

SQUARING RECTANGLES FOR DUMBBELLS

J. W. CANNON, W. J. FLOYD, AND W. R. PARRY

ABSTRACT. The theorem on squaring a rectangle (see Schramm [Israel J. Math. **84** (1993)] and Cannon-Floyd-Parry [Contemp. Math. 169, AMS, 1994]) gives a combinatorial version of the Riemann mapping theorem. We elucidate by example (the dumbbell) some of the limitations of rectangle-squaring as an approximation to the classical Riemann mapping.

1. INTRODUCTION

A *quadrilateral* is a planar disk D with four distinguished boundary points a , b , c , and d that appear in clockwise order on the boundary. These four points determine a *top edge* ab , a *right edge* bc , a *bottom edge* cd , and a *left edge* da . A *tiling* of D is a finite collection of disks T_i , called tiles, whose union fills D , whose interiors $\text{Int } T_i$ are disjoint, and whose boundaries form a finite graph Γ in D that contains ∂D .

The *rectangle-squaring theorem* (see Schramm [6] and Cannon-Floyd-Parry [1]) states that there are integers $p_i \geq 0$ parametrized by the tiles (for a tile T_i , p_i is called the *weight* of T_i), not all equal to 0, unique up to scaling, such that squares S_i of edge length p_i can be assembled with essentially the same adjacencies as the tiles T_i to form a geometric rectangle with top, bottom, and sides corresponding to the edges of D . Some combinatorial distortions are inevitable. For example, in a squared rectangle, at most four tiles can meet at a point. The distortions allowed are these: (1) a vertex may expand into a vertical segment and (2) a tile may collapse to a point. No other distortions are required.

The rectangle-squaring theorem is essentially a combinatorial version of the Riemann Mapping Theorem. It has the advantage over other versions of the Riemann Mapping Theorem that the integers p_i can be calculated by a terminating algorithm and can be approximated rapidly by various simple procedures. As a consequence, rectangle-squaring can be used as a rapid preprocessor for other computational methods of approximating the Riemann mapping.

Schramm [6] has pointed out that tilings given by the simplest subdivision rules give squarings that need not converge to the classical Riemann mapping. The purpose of this paper is to further elucidate the limitations of the method. In the classical Riemann mapping, changing the domain of the mapping anywhere typically changes the mapping everywhere. Our main result shows that this is not true for rectangle squaring. A *dumbbell* is a planar quadrilateral, constructed from squares of equal size, that consists of two blobs (the left ball and the right ball)

Received by the editors June 19, 2008.

2000 *Mathematics Subject Classification*. Primary 52C20, 52C26.

Key words and phrases. Finite subdivision rules, combinatorial moduli, squaring rectangles.

at the end joined by a relatively narrow bar of uniform height in the middle. We show that for any choices of the left and right balls of a dumbbell, the weights p_i associated with the squares in the middle of the bar are constant provided that the bar is sufficiently long and narrow. In particular, subdivision cannot lead to tilings whose squarings converge to the classical Riemann mapping.

In order to state our main theorem, we precisely define what we mean by a dumbbell. A *dumbbell* D is a special sort of conformal quadrilateral which is a subcomplex of the square tiling of the plane. It consists of a *left ball*, a *right ball*, and a *bar*. The left ball, the bar, and the right ball are all subcomplexes of D . The bar is a rectangle at least six times as wide as high. It meets the left ball in a connected subcomplex of the left side of the bar, and it meets the right ball in a connected subcomplex of the right side of the bar. The bar of a dumbbell is never empty, but we allow the balls to be empty. The *bar height* of D is the number of squares in each column of squares of the bar of D . Figure 1 shows a dumbbell with bar height 1.




FIGURE 1. A dumbbell D

If D is a dumbbell with bar height n , then a weight function ρ for D is *virtually bar uniform* if $\rho(t) = \frac{1}{n}H_\rho$ for any tile t in the bar of D whose skinny path distance to the balls of D is at least $3n$. (The skinny path distance between two subsets of D is one less than the minimum length of a chain of tiles which joins a point in one subset to a point in the other subset. The skinny path distance is used extensively in [3].) We are now ready to state our main theorem.

Dumbbell Theorem. *Every fat flow optimal weight function for a dumbbell is virtually bar uniform.*

As indicated above, the theorem has consequences for Riemann mappings. Let D be a dumbbell with bar height n , and let ρ be a fat flow optimal weight function for D . If t is a tile in the bar of D whose skinny path distance to the balls of D is at least $3n$, then the dumbbell theorem implies that $\rho(t) = \frac{1}{n}H_\rho$. Now we subdivide D using the binary square finite subdivision rule \mathcal{R} , which subdivides each square into four subsquares. We see that $\mathcal{R}(D)$ is a dumbbell with bar height $2n$. Let s be a tile of $\mathcal{R}(D)$ contained in t . Then the skinny path distance from s to the balls of $\mathcal{R}(D)$ is at least $6n$. If σ is a fat flow optimal weight function for $\mathcal{R}(D)$, then the dumbbell theorem implies that $\sigma(s) = \frac{1}{2n}H_\sigma$. It follows that as we repeatedly subdivide D using the binary square finite subdivision rule and normalize the optimal weight functions so that they have the same height, if they converge, then they converge

to the weight function of an affine function in the middle of the bar of D . The only way that a Riemann mapping of D can be affine in the middle of the bar of D is for D to be just a rectangle. Thus our sequence of optimal weight functions almost never converges to the weight function of a Riemann mapping of D . We formally state this result as a corollary to the dumbbell theorem.

Corollary. *When we repeatedly subdivide a dumbbell using the binary square subdivision rule, if the resulting sequence of fat flow optimal weight functions converges to the weight function of a Riemann mapping, then the dumbbell is just a rectangle.*

Here is a brief outline of the proof of the dumbbell theorem. Let D be a dumbbell, and let ρ be a fat flow optimal weight function for D . From ρ we construct a new weight function σ for D . The weight function σ is constructed so that it is virtually bar uniform and if t is a tile of D not in the bar of D , then $\sigma(t) = \rho(t)$. Much effort shows that $H_\sigma = H_\rho$. Because σ is a weight function for D with $H_\sigma = H_\rho$ and ρ is optimal, $A_\rho \leq A_\sigma$. Because σ agrees with ρ outside the bar of D , it follows that there exists a column of squares in the bar of D whose ρ -area is at most its σ -area. The main difficulty in the proof lies in controlling the σ -areas of such columns of squares. The key result in this regard is Theorem 5.7. Most of this proof is devoted to proving Theorem 5.7. It gives us enough information about the σ -areas of such columns of squares to conclude that $H_\sigma = H_\rho$ and eventually that $\sigma = \rho$. Thus ρ is virtually bar uniform.

We give some simple examples in Section 2. Sections 3 and 4 are relatively easy. The skinny cut function Φ is defined in the first paragraph of Section 5. Once the reader understands the definition of Φ , the statement of Theorem 5.7 can be understood. After understanding the statement of Theorem 5.7, the reader can read Section 6 to get a better grasp of the proof of the dumbbell theorem outlined in the previous paragraph. Theorem 5.7 is the key ingredient, and its proof presents the greatest difficulties in our argument.

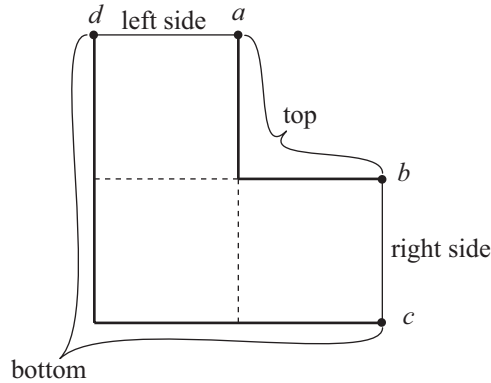

In Section 7 we discuss without proofs an assortment of results which are related to (but not used in) our proof of the dumbbell theorem.

2. EXAMPLES

We give some simple examples here to illustrate the theorem.

Example 2.1. We begin with an example that motivated this work. Consider the topological disk D_1 shown in Figure 2. The left ball of D_1 is a union of two tiles, the bar of D_1 is a single tile, and the right ball of D_1 is empty. D_1 is not a dumbbell because the bar isn't wide enough, but since it has so few tiles it is easier to analyze. Figure 3 shows the first three subdivisions of D_1 with respect to the binary square subdivision rule; the tiling $\mathcal{R}^{n+1}(D_1)$ is obtained from $\mathcal{R}^n(D_1)$ by subdividing each square into four subsquares. We consider each subdivision as a conformal quadrilateral by choosing the same four points as vertices and the same labeling of the edges.

We considered this example because we were interested in the squared rectangles corresponding to the optimal weight functions for the fat flow moduli of this sequence of quadrilaterals. Given a tiling \mathcal{T} (or, more generally, a shingling) of a conformal quadrilateral Q , a *weight function* on \mathcal{T} is a non-negative real-valued function ρ on the set of tiles of \mathcal{T} . If ρ is a weight function on \mathcal{T} and t is a tile of \mathcal{T} ,

FIGURE 2. The quadrilateral D_1 FIGURE 3. The first three subdivisions of D_1

then $\rho(t)$ is the *weight* of t . One can use a weight function to assign “lengths” to paths in Q and “areas” to subsets of Q . The ρ -*length* of a path α in Q is the sum of the weights of the tiles that the image of α intersects, and the ρ -*area* $A_\rho(W)$ of a subset W of Q is the sum of the squares of the weights of the tiles that intersect W . One then defines the ρ -*height* H_ρ of Q to be the minimum ρ -length of a path in Q that joins the top and bottom of Q , and one defines the ρ -*area* A_ρ of Q to be $A_\rho(Q)$. The *fat flow modulus* of Q with respect to ρ is H_ρ^2/A_ρ . The *optimal weight function* is a weight function whose fat flow modulus is the supremum of the fat flow moduli of weight functions; it exists and is unique up to scaling. There is a squared rectangle corresponding to Q whose squares correspond to the tiles of \mathcal{T} with non-zero weights under the optimal weight function; furthermore, the side length of one of these squares is the weight of the corresponding tile.

