Discrete Riemann Maps and the Parabolicity of Tilings

Andrew S. Repp

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Mathematics

William J. Floyd, Chair
Peter E. Haskell
Peter A. Linnell
Robert A. McCoy
James E. Thomson

May 4, 1998
Blacksburg, Virginia

Keywords: Tilings, Parabolic, Modulus, Riemann Map

Discrete Riemann Maps and the Parabolicity of Tilings

Andrew Repp

(ABSTRACT)

The classical Riemann Mapping Theorem has many discrete analogues. One of these, the Finite Riemann Mapping Theorem of Cannon, Floyd, Parry, and others, describes finite tilings of quadrilaterals and annuli. It relates to several combinatorial moduli, similar in nature to the classical modulus. The first chapter surveys some of these discrete analogues. The next chapter considers appropriate extensions to infinite tilings of half-open quadrilaterals and annuli. In this chapter we prove some results about combinatorial moduli for such tilings. The final chapter considers triangulations of open topological disks. It has been shown that one can classify such triangulations as either parabolic or hyperbolic, depending on whether an associated combinatorial modulus is infinite or finite. We obtain a criterion for parabolicity in terms of the degrees of vertices that lie within a specified distance of a given base vertex.

Dedication

To my Lord and Savior

Jesus Christ

To Whom I owe all things

Acknowledgments

It would be impossible to list all of the people who were instrumental in bringing me to this point. However, I would in particular like to thank my advisor, Bill Floyd, for the many hours he has spent teaching me and discussing these topics, providing excellent ideas, reading several extremely rough drafts of this material, and giving plenty of good advice. A large portion of the credit for this achievement belongs to him. In addition, I would like to thank all of my teachers, both here and at Bob Jones, for the lasting investment they made in my life. Finally, I would like to thank my family members (Russell, Gloria, Jon, Krista, and Janelle Repp; Pearl and John Sprigg), as well as my friends, for their continual encouragement, prayers, and support.

Contents

List of Figures

Chapter 1

A Survey of the Finite Case

1.1 Introduction

Tilings have fascinated children, artisans, and mathematicians for centuries. Recently, some mathematicians who study tilings and related structures (such as circle packings) have discovered that the combinatorics of these tilings provide discrete analogues of various results from conformal geometry.

In this first chapter, we shall survey some known results for finite tilings of quadrilaterals and annuli. One prominent subject of these results is a combinatorial modulus, which, as we shall see, gives us a discrete analogue of the Riemann Mapping Theorem. In Chapter 2, we shall extend the appropriate definitions to infinite tilings and prove some new results for such tilings. In Chapter 3, we shall consider triangulations of the open unit disk. It is known that one can classify these tilings as either parabolic or hyperbolic. We shall provide a criterion for parabolicity in terms of the valences of vertices within a specified distance of a given base vertex. In order to do so, we shall determine the combinatorial modulus of various subannuli and then use these moduli to estimate the modulus of the entire tiling.

We begin by taking a closer look at the Riemann Mapping Theorem.

1.2 Perspectives on the Riemann Mapping Theorem

The classical Riemann Mapping Theorem is well-known to students of complex analysis. One statement of this remarkable theorem is the following:

Theorem 1.1 (Riemann Mapping Theorem) *Let G be an open, simply connected, and proper subset of the complex plane; let x be a point in G. Then there is a one-to-one conformal* *map f from G onto the interior of the unit disk such that $f(x) = 0$. Furthermore, this map is unique up to rotations of the disk.*

The surprising feature of this theorem is that one can make the map f conformal. One way to understand a conformal map is to think of it as a scaling at each point, such that all of the scalings fit together into a smooth map. To be more precise, we multiply the metric at each point z by a positive function $\rho(z)$, so that $|dz|$ in the preimage becomes $\rho(z)|dz|$ in the image. Hence, conformal maps take infinitesimal circles to infinitesimal circles. Indeed, the ratio of the radii of the two infinitesimal circles (one being in the preimage, the other being the corresponding circle in the image) is the value of $|f'(z)|$ at the center of the preimage circle.

Since the Riemann map is conformal, it seems plausible that one could approximate it by a discrete map involving circles. In particular, one would hope to be able to approximate the Riemann map by a mapping from a circle-packing of the original region G to a circle-packing of the unit disk. Thurston made this conjecture in 1985, and Rodin and Sullivan proved it in 1986 (see [11]), yielding a discrete Riemann mapping theorem for circle packings. Without rigorously defining all of the terms involved, Rodin and Sullivan's theorem is the following:

Theorem 1.2 (Discrete Riemann Mapping, Circle Packings) *Let G be a simply connected bounded region of the plane. For $\epsilon > 0$, consider a regular hexagonal circle packing C_ϵ of G. There exists a circle packing C'_ϵ of the unit disk such that C_ϵ is isomorphic to C'_ϵ. Then the isomorphism $C_\epsilon \to C'_\epsilon$ of circle packings determines an approximate mapping which, as ϵ approaches 0, converges to a a Riemann map from G to the unit disk.*

Note that Rodin and Sullivan do not prove the existence of the packing C'_ϵ and the isomorphism $C_\epsilon \to C'_\epsilon$ as part of their theorem; these two facts are consequences of a theorem by Koebe, Andreev, and Thurston. What Rodin and Sullivan proved was that this isomorphism determines mappings which converge to the Riemann map. See Figure 1.1 for an illustration of such an approximate map, determined by circle packings. (This figure was prepared using CirclePack, a program written by Ken Stephenson at the University of Tennessee.)

One can restate the classical Riemann Mapping Theorem in terms of quadrilaterals and annuli. This formulation is the one which shall receive the most attention in this paper.

Theorem 1.3 (Riemann Mapping Theorem) *Suppose we are given a quadrilateral Q, that is, a topological disk in the complex plane with four distinguished points on its boundary. Then there is a conformal mapping taking the interior of the quadrilateral onto the interior of a rectangle in such a way that the induced boundary map takes the four distinguished boundary points to the four corners of the rectangle.*

So the Riemann map takes a quadrilateral to a rectangle. Alternatively, one can restate the theorem so that topological annuli map conformally to right circular cylinders.

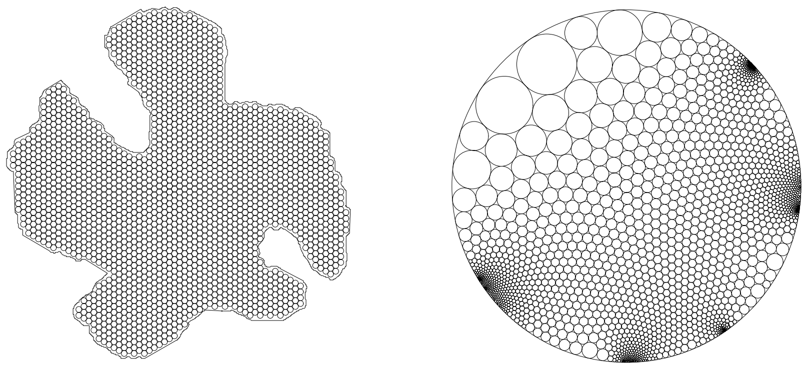

Figure 1.1: Approximating the Riemann Map by Circle Packings

Let us consider more closely the image of Q under this mapping. We have already seen that the Riemann map multiplies each differential dz by a positive local scale factor $\rho(z)$. If we consider any local scaling function $\rho(z)$ (not necessarily the one determined by the Riemann map), then it is possible to define the *height* and *area* of Q under ρ. If Q is a quadrilateral, we choose one of its four boundary segments to be the top of the quadrilateral, and we choose the opposite boundary segment to be its bottom. (If Q is an annulus, we choose one of the boundary components to be the top and the other to be the bottom.) The height H_ρ is the infimum of $\int_\alpha \rho(z)|dz|$, where the infimum is taken over all paths α connecting the bottom of Q to the top of Q. The area A_ρ is the area of Q in the new metric: $A_\rho = \int_Q \rho(z)^2 dA$. Looking at all possible local scaling functions ρ, we may define the classical *modulus* of Q to be

$$M(Q) = \sup_\rho \frac{H_\rho^2}{A_\rho}.$$

Note that if Q, after scaling by ρ, is a rectangle, then H_ρ^2/A_ρ is simply the ratio of the height of the rectangle to its width. (Similarly, if Q is isometric to a right circular cylinder under the scaling, then H_ρ^2/A_ρ is the ratio of the height to the circumference.) So, in some sense, the modulus $M(Q)$ is the supremum of the height-width ratio of Q under local scalings.

