ISRAEL JOURNAL OF MATHEMATICS 84 (1993), 97–118

SQUARE TILINGS WITH PRESCRIBED COMBINATORICS

BY

ODED SCHRAMM*

Department of Mathematics
The Weizmann Institute, Rehovot 76100, Israel
E-mail: schramm@wisdom.weizmann.ac.il

ABSTRACT

Let T be a triangulation of a quadrilateral Q, and let V be the set of vertices of T. Then there is an essentially unique tiling $Z = (Z_v\colon v \in V)$ of a rectangle R by squares such that for every edge $\langle u, v\rangle$ of T the corresponding two squares Z_u, Z_v are in contact and such that the vertices corresponding to squares at corners of R are at the corners of Q.

It is also shown that the sizes of the squares are obtained as a solution of an extremal problem which is a discrete version of the concept of extremal length from conformal function theory. In this discrete version of extremal length, the metrics assign lengths to the vertices, not the edges.

A practical algorithm for computing these tilings is presented and analyzed.

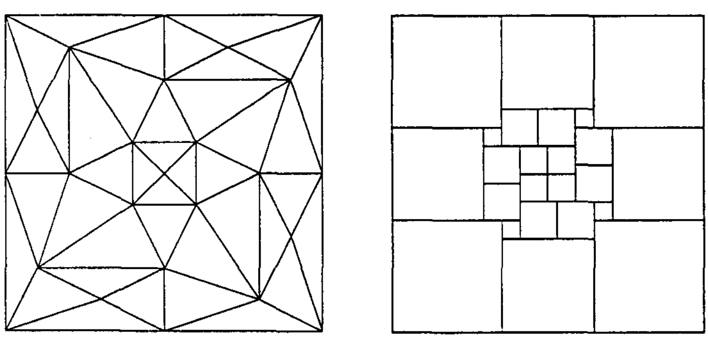

Figure 1: A triangulation and the corresponding tiling.

* The author thankfully acknowledges support of NSF grant DMS-9112150.
Received December 23, 1992

1. Introduction

To give the background and motivation for this work we need to discuss packings first.

PACKINGS. How does one describe the combinatorics of a packing? A combinatorial object associated with a packing is its **contacts graph**, defined as follows. Let the sets in the packing be indexed by a set V, $P = (P_v\colon v \in V)$, then the contacts graph (or **nerve**) of P is the graph $G = G(V, E)$ whose vertex set is V and a pair $v \neq w \in V$ are joined by an edge in E whenever $P_v \cap P_w \neq \emptyset$.

With this definition, the following question seems very natural. Given some graph G, and some geometric requirements on the packed sets, does there exists a packing with nerve G satisfying the requirements? The canonical example is when all the packed sets are required to be geometric disks (of arbitrary sizes) in $\mathbb{R}^2$. One sees immediately that the contacts graph of such a packing is a planar graph. The circle packing theorem tells us that the converse also holds: if G is a finite planar graph, then there is a packing of disks in $\mathbb{R}^2$ whose contacts graph is G. This elegant result was first proved by P. Koebe [6] as a consequence of his theorem that every finitely connected planar domain is conformally equivalent to a circle domain. Later, the circle packing theorem was rediscovered by W. Thurston [16], who also conjectured that a sequence of maps which he constructed using circle packings converges to the Riemann map from a given simply connected domain to the unit disk. This conjecture was proved by B. Rodin and D. Sullivan [10], giving a second connection between combinatorially specified circle packings and conformal maps.

In previous works, the present author has generalized the circle packing theorem in various ways ([11], [12], [13]), and the following theorem is a special case of a result in [13].

1.1 PACKING THEOREM: *Let $G = (V, E)$ be a finite planar graph, and let $(P_v\colon v \in V)$ be a collection of smooth (C^1) closed topological disks in $\mathbb{R}^2$ indexed by the vertices of G. Then there is a packing $Q = (Q_v\colon v \in V)$ with contacts graph G such that each Q_v is positively homothetic to P_v ($v \in V$); that is $Q_v = a_v P_v + b_v$, where $a_v > 0$ and $b_v \in \mathbb{R}^2$.*

The proof of this result is based on a conformal uniformization theorem of M. Brandt [1] and A. Harrington [5].

When one takes all the sets P_v to be geometric disks one obtains the circle packing theorem. The case of circles being well studied (for some references to the growing bibliography see, for example, [13]), it seems to be of some interest to investigate other special situations. It was Thurston who suggested to the author to investigate the case where the sets P_v are squares.

At first sight, it may seem that Theorem 1.1 does not apply here, since squares are usually not smooth. But one may take smooth sets which approximate the squares, obtain for them the desired packing, and then try to take the limit. There are basically two things which can go wrong in the limit. First, some new contacts may occur, and then the contacts graph will include G, rather than be equal to G. Second, some of the sets may degenerate to points. There are some tools, namely boundary conditions, that enable one to make sure that "enough" of the sets do not degenerate so that the limit packing will not be "trivial". An example of such a limiting argument can be found in [11].

What makes the case of squares especially interesting? If P is a packing of squares with edges parallel to the coordinate axis, and if the contacts graph of P is a triangulation, then the packing is actually a tiling. This follows from the following easy observation which shows that there will be no "gaps" between the squares.

1.2 OBSERVATION: *Let P_a, P_b, P_c be three rectangles whose edges are parallel to the coordinate axis. Suppose that the intersection of every two of these rectangles is nonempty. Then $P_a \cap P_b \cap P_c \neq \emptyset$.*

SQUARE TILINGS. The famous paper by Brooks, Smith, Stone and Tutte [2] studied square tilings of rectangles. Their emphasis was on *perfect tilings*, in which the sizes of the squares are all distinct, and they managed to construct nontrivial perfect square tilings of a square. They define a correspondence between square tilings of rectangles and planar multigraphs (graphs with possible multiple edges) with two poles, a source and a sink, and view the multigraph as a network of resistors in which electricity is flowing. In their setup, a vertex of the graph corresponds to a connected component of the union of the horizontal edges of the squares in the tiling, and one edge appears between two such vertices for each square whose horizontal edges lie in the corresponding connected components. Thus the graph only describes the down-up contacts between the squares, in contrast with the contacts graph. Our correspondence below graph $\to$ square tiling is very different.