Figure 4 gives optimal weight functions for the first two subdivisions of the quadrilateral D_1 , and the left side of Figure 5 gives the squared rectangles for the first three subdivisions of D_1 . (See [1] for detailed information about optimal weight functions for tilings.) For each of these tilings of D_1 , one can define a circle packing whose carrier complex is obtained from the tiling by adding a barycenter to each tile and then subdividing each tile into triangles by adding an edge from each vertex of the tile to the barycenter. By only drawing the edges of the packed carrier complex that correspond to edges of the tiling, one can use the circle packing to draw the tilings. On the right side of Figure 5, this is done for the first three subdivisions of D_1 using Stephenson’s program CirclePack [7].




FIGURE 4. Optimal weight functions for the first two subdivisions of D_1

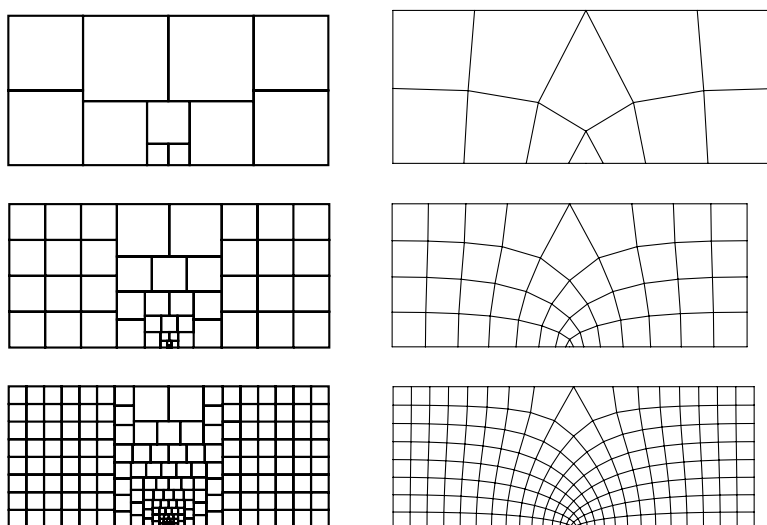


FIGURE 5. The squared rectangles and circle packings for the first three subdivisions of D_1

While one cannot expect discrete functions from squared rectangles for tilings to converge to Riemann maps, He and Schramm showed in [5] that under general hypotheses discrete functions associated to circle packings converge to Riemann maps. It may seem that the squared rectangles and the circle packings are not closely related, but in some computational examples with the pentagonal and dodecahedral subdivision rules we have found that using our squared rectangle software [4] in conjunction with CirclePack [7] can lead to dramatic reductions in the total computation time for producing the packings. (In an example with over 1,600,000 vertices, using both programs led to a reduction in the computation time from almost 38 hours to under six hours.) We wondered whether the squared rectangles for the

subdivisions of D_1 would define discrete functions which were closely related to the Riemann map for D_1 .

But in this example the optimal weight functions appeared to be constant near the two sides of the quadrilateral. As we described above, if they really are constant, then they couldn't possibly give discrete approximations to the Riemann map. We found this "apparent" constancy of weights near the sides to be very surprising, and decided to look more closely at optimal weight functions for quadrilaterals made out of square tiles.

Example 2.2. Let D be the dumbbell drawn in Figure 1. The left edge of D is the left side of the left ball, and the right edge of D is the right side of the right ball. The bar of D has height 1 and width 8. Figure 6 gives squared rectangles for the first three subdivisions of D . For these subdivisions, the weights p_i are constant on the entire bar and not just on the middle fourth of the bar.




FIGURE 6. Squared rectangles for the first three subdivisions of D

Example 2.3. The dumbbell D_2 shown in Figure 7 is similar to the dumbbell D . The two dumbbells have the same bars, but for D_2 the left and right balls are smaller. The left edge of D_2 is the top side of the left ball and the right edge of D_2 is the top side of the right ball. Squared rectangles for the first three subdivisions of D_2 are shown in Figure 8. The dumbbell theorem guarantees that the weights are constant in the middle fourth of the bar, but here they are constant on most of the bar. The quadrilateral D_1 is a subcomplex of D_2 .




FIGURE 7. The dumbbell D_2




FIGURE 8. Squared rectangles for the first three subdivisions of D_2

3. WEIGHT VECTORS

This section deals with basic properties of vectors in Euclidean space.

We fix a positive integer n , which will be the bar height of a dumbbell under consideration. A *weight vector* is an element of \mathbf{R}^n whose components are all nonnegative and not all 0. We denote by \mathcal{W} the set of all weight vectors in \mathbf{R}^n . The *height* $H(x)$ and *area* $A(x)$ of a weight vector $x = (x_1, \dots, x_n)$ are defined by

$$H(x) = \sum_{i=1}^n x_i \quad \text{and} \quad A(x) = \sum_{i=1}^n x_i^2.$$

Given a positive real number h , we let $w_h = (\frac{h}{n}, \dots, \frac{h}{n})$ and

$$\mathcal{W}_h = \{x \in \mathcal{W} : H(x) = h\}.$$

We see that $w_h \in \mathcal{W}_h$ and that $A(w_h) = \frac{h^2}{n}$.

Lemma 3.1. *Let h be a positive real number, let $x_1, x_2 \in \mathcal{W}_h$, and let d_1 and d_2 be the distances from x_1 and x_2 to w_h . Then $A(x_1) \leq A(x_2)$ if and only if $d_1 \leq d_2$.*

Proof. Let $x \in \mathcal{W}_h$. Since $H(x)$ is the dot product of x with the vector whose components are all 1, we have that $x \cdot (\frac{n}{h}w_h) = h$. Hence $x \cdot w_h = \frac{h^2}{n} = w_h \cdot w_h$. So the square of the distance from x to w_h is

$$(x - w_h) \cdot (x - w_h) = x \cdot x - 2x \cdot w_h + w_h \cdot w_h = x \cdot x - w_h \cdot w_h.$$

Thus decreasing $A(x)$ is equivalent to decreasing the distance from x to w_h . This proves the lemma. \square

Corollary 3.2. *The weight vector w_h is the unique element of \mathcal{W}_h with least area.*

Lemma 3.3. *Let x and y be distinct weight vectors such that $A(x) \geq A(y)$. Then the area function restricted to the line segment from x to y is strictly decreasing at x .*

Proof. The line segment from x to y is traversed by $(y - x)t + x$ as the parameter t varies from 0 to 1. We have that

$$A((y-x)t+x) = ((y-x)t+x) \cdot ((y-x)t+x) = (y-x) \cdot (y-x)t^2 + 2x \cdot (y-x)t + x \cdot x.$$

Viewing this as a function of t , its derivative at $t = 0$ is $2x \cdot (y - x)$. But since

$$x \cdot x - 2x \cdot y + y \cdot y = (x - y) \cdot (x - y) > 0,$$

we have that $2x \cdot y < x \cdot x + y \cdot y$. So

$$2x \cdot (y - x) = 2x \cdot y - 2x \cdot x < y \cdot y - x \cdot x = A(y) - A(x) \leq 0.$$

This proves the lemma. \square

Lemma 3.4. *Every nonempty compact convex set of weight vectors in \mathbf{R}^n contains a unique weight vector with minimal area.*

Proof. Let C be a nonempty compact convex set of weight vectors in \mathbf{R}^n . Since the area function is continuous, it has a minimum on C . Let x and y be weight vectors in C with minimal area. Lemma 3.3 implies that if $x \neq y$, then the area function restricted to the line segment from x to y is decreasing at x . This is impossible because x and y have minimal area, and so $x = y$.

This proves Lemma 3.4. \square

4. WEIGHT FUNCTIONS WHICH ARE SUMS OF STRICTLY MONOTONIC CUTS

This section deals with weight functions on rectangles that are viewed as conformal quadrilaterals. Eventually such a rectangle will be chosen to be contained in the bar of a dumbbell with the top of the rectangle in the top of the bar and the bottom of the rectangle in the bottom of the bar.

Let R be a rectangle in the plane tiled in the straightforward way by squares with n rows and m columns of squares. Let T_{ij} be the square in row i and column j for $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. With an eye toward combinatorial conformal moduli, we view R as a quadrilateral in the straightforward way. A skinny cut for R is *strictly monotonic* if it contains exactly one tile in every column of R . A weight function ρ on R is a *sum of strictly monotonic skinny cuts* if ρ is a nonnegative linear combination of characteristic functions of strictly monotonic skinny cuts.

Lemma 4.1. *In the situation of the previous paragraph, the weight function ρ is a sum of strictly monotonic skinny cuts if and only if for every $j \in \{1, \dots, m-1\}$ we have that*

$$\sum_{i=1}^n \rho(T_{ij}) = \sum_{i=1}^n \rho(T_{i,j+1})$$

and

$$\sum_{i=1}^k \rho(T_{ij}) \leq \sum_{i=1}^{k+1} \rho(T_{i,j+1}) \text{ and } \sum_{i=1}^k \rho(T_{i,j+1}) \leq \sum_{i=1}^{k+1} \rho(T_{ij}) \text{ for every } k \in \{1, \dots, n-1\}.$$

Proof. Suppose that ρ is a sum of strictly monotonic skinny cuts. Let $j \in \{1, \dots, m-1\}$, and let $k \in \{1, \dots, n-1\}$. Then every skinny cut defining ρ which contains one of T_{1j}, \dots, T_{kj} also contains one of $T_{1,j+1}, \dots, T_{k+1,j+1}$. Since ρ is a nonnegative linear combination of such skinny cuts, it follows that

$$\sum_{i=1}^k \rho(T_{ij}) \leq \sum_{i=1}^{k+1} \rho(T_{i,j+1}).$$

Likewise

$$\sum_{i=1}^k \rho(T_{i,j+1}) \leq \sum_{i=1}^{k+1} \rho(T_{ij}).$$

In the case $k = n$, every skinny cut defining ρ which contains one of T_{1j}, \dots, T_{nj} also contains one of $T_{1j+1}, \dots, T_{nj+1}$. So

$$\sum_{i=1}^n \rho(T_{ij}) \leq \sum_{i=1}^n \rho(T_{i,j+1}).$$

The opposite inequality holds by symmetry, and so the inequality is actually equality. This proves the forward implication of the lemma.