However, we know that the Riemann map gives us a local scaling ρ_0 such that ρ_0 turns Q into a rectangle (or cylinder). According to a result by Ahlfors and Beurling, this scaling ρ_0 maximizes the ratio H_ρ^2/A_ρ. Hence the supremum in the definition of $M(Q)$ is actually a maximum, and $M(Q)$ is equal to the height-width ratio of the scaled quadrilateral (or annulus) Q.

Therefore, the Riemann Mapping Theorem tells us that there is a conformal change of metric that converts a quadrilateral (or annulus) into a rectangle (or cylinder), and that this change of metric $\rho_0(z)$ optimizes the ratio H_ρ^2/A_ρ over all appropriate functions ρ. The Riemann map is implicit in this function ρ_0. In fact, given the existence of this optimal ρ—even with no other geometric information about it or about Q—one can show that it transforms the quadrilateral (or annulus) into a rectangle (or cylinder). See [3] and [6] for the outline of such a method. Essentially, sets at a constant distance from the bottom (under ρ_0) become horizontal lines in the rectangle (or cylinder), and lines of minimal length (under ρ_0) that connect top and bottom become vertical lines in the rectangle (or cylinder). Thus, this optimal local scaling function ρ_0 induces a grid on Q which become horizontal and vertical curves in the resulting rectangle or cylinder.

Summarizing the above discussion, we may say that the classical Riemann Mapping Theorem shows us how to take a function (ρ_0) that optimizes a length-area ratio and then extract complex coordinates (on the resulting rectangle or cylinder) from that function.

1.3 A Discrete Riemann Mapping for Finite Tilings

1.3.1 Extracting Geometry from Combinatorics

Now let us turn to the geometry of 3-manifolds, a field which might at first glance seem quite distant from conformal geometry. In 1982, Thurston conjectured that every closed compact 3-manifold has a canonical decomposition into pieces with structures based on the eight model 3-dimensional geometries. This Geometrization Conjecture (as it is called) contains within itself several other conjectures. One of these conjectures is the Poincaré Conjecture; another, known as Thurston's Hyperbolization Conjecture, deals with the geometry of hyperbolic 3-space, one of the most interesting of the eight model geometries. This conjecture reads as follows:

> Let M be a closed 3-manifold. Suppose that $\pi_1(M)$ is infinite, that $\pi_1(M)$ does not contain a subgroup isomorphic to $\mathbf{Z}\oplus\mathbf{Z}$, and that $\pi_1(M)$ is not a free product. Then M has a hyperbolic structure.

The Cannon-Floyd-Parry approach (focusing on group actions on the sphere at infinity) is part of one attempted program of proof of the Hyperbolization Conjecture. The following conjecture is in turn part of this approach:

> Suppose G is a negatively curved group whose space at infinity is the 2-sphere. Then G acts conformally on the 2-sphere (its space at infinity) and hence acts isometrically on hyperbolic 3-space.

By *negatively curved*, we mean that there exists a number δ such that for any geodesic triangle in the Cayley graph of G, any point on a given side of the triangle is within δ of some point on the union of the other two sides. In other words, triangles in the Cayley graph of G are uniformly thin, just as they are in standard hyperbolic space. (Some authors use the phrases *word hyperbolic* or *Gromov hyperbolic* to denote the same condition.) We shall define what we mean by the space at infinity in the next paragraph.

First consider what would be required to prove this conjecture. In order to demonstrate that G acts conformally on ∂G (its space at infinity), one needs to impose an appropriate set of complex coordinates on ∂G. What sort of structure do we have? By hypothesis, ∂G has a topological structure making it a 2-sphere; let us consider how we obtain this structure. Let Γ be the Cayley graph for G with respect to some finite generating set. Choose some base vertex v in Γ. Then for every geodesic ray $R : [0, \infty) \longrightarrow \Gamma$ such that $R(0) = v$, and for every positive integer n, we define the combinatorial half-space

$$H(R, n) = \{x \in \Gamma \mid d(x, R[n, \infty]) \leq d(x, R[0, n])\},$$

where $d(\cdot, \cdot)$ denotes distance in the Cayley graph. Now the space at infinity ∂G is defined to be the set of equivalence classes of geodesic rays in Γ, two rays being equivalent if and only if they remain a bounded distance apart. We can define the *combinatorial disk* $D(R, n)$ to be the set of all points x in ∂G such that, if a ray $S : [0, \infty) \longrightarrow \Gamma$ represents x, then $S(k)$ lies arbitrarily far inside $H(R, n)$ as k approaches infinity. In essence, $D(R, n)$ is the shadow of $H(R, n)$ on ∂G. The sets

$$\{D(R, n) \mid R : [0, \infty) \longrightarrow \Gamma \text{ a geodesic ray, } n > 0\}$$

form the basis for a topology on ∂G. (See Gromov in [9] and Swenson in [14] for more on the above.) It is with regard to this topology that we say ∂G is homeomorphic to S^2.

We also have a combinatorial structure on ∂G. If we define

$$\mathcal{D}(n) = \{D(R, n) \mid R : [0, \infty) \longrightarrow \Gamma \text{ a geodesic ray}\},$$

then each $\mathcal{D}(n)$ is a finite cover of ∂G. According to [7], we can obtain the disk cover $\mathcal{D}(n+1)$ from $\mathcal{D}(n)$ recursively for every $n \geq 1$. This recursively-refined sequence of disk structures on the space at infinity provides us with a combinatorial structure on ∂G.

Recall the problem which we were considering. We are given this combinatorial information (consisting of sequences of covers on S^2), and we want to find complex coordinates on S^2 which are in some way compatible with this combinatorial structure. When discussing the Riemann map, we noted that it takes a function that optimizes a length-area ratio and constructs complex coordinates from that function. We shall apply the same approach here. Given a finite cover of a space, we want to be able to calculate lengths and areas and optimize a similar ratio. Then, in a manner analogous to the Riemann Mapping Theorem, we can attempt to transform this optimum into a set of coordinates. The following work, culminating in a Finite Riemann Mapping Theorem, was done independently by Robertson; by Schramm ([13]); and by Cannon, Floyd, and Parry ([5]).

1.3.2 Combinatorial Moduli

We shall simplify our task by passing from arbitrary finite covers to tilings by topological disks. Suppose X is a topological quadrilateral or annulus tiled by finitely many topological disks. Our goal is to impose complex coordinates on this figure by optimizing a combinatorial modulus which we shall define below. First we must determine the functions for which to optimize the ratios. In the classical case, these functions ρ were non-negative functions on the quadrilateral or annulus. Now, however, we have a discrete case in which we can only work with a finite number of tiles. So we shall want our function to assign a non-negative number to each of these tiles. Thus, we define a *weight function* on X to be a function which assigns to every tile in X a non-negative number, such that not all tiles have zero weight. We shall attempt to optimize length-area ratios over the set of all possible weight functions.

