.2. In this section, we shall prove Lemma 5.2. In proving this bound, it will be convenient to consider the following weighted form of the Laplacian.

DEFINITION 6.1. *The weighted Laplacian $\mathcal{L}^W(G)$ of a graph G is the matrix*

$$\mathcal{L}^W(G) = W^{-1/2} L(G) W^{-1/2},$$

where $L(G)$ is the Laplacian of G , and W is a diagonal matrix whose i th diagonal entry w_i is strictly positive for all i .

We shall denote the eigenvalues of $\mathcal{L}^W(G)$ by $\tilde{\lambda}_1^W(G) \leq \dots \leq \tilde{\lambda}_n^W(G)$ and the corresponding eigenvectors by $\tilde{v}_1^W(G) \dots \tilde{v}_n^W(G)$. A straightforward calculation shows

that the weighted Laplacian has $\tilde{\lambda}_1^W = 0$ and $\tilde{v}_1^W = W^{1/2}\mathbf{1}$. Our main quantity of interest will be $\tilde{\lambda}_2^W(G)$, which we can compute using a weighted analogue of the Rayleigh quotient:

$$(6.1) \quad \tilde{\lambda}_2^W = \min_{x \perp W\mathbf{1}} \frac{\sum_{(i,j) \in E} (x_i - x_j)^2}{\sum_i x_i^2 w_i}.$$

The second eigenvector $\tilde{v}_2^W(G)$ equals $W^{1/2}x$, where x is the vector that achieves the minimum in (6.1).

If all of the weights are $\Theta(1)$, standard linear algebra shows that $\lambda_2(G)$ and $\tilde{\lambda}_2^W(G)$ differ by at most a constant factor, so proving a bound on one implies a bound on the other. (See Chung’s book [9] for detailed proofs of the above facts and for other foundational information about the weighted Laplacian.)

Before we can proceed to the body of the proof of Lemma 5.2, we require two fairly general technical lemmas about independent random variables.

LEMMA 6.2. *Let a_1, \dots, a_n be independent real-valued random variables, possibly drawn from different probability distributions. Let $w_1, \dots, w_n \in \mathbb{R}^+$ be strictly positive constants. If the expectation $\mathbb{E}[\sum_i w_i a_i] = 0$, then*

$$\mathbb{E} \left[\left(\sum_j w_j a_j \right)^2 \right] \leq \mathbb{E} \left[\sum_j w_j^2 a_j^2 \right].$$

Proof. This follows by expanding the left-hand side:

$$\begin{aligned} \mathbb{E} \left[\left(\sum_j w_j a_j \right)^2 \right] &= \mathbb{E} \left[\sum_i w_i^2 a_i^2 \right] + \mathbb{E} \left[\sum_i w_i a_i \left(\sum_{j \neq i} w_j a_j \right) \right] \\ &= \mathbb{E} \left[\sum_i w_i^2 a_i^2 \right] + \sum_i -(\mathbb{E}[w_i a_i])^2 \\ &\leq \mathbb{E} \left[\sum_j w_j^2 a_j^2 \right], \end{aligned}$$

where the second equality follows from the independence of the variables and the fact that the sum of their expectations is zero. \square

We shall now use this lemma to establish our second lemma, which is the one that will actually appear in our main proof.

LEMMA 6.3. *Let a_1, \dots, a_n be independent real-valued random variables, possibly drawn from different probability distributions, and let $w_1, \dots, w_n \in \mathbb{R}^+$ be strictly positive constants such that $\mathbb{E}[\sum_i w_i a_i] = 0$. Let $a = (a_1, \dots, a_n)$, and let $w_{\max} = \max_i w_i$. Further let*

$$b = \left(\frac{1}{\sum_i w_i} \sum_i w_i a_i \right) \mathbf{1} \quad \text{and} \quad c = a - b.$$

Then

$$\mathbb{E} \left[\sum_i w_i c_i^2 \right] \geq \left(1 - \frac{w_{\max}}{\sum_i w_i} \right) \mathbb{E} \left[\sum_i w_i a_i^2 \right].$$