For the backward implication, suppose that these equalities and inequalities are satisfied. We argue by induction on the number N of tiles T_{ij} such that $\rho(T_{ij}) \neq 0$. The statement to be proved is vacuously true if $N = 0$. So suppose that $N > 0$ and that the statement is true for smaller values of N . Let $j \in \{1, \dots, m\}$. Since the zero function is not an allowable weight function and since column sums of ρ -weights are equal, it follows that the sum of the ρ -weights of the tiles in column j is not 0. Let S_j be the tile T_{ij} with the smallest value of i such that $\rho(T_{ij}) \neq 0$. The inequalities imply that $\{S_1, \dots, S_m\}$ is a strictly monotonic skinny cut. Let $c = \min\{\rho(S_1), \dots, \rho(S_m)\}$. Let ρ' be the function on the tiles of R obtained from ρ by subtracting c times the characteristic function of $\{S_1, \dots, S_m\}$. If ρ' is the zero function, then ρ is just c times the characteristic function of $\{S_1, \dots, S_m\}$. Otherwise ρ' is a weight function which satisfies the inequalities and it has value 0 at more tiles than does ρ . So by induction ρ' is a sum of strictly monotonic skinny cuts. It follows that ρ is too.

This completes the proof of Lemma 4.1. \square

Lemma 4.2. *Let R be a rectangle as above tiled by n rows and m columns of squares T_{ij} . Let $x = (x_1, \dots, x_n)$ be a weight vector in \mathbf{R}^n . Then there exists a unique weight function ρ for R with minimal area subject to the conditions that ρ is a sum of strictly monotonic skinny cuts and $\rho(T_{i1}) = x_i$ for every $i \in \{1, \dots, n\}$.*

Proof. If ρ is a weight function for R which is a sum of strictly monotonic skinny cuts such that $\rho(T_{i1}) = x_i$ for $i \in \{1, \dots, n\}$, then the ρ -height of R is $H(x)$. Thus to minimize area, we may restrict to weight functions whose weights are all at most $H(x)$. Lemma 4.1 shows that we are minimizing area over a compact convex subset of \mathbf{R}^{nm} . This subset of \mathbf{R}^{nm} is nonempty because it contains the weight vector corresponding to the weight function ρ for R for which $\rho(T_{ij}) = x_i$ for every i and j . Now Lemma 3.4 completes the proof of Lemma 4.2. \square

5. THE SKINNY CUT FUNCTION

We maintain the setting of Section 4, and continue to let \mathcal{W} denote the set of all weight vectors in \mathbf{R}^n . Lemma 4.2 allows us to define a function $\Phi : \mathcal{W} \rightarrow \mathcal{W}$ as follows. Let $x = (x_1, \dots, x_n) \in \mathcal{W}$. We apply Lemma 4.2 in the case where $m = 2$, so that the rectangle R has n rows and only two columns. Then there exists a unique weight function ρ for R with minimal area subject to the conditions that ρ is a sum of strictly monotonic skinny cuts and $\rho(T_{i1}) = x_i$ for $i \in \{1, \dots, n\}$. Let $y = (y_1, \dots, y_n)$ be the weight vector such that $y_i = \rho(T_{i2})$ for $i \in \{1, \dots, n\}$. We set $\Phi(x) = y$. This defines Φ , which we call the *skinny cut function*. This section is devoted to the investigation of Φ .

Let $x \in \mathcal{W}$, and let $h = H(x)$. Lemma 4.1 implies that $H(\Phi(x)) = h$, so Φ maps \mathcal{W}_h into \mathcal{W}_h . Moreover it is not difficult to see that if r is a positive real number, then $\Phi(rx) = r\Phi(x)$.

Let h be a positive real number, and let $x \in \mathcal{W}_h$. From x we obtain real numbers $p_0 = 0 \leq p_1 \leq p_2 \leq \dots \leq p_n = h$ by setting $p_k = \sum_{i=1}^k x_i$ for $k \in \{0, \dots, n\}$. We denote (p_0, \dots, p_n) by $\pi(x)$ and we call it a *weak partition* of the closed interval $[0, h]$ as opposed to a *strict partition* in which the numbers p_0, \dots, p_n are required to be distinct. We call p_k the k th *partition point* of x . This correspondence gives a bijection between \mathcal{W}_h and the set of all weak partitions of $[0, h]$.

The inequalities of Lemma 4.1 can be easily interpreted in terms of partition points. Let R be a rectangle as before with n rows and two columns of tiles. Let ρ be a weight function on R which is a sum of strictly monotonic skinny cuts. Let $x = (x_1, \dots, x_n)$ and let $y = (y_1, \dots, y_n)$, where $x_i = \rho(T_{i1})$ and $y_i = \rho(T_{i2})$ for each i . Let $(p_0, \dots, p_n) = \pi(x)$, and let $(q_0, \dots, q_n) = \pi(y)$. Lemma 4.1 implies that $H(x) = H(y)$. Then, assuming that $H(x) = H(y)$, the inequalities of Lemma 4.1 are equivalent to the inequalities $q_{k-1} \leq p_k \leq q_{k+1}$ for every $k \in \{1, \dots, n-1\}$. Likewise, assuming that $H(x) = H(y)$, the inequalities of Lemma 4.1 are equivalent to the inequalities $p_{k-1} \leq q_k \leq p_{k+1}$ for every $i \in \{1, \dots, n-1\}$. We say that two weight vectors x and y are *compatible* if $H(x) = H(y)$ and $q_{k-1} \leq p_k \leq q_{k+1}$, equivalently $p_{k-1} \leq q_k \leq p_{k+1}$, for every $k \in \{1, \dots, n-1\}$, where $(p_0, \dots, p_n) = \pi(x)$ and $(q_0, \dots, q_n) = \pi(y)$. This gives us the following reformulation of the definition of Φ , formally stated as a lemma.

Lemma 5.1. *If x is a weight vector, then $\Phi(x)$ is the weight vector with minimal area which is compatible with x .*

Let $x = (x_1, \dots, x_n) \in \mathcal{W}$, and let $(p_0, \dots, p_n) = \pi(x)$. Let $k \in \{1, \dots, n-1\}$. We call p_k a *left leaner* for x if $x_k > x_{k+1}$, and we call p_k a *right leaner* for x if $x_k < x_{k+1}$. The rest of this paragraph explains this terminology. Suppose that p_k is a left leaner for x for some $k \in \{1, \dots, n-1\}$. Then $x_k > x_{k+1}$. The function $f(t) = (t - p_{k-1})^2 + (p_{k+1} - t)^2$ is convex on the closed interval $[p_{k-1}, p_{k+1}]$ with a minimum at the interval's midpoint. This and the inequality $x_k > x_{k+1}$ imply that decreasing p_k while fixing all other partition points of x decreases the area of x . So we view p_k as leaning left toward a position of less area for x . A similar discussion holds for right leaners.

Lemma 5.1 and the previous paragraph imply for every $x \in \mathcal{W}$ that $A(\Phi(x)) < A(x)$ unless x has no leaners. But if x has no leaners, then $x = w_h$, where $h = H(x)$. Corollary 3.2 implies that $\Phi(w_h) = w_h$. Thus for every positive real number h the restriction of Φ to \mathcal{W}_h has a unique fixed point, namely, w_h . We formally state some of the results just proved in the following lemma.

Lemma 5.2. *For every positive real number h the restriction of Φ to \mathcal{W}_h is a function $\Phi|_{\mathcal{W}_h} : \mathcal{W}_h \rightarrow \mathcal{W}_h$ which reduces areas of weight vectors except for at the unique fixed point, w_h .*

Now let x and y be compatible weight vectors with $(p_0, \dots, p_n) = \pi(x)$ and $(q_0, \dots, q_n) = \pi(y)$. Suppose that q_k is a left leaner for y for some $k \in \{1, \dots, n-1\}$. We say that x *blocks* q_k if $q_k = p_{k-1}$. The motivation behind this terminology is that for x and y to be compatible the inequality $q_k \geq p_{k-1}$ must be satisfied, and so even though q_k is leaning left, x prevents us from decreasing q_k to decrease the area of y . Similarly, if q_k is a right leaner for y , then we say that x blocks q_k if $q_k = p_{k+1}$.

Lemma 5.3. *Let x and y be compatible weight vectors. Then every left leaner for y which is blocked by x is a left leaner for x . Similarly, every right leaner for y which is blocked by x is a right leaner for x .*

Proof. Suppose that $x = (x_1, \dots, x_n)$, $y = (y_1, \dots, y_n)$, $\pi(x) = (p_0, \dots, p_n)$, and $\pi(y) = (q_0, \dots, q_n)$. Suppose that q_k is a left leaner for y blocked by x for some $k \in \{1, \dots, n-1\}$. This means that $y_k > y_{k+1}$ and that $q_k = p_{k-1}$. Since $q_{k-1} \geq p_{k-2}$, we have $y_k \leq x_{k-1}$. Since $q_{k+1} \geq p_k$, we have $y_{k+1} \geq x_k$. Combining these inequalities, we obtain that $x_{k-1} \geq y_k > y_{k+1} \geq x_k$, and so p_{k-1} is a left leaner for x . Hence every left leaner for y which is blocked by x is a left leaner for x . A similar argument holds for right leaners.