Next, we must define combinatorial analogues for length and area. In the continuous case, the length of a path α was $\int_\alpha \rho(z)\,|dz|$, and the area of a set Q was $\int_Q \rho(z)^2\,dA$. We shall proceed similarly in the discrete case. Given a weight function w and α, a subset of X, we define the *w-length* of α to be $\sum w(t)$, taking the sum over all tiles t in X which intersect α. If Q is a subset of X, we define the *w-area* of Q to be $A_w(Q) = \sum w(t)^2$, where again we take the sum over all tiles t in X which intersect Q.

Now if we wish to define a length-area ratio for X, we must specify the paths to consider when determining height. We shall be interested in four types of paths, and thus we shall make the following definitions. Following Cannon, Floyd, and Parry in [5], we shall define a *path* to be a non-empty set of tiles whose union is connected. A *fat path* is the set of all tiles meeting a given topological path in X. A *skinny path* is a nonempty set of tiles which can be ordered $t_1, \dots, t_m$, such that $t_i \cap t_{i+1} \neq \emptyset$ for $i = 1, \dots, m-1$. The skinny path is *closed* if $t_1 \cap t_m \neq \emptyset$. A fat path is *closed* if it has an underlying topological path that is closed.

A *flow* is a path which meets both the top and bottom of X. We will require a *fat flow* to have an underlying topological path from the bottom to the top of X, and we will require a *skinny flow* to have tiles which can be ordered such that the first tile meets the bottom of X and the last tile meets the top of X. A *cut* is a path which separates the bottom of X from the top of X. A *fat cut* will have an underlying topological path separating the bottom and top of X. If X is a quadrilateral, we shall require a *skinny cut* to have tiles which can be ordered such that the first tile meets a boundary segment of X which is not the bottom or the top, and such that the last tile meets the other boundary segment which is not the bottom or the top. If X is an annulus, we require a *skinny cut* to be a closed skinny path which separates the bottom and the top of X. Note that every flow contains a subflow which is a skinny flow, and every cut contains a subcut which is a skinny cut.

We now have four types of paths which shall be useful: fat flows, fat cuts, skinny flows, and skinny cuts. At this point we may define the lengths which we shall put into the length-area ratios. If w is a weight function on X, then by the finiteness of the tiling there will be fat flows, fat cuts, skinny flows, and skinny cuts with minimal w-length. We call such a fat flow a *w-minimal fat flow*. Likewise, we have w-minimal fat cuts, skinny flows, and skinny cuts. Naturally, the minimal w-lengths will probably not be the same for the four categories. So we define the *fat w-height* of X, denoted $H_{w,f}(X)$, to be the w-length of a w-minimal fat flow; and we define the *skinny w-height* of X, denoted $H_{w,s}(X)$, to be the w-length of a w-minimal skinny flow. For cuts, we talk about circumferences instead of heights. (We use this terminology even if X is a quadrilateral.) The *fat w-circumference* of X, denoted $C_{w,f}(X)$, is the w-length of a minimal fat cut; and the *skinny w-circumference* of X, denoted $C_{w,s}(X)$, is the w-length of a minimal skinny cut. As before, we define the w-area of X, denoted $A_w(X)$, to be the sum of the squares of the weights of the tiles in X. If there is no ambiguity involved, we shall often suppress the (X) in the notation and talk about $H_{w,f}$, $C_{w,f}$, etc.

Finally we come to the length-area ratios that will give us combinatorial moduli. Let the *fat flow modulus* and the *fat cut modulus* of X be defined as

$$M_f = \sup_w \frac{H_{w,f}^2}{A_w}, \qquad \text{and} \qquad m_f = \inf_w \frac{A_w}{C_{w,f}^2},$$

respectively, where the supremum (respectively, the infimum) is taken over all possible weight functions on X. We may also define the *skinny flow modulus* and the *skinny cut modulus* of X as

$$M_s = \sup_w \frac{H_{w,s}^2}{A_w}, \qquad \text{and} \qquad m_s = \inf_w \frac{A_w}{C_{w,s}^2},$$

respectively; again, the supremum (respectively, infimum) is taken over all possible weight functions on X. Thus, instead of one modulus (as in the classical case), we obtain four combinatorial moduli. There is no a priori reason to expect them to be the same, nor, in fact, do all four turn out to be the same.

1.3.3 The Vector Space Formulation

In considering paths, weight functions, and combinatorial moduli, it is useful to represent the problem in terms of a finite dimensional vector space. Suppose our tiling X has n tiles. We then let $\mathcal{W}$ be the n-dimensional vector space $\mathbf{R}^n$. Each dimension in this vector space corresponds to one of the tiles in the tiling. We define a *weight vector* to be a non-zero vector $w = (w_1, \ldots, w_n)$ in $\mathcal{W}$, such that $w_i \geq 0$ for $i = 1, \ldots, n$. Weight vectors are natural representations of weight functions; the weights of the tiles become components of the vector. We shall define P to be a nonempty finite subset of $\mathbf{N}^n \setminus \{0\}$, where $\mathbf{N}$ is the set of non-negative integers. A *path vector* p is an element of P. Note that path vectors correspond naturally to sets of tiles, where the ith component of p denotes how many times the tile corresponding to i is included in the set. A path vector p in $\mathcal{W}$ represents a path in the sense defined above if the union of the tiles corresponding to non-zero components of p is connected. In most of the rest of this paper, P will represent either the set of fat flows, the set of fat cuts, the set of skinny flows, or the set of skinny cuts.

The advantage of the vector space notation is that it allows us to make use of the inner product structure on the space. We can define the *w-length* of a path vector p to be $\langle p, w\rangle$, where $\langle \cdot, \cdot \rangle$ denotes the standard inner product on $\mathbf{R}^n$. Note that this definition coincides precisely with the previous definition of the w-length of a set of tiles. Given the set P of path vectors, we may define the *w-height* $H_w(P)$ to be the minimum of the w-lengths of all elements p of P. A path vector p such that $\langle p, w \rangle = H_w(P)$ is called a *w-minimal path vector*. We define the *w-area* A_w to be $\langle w, w \rangle = \|w\|^2$. Note that this would correspond to the w-area of the entire tiling X in our previous formulation. Accordingly, we define the *modulus* M with respect to P to be

$$M = \sup_w \frac{H_w(P)^2}{A_w},$$

where we take the supremum over all weight vectors w.

Consider how the modulus as defined in this vector space setting relates to the four combinatorial moduli defined earlier. If P corresponds to the set of fat flows or the set of skinny flows, then $H_w(P)$ will equal $H_{w,f}$ or $H_{w,s}$, respectively. The modulus M will in turn equal M_f or M_s, respectively. On the other hand, if P corresponds to the set of fat cuts or the set of skinny cuts, then $H_w(P)$ will equal $C_{w,f}$ or $C_{w,s}$, respectively. Comparing the definitions, it is apparent that the modulus M will then equal $1/m_f$ or $1/m_s$, respectively. Thus, one can express each of the combinatorial moduli in terms of this vector space modulus. One can apply results from the vector space setting to the combinatorial moduli by using the appropriate sets for P, the set of path vectors. We shall pass freely between the tiling and vector space representations.