We now prepare to state our basic theorem. Let D be a closed triangulated topological disk, and denote the sets of vertices, edges, and faces of the triangulation by V, E and F, respectively. Let $\partial D = B_1 \cup B_2 \cup B_3 \cup B_4$ be a decomposition of ∂D into 4 nontrivial arcs of the triangulation, in cyclic order. That is, each B_j is a nonempty connected union of edges of the triangulation, and $B_j \cap B_k$ is empty if $k = j + 2$ and consists of a vertex of the triangulation if $k = j + 1$ or $j = 1, k = 4$. The collection $T = (V, E, F; B_1, B_2, B_3, B_4)$ will be called a **triangulation of a quadrilateral.**

We will prove the following theorem.

1.3 THEOREM: *Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be a triangulation of a quadrilateral. Then there is an $h > 0$ and a square tiling $Z = (Z_v\colon v \in V)$ of the rectangle $R = [0, h^{-1}] \times [0, h]$ such that*

$$Z_v \cap Z_u \neq \emptyset \quad \text{whenever } \langle v, u \rangle \in E. \tag{1.1}$$

Moreover, let R_1, R_2, R_3, R_4 be the bottom, left, top, and right edges of R, respectively. Then it can also be required that for each $j = 1, 2, 3, 4$ we have

$$Z_v \cap R_j \neq \emptyset \quad \text{whenever } v \in B_j. \tag{1.2}$$

Under these conditions, the number h and the tiling Z are uniquely determined.

Note that we allow the possibility that some of the squares degenerate to points, as is the case for the tiling in Figure 1. (There are some situations where this can be ruled out; see 10.1.) Also, the contacts graph of the tiling may contain some edges in addition to those in E, but the corresponding intersections of squares will consist of only single points. These 'unpleasant degenerations' are consequences of the fact that squares are not smooth.

The existence part of Theorem 1.3 can be obtained from [11], [12], or [13], and uniqueness can be proven with some modification of the method of [8]. Our intention here is to give a new elementary proof. This approach also yields a practical algorithm for computing the tilings, and, in the author's view, gives insight into their structure.

The proof shows that the sizes of the squares are solutions to a certain extremal problem, 'discrete extremal length'. Extremal length is a central tool in the study of conformal and quasiconformal maps [7]. R. Duffin [4] has introduced a discrete notion of extremal length on graphs. However, the notion we use is due to J. W. Cannon [3], and is distinct from Duffin's in that lengths are assigned to vertices, rather than to edges.

Theorem 1.3 and some of the results below were independently proved by Walter Parry, and follow from his [9]. Parry also uses Cannon's discrete extremal length.

In Section 2 we will introduce the continuous notion of extremal length, for background. Section 3 defines Cannon's discrete extremal length, and gives a simple existence and uniqueness result for extremal metrics. Section 4 will show that square tilings yield extremal metrics, while Section 5 will do the converse. In Section 6 the duality properties for extremal length are indicated. An algorithm for computing square tilings is presented and analyzed in Section 7. Section 8 will discuss tilings by rectangles of specified aspect ratios; most of the theory carries through. In Section 9 we consider extremal length in arbitrary connected graphs, not necessarily planar. Finally, Section 10 consists of some concluding remarks, about degeneracies in tilings, about periodic tilings, and about non convergence to the Riemann map.

2. Extremal length

Our basic tool is a discrete version of extremal length. We now recall the continuous notion (see also [7]), not because we rely on it in the following, but only to put our arguments in perspective.

Let $Q = Q(B_1, B_2, B_3, B_4)$ be a quadrilateral in $\mathbb{R}^2$. This means that Q is a closed topological disk, and that ∂Q consists of the 4 distinguished interiorwise disjoint arcs B_1, B_2, B_3, B_4, in positive order. Let Γ be the collection of all rectifiable paths in Q which connect B_1 to B_3. Define a (conformal) metric m to be a (Borel measurable) nonnegative function on Q. The area of m is defined by

$$\text{area}(m) = \int_Q m^2\, dx\, dy = \|m\|^2.$$

(Wherever we use $\|\cdot\|$, we will mean the L^2 norm.) Let γ be a rectifiable curve in Q, then its m-length is defined as

$$l_m(\gamma) = \int_\gamma m\, |d\gamma|,$$

and l_m, the length of m, is defined as the m-distance from B_1 to B_3; that is,

$$l_m = \inf_{\gamma\in\Gamma} l_m(\gamma).$$

The **extremal length** of Q is then defined as

$$L(Q) = \sup_m \frac{l_m^2}{\|m\|^2},$$

where the supremum is taken over all conformal metrics with nonzero area. A metric m realizing the supremum is called an extremal metric.

The main point about the extremal length is that it is invariant under maps which are conformal in interior(Q). More specifically, let $f\colon Q \to Q^*$ be a conformal map between quadrilaterals, and let m^* be a conformal metric on Q^*, then $l_m = l_{m^*}$ and $\|m\| = \|m^*\|$, where $m(z) = |f'(z)|m^*(f(z))$. (Here, of course, l_{m^*} and $\|m^*\|$ refer to the quadrilateral Q^*.) It turns out that when Q^* is a rectangle the extremal metrics are just the positive constants. Since every quadrilateral Q is conformally equivalent to some rectangle, this means that the extremal metrics for Q are of the form $m(z) = |f'(z)|$, where f is a conformal homeomorphism to a rectangle. To put it differently, $|f'(z)|$ is given as a solution of an extremum problem.

We shall see that the problem of tiling a rectangle by squares with prescribed combinatorics is in a limited sense a discrete analog of the problem of finding a conformal map from a given quadrilateral to a rectangle: the sizes of the squares will be obtained from a "discrete extremal metric". In another paper [14] the author applies similar ideas, but in the continuous setting, to the study of conformal uniformization of multiply connected planar domain.

3. Discrete extremal length

We now introduce Cannon's concept of extremal length on a graph. As mentioned in the introduction, this notion of discrete extremal length is special in that the metrics give sizes to vertices, rather than to edges.

Let $G = (V, E)$ be a finite connected graph, and let A, B be subsets of V. A **path** in the graph is a sequence $(v_0, v_1, \ldots, v_n)$ of vertices such that $\langle v_{j-1}, v_j\rangle \in E$ for $j = 1, 2, \ldots, n$. A nonnegative function $m\colon V \to [0, \infty)$ will be called a (discrete) **metric** on G. Given a path $\gamma = (v_0, v_1, \ldots, v_n)$ and a metric m, we define the m-length of γ to be

$$l_m(\gamma) = \sum_{j=0}^{n} m(v_j).$$

Note that a shortest path from a vertex v to itself is the path (v), and its length is $m(v)$.