This proves Lemma 5.3. \square

Lemma 5.4. *Let $x \in \mathcal{W}$. Then every left leaner for $\Phi(x)$ is a left leaner for x blocked by x , and every right leaner for $\Phi(x)$ is a right leaner for x blocked by x . Conversely, if y is a weight vector compatible with x such that every leaner for y is blocked by x , then $y = \Phi(x)$.*

Proof. Let q be a left leaner for $\Phi(x)$. Then q is blocked by x , for otherwise it is possible to decrease $A(\Phi(x))$ by decreasing q . Thus every left leaner for $\Phi(x)$ is blocked by x . Now Lemma 5.3 implies that every left leaner for $\Phi(x)$ is a left leaner for x . Similarly, every right leaner for $\Phi(x)$ is a right leaner for x blocked by x .

Now suppose that y is a weight vector compatible with x such that every leaner for y is blocked by x . Suppose that $y \neq \Phi(x)$. All weight vectors on the line segment from y to $\Phi(x)$ are compatible with x . Lemma 3.3 implies that the area function restricted to this line segment is strictly decreasing at y . But as we move the partition points of y linearly toward the partition points of $\Phi(x)$ we either move leaners away from blocked positions, which increases area, or we move nonleaners, which also increases area. Thus $y = \Phi(x)$.

This proves Lemma 5.4. \square

We next introduce the notion of segments. Let $x \in \mathcal{W}_h$ for some positive real number h , and let $\pi(x) = (p_0, \dots, p_n)$. Let i and j be indices such that p_i is either 0 or a leaner of x , p_j is either h or a leaner of x , $i < j$ and p_k is not a leaner of x if $i < k < j$. Then $x_{i+1} = \dots = x_j$, $x_i \neq x_{i+1}$ if $i > 0$, and $x_{j+1} \neq x_j$ if $j < n$. We call (x_{i+1}, \dots, x_j) a *segment* of x . In other words, the leaners of x parse x into segments such that the coordinates of x in every segment are equal, and segments are maximal with respect to this property. We call p_i and p_j the *endpoints* of the segment (x_{i+1}, \dots, x_j) . By the *value* of a segment we mean the value of any component of x in that segment. By the *dimension* of a segment we mean the number of components in it. By the *height* of a segment we mean the sum of its components. Each endpoint of a segment is either a leaner or an endpoint of the interval $[0, h]$. If the left endpoint of a segment is a left leaner, then we say that it *leans away from* the segment, and if it is a right leaner, then we say that it *leans toward* the segment. The situation is similar for right endpoints.

Let y be a weight vector in the image of Φ . A weight vector x is called a *minimal preimage* of y if it satisfies the following:

- (i) $\Phi(x) = y$.
- (ii) If x' is a weight vector with $\Phi(x') = y$, then $A(x) \leq A(x')$, with equality if and only if $x = x'$.

Lemma 5.5. *Let y be a weight vector, and let t_1, \dots, t_m be the segments of y in order. Then y is in the image of Φ if and only if there is a weight vector x with segments s_1, \dots, s_m in order which satisfies the following:*

- (1) *The height of s_i equals the height of t_i for $i \in \{1, \dots, m\}$.*
- (2) *For every $i \in \{1, \dots, m\}$ the dimension of s_i is the dimension of t_i plus $\alpha_i - \tau_i$ where α_i is the number of endpoints of t_i which lean away from t_i and τ_i is the number of endpoints of t_i which lean toward t_i .*

Furthermore, if y is in the image of Φ , then this vector x is unique and is the minimal preimage of y .

Proof. We first suppose that y is in the image of Φ , and let w be a weight vector such that $\Phi(w) = y$. Suppose that the right endpoint of t_1 leans toward t_1 , namely, that it is a left leaner. Suppose that the right endpoint of t_1 is partition point k of y . Lemma 5.4 implies that this left leaner is blocked by w , and so it is partition point number $k - 1$ of w . This proves that if the right endpoint of t_1 leans toward t_1 , then the dimension of t_1 is at least 2. Similarly, if the left endpoint of the last segment t_m leans toward t_m , then the dimension of t_m is at least 2. More generally, if both endpoints of some segment of y lean toward that segment, then the dimension of that segment is at least 3. With the results of this paragraph and the fact that $\sum_{i=1}^m (\alpha_i - \tau_i) = 0$, we see that conditions 1 and 2 in the statement of Lemma 5.5 uniquely determine a weight vector x .

For the converse, suppose that x is a weight vector with segments s_1, \dots, s_m in order which satisfies conditions 1 and 2 of the lemma. An induction argument on i shows that the endpoints of t_i other than 0 and h are blocked by x and that x is compatible with y . By Lemma 5.4, $\Phi(x) = y$ and so y is in the image of Φ .

For the last statement, suppose that y is in the image of Φ . Let x be the weight vector which satisfies conditions 1 and 2 of the lemma, and let x' be a weight vector with $\Phi(x') = y$. Let $i \in \{1, \dots, m\}$. Because the endpoints of t_i other than 0 and h are blocked by x' , the number of components of x' between these endpoints equals the number of components of x between these endpoints, namely, the dimension of s_i . But since the components of s_i are all equal, Corollary 3.2 implies that the area of s_i is at most the sum of the squares of the corresponding components of x' . Thus $A(x) \leq A(x')$, and equality holds if and only if $x = x'$.

This proves Lemma 5.5. □

This section has thus far been concerned with the definition of Φ and basic relationships between a weight vector and its image under Φ . We now turn to area estimates. The next result, Theorem 5.6, prepares for Theorem 5.7, which is the key ingredient in our proof of the dumbbell theorem.

Theorem 5.6. *Let h be a positive real number, and let $x \in \mathcal{W}_h$. Then there exists a weight vector $w \in \mathcal{W}_h$ satisfying the following conditions.*

- (1) *Both w and $\Phi(w)$ have at most two segments with nonzero values. Furthermore, if they have two segments with nonzero values, then these segments are consecutive.*
- (2) *The weight vector w is the minimal preimage of $\Phi(w)$.*
- (3) *$A(w) \leq A(x)$.*
- (4) *$A(\Phi(w)) \geq A(\Phi(x))$.*

Proof. Set

$$W = \{z \in \mathcal{W}_h : A(z) \leq A(x) \text{ and } A(\Phi(z)) \geq A(\Phi(x))\}.$$

Since every element of W satisfies conditions 3 and 4, it suffices to prove that W contains a weight vector satisfying conditions 1 and 2. We will do this by first proving that W contains an element with minimal area. It follows that w is the minimal preimage of $\Phi(w)$, giving condition 2. We then prove that w satisfies condition 1.

In this paragraph we prove that W contains an element with minimal area. Since $x \in W$, W is not empty. The set \mathcal{W}_h is compact, and so the closure \overline{W} of W in \mathcal{W}_h is compact. Since the area function is continuous, there exists a weight vector w in \overline{W} with minimal area. Because the area function is continuous and $w \in \overline{W}$, we have that $A(w) \leq A(x)$. It remains to prove that $A(\Phi(w)) \geq A(\Phi(x))$. For this let w' be a weight vector in W near w . Lemma 5.1 shows that $\Phi(w)$ is compatible with w . Because w' is near w , there exists a weight vector y near $\Phi(w)$ which is compatible with w' . Because y is near $\Phi(w)$ the area of y is not much larger than the area of $\Phi(w)$. Now Lemma 5.1 shows that the area of $\Phi(w')$ is not much larger than the area of $\Phi(w)$. Since $A(\Phi(w')) \geq A(\Phi(x))$, it follows that $A(\Phi(w)) \geq A(\Phi(x))$. Thus $w \in W$. This proves that W contains an element with minimal area.

We fix an element w in W with minimal area. Lemma 5.5 implies that w is the minimal preimage of $\Phi(w)$. To prove Theorem 5.6 it suffices to prove that w cannot have nonconsecutive segments with nonzero values. We do this by contradiction. So suppose that w has nonconsecutive segments with nonzero values. Let s_1, \dots, s_m be the segments of w in order with dimensions d_1, \dots, d_m , values v_1, \dots, v_m , and heights h_1, \dots, h_m , so that $h_i = d_i v_i$ for every $i \in \{1, \dots, m\}$.

In this paragraph we show that w has two consecutive segments with nonzero values. We proceed by contradiction: suppose that w does not have two consecutive segments with nonzero values. Then every other value of w is 0. Let $i \in \{1, \dots, m\}$ such that $v_i \neq 0$, $v_{i+1} = 0$, and $v_{i+2} \neq 0$. Because of the symmetry with respect to the order of the components of w , we may, and do, assume that $v_i \leq v_{i+2}$. Let A_0 be the sum of the areas of the segments of w other than s_i , s_{i+1} , and s_{i+2} . Then

$$A(w) = A_0 + \frac{h_i^2}{d_i} + \frac{h_{i+2}^2}{d_{i+2}}.$$

Similarly, if B_0 is the sum of the areas of the segments of $\Phi(w)$ other than those corresponding to s_i , s_{i+1} , and s_{i+2} , then

$$A(\Phi(w)) = B_0 + \frac{h_i^2}{e_i} + \frac{h_{i+2}^2}{e_{i+2}},$$

where

$$e_i = \begin{cases} d_i + 1, & \text{if } i = 1, \\ d_i + 2, & \text{if } i > 1, \end{cases}$$

and

$$e_{i+2} = \begin{cases} d_{i+2} + 1, & \text{if } i = m, \\ d_{i+2} + 2, & \text{if } i < m. \end{cases}$$

We construct a new weight vector w' by modifying w as follows. Where w has segments s_i , s_{i+1} , s_{i+2} with heights h_i , 0, h_{i+2} and dimensions d_i , d_{i+1} , d_{i+2} , the weight vector w' has segments with heights 0, h_i , h_{i+2} and dimensions d'_i , $d_i + 1$, d_{i+2} . The heights and dimensions of all other segments of w' equal the

corresponding heights and dimensions of w except that if $i > 1$, then w' has one fewer segment than w because segment s_{i-1} of w has value 0 and so the components of w in s_{i+1} must be adjoined to s_{i-1} . So

$$A(w') = A_0 + \frac{h_i^2}{d_i + 1} + \frac{h_{i+2}^2}{d_{i+2}} < A(w) \leq A(x).$$

Using Lemma 5.5 we see that

$$A(\Phi(w')) = B_0 + \frac{h_i^2}{d_i + 1} + \frac{h_{i+2}^2}{e_{i+2}} \geq A(\Phi(w)) \geq A(\Phi(x)).$$

This contradicts the fact that w is an element in W with minimal area. Thus w has two consecutive segments with nonzero values. If w only has two segments with nonzero values, then we are done. Hence we may assume that w has at least three segments with nonzero values.