Proof. This follows by direct calculation:

$$\begin{aligned}
\mathbb{E} \left[\sum_i w_i c_i^2 \right] &= \mathbb{E} \left[\sum_i w_i \left(a_i - \frac{1}{\sum_j w_j} \left(\sum_j w_j a_j \right) \right)^2 \right] \\
&= \mathbb{E} \left[\sum_i w_i a_i^2 \right] + \frac{1}{(\sum_i w_i)^2} \mathbb{E} \left[\sum_i w_i \left(\sum_j w_j a_j \right)^2 \right] \\
&\quad - \frac{2}{\sum_i w_i} \mathbb{E} \left[\sum_i w_i a_i \left(\sum_j w_j a_j \right) \right] \\
&= \mathbb{E} \left[\sum_i w_i a_i^2 \right] + \frac{1}{\sum_i w_i} \mathbb{E} \left[\left(\sum_j w_j a_j \right)^2 \right] - \frac{2}{\sum_i w_i} \mathbb{E} \left[\left(\sum_j w_j a_j \right)^2 \right] \\
&= \mathbb{E} \left[\sum_i w_i a_i^2 \right] - \frac{1}{\sum_i w_i} \mathbb{E} \left[\left(\sum_j w_j a_j \right)^2 \right] \\
&\geq \mathbb{E} \left[\sum_i w_i a_i^2 \right] - \frac{1}{\sum_i w_i} \mathbb{E} \left[\sum_j w_j^2 a_j^2 \right] \\
&= \mathbb{E} \left[\sum_i \left(1 - \frac{w_i}{\sum_j w_j} \right) w_i a_i^2 \right] \\
&\geq \left(1 - \frac{w_{\max}}{\sum_i w_i} \right) \mathbb{E} \left[\sum_i w_i a_i^2 \right],
\end{aligned}$$

where second-to-last inequality follows from Lemma 6.2. \square

We are now prepared to prove Lemma 5.2.

Proof of Lemma 5.2. Let $G = (V_G, E_G)$ be the original graph, and let $G' = (V_{G'}, E_{G'})$ be the graph that results from performing k successive hexagonal refinements on G . The embeddings into surfaces endow both G and G' with triangulations; let T_G and $T_{G'}$ be the respective sets of triangles in these triangulations. There is a natural inclusion $\iota : V_G \hookrightarrow V_{G'}$, since the subdivision procedure only adds vertices to the original set. There is also a map $\eta : T_{G'} \rightarrow T_G$ that takes a triangle from the subdivided graph to the one in the original graph from which it arose. For a vertex v in either graph, let $N(v)$ be the set of triangles containing it. For a vertex $w \in V_G$, let $P(w) = \eta^{-1}(N(w))$ be the set of triangles in $T(G')$ taken by η to elements of $N(w)$. (See Figure 6.1.)

Our proof will proceed by producing a randomized construction of a subgraph H of G' . Given a vector that assigns a value to every vertex of G' , we can obtain such a vector on H by restriction. We shall also show how to use such a vector on H to construct such a vector on G . The vectors on the different graphs will give rise to Rayleigh quotients on the graphs (some of which will be weighted), where the Rayleigh quotients for G and H will depend on the random choices made in the construction of H . By relating the terms in the different Rayleigh quotients, we shall then provide a probabilistic proof that there exists an H that gives rise to a small Rayleigh quotient




FIG. 6.1. A subdivided graph, with $P(w)$ and $N(w)$ shaded for a vertex w .

on G , which will suffice to prove our desired bound.

H will be produced by randomly choosing a representative in $V_{G'}$ for each vertex in V_G and representing every edge in E_G by a randomly chosen path in G' between the representatives of its endpoints.

We first construct the map $\pi_V : V_G \rightarrow V_{G'}$ that chooses the representatives of the vertices. For each $v \in V_G$ we choose $\pi_V(v)$ uniformly at random from the vertices contained in $P(v)$ that are at least as close to $\iota(v)$ as to $\iota(w)$ for any other $w \in V_G$. Vertices in $P(v)$ that are equally close to $\iota(v)$ and $\iota(w)$ should be arbitrarily assigned to either v or w , but not both.

We now construct π_E , which maps edges in E_G to paths in G' . Let $e = (v_1, v_2)$ be an edge in G , and let w_1 and w_2 equal $\pi_V(v_1)$ and $\pi_V(v_2)$, respectively. The two neighborhoods in G , $N(v_1)$ and $N(v_2)$, share exactly two triangles, t_1 and t_2 . Let x be a vertex randomly chosen from the vertices in $\eta^{-1}(t_1 \cup t_2)$. We shall construct a path from each w_i ($i = 1, 2$) to x , so that their composition gives a path from w_1 to w_2 . We shall use the same construction for each, so, without loss of generality, we shall just construct the path from w_1 to x .