The (A, B) **length** of a metric m on G is defined as

$$l_m = \inf_\gamma l_m(\gamma),$$

where the infimum is over all paths γ which start at A and end at B. The **area** of a metric m is just the square of its 2-norm:

$$\text{area}(m) = \|m\|^2 = \sum_{v \in V} m(v)^2,$$

and the **normalized length** of m is defined as

$$\hat{l}(m) = \frac{l_m^2}{\|m\|^2}.$$

Now, the **extremal length** of $(G; A, B)$ is

$$L(G; A, B) = \sup_m \hat{l}(m),$$

where the supremum is over all metrics m of positive area. An **extremal metric** for $(G; A, B)$ is one which realizes this supremum.

Note that $l_{am} = a l_m$ when a is a positive constant. Therefore, $\hat{l}(m)$ does not change if we scale m by a positive constant factor.

3.1 LEMMA: *Let G, A, B be as above. Then there is an extremal metric m_0 for $(G; A, B)$. It is unique, up to scaling.*

Proof: Let M_1 be the set of all metrics m such that $l_m \geqslant 1$. Observe that M_1 is a nonempty closed convex set. Therefore M_1 has a unique element m_0 of least norm. Let m_1 be any metric with $l_{m_1} > 0$. Then $m_1/l_{m_1} \in M_1$, and the lemma follows by the scaling invariance of $l_m^2/\|m\|^2$. ∎

4. Square tilings give extremal metrics

Lemma 3.1 and the following lemma will imply the uniqueness part of Theorem 1.3.

4.1 LEMMA: *Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be as in Theorem 1.3, and suppose that h and Z satisfy all the requirements there. Let $G = (V, E)$ be the 1-skeleton of T, and let $s(v)$ denote the edge length of the square Z_v, for $v \in V$. Then s is an extremal metric for $(G; B_1, B_3)$.*

Proof: Let m be some metric on G of positive area. Let $t \in [0, h^{-1}]$, let β_t be the line $\{t\} \times \mathbb{R}$, and let $\gamma_t = \{v \in V \colon \beta_t \cap Z_v \neq \emptyset\}$. Then, obviously, γ_t contains the set of vertices of a simple path from B_1 to B_3. Consequently,

$$l_m \leqslant \sum_{v \in \gamma_t} m(v).$$

We integrate this inequality over t, to obtain,

$$h^{-1} l_m \leqslant \int_{t=0}^{h^{-1}} \sum_{v \in \gamma_t} m(v)\, dt.$$

Each $v \in V$ is in γ_t for t in an interval of length $s(v)$. This enables us to rewrite the right hand side, and get

$$h^{-1} l_m \leqslant \sum_{v \in V} s(v) m(v) \leqslant \|s\| \|m\|.$$

Since the area of R is 1, it follows that $\|s\| = 1$, and it is clear that $l_s = h$. Therefore,

$$\hat{l}(m) = \frac{l_m^2}{\|m\|^2} \leqslant \frac{l_s^2}{\|s\|^2} = \hat{l}(s),$$

as required. ∎

5. Extremal metrics give square tilings

5.1 THEOREM: *Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be a triangulation of a quadrilateral, and let $G = (V, E)$ be the 1-skeleton of T. Let m be the extremal metric for $(G; B_1, B_3)$ that satisfies $\|m\| = 1$. Set $h = l_m$, $R = [0, h^{-1}] \times [0, h]$, and for each $v \in V$ let*

$$Z_v = [x(v) - m(v), x(v)] \times [y(v) - m(v), y(v)], \tag{5.1}$$

where $x(v)$ [respectively $y(v)$] is the least m-length of a path from B_2 [respectively B_1] to v. Then $Z = (Z_v \colon v \in V)$ is a square tiling of the rectangle R which satisfies the contact requirements (1.1) and (1.2) of Theorem 1.3.

Note that each Z_v is a square of side length $m(v)$.

Proof: Suppose that $\langle u,v\rangle$ is an edge in T. Since a path from B_1 to v can be obtained by appending the edge $\langle u,v\rangle$ to any path from B_1 to u, it is clear that $y(v)-m(v)\leqslant y(u)$. By symmetry, we also have $y(u)-m(u)\leqslant y(v)$, and these two inequalities imply that $[y(v)-m(v),y(v)]\cap[y(u)-m(u),y(u)]\neq\emptyset$. Since a similar argument shows that $[x(v)-m(v),x(v)]\cap[x(u)-m(u),x(u)]\neq\emptyset$, we conclude that $Z_v\cap Z_u\neq\emptyset$.

Now set $\hat{R}_1=\{(x,0)\colon x\geqslant 0\},\hat{R}_2=\{(0,y)\colon y\geqslant 0\},\hat{R}_3=\{(x,y)\colon x\geqslant 0,y\geqslant h\}$, and $\hat{R}_4=\{(x,y)\colon x\geqslant h^{-1},y\geqslant 0\}$. It is clear from the definition of Z_v that $Z_v\cap\hat{R}_j\neq\emptyset$ when $j=1,2$ and $v\in B_j$. The same also holds for $j=3, v\in B_3$, by the definition of h. We now need to work a bit to prove the same for $j=4$.

Let γ be a path of least m-length which connects B_2 and B_4. For $t\geqslant 0$, set $m_t(v)=m(v)+t$ for v in γ, and $m_t(v)=m(v)$, otherwise. Since every path from B_1 to B_3 must intersect γ, we have $l_{m_t}\geqslant l_m+t$. So $D_+(l_{m_t})\geqslant 1$, where D_+ denotes the one sided derivative with respect to t, as $t\downarrow 0$. It is also easy to compute $D_+(\|m_t\|^2)$:

$$D_+(\|m_t\|^2)=\sum_{v\in\gamma}D_+\left((m(v)+t)^2\right)=\sum_{v\in\gamma}2m(v)=2l_m(\gamma).$$

Since $m=m_0$ is the extremal metric, we have

$$0\geqslant D_+\hat{l}(m_t)=D_+\left(\frac{l_{m_t}^2}{\|m_t\|^2}\right)=\frac{\|m\|^2D_+\left(l_{m_t}^2\right)-l_m^2D_+\left(\|m_t\|^2\right)}{\|m\|^4}.$$

We use our previous computations, and the normalization $\|m\|=1$, to get

$$0\geqslant D_+\left(l_{m_t}^2\right)-l_m^2D_+\left(\|m_t\|^2\right)\geqslant 2l_m-2l_m^2l_m(\gamma).$$

Since $l_m=h$, this gives $l_m(\gamma)\geqslant h^{-1}$. By the choice of γ, this implies that $Z_v\cap\hat{R}_4\neq\emptyset$ when $v\in B_4$.