Next suppose that w' is a weight vector with the same number of segments as w , the segments of w' have the same dimensions as the corresponding segments of w but the leaners of w' are obtained by perturbing the leaners of w , equivalently, the heights of the segments of w' are obtained by perturbing the heights of the segments of w . If this perturbation is small enough, then Lemma 5.4 implies that $\Phi(w')$ is obtained from $\Phi(w)$ by exactly the same perturbation. We will prove that there exists such a perturbation such that $A(w') < A(w)$ and $A(\Phi(w')) = A(\Phi(w))$. Thus w' is an element of W with smaller area than w , a contradiction which will complete the proof of Theorem 5.6.

To begin the construction of such a perturbation, we note that the area of w is

$$\sum_{i=1}^m d_i v_i^2 = \sum_{i=1}^m d_i \left(\frac{h_i}{d_i} \right)^2 = \sum_{i=1}^m \frac{1}{d_i} h_i^2 = \sum_{i=1}^m a_i h_i^2,$$

where $a_i = d_i^{-1}$ for every $i \in \{1, \dots, m\}$. Recall that w is the minimal preimage of $\Phi(w)$. As above, if t_1, \dots, t_m are the segments of $\Phi(w)$ in order with dimensions $A_1^{-1}, \dots, A_m^{-1}$, then the area of $\Phi(w)$ is

$$\sum_{i=1}^m A_i h_i^2.$$

This leads us to define functions $f: \mathbf{R}^m \rightarrow \mathbf{R}$ and $F: \mathbf{R}^m \rightarrow \mathbf{R}$ so that

$$f(x_1, \dots, x_m) = \sum_{i=1}^m a_i (x_i + h_i)^2$$

and

$$F(x_1, \dots, x_m) = \sum_{i=1}^m A_i (x_i + h_i)^2.$$

To prove Theorem 5.6, it suffices to find points $x = (x_1, \dots, x_m) \in \mathbf{R}^m$ arbitrarily near the origin 0 such that $\sum_{i=1}^m x_i = 0$, $f(x) < f(0)$, $F(x) = F(0)$, and $x_i \geq 0$ if $h_i = 0$.

We do this while fixing all but three components of (x_1, \dots, x_m) . We specify these components later in this paragraph. Since only three components of (x_1, \dots, x_m) are allowed to vary, the functions f and F are really functions of

three variables. To simplify notation, we now view f and F as functions from \mathbf{R}^3 to \mathbf{R} with

$$f(x, y, z) = a(x + x_0)^2 + b(y + y_0)^2 + c(z + z_0)^2 + d$$

and

$$F(x, y, z) = A(x + x_0)^2 + B(y + y_0)^2 + C(z + z_0)^2 + D.$$

Since w has two consecutive segments with nonzero values, by Lemma 5.5, $\Phi(w)$ has two consecutive segments with nonzero values. Let the variable x correspond to a segment of $\Phi(w)$ with nonzero value which is adjacent to a segment with larger value. Let z correspond to a segment of $\Phi(w)$ with maximal value. Let y correspond to any segment of $\Phi(w)$ with nonzero value not already chosen. Note that $x_0, y_0, z_0 > 0$. Because the segment of $\Phi(w)$ corresponding to x has an endpoint leaning away from it, its dimension A^{-1} is at most as large as the dimension a^{-1} of the corresponding segment of w . Thus $A^{-1} \leq a^{-1}$, and so $A \geq a$. Because the segment of $\Phi(w)$ corresponding to z has no endpoint leaning away from it, $C^{-1} > c^{-1}$, and so $C < c$. Because the value of the segment of $\Phi(w)$ corresponding to z is larger than the value of the segment corresponding to x , we have $Cz_0 > Ax_0$. To prove Theorem 5.6 it suffices to find points $(x, y, z) \in \mathbf{R}^3$ arbitrarily near $(0, 0, 0)$ such that $x + y + z = 0$, $f(x, y, z) < f(0, 0, 0)$, and $F(x, y, z) = F(0, 0, 0)$.

The set of all solutions (x, y, z) to the equation $f(x, y, z) = f(0, 0, 0)$ is an ellipsoid containing the point $(0, 0, 0)$. The gradient of f at $(0, 0, 0)$ is $(2ax_0, 2by_0, 2cz_0)$. Because ax_0, by_0, cz_0 are not all equal, the plane given by $x + y + z = 0$ contains the point $(0, 0, 0)$, but it is not tangent to the ellipsoid. Thus this plane intersects the ellipsoid in an ellipse. It likewise intersects the ellipsoid given by $F(x, y, z) = F(0, 0, 0)$ in an ellipse. These ellipses both lie in this plane and they have a point in common. If they intersect transversely at $(0, 0, 0)$, then it is easy to find points arbitrarily near $(0, 0, 0)$ which lie within the first ellipse and lie on the second one. In other words, Theorem 5.6 is true if the ellipses intersect transversely.

Finally, we assume that the ellipses are tangent at $(0, 0, 0)$. Let L be the line in \mathbf{R}^3 tangent to the ellipses at $(0, 0, 0)$. Then L lies in the plane given by $x + y + z = 0$ and it lies in the tangent planes to both of the ellipsoids. These three planes have normal vectors given by $(1, 1, 1)$, (ax_0, by_0, cz_0) and (Ax_0, By_0, Cz_0) . Thus these three vectors are linearly dependent. Thus the same is true for the columns of the matrix

$$\begin{bmatrix} 1 & ax_0 & Ax_0 \\ 1 & by_0 & By_0 \\ 1 & cz_0 & Cz_0 \end{bmatrix}.$$

The cross product of any two rows is orthogonal to all three rows. The cross product of row 1 and row 3 is

$$v = ((aC - Ac)x_0z_0, Ax_0 - Cz_0, cz_0 - ax_0).$$

Since $a \leq A$ and $C < c$, we have $(aC - Ac)x_0z_0 < 0$. Since $Cz_0 > Ax_0$, we have $Ax_0 - Cz_0 < 0$. Since v is orthogonal to the rows of the preceding matrix, it is also orthogonal to the rows of

$$\begin{bmatrix} \frac{1}{x_0} & a & A \\ \frac{1}{y_0} & b & B \\ \frac{1}{z_0} & c & C \end{bmatrix}.$$

So

$$(aC - Ac)x_0z_0 \left(\frac{1}{x_0}, \frac{1}{y_0}, \frac{1}{z_0} \right) + (Ax_0 - Cz_0)(a, b, c) + (cz_0 - ax_0)(A, B, C) = (0, 0, 0).$$

Solving for (a, b, c) , we find that

$$(a, b, c) = r \left(\frac{1}{x_0}, \frac{1}{y_0}, \frac{1}{z_0} \right) + s(A, B, C)$$

for some real numbers r and s with $r < 0$. Viewing $f(x, y, z)$ and $F(x, y, z)$ as dot products of vectors except for the final additive constants, we have that

$$\begin{aligned} f(x, y, z) - f(0, 0, 0) &= (a, b, c) \cdot ((x + x_0)^2, (y + y_0)^2, (z + z_0)^2) - (a, b, c) \cdot (x_0^2, y_0^2, z_0^2) \\ &= (a, b, c) \cdot (x^2 + 2x_0x, y^2 + 2y_0y, z^2 + 2z_0z) \\ &= r \left(\frac{1}{x_0}, \frac{1}{y_0}, \frac{1}{z_0} \right) \cdot (x^2 + 2x_0x, y^2 + 2y_0y, z^2 + 2z_0z) \\ &\quad + s(A, B, C) \cdot (x^2 + 2x_0x, y^2 + 2y_0y, z^2 + 2z_0z) \\ &= r \left[\left(\frac{1}{x_0}x^2 + \frac{1}{y_0}y^2 + \frac{1}{z_0}z^2 \right) + 2(x + y + z) \right] + s[F(x, y, z) - F(0, 0, 0)] \\ &= r \left(\frac{1}{x_0}x^2 + \frac{1}{y_0}y^2 + \frac{1}{z_0}z^2 \right) + s[F(x, y, z) - F(0, 0, 0)]. \end{aligned}$$

So if (x, y, z) is any point on the ellipse corresponding to F , then $F(x, y, z) - F(0, 0, 0) = 0$ and so $f(x, y, z) - f(0, 0, 0) < 0$, as desired.

This proves Theorem 5.6. \square

Theorem 5.7. *Let h be a positive real number. Let x_1, \dots, x_{3n} be weight vectors in \mathcal{W}_h such that $A(x_{i+1}) \leq A(\Phi(x_i))$ for every $i \in \{1, \dots, 3n-1\}$. Then $x_{3n} = w_h$.*

Proof. Lemma 5.2 implies that Φ does not increase area, and so

$$A(x_1) \geq \dots \geq A(x_{3n}).$$

Lemma 5.2 and Corollary 3.2 furthermore imply that if two of x_1, \dots, x_{3n} have the same area, then $x_{3n} = w_h$. So we may, and do, assume that x_1, \dots, x_{3n} are distinct.