Both w_1 and x are in $P(v_1)$, and we give a general procedure for constructing a path between any two such vertices. The images under the inclusion ι of the triangles in $N(v_1)$ encircle $\iota(v_1)$. Suppose w_1 is contained in T_1 , and x is contained in T_2 . Traversing the triangles in a clockwise order from T_1 to T_2 gives one list of triangles, and traversing in a counterclockwise order gives another. Let $T_1, Q_1, \dots, Q_\ell, T_2$ be the shorter of these two lists, with a random choice made if the two lists are the same length. Choose a random vertex a_i in each Q_i , and let $a_0 = w_1$ and $a_{\ell+1} = x$. We thus have a vertex representing each triangle in the list. Our path will consist of a sequence of segments from each representative to the next.

Note that all of the triangles are distinct, except if $T_1 = T_2$ and the list is of length 2. We suppose for now that we have two vertices a_i and a_{i+1} in distinct triangles, and we deal with the degenerate case later. The two triangles in question




FIG. 6.2. An illustration of how the grid graph exists as a subgraph of the union of two adjacent subdivided triangles.




FIG. 6.3. The entire construction illustrated for a given edge of the original graph.

are adjacent, and their union contains a grid graph as a subgraph. (See Figure 6.2.) Given two vertices in a grid, there is a unique path between them that one obtains by first moving horizontally and then vertically, and another that one obtains by moving vertically and then horizontally. (These two coincide if there is a line connecting the two points.) Randomly choose one of these two paths. This is the path connecting a_i to a_{i+1} . If a_i and a_{i+1} lie in the same triangle, randomly choose one of the two adjacent triangles to form a grid, and then use the above construction. Composing the paths between each a_i and a_{i+1} completes the construction of π_E . The entire construction is illustrated in Figure 6.3.

We now consider the Rayleigh quotients for the three graphs that we have con-

structed. After k hexagonal refinements, every edge in G is split into $r = 2^k$ pieces, every triangle gets replaced with r^2 smaller triangles, and the number of vertices grows quadratically in r . A vector $y \in \mathbb{R}^{|V_{G'}|}$ that assigns a value to each vertex in G' gives the Rayleigh quotient

$$R(G') = \frac{\sum_{(i,j) \in E_{G'}} (y_i - y_j)^2}{y^T y}.$$

This induces a vector on the vertices of H by restriction. The probability, taken over the random choices in the construction of π_V and π_E , that a given edge of G' appears on the path representing a given edge e of G is zero if it is not in $P(\alpha)$ with α equal to one of the endpoints of e , and at most $O(1/r)$ otherwise. Since the maximum degree of a vertex in G is assumed constant, the expected number of times that a given edge of G' occurs in H is $O(1/r)$. Every vertex in G' is selected as a representative of a vertex in G with probability $\Theta(1/r^2)$. It thus follows that

$$(6.2) \quad \mathbb{E} \left[\sum_{(i,j) \in E_H} (y_i - y_j)^2 \right] \leq O\left(\frac{1}{r}\right) \sum_{(i,j) \in E_{G'}} (y_i - y_j)^2$$

and

$$(6.3) \quad \mathbb{E} \left[\sum_{i \in V_G} w_i y_{\pi_V(i)}^2 \right] = \Theta\left(\frac{1}{r^2}\right) \sum_{i \in V_{G'}} y_i^2,$$

where the expectations are taken over the random choices in the construction of (π_V, π_E) , and the w_i are any weights that are bounded above and below by positive constants.

Let \bar{y} be the vector in $\mathbb{R}^{|V_G|}$ whose i th coordinate is $y_{\pi_V(i)}$. Each coordinate \bar{y}_i of \bar{y} is chosen independently from a distinct set S_i of the coordinates of y , and every coordinate is contained in one of these sets. Let $s_i = |S_i|$, let $s_{\min} = \min_i s_i$, and take W to be the diagonal matrix whose i th diagonal entry w_i equals s_i/s_{\min} . The probability that a given vertex in S_i is selected equals $1/s_i$, so we have that

$$\mathbb{E} \left[\sum_{j \in V_G} w_j \bar{y}_j \right] = \sum_{k \in V_{G'}} y_k = 0.$$

(The need to weight the terms on the left-hand side of this expression by the w_i is what will necessitate the use of the weighted Laplacian in our proof.) The size of each S_i is approximately proportional to the degree of the i th vertex of G , so the w_i are all bounded above by a constant, and they are all at least one by definition. The eigenvalue $\tilde{\lambda}_2^W(G)$ of the weighted Laplacian is thus within a constant factor of the standard Fiedler value $\lambda_2(G)$.