1.3.4 Optimal Weight Functions

We define an *optimal weight vector* to be a weight vector which achieves the supremum of $H_w(P)^2/A_w$. Thus, if w is an optimal weight vector, $M = H_w(P)^2/A_w$. Note that if w is an optimal weight vector and λ is a constant, then

$$\frac{H_{\lambda w}(P)^2}{A_{\lambda w}} = \frac{(\inf_p \langle p, \lambda w \rangle)^2}{\langle \lambda w, \lambda w \rangle} = \frac{\lambda^2 (\inf_p \langle p, w \rangle)^2}{\lambda^2 \langle w, w \rangle} = \frac{\lambda^2 H_w^2(P)}{\lambda^2 A_w} = M.$$

So any optimal weight vector remains an optimal weight vector under scaling.

The following theorem, proved in [5], addresses the existence and uniqueness of optimal weight vectors.

Theorem 1.4 *Let $\mathcal{W}$ and P be as described above. Then there exists an optimal weight vector with $A_w = 1$. All other optimal weight vectors are scalar multiples of this one.*

Proof: We have seen that scaling a weight vector w does not change the ratio $H_w(P)^2/A_w$; hence, when defining M, we need only take the supremum of this ratio over all vectors on the unit $(n-1)$-sphere S^{n-1}. If w is on this sphere, then $H_w(P)^2/A_w = H_w(P)^2$. Now the function which maps w to $H_w(P)^2$ is continuous, being the infimum of a finite number of continuous functions. Since S^{n-1} is compact, there is a vector w_0 on S^{n-1} such that $H_w(P)^2$ is at a maximum. Then $H_{w_0}(P)^2/A_{w_0} = M$, and w_0 is an optimal weight vector.

The uniqueness of w_0 follows from convexity of S^{n-1}. If w_0 and w_1 are distinct optimal weight vectors on S^{n-1}, then let $w = (1-t)w_0 + tw_1$ for some t, $0 < t < 1$. If p is a path vector, then we have

$$\langle p, w\rangle = (1-t)\langle p, w_0\rangle + t\langle p, w_1\rangle \geq (1-t)H_{w_0} + tH_{w_1} = \sqrt{M}.$$

Since $0 < \|w\| < 1$, $(1/\|w\|)w$ lies on S^{n-1}, and $\langle p, \frac{1}{\|w\|}w\rangle > \langle p, w\rangle \geq \sqrt{M}$ for all paths p, then we know that

$$H_{\frac{1}{\|w\|}w}(P) > \sqrt{M},$$

which cannot be. $\square$

Note that the finite-dimensionality of $\mathcal{W}$ played a key role in demonstrating the existence of optimal weight vectors.

Since optimal weight vectors exist, and since the ratio of height squared to area is invariant under scaling, we may find the modulus by taking the maximum over the unit sphere:

$$M = \max_{\|w\|=1} \frac{H_w^2(P)}{A_w} = \max_{\|w\|=1} H_w^2(P).$$

It turns out that one can characterize optimal weight vectors as the weighted sum of their minimal paths, as stated in the following theorem:

Theorem 1.5 *Let w be a weight vector, and let $p_1, \ldots, p_k$ be the set of vectors in P whose w-length is minimal. I.e., $\langle p_i, w\rangle = H_w(P)$ for $i = 1, \ldots, k$. Then w is an optimal weight vector if and only if there exist non-negative real numbers $a_1, \ldots, a_k$ such that*

$$w = \sum_{i=1}^{k} a_i p_i.$$

Furthermore, for every optimal weight vector w, there is a scalar multiple λw of w ($\lambda > 0$) such that the coefficients $a_1, \ldots, a_k$ for λw are non-negative integers.

See [5] for proofs of the various statements of this theorem.

1.3.5 Fat and Skinny Cuts and Flows

Now that we have characterized optimal weight vectors in the vector space setting, let us return to the four specific combinatorial moduli previously defined. Theorem 1.4 implies the existence of an optimal weight function w_{M_f} for fat flows such that

$$M_f = \frac{H^2_{w_{M_f},f}}{A_{w_{M_f}}}.$$

It also tells us that this weight function is unique up to scaling. Furthermore, Theorem 1.5 implies that w_{M_f} is the weighted sum of w_{M_f}-minimal fat flows.

Likewise, there exists an optimal weight function w_{m_f} for fat cuts such that

$$m_f = \frac{A_{w_{m_f}}}{C_{w_{m_f},f}},$$

and w_{m_f} is unique up to scaling. This weight function w_{m_f} is the weighted sum of w_{m_f}-minimal fat cuts. Similarly, we can find optimal weight functions w_{M_s} and w_{m_s} for skinny flows and cuts, respectively; w_{M_s} will be the weighted sum of its minimal skinny flows, and w_{m_s} is the weighted sum of its minimal skinny cuts.

As we noted before, there is no a priori reason to expect the four combinatorial moduli to be the same, nor to expect the corresponding optimal weight functions to be comparable to each other. However, Cannon, Floyd, and Parry have shown in [5] that the moduli and optimal weight functions for fat flows and skinny cuts are the same and that a similar relationship obtains between skinny flows and fat cuts:

Theorem 1.6 *Let X be a topological quadrilateral or annulus tiled by finitely many tiles. Then $M_f = m_s$ and $M_s = m_f$. If one normalizes the optimal weight vectors to have unit area, then $w_{M_f} = w_{m_s}$ and $w_{M_s} = w_{m_f}$.*

So, up to scaling, there are only two optimal weight vectors on any finite tiling of a quadrilateral or annulus. One of them optimizes fat flows and skinny cuts; the other optimizes fat cuts and skinny flows. Cannon, Floyd, and Parry present an algorithm for finding these optimal weight functions in [5].

The proof of Theorem 1.6 in [5] constructs a sequence of skinny cuts (called *level curves*) at integral distances from the bottom; it then shows that one can choose underlying topological paths for these level curves which are disjoint, except possibly at points where four or more tiles come together. There is a similar disjoint family of paths underlying fat flows. For any given tile t, the number of level curves containing t is equal to the number of paths underlying fat flows which pass through t. Similar results hold for fat cuts and skinny flows.

Let us restrict our attention to fat flows and skinny cuts; for now let us also suppose that X is a topological quadrilateral. We have a family of fat flows with disjoint underlying paths and a family of skinny cuts with essentially disjoint underlying paths. If we identify these cuts and flows (or rather, the underlying paths) with a grid on graph paper, then the quadrilateral becomes a rectangle. The width of the rectangle will be the number of paths underlying fat flows, and the height of the rectangle will be the number of underlying skinny cuts. The height-width ratio of the rectangle will be the combinatorial modulus $M_f = m_s$. Each tile will become a square, since the number of fat flows passing through it and the number of skinny cuts passing through it will be equal. Thus, the tiled quadrilateral, containing only combinatorial information, has become a rectangle tiled by squares. We have succeeded in extracting geometric information from combinatorics, using an analogue of classical modulus. The Finite Riemann Mapping Theorem summarizes the above discussion and states that similar results hold under fat cut optimization and for tiled annuli. It was proved by Cannon, Floyd, and Parry ([5]); Oded Schramm ([13]); and John Robertson.

Theorem 1.7 (Finite Riemann Mapping Theorem) *Every tiled quadrilateral X corresponds under fat flow optimization to a rectangle tiled by squares. It corresponds under fat cut optimization to another rectangle tiled by squares. Analogous results hold for tiled annuli.*

1.4 An Example

Let us apply the Finite Riemann Mapping Theorem to a specific tiled quadrilateral. This quadrilateral will consist of the European landmass, excluding Scandinavia and the regions included in the former Soviet Union. We must distinguish four points on the boundary; the points we choose are Gibraltar, the northwestern tip of Brittany, the northeastern corner of Poland, and the southern tip of Greece. The tiles will, of course, be the national territories of Europe. We exclude San Marino and the Vatican City (independent territories surrounded by Italy) since each tile must be simply connected.