Since $\Phi(hx) = h\Phi(x)$ for every weight vector $x \in \mathcal{W}_1$, we may, and do, assume that $h = 1$.

Let $i \in \{1, \dots, 3n-1\}$. Theorem 5.6 implies that there exists a weight vector $x'_i \in \mathcal{W}_1$ such that both x'_i and $\Phi(x'_i)$ have at most two segments with nonzero values, the weight vector x'_i is the minimal preimage of $\Phi(x'_i)$, $A(x'_i) \leq A(x_i)$ and $A(\Phi(x'_i)) \geq A(\Phi(x_i))$. Set $x'_{3n} = x_{3n}$. Then

$$A(x'_{i+1}) \leq A(x_{i+1}) \leq A(\Phi(x_i)) \leq A(\Phi(x'_i))$$

for every $i \in \{1, \dots, 3n-1\}$. Thus to prove Theorem 5.7, we may replace x_1, \dots, x_{3n} by x'_1, \dots, x'_{3n} . In other words we may, and do, assume for every $i \in \{1, \dots, 3n-1\}$ that x_i and $\Phi(x_i)$ do not have nonconsecutive segments with nonzero values and that x_i is the minimal preimage of $\Phi(x_i)$.

At this point the argument splits into two cases which depend on how many of the weight vectors x_1, \dots, x_{3n} have segments with value 0. In Case 1 we prove Theorem 5.7 under the assumption that at least n of the weight vectors x_1, \dots, x_{3n}

have a segment with value 0. In Case 2 we prove Theorem 5.7 under the assumption that fewer than n of the weight vectors x_1, \dots, x_{3n} have a segment with value 0.

Case 1. Assume that at least n of the weight vectors x_1, \dots, x_{3n} have a segment with value 0. Let z_1, \dots, z_n be distinct elements of $\{x_1, \dots, x_{3n}\}$ in order each of which has a segment with value 0. Then $A(z_{i+1}) \leq A(\Phi(z_i))$ for every $i \in \{1, \dots, n-1\}$.

In this paragraph we show that in Case 1 it suffices to prove for every $i \in \{1, \dots, n-1\}$ that

$$(5.8) \quad A(z_{i+1})(A(z_i) + 1) - A(z_i) \leq 0.$$

Indeed, this inequality is equivalent to the inequality

$$A(z_{i+1}) \leq \frac{1}{A(z_i)^{-1} + 1}.$$

Since $H(z_1) = 1$, we have $A(z_1) \leq 1$. A straightforward induction argument based on the preceding displayed inequality shows that $A(z_i) \leq 1/i$ for every $i \in \{1, \dots, n\}$. In particular $A(z_n) \leq 1/n$. But Corollary 3.2 implies that w_1 is the unique element in \mathcal{W}_1 with area at most $1/n$. Hence $z_n = x_{3n} = w_1$. Thus to prove Theorem 5.7 in Case 1 it suffices to prove line (5.8) for every $i \in \{1, \dots, n-1\}$.

In this paragraph we assume that there exists $i \in \{1, \dots, n-1\}$ such that z_i has only one segment with nonzero value. Let d be the dimension of this segment. Then $A(z_i) = 1/d$. Moreover $\Phi(z_i)$ also has just one segment with nonzero value with dimension either $d+1$ or $d+2$. Hence $A(\Phi(z_i))$ is either $1/(d+1)$ or $1/(d+2)$. So line (5.8) holds. Thus to prove Theorem 5.7 in Case 1 it suffices to prove line (5.8) under the assumption that z_i has two segments with nonzero value.

So let $i \in \{1, \dots, n-1\}$, and suppose that z_i has two consecutive segments s and t in order with nonzero value. Let p be the right endpoint of s , equivalently, the left endpoint of t . Let d be the dimension of s , and let e be the dimension of t . The value of s is p/d , and the value of t is $(1-p)/e$. It follows that

$$(5.9) \quad A(z_i) = d^{-1}p^2 + e^{-1}(1-p)^2.$$

Let d' and e' be the dimensions of the segments s' and t' of $\Phi(z_i)$ corresponding to s and t . Then

$$A(\Phi(z_i)) = (d')^{-1}p^2 + (e')^{-1}(1-p)^2.$$

By symmetry we may assume that

$$\frac{p}{d'} \leq \frac{1-p}{e'}.$$

The last inequality implies that p is a right leaner for $\Phi(z_i)$. Hence Lemma 5.4 implies that p is a right leaner for z_i and so

$$\frac{p}{d} \leq \frac{1-p}{e}.$$

Combining this with line (5.9) yields that

$$A(z_i) = d^{-1}p^2 + e^{-1}(1-p)^2 \leq e^{-1}p(1-p) + e^{-1}(1-p)^2 = e^{-1}(1-p).$$

Similarly, we have that

$$(5.10) \quad A(\Phi(z_i)) \leq (e')^{-1}(1-p).$$

Because z_i has a segment with value 0 and z_i is the minimal preimage of $\Phi(z_i)$, it is also true that $\Phi(z_i)$ has a segment with value 0. These facts and Lemma 5.5 imply that if $\Phi(z_i)$ has a segment with value 0 preceding s' , then $d' = d$ and $e' \geq e + 1$. If $\Phi(z_i)$ does not have a segment with value 0 preceding s' , then $d' = d - 1$ and $e' = e + 2$.

Suppose that $\Phi(z_i)$ has a segment with value 0 preceding s' . The previous paragraph shows that $d' = d$ and $e' \geq e + 1$. Combining this with line (5.10) yields that $A(\Phi(z_i)) \leq (e + 1)^{-1}(1 - p)$. Hence

$$\begin{aligned}
& A(z_{i+1})(A(z_i) + 1) - A(z_i) \\
& \leq A(\Phi(z_i))(A(z_i) + 1) - A(z_i) \\
& \leq A(\Phi(z_i))A(z_i) + A(\Phi(z_i)) - A(z_i) \\
& \leq \frac{1}{e(e+1)}(1-p)^2 + \frac{1}{d'}p^2 + \frac{1}{e'}(1-p)^2 - \frac{1}{d}p^2 - \frac{1}{e}(1-p)^2 \\
& \leq \frac{1}{e(e+1)}(1-p)^2 + \frac{1}{d}p^2 + \frac{1}{e+1}(1-p)^2 - \frac{1}{d}p^2 - \frac{1}{e}(1-p)^2 \\
& \leq \frac{1}{e(e+1)}(1-p)^2 - \frac{1}{e(e+1)}(1-p)^2 \\
& \leq 0.
\end{aligned}$$

Thus line (5.8) is true if $\Phi(z_i)$ has a segment with value 0 preceding s' .

Next assume that $\Phi(z_i)$ does not have a segment with value 0 preceding s' . Then the next-to-last paragraph shows that $d' = d - 1$ and $e' = e + 2$. Combining this with line (5.10) yields that $A(\Phi(z_i)) \leq (e + 2)^{-1}(1 - p)$. We also have that

$$\frac{p}{d-1} = \frac{p}{d'} \leq \frac{1-p}{e'} = \frac{1-p}{e+2}.$$

Hence

$$\begin{aligned}
& A(z_{i+1})(A(z_i) + 1) - A(z_i) \\
& \leq A(\Phi(z_i))(A(z_i) + 1) - A(z_i) \\
& \leq A(\Phi(z_i))A(z_i) + A(\Phi(z_i)) - A(z_i) \\
& \leq \frac{1}{e(e+2)}(1-p)^2 + \frac{1}{d'}p^2 + \frac{1}{e'}(1-p)^2 - \frac{1}{d}p^2 - \frac{1}{e}(1-p)^2 \\
& \leq \frac{1}{e(e+2)}(1-p)^2 + \frac{1}{d-1}p^2 + \frac{1}{e+2}(1-p)^2 - \frac{1}{d}p^2 - \frac{1}{e}(1-p)^2 \\
& \leq \frac{1}{e(e+2)}(1-p)^2 + \frac{1}{(d-1)d}p^2 - \frac{2}{e(e+2)}(1-p)^2 \\
& \leq \frac{1}{e(e+2)}(1-p)^2 + \frac{1}{e(e+2)}(1-p)^2 - \frac{2}{e(e+2)}(1-p)^2 \\
& \leq 0.
\end{aligned}$$

Thus line (5.8) is true if $\Phi(z_i)$ does not have a segment with value 0 preceding s' .

The proof of Theorem 5.7 is now complete in Case 1, namely, Theorem 5.7 is proved if at least n of the weight vectors x_1, \dots, x_{3n} have a segment with value 0.

Case 2. Suppose that fewer than n of the weight vectors x_1, \dots, x_{3n} have a segment with value 0. Then at least $2n$ of the weight vectors x_1, \dots, x_{3n} do not have a segment with value 0. Let z_1, \dots, z_{2n} be distinct elements of $\{x_1, \dots, x_{3n}\}$ in order

none of which has a segment with value 0. Then $A(z_{i+1}) \leq A(\Phi(z_i))$ for every $i \in \{1, \dots, 2n-1\}$. As we have seen, to finish the proof of Theorem 5.7, it suffices to prove that $z_{2n} = w_1$. We do this by contradiction. Suppose that $z_{2n} \neq w_1$. As we have seen, it follows that $z_i \neq w_1$ for $i \in \{1, \dots, 2n\}$.

Let $i \in \{1, \dots, 2n-1\}$. Then z_i has at most two segments, z_i has no segment with value 0, and z_i is the minimal preimage of $\Phi(z_i)$. The only weight vector in \mathcal{W}_1 with one segment is w_1 , so z_i has two segments neither of which has value 0.