Let z be the vector

$$z = \bar{y} - \left(\frac{\sum_i w_i \bar{y}_i}{\sum_i w_i} \right) \mathbf{1},$$

so that z differs from \bar{y} by a multiple of the all-ones vector and is orthogonal to $W\mathbf{1}$. By applying Lemma 6.3 to (6.3), we obtain

$$(6.4) \quad \mathbb{E} \left[\sum_{i \in V_G} w_i z_i^2 \right] \geq \left(1 - \frac{w_{\max}}{\sum_i w_i} \right) \mathbb{E} \left[\sum_{i \in V_G} w_i \bar{y}_i^2 \right] = \Theta\left(\frac{1}{r^2}\right) \sum_{i \in V_{G'}} y_i^2.$$

Multiplying the inequalities in (6.2) and (6.4) by the appropriate factors and combining them yields

$$(6.5) \quad O(r) \left(\sum_{(i,j) \in E_{G'}} (y_i - y_j)^2 \right) \cdot \mathbb{E} \left[\sum_{i \in V_G} w_i z_i^2 \right] \geq \left(\sum_{i \in V_{G'}} y_i^2 \right) \cdot \mathbb{E} \left[\sum_{(i,j) \in E_H} (y_i - y_j)^2 \right].$$

This implies that there exists some choice of (π_V, π_E) for which the left-hand side of (6.5) is greater than or equal to the right-hand side, in which case we would have

$$(6.6) \quad \frac{\sum_{(i,j) \in E_H} (y_i - y_j)^2}{\sum_{i \in V_G} w_i z_i^2} \leq O(r) \frac{\sum_{(i,j) \in E_{G'}} (y_i - y_j)^2}{\sum_{i \in V_{G'}} y_i^2} = O(r)R(G').$$

Now suppose that we assign to each vertex $v \in V_G$ the value assumed by y at $\pi_V(v)$. Using the fact that the maximum degree of a vertex is bounded, so that there are $O(1)$ triangles surrounding any vertex in G , we see that every path representing an edge is of length $O(r)$. We note that if i_1, \dots, i_s is a sequence of vertices,

$$(y_{i_s} - y_{i_1})^2 \leq s \sum_{a=1}^{s-1} (y_{i_{a+1}} - y_{i_a})^2.$$

As such, we have

$$(6.7) \quad \sum_{(i,j) \in E_G} (y_{\pi_V(i)} - y_{\pi_V(j)})^2 \leq O(r) \sum_{(i,j) \in E_H} (y_i - y_j)^2.$$

Since z is obtained from \bar{y} by subtracting a multiple of the all-ones vector,

$$z_i - z_j = y_{\pi_V(i)} - y_{\pi_V(j)}$$

for any i and j . Plugging this into (6.7) gives

$$\sum_{(i,j) \in E_G} (z_i - z_j)^2 \leq O(r) \sum_{(i,j) \in E_H} (y_i - y_j)^2,$$

and applying this to the inequality in (6.6) yields

$$\frac{\sum_{(i,j) \in E_G} (z_i - z_j)^2}{\sum_{i \in V_G} w_i z_i^2} \leq O(r^2)R(G').$$

We have thus constructed an assignment of values to the vertices of G that is orthogonal to the vector $W\mathbf{1}$ and produces a weighted Rayleigh quotient of $O(r^2)R(G')$. If we choose the y_i to be the values that give the Fiedler value of G' , we thus obtain, by (6.1) and the fact that the w_i are $\Theta(1)$,

$$\lambda_2(G) = \Theta(1)\tilde{\lambda}_2^W(G) \leq O(r^2)\lambda_2(G').$$

Since the number of vertices in G' grows as r^2 times the number of vertices in G , this completes the proof of Lemma 5.2. \square

Acknowledgments. I would like to thank Dan Spielman and Shang-Hua Teng for introducing me to this problem and for several extremely helpful conversations. I would also like to thank Nathan Dunfield, Phil Bowers, and Ken Stephenson for their guidance on the theory of circle packings, and Christopher Mihelich for his help in simplifying the proof of Lemma 6.3. I would also like to thank the anonymous referees for their extremely helpful comments and suggestions.