We will now demonstrate that $\cup_{v\in V}Z_v\supset R$, by showing that ∂R is homotopic to a constant in $\cup_{v\in V}Z_v\cup\left(\mathbb{R}^2-\text{interior}(R)\right)$. For convenience of notation, we assume that each triangular face $\langle u,v,w\rangle$ of T is parameterized by an equilateral triangle of side length 1, and that these parameterizations are compatible along the edges, in the obvious manner. (That is, we have a piecewise linear structure on T.) Therefore, for example, the notion of the center of a face or an edge is well defined.

Define a map $f\colon T\to\bigcup_{v\in V}Z_v$, as follows. For each vertex $v\in V$ choose $f(v)$ to be some point in Z_v, such that $f(v)\in\hat{R}_j$ whenever $j=1,2,3,4$ and $v\in B_j$.

Here we use 1.2 for the vertices v at the corners, $v \in B_j \cap B_k, k - j = \pm 1 \bmod 4$. (It is clear that 1.2 applies, even though the "rectangles" $\hat{R}_j$ are not compact and some are degenerated.) We'll implicitly use 1.2 below. For each edge $\langle u, v\rangle \in E$, let $p_{\langle u,v\rangle}$ be the midpoint of the edge, and choose $f(p_{\langle u,v\rangle})$ to be some point in the intersection $Z_u \cap Z_v$. We also require that $f(p_{\langle u,v\rangle}) \in \hat{R}_j$ if $v, u \in B_j$. For each triangular face $\langle u, v, w\rangle$ in T, let $p_{\langle u,v,w\rangle}$ be the center of $\langle u, v, w\rangle$, and choose $f(p_{\langle u,v,w\rangle})$ to be some point in the intersection $Z_u \cap Z_v \cap Z_w$. Let T^* be the first barycentric subdivision of T; the triangular faces of T^* have the form $\langle u, p_{\langle u,v\rangle}, p_{\langle u,v,w\rangle}\rangle$. Now extend the map f to T by requiring it to be affine on each triangular face of T^*. It is clear that there are no conflicts in the definition of f on the vertices and edges. It is also clear that $f(T) \subset \cup_{v\in V} Z_v$, since for each face $\langle u, p_{\langle u,v\rangle}, p_{\langle u,v,w\rangle}\rangle$ of T^* the three points $f(u), f(p_{\langle u,v\rangle}), f(p_{\langle u,v,w\rangle})$ are in Z_u, and Z_u is convex.

From the fact that $f(u), f(p_{\langle u,v\rangle}) \in \hat{R}_j$ whenever $u, v \in B_j$ it follows that the restriction of f to $\partial T = B_1 \cup B_2 \cup B_3 \cup B_4$ is homotopic to ∂R in $\mathbb{R}^2 - \text{interior}(R)$. (Just take $H(p,t) = k(f(p), t)$, where $k((x,y),t) = (\min(x, h^{-1}+t), \min(y, h+t))$, as a first homotopy, and then "stretch" $H(p, 0)$ over ∂R.) Since ∂T is homotopic to a constant in T, we see that ∂R is homotopic to a constant in $\bigcup_{v\in V} Z_v \cup (\mathbb{R}^2 - \text{interior}(R))$. This clearly implies that $\bigcup_{v\in V} Z_v \supset R$.

Now that we have established $\bigcup_{v\in V} Z_v \supset R$, we recall that $1 = \|m\|^2$ is the sum of the areas of the squares in Z. Since the area of R is 1, this implies that actually $\bigcup_{v\in V} Z_v = R$, that there are no overlaps of positive area among the squares, and that no square protrudes out of R. Thus, the proof of the theorem is complete. ∎

Proof of Theorem 1.3: Existence follows from Lemma 3.1 and Theorem 5.1, and uniqueness follows from Lemma 4.1 and Lemma 3.1. ∎

6. Duality

Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be a triangulation of a quadrilateral, and let m be a metric on the 1-skeleton $G = (V, E)$ of T. By taking the tiling in Theorem 5.1 and rotating it by $\pi/2$ we see that an extremal metric for $(G; B_1, B_3)$ is also an extremal metric for $(G; B_2, B_4)$. Moreover, we have

$$L(G; B_1, B_3) = L(G; B_2, B_4)^{-1}.$$

We like to think of this phenomenon as "duality". Note that duality is not only a consequence of our results, but also is evident in the proof of Theorem 5.1, at the point where it was shown that the m-distance from B_1 to B_3 is at least h^{-1}. Actually, that was the only place where the fact that m is an extremal metric was used.

Since the same metric is extremal for $(G; B_1, B_3)$ and for $(G; B_2, B_4)$, it must also achieve the supremum for

$$1 = \sup_m \frac{l_m w_m}{\|m\|^2}, \tag{6.1}$$

where l_m [respectively w_m] denotes the least m-length of a path from B_1 to B_3 [respectively B_2 to B_4].

Duality enables us to easily get estimates for $L(G; B_1, B_3)$. Suppose that we have some metric m on G. Then the definitions give $L(G; B_1, B_3) \geqslant \hat{l}(m)$. On the other hand, we also have

$$L(G; B_1, B_3) = L(G; B_2, B_4)^{-1} \leqslant \frac{\|m\|^2}{w_m^2}. \tag{6.2}$$

7. The algorithm

We now give an algorithm for computing an extremal metric for a triangulation of a quadrilateral $T = (V, E, F; B_1, B_2, B_3, B_4)$. It is clear that the square tiling can be easily obtained from the extremal metric. In general, the algorithm will only converge to an extremal metric, and will not halt.

7.1 Algorithm:

(1) Let γ_0 be a path connecting B_2 and B_4 with the fewest possible number of vertices, and let n_0 be the number of vertices in γ_0. Set $m(v) = 1/\sqrt{n_0}$ for vertices v in γ_0 and $m(v) = 0$, otherwise.

(2) Find a path γ of least m-length from B_2 to B_4, let $w_m = l_m(\gamma)$, and let n be the number of vertices in γ.

(3) If $w_m l_m = \|m\|^2$, then stop.

(4) Let $m^*(v) = m(v) + \delta$ for vertices v in γ and $m^*(v) = m(v)$ otherwise, where

$$\delta = \frac{\|m\|^2 - l_m w_m}{n l_m - w_m}.$$

(5) Replace m by $m^*/\|m^*\|$ and go to step (2). ∎

It is first necessary to check that the δ in step (4) is not negative (to make sure that m^* is a metric). This and other useful inequalities are established in the following lemma.