Let the first segment of z_i have right endpoint p and dimension d . Then the second segment of z_i has left endpoint p and dimension $n-d$. So

$$\begin{aligned} A(z_i) &= \frac{1}{d}p^2 + \frac{1}{(n-d)}(1-p)^2 = \frac{(n-d)p^2 + d(1-p)^2}{d(n-d)} \\ &= \frac{np^2 - dp^2 + d - 2dp + dp^2}{d(n-d)} = \frac{n^2p^2 - 2dnp + dn}{d(n-d)n} \\ &= \frac{(np-d)^2 - d^2 + dn}{d(n-d)n} = \frac{(d-np)^2}{d(n-d)n} + \frac{1}{n}. \end{aligned}$$

Set

$$B(z_i) = \frac{(d-np)^2}{d(n-d)}.$$

Then

$$A(z_i) = \frac{1}{n}B(z_i) + \frac{1}{n}.$$

Since z_i has two segments neither of which has value 0, the same is true of $\Phi(z_i)$. We define $B(\Phi(z_i))$ in the same way that we define $B(z_i)$. Then the given inequality $A(z_{i+1}) \leq A(\Phi(z_i))$ is equivalent to the inequality $B(z_{i+1}) \leq B(\Phi(z_i))$. We focus on this latter inequality.

Let a be the positive real number such that $B(z_i) = a^2$. Then

$$\begin{aligned} \frac{(d-np)^2}{d(n-d)} &= a^2, \\ (d-np)^2 &= a^2d(n-d), \\ d^2 - nd + a^{-2}(d-np)^2 &= 0, \\ \left(d - \frac{n}{2}\right)^2 + a^{-2}(d-np)^2 &= \frac{n^2}{4}, \\ \left(\frac{2d}{n} - 1\right)^2 + a^{-2}\left(\frac{2d}{n} - 2p\right)^2 &= 1. \end{aligned}$$

We are led to consider the family \mathcal{E} of all ellipses of the form $x^2 + a^{-2}y^2 = 1$, where a is a positive real number. The open upper halves of the ellipses in \mathcal{E} fill the open half infinite strip $S = \{(x, y) \in \mathbf{R}^2 : |x| < 1, y > 0\}$.

The weight vector z_i determines the point

$$P(z_i) = (x_i, y_i) = \left(\frac{2d}{n} - 1, \frac{2d}{n} - 2p\right)$$

in \mathbf{R}^2 , and we next show that we may assume that $P(z_i) \in S$. It is easy to see that $|x_i| < 1$, so what is needed is the inequality $y_i > 0$. Since z_i and $\Phi(z_i)$ have exactly

one leaner, by symmetry we may, and do, assume that this leaner is a right leaner. Hence

$$\begin{aligned}\frac{1}{d}p &< \frac{1}{n-d}(1-p), \\ p(n-d) &< d(1-p), \\ pn &< d.\end{aligned}$$

It follows that $y_i > 0$, and so $P(z_i) \in S$. It is even true that $y_i < x_i + 1$, and so $P(z_i)$ lies in the triangle $T = \{(x, y) \in \mathbf{R}^2 : |x| < 1, 0 < y < x + 1\}$.

The point of \mathbf{R}^2 corresponding to $\Phi(z_i)$ likewise lies in T . Moreover since the first segment of z_i has right endpoint p and dimension d , it follows that the first segment of $\Phi(z_i)$ has right endpoint p and dimension $d - 1$. In other words, the point of T corresponding to $\Phi(z_i)$ is

$$\left(\frac{2(d-1)}{n} - 1, \frac{2(d-1)}{n} - 2p\right) = \left(x_i - \frac{2}{n}, y_i - \frac{2}{n}\right).$$

Now we can reformulate our problem as follows. Let $\tau: \mathbf{R}^2 \rightarrow \mathbf{R}^2$ be the translation defined by $\tau(x, y) = (x - \frac{2}{n}, y - \frac{2}{n})$. We begin with a point $t_1 \in T$. Let E_1 be the ellipse in \mathcal{E} containing t_1 . We have that $\tau(t_1) \in T$ and $\tau(t_1)$ lie within E_1 . Let E'_1 be the ellipse in \mathcal{E} containing $\tau(t_1)$. Let t_2 be a point of T either on or within E'_1 . We iterate this construction to obtain points $t_1, \dots, t_{2n} \in T$. To complete the proof of Theorem 5.7 it suffices to prove that it is impossible to construct points t_1, \dots, t_{2n} in this way.

To this end, let A be the subset of \mathbf{R}^2 which is the union of the closed line segment joining $(0, 0)$ with $(0, 1)$ and the closed line segment joining $(0, 1)$ with $(1, 2)$ (see Figure 9). We define a function $\alpha: T \rightarrow \mathbf{R}$ as follows. Let $t \in T$. There exists a unique ellipse $E \in \mathcal{E}$ which contains t , and $A \cap E$ consists of a single point. Let $\alpha(t)$ be the y -coordinate of this point.




FIGURE 9. Defining $\alpha(t)$.

We will prove that if t_1, \dots, t_{2n} are points in T as in the next-to-last paragraph, then $\alpha(t_{i+1}) < \alpha(t_i) - 1/n$ for every $i \in \{1, \dots, 2n - 1\}$. This suffices to complete

the proof of Theorem 5.7. Because t_{i+1} lies either on or within the ellipse in \mathcal{E} containing $\tau(t_i)$, it actually suffices to prove that $\alpha(\tau(t_i)) < \alpha(t_i) - 1/n$ for every $i \in \{1, \dots, 2n - 1\}$, which is what we do.

So let $i \in \{1, \dots, 2n - 1\}$. Let E_i and E'_i be the ellipses in \mathcal{E} which contain t_i and $\tau(t_i)$.

First suppose that the x -coordinates of both t_i and $\tau(t_i)$ are nonnegative (see Figure 10). In this case it is clear that we actually have that $\alpha(\tau(t_i)) < \alpha(t_i) - 2/n$, as desired.




FIGURE 10. When the x -coordinates of both t_i and $\tau(t_i)$ are nonnegative.

Next suppose that the x -coordinates of both t_i and $\tau(t_i)$ are nonpositive (see Figure 11). Since $x^2 + a^{-2}y^2 = 1$, we have that $y^2 = a^2(1 - x^2)$. Differentiating




FIGURE 11. When the x -coordinates of both t_i and $\tau(t_i)$ are nonpositive.

this with respect to x yields that $2yy' = -2a^2x$. Solving for y' and replacing a^2 with $y^2/(1 - x^2)$ yields that

$$(5.11) \quad y' = \frac{xy}{x^2 - 1}.$$

Since $y < x + 1$ and $x/(x^2 - 1) \geq 0$ if $-1 < x \leq 0$, we have that $y' < x/(x - 1)$. It follows that $y' < 1/2$ throughout T . This implies that the vertical distance from t_i to E'_i is at least $1/n$. Line (5.11) implies that E_i increases more than E'_i as they approach A from the left. It follows that $\alpha(\tau(t_i)) < \alpha(t_i) - 1/n$, as desired.

Finally, if the x -coordinate of t_i is positive and the x -coordinate of $\tau(t_i)$ is negative, then an argument similar to that in the previous paragraph shows again that $\alpha(\tau(t_i)) < \alpha(t_i) - 1/n$.

This completes the proof of Theorem 5.7. \square

6. THE PROOF OF THE DUMBBELL THEOREM

We prove the dumbbell theorem in this section.

Let D be a dumbbell. Let ρ be a fat flow optimal weight function for D . Let $h = H_\rho$. We scale ρ so that it is a sum of minimal skinny cuts, making h an integer. Suppose that D has bar height n . Let C_1, \dots, C_{3n} be the leftmost $3n$ columns in B in left-to-right order, and let D_1, \dots, D_{3n} be the rightmost $3n$ columns in B in right-to-left order.

For every $i \in \{1, \dots, 3n\}$, applying ρ to the tiles in C_i obtains a weight vector $y_i \in \mathbf{R}^n$ and applying ρ to the tiles in D_i obtains a weight vector $z_i \in \mathbf{R}^n$. From the assumption that $H_\rho = h$, it follows that $H(y_1) \geq h$, and so by scaling y_1 we can obtain a weight vector $x_1 \in \mathcal{W}_h$ such that $A(x_1) \leq A(y_1)$. Let j be the largest element of $\{1, \dots, 3n\}$ such that there exist weight vectors $x_1, \dots, x_j \in \mathcal{W}_h$ with $A(x_1) \leq A(y_1)$ and if $j > 1$, then $A(x_j) \geq A(y_j)$ and $A(x_{i+1}) \leq A(\Phi(x_i))$ for every $i \in \{1, \dots, j-1\}$. Similarly, let k be the largest element of $\{1, \dots, 3n\}$ such that there exist weight vectors $x_1, \dots, x_k \in \mathcal{W}_h$ with $A(x_1) \leq A(z_1)$ and if $k > 1$, then $A(x_k) \geq A(z_k)$ and $A(x_{i+1}) \leq A(\Phi(x_i))$ for every $i \in \{1, \dots, k-1\}$.

In this paragraph we define a new weight function σ on D . Let t be a tile of D . If t is contained in either one of the balls of D or one of the columns $C_1, \dots, C_{j-1}, D_{k-1}, \dots, D_1$, then $\sigma(t) = \rho(t)$. If t is in B and strictly between C_{3n} and D_{3n} , then $\sigma(t) = h/n$. Now suppose that t is contained in C_j . We have that ρ is a sum of h minimal skinny cuts. We define $\sigma(t)$ to be the number of these skinny cuts c with the property that as c proceeds from the left side of D to C_j , the tile t is the first tile of C_j in c . This defines $\sigma(t)$ for every tile t in C_j . These values determine a weight vector $\bar{x}_j \in \mathcal{W}_h$ in the straightforward way. If $j < 3n$, then we use the weight vector $x_{j+1} = \Phi(\bar{x}_j)$ in the straightforward way to define $\sigma(t)$ for every tile t in C_{j+1} . We inductively set $x_{i+1} = \Phi(x_i)$ and use this weight vector in the straightforward way to define $\sigma(t)$ for every tile t in C_{i+1} for every $i \in \{j+1, \dots, 3n-1\}$. This defines $\sigma(t)$ for every tile t in C_j, \dots, C_{3n} . We define $\sigma(t)$ for every tile t in D_k, \dots, D_{3n} analogously. The definition of σ is now complete.