REFERENCES

- [1] E. M. ANDREEV, *Convex polyhedra in Lobachevskii space*, Math. USSR Sbornik, 10 (1970), pp. 413–440.
- [2] N. ALON, P. SEYMOUR, AND R. THOMAS, *A separator theorem for nonplanar graphs*, J. Amer. Math. Soc., 3 (1990), pp. 801–808.
- [3] C. J. ALPERT AND A. B. KAHNG, *Recent directions in netlist partitioning: A survey*, Integration: The VLSI Journal, 19 (1995), pp. 1–81.
- [4] A. F. BEARDON AND K. STEPHENSON, *The uniformization theorem for circle packings*, Indiana Univ. Math. J., 39 (1990), pp. 1383–1425.
- [5] P. BOWERS AND K. STEPHENSON, *Uniformizing dessins and Belyi maps via circle packing*, Mem. Amer. Math. Soc., 170 (2004).
- [6] P. CHAN, M. SCHLAG, AND J. ZIEN, *Spectral k -way ratio-cut partitioning and clustering*, IEEE Trans. Computer-Aided Design Integrated Circuits Systems, 13 (1994), pp. 1088–1096.
- [7] T. CHAN AND D. C. RESASCO, *A framework for the analysis and construction of domain decomposition preconditioners*, in Proceedings of the First International Symposium on Domain Decomposition Methods, Paris, 1987, R. Glowinski et al., eds., SIAM, Philadelphia, 1988, pp. 217–230.
- [8] T. CHAN AND B. SMITH, *Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes*, Contemp. Math., (1993), pp. 1–14.
- [9] F. R. K. CHUNG, *Spectral Graph Theory*, AMS, Providence, RI, 1997.
- [10] H. DJIDJEV, *A linear algorithm for partitioning graphs*, C. R. Acad. Bulgare Sci., 35 (1982), pp. 1053–1056.
- [11] H. DJIDJEV AND J. REIF, *An efficient algorithm for the genus problem with explicit construction of forbidden subgraphs*, in Proceedings of the 23rd Annual ACM Symposium on the Theory of Computing, New Orleans, LA, 1991, ACM, New York, pp. 337–348.
- [12] I. FILOTTI, G. MILLER, AND J. REIF, *On determining the genus of a graph in $o(v^{O(g)})$ steps*, in Proceedings of the 11th Annual ACM Symposium on the Theory of Computing, Atlanta, GA, 1979, ACM, New York, pp. 27–37.
- [13] O. FORSTER, *Lectures on Riemann Surfaces*, Springer-Verlag, New York, 1981.
- [14] J. GILBERT, J. HUTCHINSON, AND R. TARJAN, *A separation theorem for graphs of bounded genus*, J. Algorithms, 5 (1984), pp. 391–407.
- [15] S. GUATTERY AND G. L. MILLER, *On the performance of spectral graph partitioning methods*, in Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, San Francisco, CA, 1995, SIAM, Philadelphia, 1995, pp. 233–242.
- [16] 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.
- [17] Z.-X. HE AND O. SCHRAMM, *Fixed points, Koebe uniformization, and circle packings*, Ann. Math., 137 (1993), pp. 369–406.
- [18] P. KOEBE, *Kontaktprobleme der Konformen Abbildung*, Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl., 88 (1936), pp. 141–164.
- [19] R. J. LIPTON AND R. E. TARJAN, *A separator theorem for planar graphs*, SIAM J. Appl. Math., 36 (1979), pp. 177–189.
- [20] G. MCCAUGHAN, *A recurrence/transience result for circle packings*, Proc. Amer. Math. Soc., 126 (1998), pp. 3647–3656.
- [21] B. MOHAR, *A linear time algorithm for embedding graphs in an arbitrary surface*, SIAM J. Discrete Math., 12 (1999), pp. 6–26.
- [22] G. MIHAIL, *Conductance and convergence of Markov chains—A combinatorial treatment of expanders*, in Proceedings of the 29th Annual IEEE Conference on Foundations of Computer Science, 1989, IEEE Press, Piscataway, NJ, pp. 526–531.
- [23] B. RODIN AND D. SULLIVAN, *The convergence of circle packings to the Riemann mapping*, J. Differential Geometry, 26 (1987), pp. 349–360.
- [24] H. SIMON, *Partitioning of unstructured problems for parallel processing*, Comput. Systems Engineering, 2 (1991), pp. 135–148.

- [25] D. A. SPIELMAN AND S.-H. TENG, *Spectral partitioning works: Planar graphs and finite element meshes*, in Proceedings of the 37th Annual IEEE Conference on Foundations of Computer Science, Burlington, VT, 1996, IEEE Press, Piscataway, NJ, 1996, pp. 96–105.
- [26] K. STEPHENSON, *Circle packing and discrete analytic function theory*, in Handbook of Complex Analysis, Vol. 1: Geometric Function Theory, R. Kühnau, ed., Elsevier, New York, Chapter 11.
- [27] C. THOMASSEN, *The graph genus problem is NP-complete*, J. Algorithms, 10 (1989), pp. 568–576.
- [28] W. THURSTON, *The Geometry and Topology of 3-Manifolds*, Princeton University Notes, preprint.
- [29] R. D. WILLIAMS, *Performance of dynamic load balancing algorithms for unstructured mesh calculations*, Concurrency: Pract. Exp., 3 (1991), pp. 457–481.