7.2 LEMMA: *Each time we reach step (4) in the algorithm we have*

$$\hat{l}(m) \geqslant n_0^{-1}, \tag{7.1}$$

$$\delta \geqslant 0, \tag{7.2}$$

$$\hat{l}(m^*) - \hat{l}(m) \geqslant \frac{(\|m\|^2 - l_m w_m)^2}{\|m\|^2(n\|m\|^2 - w_m^2)} \geqslant 0. \tag{7.3}$$

Proof: Assuming for the moment that (7.1) holds, we will demonstrate that (7.2) and (7.3) follow. By (6.1), the numerator in the formula for δ is nonnegative. Using (7.1) and (6.1), we have $l_m^2 \geqslant n_0^{-1}\|m\|^2 \geqslant n^{-1} l_m w_m$. The last inequality must actually be strict, since we just passes step (3). This then establishes (7.2).

As in the proof of Theorem 5.1, since $\delta \geqslant 0$ it is clear that $l_{m^*} \geqslant l_m + \delta$. On the other hand, we have $\|m^*\|^2 = \|m\|^2 + 2w_m\delta + n\delta^2$. Thus

$$\hat{l}(m^*) - \hat{l}(m) \geqslant \frac{(l_m + \delta)^2}{\|m\|^2 + 2w_m\delta + n\delta^2} - \frac{l_m^2}{\|m\|^2},$$

and after substituting in the value of δ and simplifying, we get

$$\hat{l}(m^*) - \hat{l}(m) \geqslant \frac{(\|m\|^2 - l_m w_m)^2}{\|m\|^2(n\|m\|^2 - w_m^2)}.$$

To check that this is nonnegative, it must be verified that $n\|m\|^2 \geqslant w_m^2$. Using (7.1) and (6.1), we get

$$w_m^2 \leqslant n_0 w_m^2 \hat{l}(m) \leqslant n w_m^2 l_m^2 / \|m\|^2 \leqslant n\|m\|^2.$$

This shows then that (7.1) implies (7.3).

The proof of (7.1) will proceed inductively. For the base of the induction, when we reach step (4) for the first time, we have $l_m = n_0^{-1/2}$, $\|m\|^2 = 1$, and therefore (7.1) holds. In the inductive step, since $\hat{l}(m)$ is the same as $\hat{l}(m^*)$ of the previous visit to step (4), the inductive assumption together with (7.3) of the last visit yield (7.1). This establishes the lemma. ∎

A metric which satisfies $l_m w_m = \|m\|^2$ is extremal, and therefore, if the algorithm stops, it stops with an extremal metric. Typically, however, it will not halt.

We now investigate the convergence of the algorithm to the extremal metric. Let M be the extremal metric for T with $\|M\| = 1$, and let $L = L(T) = \hat{l}(M)$. Let $m(j)$ be the metric m at the j-th iteration of the loop in the algorithm (or the metric when the algorithm stops, if it stops before iteration j). We shall prove the following.

7.3 THEOREM: *The algorithm converges to M, namely $m(j) \to M$ as $j \to \infty$. Moreover, the following estimates hold for $j = 1, 2, \ldots$.*

$$\|M - m(j)\| \leqslant 4\sqrt{L|V|/j}, \tag{7.4}$$

$$L - \hat{l}(m(j)) \leqslant 8L^2|V|/j, \tag{7.5}$$

$$L^{-1/2} - w_{m(j)} \leqslant 4|V|\sqrt{L/j}, \tag{7.6}$$

$$\min\left\{1 - \frac{w_{m(i)} l_{m(i)}}{\|m(i)\|^2} \;\middle|\; i = j+1, j+2, \ldots, 2j\right\} \leqslant 3L|V|/j. \tag{7.7}$$

Here $|V|$ denotes the cardinality of V.

These four inequalities are various measures of the quality of convergence of the algorithm. Of these, the estimate (7.7) is the most geometric one: if one defines squares as in (5.1), but with a metric m which is not extremal, then $1 - w_m l_m/\|m\|^2$ measures the failure of the squares to tile the rectangle $[0, l_m] \times [0, w_m]$. It is the proportion of the sum of the areas of the squares which is in overlaps or falls outside of that rectangle. Estimate (7.7) then means that the "best" metric among the metrics encountered in the first $2j$ steps will have a proportion of at most $O(j^{-1})$ "wasted" area.

Inequality (7.5) estimates how fast the normalized length of $m(j)$ converges to the extremal length L, and inequality (7.6) estimates the dual convergence of $w_{m(j)}$ to the extremal length of the "rotated triangulation".

It is clear that the left hand sides of the inequalities in 7.3 are all nonnegative.

Proof: Set $r(j) = \hat{l}(m(j))/L$. Since $w_m^2/\|m\|^2 \leqslant L^{-1}$, inequality (7.3) gives

$$r(j+1) - r(j) \geqslant \frac{\left(1 - \frac{l_{m(j)} w_{m(j)}}{\|m(j)\|^2}\right)^2}{nL} \geqslant \frac{\left(1 - \frac{l_{m(j)} w_{m(j)}}{\|m(j)\|^2}\right)^2}{|V|L} \geqslant \frac{\left(1 - \sqrt{r(j)}\right)^2}{|V|L}. \tag{7.8}$$

Since $r(j) \leqslant 1$, using (7.8) and $\sqrt{1-\epsilon} \leqslant 1 - \epsilon/2$ it is easy to prove inductively that for every j

$$r(j) \geqslant 1 - 8L|V|j^{-1}. \tag{7.9}$$

This establishes (7.5).

To prove (7.4), we note that M/l_M and $m(j)/l_{m(j)}$ both belong to the convex set of all metrics m such that $l_m \geqslant 1$, and that M/l_M is of least norm among that set. This implies that

$$\left(\frac{m(j)}{l_{m(j)}} - \frac{M}{l_M}\right) \cdot \frac{M}{l_M} \geqslant 0,$$

where $\cdot$ denotes the inner product. We therefore get

$$m(j) \cdot M \geqslant l_{m(j)}/l_M.$$

Since $\|M\| = \|m(j)\| = 1$, using (7.9), this gives

$$\begin{aligned}\|M - m(j)\|^2 = 1 - 2m(j) \cdot M + 1 \leqslant 2(1 - l_{m(j)}/l_M) = 2\left(1 - \sqrt{r(j)}\right)\\ \leqslant 2\left(1 - \sqrt{1 - 8L|V|/j}\right) \leqslant 16L|V|/j,\end{aligned}$$

and proves (7.4).