Figure 1.2 shows the original quadrilateral and the image under the Finite Riemann Mapping Theorem. The optimization was performed using programs written by Bill Floyd and Jim Cannon. The squares of the image are labeled with the name of the countries they represent, and the tiles in the original quadrilateral are labeled with their weight under fat flow optimization.

The most noticeable feature of this figure is the great difference between geopolitical and combinatorial importance. Napoleon Bonaparte would have been proud to know that France is (combinatorially) greater than the rest of Europe combined. On the other hand, Spain, Portugal, and Denmark—among other countries—receive a weight of zero under fat flow optimization and thus do not even appear in the combinatorial landscape. It is easy to see why this occurs, since the optimal weight function w must be the sum of w-minimal fat flows. If Spain, for instance, had non-zero weight under w, then there would be a w-minimal fat flow passing through it. Any fat flow passing through Spain must also pass through France. However, France has a border on the Mediterranean Sea (the bottom of our quadrilateral), and hence any fat flow continuing on through Spain could not be minimal. In essence,

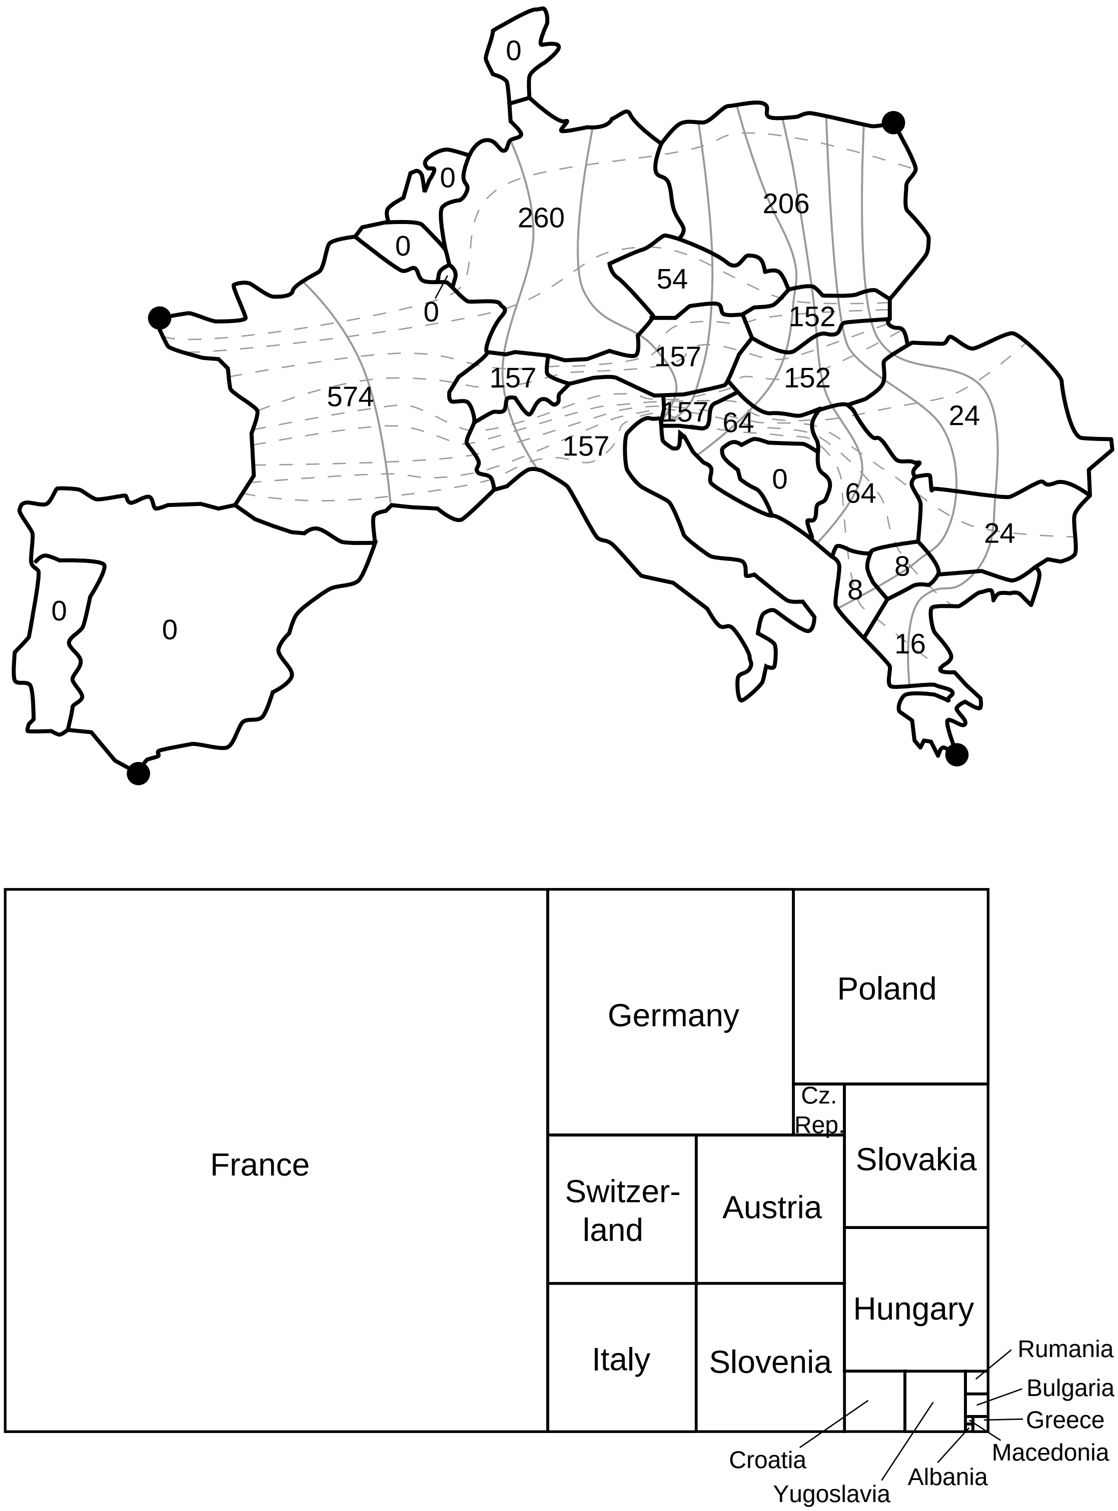

Figure 1.2: Optimal Weights for Europe

France intercepts any minimal fat flow that would try to pass through Spain. Hence, it (and Portugal) receive a weight of zero. Likewise, Germany cuts off Denmark; France and Germany together cut off the Benelux countries; and Croatia and Yugoslavia conspire to cut off Bosnia/Herzogovina. The same would have happened to the other small countries of Europe (such as Andorra and Liechtenstein) had they been included in the calculations. Surprisingly, the small country of Slovenia receives a large weight because it borders on the Mediterranean to one side and Hungary to the other. The fact that its seafront is miniscule does not reduce its combinatorial importance.

Let us consider more closely the mathematical aspects of this example. The light vertical lines on the map of Europe in Figure 1.2 are paths underlying minimal fat flows. They run roughly south to north, and they correspond to vertical lines in the squared rectangle. If we label these fat flows $f_1, \ldots, f_8$ (from west to east) and call the optimal weight function w, then

$$w = 574f_1 + 157f_2 + 103f_3 + 54f_4 + 64f_5 + 64f_6 + 8f_7 + 16f_8,$$

the weighted sum of minimal fat flows. Each of these flows has a length equal to the height $H_{w,f}$ of the quadrilateral, which is 574.