By assumption there exist weight vectors $x_1, \dots, x_j \in \mathcal{W}_h$ with $A(x_1) \leq A(y_1)$ and if $j > 1$, then $A(x_j) \geq A(y_j)$ and $A(x_{i+1}) \leq A(\Phi(x_i))$ for every $i \in \{1, \dots, j-1\}$. In defining σ we constructed weight vectors $\bar{x}_j, x_{j+1}, \dots, x_{3n} \in \mathcal{W}_h$ with $x_{j+1} = \Phi(\bar{x}_j)$ and $x_{i+1} = \Phi(x_i)$ for every $i \in \{j+1, \dots, 3n-1\}$. We redefine x_j to be \bar{x}_j . Then x_1, \dots, x_{3n} are weight vectors in \mathcal{W}_h such that $A(x_{i+1}) \leq A(\Phi(x_i))$ for every $i \in \{1, \dots, 3n-1\}$. Theorem 5.7 implies that $x_{3n} = w_h$. The situation is analogous at the right end of B .

In this paragraph we prove that $H_\sigma = h$. Because σ is constant on the union of the columns of B between C_{3n-1} and D_{3n-1} , the restriction of σ to these columns is a sum of skinny cuts. The definition of Φ implies that these skinny cuts can be

extended to the columns of B from C_j to D_k so that the restriction of σ to the columns from C_j to D_k is a sum of skinny cuts. The definition of σ implies that these skinny cuts can be extended to all of D so that σ is greater than or equal to a weight function which is a sum of h skinny cuts. We conclude that $H_\sigma \geq h$. But since the columns of D from C_j to D_k have σ -height h , it is in fact the case that $H_\sigma = h$.

In this paragraph we prove that $A_\sigma = A_\rho$. We proceed by contradiction. Suppose that $A_\sigma \neq A_\rho$. Since $H_\sigma = h = H_\rho$ and ρ is optimal, it follows that $A_\rho < A_\sigma$. By definition ρ and σ agree away from the columns of B from C_j to D_k . Hence the values of ρ on the tiles of one of the columns C from C_j to D_k determines a weight vector in \mathcal{W}_h for some $h' \geq h$ whose area is less than the weight vector in \mathcal{W}_h obtained from the values of σ on the tiles of C . If C is between C_{3n} and D_{3n} , then the weight vector determined by σ is w_h . Corollary 3.2 shows that this is impossible because the area of w_h is minimal. Thus C is either one of the columns C_j, \dots, C_{3n} or D_k, \dots, D_{3n} . By symmetry we may assume that $C = C_i$ for some $i \in \{j, \dots, 3n\}$. Then $A(y_i) < A(x_i)$. By construction we have that $A(\bar{x}_j) \leq A(y_j)$. Since we replaced the original value of x_j by \bar{x}_j , we have that $A(x_j) \leq A(y_j)$, and so $i > j$. Now we have a contradiction to the choice of j . This proves that $A_\sigma = A_\rho$.

Now we have that $H_\sigma = H_\rho$ and $A_\sigma = A_\rho$. It follows that $\sigma = \rho$. Since σ is virtually bar uniform, we see that ρ is virtually bar uniform.

This proves the dumbbell theorem.

7. NOTES

In this section we state without proof a number of results related to the dumbbell theorem.

We say that a dumbbell D is vertically convex if whenever x and y are points in D such that the line segment joining them is vertical, then this line segment lies in D . We say that the sides of D are extreme if the sides of D are contained in vertical lines such that D lies between these lines. Suppose that the vertical lines determined by our tiling of the plane intersect the x -axis in exactly the set of integers. We say that a skinny cut for D is monotonic if it has an underlying path which is the graph of a function and that the restriction of this function to the interval determined by any pair of consecutive integers is either monotonically increasing or decreasing. If the dumbbell D is vertically convex and its sides are extreme, then every optimal weight function for D is a sum of monotonic minimal skinny cuts.

We recall Lemma 4.2. Let R be a rectangle tiled by n rows and m columns of squares T_{ij} . Let $x = (x_{11}, \dots, x_{n1})$ be a weight vector in \mathbf{R}^n . Then there exists a unique weight function ρ for R with minimal area subject to the conditions that ρ is a sum of strictly monotonic skinny cuts and $\rho(T_{i1}) = x_i$ for every $i \in \{1, \dots, n\}$. The skinny cut function Φ is defined so that if $m = 2$, then $\Phi(x) = (\rho(T_{12}), \dots, \rho(T_{n2}))$. It is in fact true for every $m \geq 2$ that $\Phi^{j-1}(x) = (\rho(T_{1j}), \dots, \rho(T_{nj}))$ for every $j \in \{2, \dots, m\}$.

Let h be a positive real number, and let $x \in \mathcal{W}_h$. Let $\pi(x) = (p_0, \dots, p_n)$. We define a nonnegative integer μ_i for every $i \in \{1, \dots, n-1\}$. Let $i \in \{1, \dots, n-1\}$. If $p_i < \frac{hi}{n}$, then μ_i is the number of partition points p_i, \dots, p_{n-1} in the half-closed half-open interval $[p_i, \frac{hi}{n})$. If $p_i > \frac{hi}{n}$, then μ_i is the number of partition points p_1, \dots, p_i in the half-open half-closed interval $(\frac{hi}{n}, p_i]$. If $p_i = \frac{hi}{n}$, then $\mu_i = 0$.

Set $\mu = \max\{\mu_1, \dots, \mu_{n-1}\}$. Note that $\mu \leq n - 1$, and that $\mu = n - 1$ for the weight vector $(0, 0, 0, \dots, h)$. It turns out that $\Phi^m(x) = w_h$ for some nonnegative integer m if and only if $m \geq \mu$. This implies that $\Phi^{n-1}(x) = w_h$ for every weight vector $x \in \mathcal{W}_h$. In other words, if x_1, \dots, x_{n-1} are weight vectors in \mathcal{W}_h such that $x_{i+1} = \Phi(x_i)$ for every $i \in \{1, \dots, n-2\}$, then $x_{n-1} = w_h$. Compare this with Theorem 5.7. We see that the integer $3n$ in Theorem 5.7 cannot be replaced by an integer less than $n - 1$. It is in fact true that there exists a positive real number ϵ such that the integer $3n$ in Theorem 5.7 cannot be replaced by an integer less than $(1 + \epsilon)n$. It is probably true that $3n$ can be replaced by $3n/2$.

Not surprisingly, the skinny cut function Φ is continuous. What might be surprising is that it is in fact piecewise affine. If c_n denotes the number of affine pieces of Φ , then the sequence (c_n) satisfies a quadratic linear recursion with eigenvalue $1 + \sqrt{2}$.

There exists a theory of combinatorial moduli “with boundary conditions”, as alluded to by Lemma 4.2. For simplicity we deal with a rectangle R with n rows of tiles. Let h be a positive real number, and let x be a weight vector in \mathbf{R}^n with height h . Here are two ways to define optimal weight functions with boundary conditions. In the first way we minimize area over all weight functions on R with fat flow height h such that the values of ρ on the first column of R give x . In the second way we maximize the fat flow modulus over all weight functions on R whose values on the first column of R give a scalar multiple of x . It turns out that these constructions are equivalent. They yield a weight function which is unique up to scaling. Interpreted properly, most of the results of [1] hold in this setting. For example, the optimal weight function is a sum of minimal fat flows (modulo the first column). A straightforward modification of the algorithm for computing optimal weight functions even applies in this setting.

REFERENCES

- [1] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Squaring rectangles: the finite Riemann mapping theorem*, The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (Brooklyn, NY, 1992), Contemp. Math., 169, Amer. Math. Soc., Providence, RI, 1994, pp. 133–212. MR1292901 (95g:20045)
- [2] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Finite subdivision rules*, Conform. Geom. Dyn. **5** (2001), 153–196 (electronic). MR1875951 (2002j:52021)
- [3] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Expansion complexes for finite subdivision rules I*, Conform. Geom. Dyn. **10** (2006), 63–99 (electronic). MR2218641 (2007c:30048)
- [4] J. W. Cannon, W. J. Floyd, and W. R. Parry, *precp.p*, software, available from <http://www.math.vt.edu/people/floyd>.
- [5] Z.-x. He and O. Schramm, *On the convergence of circle packings to the Riemann map*, Invent. math. **125** (1996), 285–305. MR1395721 (97i:30009)
- [6] O. Schramm, *Square tilings with prescribed combinatorics*, Israel J. Math. **84** (1993), 97–118. MR1244661 (94h:52045)
- [7] K. Stephenson, *CirclePack*, software, available from <http://www.math.utk.edu/~kens>.

DEPARTMENT OF MATHEMATICS, BRIGHAM YOUNG UNIVERSITY, PROVO, UTAH 84602
E-mail address: cannon@math.byu.edu

DEPARTMENT OF MATHEMATICS, VIRGINIA TECH, BLACKSBURG, VIRGINIA 24061
E-mail address: floyd@math.vt.edu
URL: <http://www.math.vt.edu/people/floyd>

DEPARTMENT OF MATHEMATICS, EASTERN MICHIGAN UNIVERSITY, YPSILANTI, MICHIGAN 48197
E-mail address: walter.parry@emich.edu