The proof of (7.6) is straightforward:

$$L^{-1/2} - w_{m(j)} = w_M - w_{m(j)} \leqslant \|M - m(j)\|_1 \leqslant \sqrt{|V|}\|M - m(j)\|_2 \leqslant 4|V|\sqrt{L/j},$$

by (7.4).

Using (7.8), we have

$$1 \geqslant r(2j+1) \geqslant r(j+1) + \sum_{i=j+1}^{2j} |V|^{-1}L^{-1}\left(1 - \frac{l_{m(i)}w_{m(i)}}{\|m(i)\|^2}\right)^2,$$

which implies

$$\min\left\{\left(1 - \frac{l_{m(i)}w_{m(i)}}{\|m(i)\|^2}\right)^2 \,\middle|\, i = j+1, \ldots, 2j\right\} \leqslant \frac{L|V|(1 - r(j+1))}{j} \leqslant \frac{8L^2|V|^2}{j^2}.$$

This gives (7.7), and completes the proof. ∎

To complete the discussion of the algorithm, we need to estimate the running time of one iteration of the loop. The only substantial computations there involve finding γ and computing l_m and w_m. Thus, the question which faces us is: given a graph $G = (V, E)$, two sets of vertices $A, B \subset V$, and a metric m on G, how can we efficiently find a path from A to B of least m-length. This is very similar to the problem of determining the shortest path between two nodes in a graph with edges of various lengths. Indeed, the algorithm given in [15, Chapter 31] for the latter task applies with only minor changes to our situation, and runs in $O((|E| + |V|)\log|V|)$ computation steps. Since the graph is planar in our case, we have $|E| = O(|V|)$, and thus the loop in Algorithm 7.1 needs at most $O(|V|\log|V|)$ computation steps. We get the following

7.4 COROLLARY: *Let* $0 < \epsilon < 1$. *Algorithm* 7.1 *can be used to compute an approximate square tiling for* T *with at most* ϵ(*total area*) *in overlaps or outside the "tiled" rectangle in* $O(\epsilon^{-1}|V|^2\log|V|)$ *computation steps.*

Proof: Use the algorithm simultaneously to compute the tiling for T and for the rotated triangulation $T^* = (V, E, F; B_2, B_3, B_4, B_1)$. Since $L(T) = L(T^*)^{-1}$, the above estimate for the running time of a single loop of the algorithm and inequality (7.7) show that one of the two runs will produce the required approximate tiling in the stated time. Alternatively, one can abort the computation for T [respectively T^*] once it is clear that $L(T^*) \geqslant 1/2$ [respectively $L(T) \geqslant 1/2$]. ∎

Remark: In (7.8) we have bounded n (the length of γ) by $|V|$. It seems reasonable that "usually" n will be approximately equal to $\sqrt{|V|}$. In such situations we expect a running time not worse than $O(\epsilon^{-1}|V|^{3/2}\log|V|)$. ∎

Remark: Suppose that $Z = (Z_v\colon v \in V)$ is a tiling as in Theorem 1.3. A directed edge $\langle v, u\rangle \in E$ is **vertical** if the contact between Z_u and Z_v is along the bottom edge of Z_u and the top edge of Z_v. Let $E_\uparrow$ be the collection of vertical edges, and let $G_\uparrow = (V, E_\uparrow)$. (It is a good exercise to show that $L(G; B_1, B_3) = L(G_\uparrow; B_1, B_3)$.) Given $G_\uparrow$, the tiling Z can be computed by solving a system of linear equations, using the methods of [2]. This means that it is possible to compute Z exactly in finite time by trying all possible $G_\uparrow$, but that's clearly not practical. An efficient algorithm to calculate $G_\uparrow$ would be interesting. ∎

8. Tilings by rectangles

It turns out that all the results above can be easily generalized to tilings with rectangles of specified aspect ratios. To obtain these generalizations, it is necessary to consider a slightly more general notion of discrete extremal length.

Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be a triangulation of a quadrilateral, and let $\alpha\colon V \to (0, \infty)$ be some assignment of weights to the vertices. (These weights will become the aspect ratios.) For any metric $m\colon V \to [0,\infty)$ on T define its α-area $\|m\|_\alpha^2$ by

$$\|m\|_\alpha^2 = \sum_{v\in V} \alpha(v) m(v)^2.$$

Define the α-extremal length of T to be

$$L_\alpha(T) = \sup_m \frac{l_m^2}{\|m\|_\alpha^2},$$

where the supremum is over all metrics of positive area. An α-extremal metric is one which achieves the supremum. As in 3.1, an α-extremal metric always exists, and is unique up to a positive scaling factor.

We have the following generalization of Theorems 1.3 and 5.1.

8.1 THEOREM: *Let $T = (V, E, F; B_1, B_2, B_3, B_4)$ be a triangulation of a quadrilateral, and let $\alpha\colon V \to (0,\infty)$ be an assignment of weights to the vertices. Then there is a unique h and a unique tiling $Z = (Z_v\colon v \in V)$ of the rectangle $R = [0, h^{-1}] \times [0, h]$ such that the contact requirements (1.1) and (1.2) of Theorem 1.3 are satisfied and such that each Z_v is a "horizontal" rectangle with width to height ratio equal to $\alpha(v)$. This tiling is given by*

$$Z_v = [x(v) - \alpha(v)m(v), x(v)] \times [y(v) - m(v), y(v)], \tag{8.1}$$

where m is the α-extremal metric of T with α-area equal to 1, $x(v)$ is the least αm-length of a path from B_2 to v, and $y(v)$ is the least m-length of a path from B_1 to v.

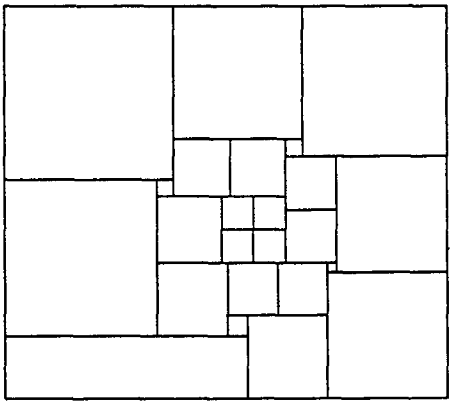

Figure 2: A tiling with one wide rectangle.

Figure 2 shows a tiling corresponding to the same triangulation as in Figure 1, but where one corner square is dictated to become a rectangle with width to height ratio equal to 4.