The same figure shows paths underlying minimal skinny cuts as light dotted lines; they run roughly east to west, and they correspond to horizontal lines in the squared rectangle. By considering the cuts, we can understand the cause of French combinatorial aggrandizement. Since France touches both the top and the bottom of the quadrilateral, every cut must pass through France. Hence, every w-minimal cut contributes to the weight of France. This is true for no other country, and hence the French combinatorial weight exceeds that of any other European country. We may, according to Theorem 1.6, also express w as the sum of w-minimal skinny cuts. Label the cuts (from south to north) $c_1, \ldots, c_9$. Then

$$w = 8c_1 + 8c_2 + 24c_3 + 24c_4 + 93c_5 + 59c_6 + 98c_7 + 54c_8 + 206c_9.$$

Each of these cuts has a length equal to the width $C_{w,s}$ of the quadrilateral, which is 1040.

Let us explicitly calculate M_f and m_s for this tiling. We already have $H_{w,f}$ and $C_{w,s}$, and we know that w is an optimal weight function. A_w, the area of Europe under w, is the sum of the squares of the tiles, which is $596,960 = 574 \cdot 1040$. Hence,

$$M_f = \frac{H_{w,f}^2}{A_w} = \frac{574^2}{596,960} = \frac{574}{1040},$$

and

$$m_s = \frac{A_w}{C_{w,f}^2} = \frac{596,960}{1040^2} = \frac{574}{1040}.$$

The two moduli are indeed the same and are equal to the height-width ratio of the rectangle.

1.5 Related Results

1.5.1 The Kirchhoff Inequalities

The word "flow" in the discussion above suggests an analogy to electrical current. In particular, we could replace each square (in the squared rectangle) with a resistor. Two resistors would be connected if and only if the corresponding squares share a horizontal edge. If one then applies a positive voltage to the entire bottom of the rectangle and a negative voltage to the entire top, the electrical current would flow precisely along the minimal fat flows described above. To what, then, would the minimal skinny cuts correspond? Their natural analogue would be equipotential curves in the circuit, which would cut across the resistors just as the cuts "cut" across the squares.

We can make the analogy more precise. Let the width of each square be the amount of current flowing through the resistor (in Amperes), and let each resistor have a resistance of 1 Ω. Ohm's Law states that the voltage drop across the resistor equals the current (width) times the resistance (one). Thus, since the squares have equal height and width, we may say that the voltage drop across a resistor equals the height of the corresponding square. So the width of the entire rectangle would be the total amount of current flowing through the circuit, and the height of the rectangle would be the total voltage drop across the circuit. The modulus (equal to height divided by width) would be the net resistance of the circuit.

To complete the analogy, it remains to show that a circuit as described above would have currents proportional to the size of the squares. Kirchhoff's Laws guarantee that it would; see [5] for a demonstration of this fact. These laws state that

1. the net current entering any resistor is equal to the net current leaving that resistor; and

2. that the net voltage drop around a cycle of resistors is zero.

So the squared rectangle corresponds to a planar circuit consisting of unit resistors.

Note that this analogy in itself is not sufficient to determine the fat flow optimal weight function for a given tiling. Recall that two resistors in the circuit are connected if and only if their corresponding squares share a horizontal edge in the squared rectangle, and two tiles share a horizontal edge in the squared rectangle if and only a minimal fat flow passes from one to the other. Given two adjacent tiles, there is no guarantee that a minimal fat flow will pass across the shared edge, and thus there is no guarantee that the corresponding resistors will be connected in the circuit. Hence, the set of adjacencies in the circuit is a subset of the adjacencies in the original tiling. The circuit that actually corresponds to the squared rectangle (produced by the optimal weight function) is only one of many possible circuits one could construct from the original tiling. We need extra conditions to determine the circuit that corresponds to the squared rectangle and, hence, to the optimal weight function.

The *Kirchhoff Inequalities* given in [5] contain the additional conditions. In this formulation, we consider the dual graph of the tiling. The tiles will become vertices in this dual graph, and any tiles meeting along an edge will have their corresponding vertices connected by an edge in the dual graph. We will add two tiles, t_0 and t_1, representing the bottom and top of the quadrilateral (or annulus) respectively. We shall use the same labels for tiles and their corresponding vertices. Now, if s and t are two tiles meeting along an edge, then there is a directed edge in the dual graph from s to t. We will label this edge (s, t). The same edge in the opposite direction would be labeled (t, s). Assign each directed edge in the graph a variable $i(s, t)$ such that $i(s, t) = -i(t, s)$. If $i(s, t)$ is positive, then $i(s, t)$ represents current flowing from s into t; if $i(s, t)$ is negative, then $i(t, s)$ is positive, and the current flows from t into s instead. The *weight* of a tile will be the sum of the currents flowing *into* the tile; we denote the weight of tile t by $w(t)$. Note that $w(t) \geq 0$.

The Kirchhoff inequalities, as stated in [5], are as follows:

1. **Vertex Equations:** For every vertex s other than t_0 and t_1, the sum of the currents flowing into s equals the sum of the currents flowing out of s.

2. **Loop Equations:** Let $s_0, \ldots, s_k$ denote a sequence of tiles forming the vertices of a loop in the dual graph. Assume in addition that each edge $(s_0, s_1), (s_1, s_2), \ldots, (s_k, s_0)$ carries nonzero current. We say that the loop is *rising* at vertex s_i if the currents (s_{i-1}, s_i) and (s_i, s_{i+1}) are both positive, and we say that the loop is *falling* at s_i if these two currents are both negative. Then if we add the weights of all vertices in the loop where the current is rising, and then subtract the weights of all vertices in the loop where the current is falling, the sum will be zero.

3. **Loop Inequalities:** We define the length of a fat path to be the sum of the weights of its vertices. Then we require that there be at least one minimal fat path from t_0 to t_1 that is (strictly) rising at each of its vertices except t_0 and t_1. Note that in order for a path to be rising at each vertex, all of its edges must carry current. Thus this condition essentially says that this rising path, being minimal, cannot be shortened by shunting the path through edges that carry no current. It is this condition that causes these three conditions together to be called the Kirchhoff *Inequalities*.

Before stating the main result dealing with the Kirchhoff Inequalities, let us examine how they correspond to the Kirchhoff Laws for electrical circuits mentioned above. We envision the tiles (vertices) as resistors, and the current-carrying edges as connections between the resistors. The Vertex Equations say that the net amount of current entering a resistor must be the net amount of current leaving the resistor, equivalent to Kirchhoff's current law. In addition, the weight of each tile is the total current entering it. If each of the resistors has unit resistance, then the voltage drop across it would be equal to the amount of current entering it and hence would be equal to the weight of the tile. The Loop Equations therefore imply that the sum of the voltage drops around a loop of resistors must be zero, equivalent to Kirchhoff's voltage law. The Loop Inequalities provide the condition which determines the optimal weight function. Because of these inequalities, all strictly rising paths will be seen to be equivalent to minimal fat flows; the sum of the rising paths through a resistor will be equal to its weight; and thus the weight determined by the current will be the weight determined by fat flow optimization.

The following result and proof appear in [5]:

Theorem 1.8 *Suppose that there exist currents $i(s,t)$, not all 0, satisfying the Kirchhoff inequalities. Assume that the values of the currents are all integers. Then the associated weights $w(s)$, $s \neq t_0, t_1$, give us a weight function which optimizes fat flows from t_0 to t_1.*

Proof: Draw $|i(s,t)|$ lines parallel to each edge (s,t) of the dual graph in the direction of current flow. By the Vertex Equations, one may extend the current lines through the vertices (so that they join each other), except at the bottom and top vertices t_0 and t_1. By definition of w (the weight function), w is the sum of the paths resulting from the joined current lines. Each current line is a path that is rising at each of its vertices (except possibly at t_0 and t_1).