The proof of 8.1 is a simple modification of our arguments in the previous sections.

The algorithm for computing square tilings can also be easily modified to compute tilings with rectangles of specified aspect ratios, as follows. The path γ_0 is chosen as a path between B_2 and B_4 of least α-length, and n_0 is replaced by $l_\alpha(\gamma_0)$, the α-length of γ_0. The path γ is chosen as a path of least αm-length from B_2 to B_4, w_m is replaced by $l_{\alpha m}(\gamma)$, n is replaced by $l_\alpha(\gamma)$, and the norm $\|\cdot\|$ is replaced by $\|\cdot\|_\alpha$. The rest stays the same, and the analysis of the algorithm proceeds similarly.

9. Extremal length in arbitrary graphs

We have been mostly interested in geometric applications, but it seems worthwhile to see how much of what we've done carries over to arbitrary graphs.

Let $G = (V, E)$ be a finite connected graph, and let A, B be subsets of V. For a metric m: $V \to [0, \infty)$ and a path γ, we denote as usual $l_m(\gamma) = \sum_{v\in\gamma} m(v)$, and let l_m be the shortest m-length of a path from A to B. When the graph G is arbitrary, there is no special reason to use the 2-norm, and we'll maintain as much generality as seems appropriate. Let $a > 0$, and let α: $V \to (0, \infty)$. Set

$$N_{a,\alpha}(m) = \sum_{v\in V} \alpha(v) m(v)^{1+a},$$

and

$$L_{a,\alpha} = L_{a,\alpha}(G; A, B) = \sup_m \frac{l_m^{1+a}}{N_{a,\alpha}(m)},$$

where the supremum extends over all metrics which are not identically zero. As usual, an extremal metric is one which achieves the supremum.

Given a metric m and a vertex $v \in V$, we let $y_m(v)$ denote the least m-length of a path from A to v. For $t \in \mathbb{R}$, let $V_m(t) = \{v \in V: y_m(v) - m(v) \leqslant t < y_m(t)\}$. It is clear that $V_m(t)$ separates A from B when $t \in [0, l_m)$. Finally, for a set of vertices $U \subset V$, let the m-**dual-length** of U be defined by

$$l_m^*(U) = \sum_{v\in U} \alpha(v) m(v)^a,$$

and let

$$l_m^* = \inf_U l_m^*(U),$$

where the infimum extends over all sets $U \subset V$ which separate A from B.

9.1 THEOREM: *Let G, A, B, a, α be as above, and let $m\colon V \to [0,\infty)$ be a metric which is not identically zero.*

(1) *The extremal metric M for $L_{a,\alpha}$ exists, and is unique up to a positive scaling factor.*
(2) *For all $t \in [0, l_M)$ we have $l_M^* = l_M^*(V_M(t))$.*
(3) *We have*

$$l_m l_m^* \leqslant N_{a,\alpha}(m) \tag{9.1}$$

$$\frac{(l_m^*)^{1+a^{-1}}}{N_{a,\alpha}(m)} \leqslant L_{a,\alpha}^{-1/a}, \tag{9.2}$$

with equality holding in both if m is extremal for $L_{a,\alpha}$. Conversely, equality in either of these implies that m is extremal.

The inequality (9.1) is a generalization of (6.1), and inequality (9.2) is a generalization of duality: The left hand side of (9.2) should be interpreted as the "normalized dual length" of m. The reader may wish to figure out what geometric phenomenon (2) generalizes.

Proof: The proof of (1) is the same as the proof of Lemma 3.1. (Note that for uniqueness it is necessary that $a > 0$.)

Let $m_1, m_2\colon V \to [0,\infty)$ be two metrics. We have

$$\begin{aligned} l_{m_1} l_{m_2}^* &\leqslant \int_{t=0}^{l_{m_1}} l_{m_2}^*(V_{m_1}(t))\,dt \\ &\leqslant \sum_{v\in V} \alpha(v) m_1(v) {m_2}^*(v)^a \\ &\leqslant \left(\sum_{v\in V} \alpha(v) m_1(v)^{1+a}\right)^{1/(1+a)} \left(\sum_{v\in V} \alpha(v) m_2(v)^{1+a}\right)^{a/(1+a)} \\ &= N_{a,\alpha}(m_1)^{1/(1+a)} N_{a,\alpha}(m_2)^{a/(1+a)}, \end{aligned} \tag{9.3}$$

where the second inequality comes from considering the contribution of each vertex $v \in V$ to the integral, and the third is an application of the Hölder inequality. Taking $m_1 = m_2 = m$ in (9.3), we get (9.1). Also (9.2) is obtained from (9.3) by setting $m_1 = M$ and $m_2 = m$.

Now let $U \subset V$ be a set which separates A from B, let $t \geqslant 0$, and let $M_t\colon V \to [0,\infty)$ be defined by $M_t(v) = M(v)+t$ for $v \in U$ and $M_t(v) = M(v)$ for $v \in V-U$. As in the computation in the proof of Theorem 5.1, it is easy to see that

$$0 \geqslant D_+\left(\frac{l_{M_t}^{1+a}}{N_{a,\alpha}(M_t)}\right)$$

implies that $l_M l_M^*(U) \geqslant N_{a,\alpha}(M)$, which means that equality holds in (9.1) when m is extremal. This also shows that equality holds in (9.2) for extremal m.

Conversely, if equality holds in (9.2) for m, then equality must hold in (9.3) when $m_2 = m$, $m_1 = M$. This implies that m is a scalar multiple of M, and is therefore extremal. Alternatively, the same conclusion can be obtained by showing that the metric which achieves the maximum of the left hand side of (9.2) is unique, up to a positive scaling factor.

Using (9.2), it is easy to see that equality in (9.1) implies that m is extremal.

Since $l_M l_M^* = N_{a,\alpha}(M)$, comparing with (9.3), we see that $l_M^*(V_M(t)) = l_M^*$ holds for almost every $t \in [0, l_M]$. But for all $t_1 \in [0, l_M)$ there is some $t_2 > t_1$ such that $V_M(t) = V_M(t_1)$ for all $t \in [t_1, t_2)$. Thus statement (2) is verified, completing the proof. ∎

10. Concluding remarks

DEGENERACIES. It is quite possible that some squares in the tiling given in Theorem 1.3 will actually degenerate to points. Let T^1 denote the 1-skeleton of T. If there is a 3-cycle in T^1 (a closed path of length 3) which is not the boundary of a triangle of T, then all the squares corresponding to vertices "inside" the cycle will degenerate to a single point. This can be shown either geometrically, using 1.2, or by considering an extremal metric.