By the Loop Equations, there cannot be any closed current paths. If there were such paths, then the path would be rising at each vertex, and the sum of the weights around the loop would be positive, instead of being zero as required. Thus each current path must join t_0 and t_1.

By the Loop Inequalities, there is a strictly rising minimal path from t_0 to t_1. Now if there is a rising current path from t_1 to t_0, then concatenate this path with the path rising from t_0 to t_1. The result will be a loop rising at every vertex besides t_0 and t_1, contradicting the Loop Equations. Hence, all current flows from t_0 to t_1.

Now suppose two of these paths, p and q, have different lengths. Suppose the length of p is greater than the length of q. Then if we let pq be the concatenation of these two paths, we have a loop that is rising along p and falling along q. Summing up the weights along p and subtracting the weights along q, we obtain a positive number. However, the Loop Equations stipulate that we should obtain zero instead. Hence all of these current paths have the same (minimal) length. So we conclude that the current paths are w-minimal fat flows in the tiling.

Hence w is the sum of w-minimal fat flows, and we conclude that w is therefore an optimal weight function by Theorem 1.5. $\square$

Note that the Kirchhoff inequalities provide a means of *proving* that a given weight function is indeed optimal, but they do not provide a means of *finding* an optimal weight function.

1.5.2 Two Theorems Dealing with Modulus

Later we shall refer to two theorems dealing with fat flow modulus. The first of these, the Layer Theorem from [3], essentially states that moduli are superadditive:

Theorem 1.9 (Layer Theorem) *Suppose a tiled quadrilateral or annulus X is divided into a family $\{X_j\}_{j=1}^n$ of tiled quadrilaterals or annuli, respectively, such that two distinct members of the family are disjoint except possibly at their boundaries. Then the fat flow modulus of X is greater than or equal to the sum of the fat flow moduli of the sets X_j.*

Proof: Let w_j be an optimal weight function for X_j, with $j = 1, \ldots, n$. Assume that the functions w_j are scaled such that $C_{w_j,s}(X_j) = 1$. Then if M_j is the fat flow modulus of X_j, we have

$$M_j = \frac{A_{w_j}(X_j)}{1} = A_{w_j},$$

and, therefore, we also have

$$M_j = \frac{H^2_{w_j,f}}{A_{w_j}} = \frac{H^2_{w_j,f}}{M_j} \implies M_j = H_{w_j,f}.$$

Now by Theorems 1.5 and 1.6, we may write

$$w_j = \sum_{k=1}^{j_n} a_{jk} \cdot c_{jk}$$

for positive numbers a_{jk} and w_j-minimal skinny cuts c_{jk} in X_j. Now $H_{w_j,f}(X_j) = \sum_{k=1}^{j_n} a_{jk}$, since each w_j-minimal fat flow in X_j will have to cross each of these cuts. Hence,

$$M_j = A_{w_j}(X_j) = H_{w_j,f}(X_j) = \sum_{k=1}^{j_n} a_{jk}.$$

Now define a weight function w on the entire tiling X by letting

$$w = \sum_{j=1}^{n} \sum_{k=1}^{j_n} a_{jk} \cdot c_{jk}.$$

Any w minimal fat flow must cross each of these skinny cuts; hence its length, $H_{w,f}(X)$, will be at least

$$\sum_{j=1}^{n} \sum_{k=1}^{j_n} a_{jk} = \sum_{j=1}^{n} M_j.$$

The area $A_w(X)$ will be exactly the sum of the areas of the sets Q_j under the weight functions w_j. So

$$A_w(X) = \sum_{j=1}^{n} A_{w_f}(X_j) = \sum_{j=1}^{n} M_j.$$

So, since we take the supremum over all possible weight functions to find the modulus M of X, we have

$$M \geq \frac{H_{w,f}(X)^2}{A_w X} \geq \sum_{j=1}^{n} M_j,$$

which was to be proven. $\square$

We shall require one other theorem, known as the Bounded Overlap Theorem. This theorem, appearing in [3], provides a way to compare the moduli of two different tilings of the same quadrilateral or annulus.

Theorem 1.10 (Bounded Overlap Theorem) *Suppose that a tiled quadrilateral or annulus X has two tilings T and T', such that no element of T intersects more than K elements of T' and such that no element of T' intersects more than K elements of T. Let $M(X,T)$ be the fat flow modulus of X with the tiling T, and let $M(X,T')$ be the fat flow modulus of X with the tiling T'. Then*

$$M(Q,T) \leq K^3 \cdot M(Q,T').$$

The proof of the Bounded Overlap Theorem (in [3]) does not depend on finiteness of the tiling. Thus, we shall later be able to extend it to deal with certain infinite tilings.

1.6 Other Riemann Mapping Theorems

1.6.1 The Combinatorial Riemann Mapping Theorem

We close this chapter by noticing some other combinatorial variations on the Riemann Mapping Theorem. We begin by referring to the motivation for the Finite Riemann Mapping Theorem given in Section 1.3.1. We had a sequence of covers $\{D(n)\}_{n=1}^{\infty}$ on the 2-sphere, and our goal was to take this combinatorial information and extract complex coordinates from it. With the Finite Riemann Mapping Theorem, we succeeded in obtaining moduli (and thus geometric information) from one member of this sequence. However, our goal was for the coordinates to reflect combinatorial information encoded in the entire sequence of successively refined covers. We could do so by taking the limit of the coordinates obtained from the covers. However, in order for this approach to work, the coordinates of the successive covers must be at least roughly compatible. In particular, since true quadrilaterals or annuli never have zero or infinite classical modulus, the sequence of moduli obtained from these covers cannot diverge or approach zero. Cannon proves the Combinatorial Riemann Mapping Theorem in [2]; it states, essentially, that one can obtain complex coordinates on the sphere such that the moduli of annuli (with respect to the coordinates) are within a global multiplicative bound of the limits of the combinatorial moduli. The two major hypotheses are the following:

1. If A is an annulus, let $M(A, D(n))$ be the fat flow modulus with respect to the cover $D(n)$, and let $m(A, D(n))$ be the fat cut modulus with respect to the cover $D(n)$. Then there must exist $K > 1$ such that for each annulus A, there exists $r > 0$ such that $m(A, D(n))$ and $M(A, D(n))$ are in $[r, Kr]$ for sufficiently large n. This condition ensures that the combinatorial moduli are compatible one with another.

2. Given a point x, a neighborhood N of x, and an integer J, there is an annulus A surrounding x in N such that $m(A, D(n))$ and $M(A, D(n))$ are greater than J for sufficiently large n.

Sequences of covers (with mesh locally approaching zero) are called *conformal* if they satisfy these two conditions. See [4] for a discussion of conformality of covers arising from finite subdivision rules.

1.6.2 Riemann Mappings on the Edges of Graphs

We have already mentioned that Oded Schramm gives an independent proof of the Finite Riemann Mapping Theorem ([13]). Rather than considering a tiling, Schramm considers a finite planar graph corresponding to the dual graph of a tiling. Thus, the vertices in the graph would correspond to the squares in the squared rectangle, and the optimal weight function assigns weights to the vertices. For this reason, He and Schramm refer to the optimal weight function in terms of *Vertex Extremal Length* (see [10]). It is also possible to assign weights to the edges (instead of to the vertices), in which case the calculations would be in terms of *Edge Extremal Length*. In [1], Benjamini and Schramm consider graphs dual to tilings of annuli. They demonstrate the existence of an optimal weight function assigning weights to the edges of the graph, and they show that this function maps the annulus to a right circular cylinder tiled by squares. Each square would correspond to an edge in the original graph. They consider both finite and infinite graphs.