We now investigate when degeneracies can occur. Suppose that some squares do degenerate. Let V_0 be a connected component of the set of vertices which correspond to degenerated squares, and let ∂V_0 be the set of all vertices which are not in V_0 but which neighbor with V_0. All the squares $Z_v, v \in V_0$ must be the same point, and Z_u must contain that point for every $u \in \partial V_0$. Since the angles of a square are $\pi/2$, it is impossible for 5 non degenerate squares with disjoint interiors to contain the same point. This means that $|\partial V_0| \leqslant 4$. ∂V_0 is either a closed path, or a path which begins and ends at the boundary, $B_1 \cup B_2 \cup B_3 \cup B_4$. If ∂V_0 is not a closed path, and it begins and ends at the boundary, then a similar consideration of angles shows that actually $|\partial V_0| \leqslant 2$, since vertices of the boundary correspond to squares on the boundary of the rectangle. Moreover, if the endpoints of ∂V_0 are not on the same B_j, then we get the ridiculous $|\partial V_0| \leqslant 1$. Thus we have the following

10.1 PROPOSITION: *In Theorem* 1.3, *let* $G = (V, E)$ *be the* 1*-skeleton of* T. *Suppose that for each* $j = 1, 2, 3, 4$ *and for every edge* $e \in E$ *whose vertices are in* B_j *we have* $e \subset B_j$. *Also suppose that every simple closed path in* G *which separates* V *contains at least* 5 *vertices. Then no square in the tiling of Theorem* 1.3 *degenerates to a point.*

PERIODIC TILINGS. If T is a triangulation of a compact cylinder, then, with the methods of this paper, it is not too difficult to obtain a square tiling of a cylinder whose contacts graph includes T^1, the 1-skeleton of T. This is analogous to having a periodic square tiling of a horizontal strip in the plane. It is worthwhile to note that the usual pattern of bricks in a wall shows that the periodic tiling is not unique, even if one normalizes the strip to have height 1 and considers two tilings to be the same if they differ by a translation. However, with this normalization, the sizes of the squares are uniquely determined, since they correspond to an extremal metric for T^1.

Regarding doubly periodic tilings, there is the following

10.2 THEOREM: *Let* T *be a triangulation of the plane* $\mathbb{C}$ *which is invariant under the translations* $z \to z + 1$ *and* $z \to z + i$, *let* V *be the set of vertices of* T, *and let* E *be the set of edges. Then there is a complex number* λ *and a tiling* $Z = (Z_v \colon v \in V)$ *of the plane by squares with edges parallel to the coordinate axes such that* (1.1) *holds and such that* $Z_v = Z_u + n + m\lambda$ *whenever* $v = u + n + mi$ *with* $n, m \in \mathbb{Z}$.

This theorem follows from [13]. It is probably possible to prove it with the methods of this paper, but the author has failed to find a relatively simple proof along this vein.

Again, the "bricks pattern" excludes uniqueness, but the sizes of the squares are probably uniquely determined.

NON-CONVERGENCE TO THE RIEMANN MAP. For anyone acquainted with [10], it is a natural question to ask whether square tilings can be used as discrete approximations for the conformal map from a simply connected domain to a rectangle. The answer is no, at least if one attempts to use the combinatorics of the hexagonal lattice. Figure 3 illustrates this: The tiling on the left is obtained as the tiling associated with the contacts graph of the tiling on the right, but the associated map is not close to a conformal map.

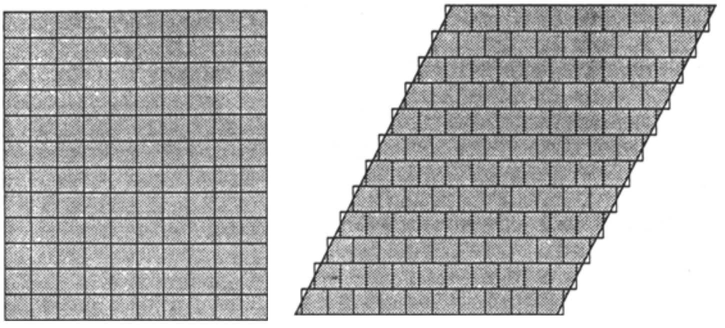

Figure 3: Non-convergence to the Riemann map.

References

[1] M. Brandt, *Ein Abbildungssatz für endlich-vielfach zusammenhängende Gebiete,* Bull. de la Soc. des Sci. et des Lettr. de Lódź, Vol. XXX, 3, 1980.

[2] R. L. Brooks, C. A. B. Smith, A. H. Stone and W. T. Tutte, *The dissection of squares into squares,* Duke Math. J. **7** (1940), 312–340.

[3] J. W. Cannon, Informal seminar, Princeton University.

[4] R. J. Duffin, *The extremal length of a network,* J. Math. Anal. Appl. **5** (1962), 200–215.

[5] A. N. Harrington, *Conformal Mappings onto domains with arbitrarily specified boundary shapes,* J. Analyse Math. **41** (1982), 39–53.

[6] P. Koebe, *Kontaktprobleme der konformen Abbildung,* Berichte Verhande. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Klasse **88** (1936), 141–164.

[7] O. Lehto and K. I. Virtanen, *Quasiconformal Mappings in the Plane,* Springer, Berlin, 1973, 258pp.

[8] M. Shiffman, *Uniqueness theorems for conformal mappings of multiply connected domains,* Proc. Natl. Acad. Sci. U.S.A. **27** (1941), 137–139.

[9] W. Parry, *Notes on optimal weight functions*, preprint.

[10] B. Rodin and D. Sullivan, *The convergence of circle packings to the Riemann mapping*, J. Differ. Geom. **26** (1987), 349–360.

[11] O. Schramm, *Packing two-dimensional bodies with prescribed combinatorics and applications to the construction of conformal and quasiconformal mappings*, Ph.D. thesis, Princeton, 1990.

[12] O. Schramm, *Existence and uniqueness of packings with specified combinatorics*, Israel J. Math. **73** (1991), 321–341.

[13] O. Schramm, *Conformal uniformization and packings*, preprint.

[14] O. Schramm, *Transboundary extremal length*, preprint.

[15] R. Sedgewick, *Algorithms*, 2nd edition, Addison-Wesley, 1988, 657pp.

[16] W. P. Thurston, *The finite Riemann mapping theorem*, invited address, International Symposium in Celebration of the Proof of the Bieberbach Conjecture, Purdue University, 1985.