Chapter 2

Some Results for Infinite Tilings

So far, we have seen that the fat flow modulus equals the skinny cut modulus for finite tilings of quadrilaterals and annuli. For such tilings we have also seen that fat flow optimal weight functions exist, that they are skinny cut optimal weight functions as well, and that they are unique up to scaling.

In chapter 3, we shall consider tilings of the open unit disk; such tilings would be infinite and have no boundary tiles. Suppose we remove one tile from such a tiling. The remaining tiles will form a half-open annulus—an annulus with one boundary component removed. The combinatorial modulus of this half-open annulus, as we shall see, will determine whether the tiling is parabolic or hyperbolic, both of which terms we shall define in that chapter.

Since such tilings do arise, it could therefore be useful to prove some results about them. Extending results from the finite case to the infinite is not a trivial task. For instance, the proof of Theorem 1.4 uses a compactness argument to guarantee the existence of optimal weight functions. This argument was valid because the sphere S^n is compact, n being the number of tiles in the tiling. The infinite dimensional unit sphere S^∞, however, is not compact. Would optimal weight functions necessarily exist for infinite tilings? What if a combinatorial modulus turns out to be infinite? We shall consider some such questions and prove some results concerning weight functions on such tilings.

2.1 Heights and Minimal Paths

We begin by defining the types of tilings under consideration. Let X be a simply connected set with two distinguished points on $\partial X \cap X$. If X is homeomorphic to $\{(x, y) \mid x \in [0, 1], y \in [0, 1)\}$, such that the homeomorphism maps the two distinguished points to $(0, 0)$ and $(1, 0)$, then we say that X is (topologically) a *half-open quadrilateral*. Fix such a homeomorphism ψ. The two distinguished points divide $\partial X \cap X$ into three curves. The homeomorphism

ψ will map one of these curves to the segment $\{(x,0)\,|\,x\in[0,1]\}$; we call this piece of the boundary the *bottom* of X. The piece which maps to $\{(0,y)\,|\,y\in[0,1)\}$ is called the *left* of X; and the piece which maps to $\{(1,y)\,|\,y\in[0,1)\}$ is called the *right* of X.

On the other hand, if X is homeomorphic to $\{(x,y)\,|\,1\le x^2+y^2<2\}$, we say that X is a *half-open annulus*. The homeomorphism ψ will map ∂X to the unit circle; we call this boundary the *bottom* of X.

Now suppose X is tiled by closed topological disks. We assume that the tiling is locally finite, but the number of tiles will be (countably) infinite. We shall define weights, lengths, and areas as in the finite case. Let w be a non-zero function that assigns to each tile t in X a non-negative *weight*, $w(t)$. We assume that w is ℓ^2; i.e.,

$$\sum_{t\in X} w(t)^2 < \infty.$$

Then the *w-area* of X, denoted A_w, is the sum of the squares of the weights of the tiles. If C is a collection of tiles, then the *w-length* of C, denoted $\mathrm{len}_w(C)$, is the sum of $w(t)$, taken over the tiles t in C.

A connected collection of tiles p is called a *fat path* if there exists a topological path α such that p includes all tiles which intersect α. On the other hand, a nonempty subset p of X is called a *skinny path* if it consists of distinct tiles $\{t_1,t_2,\ldots\}$ such that $t_{i-1}\cap t_i\neq\emptyset$ for $i\ge 2$. We also allow a skinny path to be a loop, in which case the path would consist of tiles $\{t_1,t_2,\ldots,t_k\}$, and we would require in addition that $t_1\cap t_k\neq\emptyset$. Note that each skinny path also covers an underlying topological path. If we have a designated set of paths P for X, then we can define the *height* of X with respect to P to be $H_w=\inf_P \mathrm{len}_w(p)$. A *$w$-minimal path* is a path $p\in P$ which actually achieves this infimum.

We now define flows and cuts. Since X has no top, we cannot define flows in terms of paths connecting top and bottom. Instead, let f be a connected set of tiles in X. Suppose there exists a topological path $\alpha:[0,1)\to X$ such that the tiles in f cover α. For any $r\in[0,1)$, let $F(r)$ be the union of all tiles in f which intersect $\alpha([r,1))$. Suppose α satisfies the following conditions:

1. the point $\alpha(0)$ lies on the bottom of X; and
2. for all compact subsets K of X, there exists $r\in[0,1)$ such that $F(r)\cap K=\emptyset$.

Then we say that f is a *flow*. In essence, a flow has an underlying topological path whose image under the homeomorphism ψ begins at $y=0$ and approaches arbitrarily close to $y=1$ (if X is a quadrilateral) or to $\{(x,y)\,|\,x^2+y^2=2\}$ (if X is an annulus). Note that *any* connected set of tiles which touches the bottom and contains infinitely many distinct tiles must be a flow, by local finiteness of the tiling. Clearly flows can be either *fat flows* or *skinny flows*. In particular, if f is a fat flow, then there is a topological path α such that all tiles intersecting α are in f. We will require that this path α hit the bottom. Thus, a fat flow will contain a tile t such that $\partial t \cap B$ contains an open interval.

A *cut* is a finite connected set of tiles with an underlying topological path which separates X into two components; one of these components must contain the entire bottom of X and only finitely many tiles. A cut that is a fat path is called a *fat cut*; a cut that is a skinny path is called a *skinny cut*.

Thus we can define four "heights," corresponding to the four sets of paths—fat and skinny flows, and fat and skinny cuts. For fat and skinny flows, we denote the corresponding w-heights by $H_{w,f}$ and $H_{w,s}$, respectively, as in the finite case. If we are dealing with fat or skinny cuts, we call the heights "circumferences," again as in the finite case; hence, we denote the w-heights corresponding to fat and skinny cuts by $C_{w,f}$ and $C_{w,s}$, respectively. A fat flow whose w-length is $H_{w,f}$ is called a w-minimal fat flow; we define w-minimal fat cuts, skinny flows, and skinny cuts similarly.

We may now define the four combinatorial moduli for these tilings. We define the *fat flow modulus* and *fat cut modulus* as

$$M_f = \sup_w \frac{H_{w,f}^2}{A_w} \qquad \text{and} \qquad m_f = \inf_w \frac{A_w}{C_{w,f}^2},$$

where the supremum and infimum are taken over all weight functions w on X. Likewise, we define the *skinny flow modulus* and *skinny cut modulus* as

$$M_s = \sup_w \frac{H_{w,s}^2}{A_w} \qquad \text{and} \qquad m_s = \inf_w \frac{A_w}{C_{w,s}^2}.$$

If a weight function actually achieves the supremum (respectively, infimum) in the above definitions, we call it an *optimal weight function* for that particular type of path. We often write w_{M_f}, w_{M_s}, w_{m_f}, and w_{m_s} for a fat flow optimal weight function, skinny flow optimal weight function, etc.

The following proposition proves the existence of minimal fat and skinny flows.

Proposition 2.1 *If X is a half-open quadrilateral or annulus with a locally finite tiling as described above, then minimal fat and skinny flows exist for any weight function.*

Proof: Before we start, we note that the number of tiles touching the bottom of X is finite by the local finiteness of the tiling. Also note that one can order the tiles in either a fat or a skinny flow as $\{t_1, t_2, \ldots\}$, such that t_1 is a bottom tile and $t_i \cap \left(\bigcup_{j=1}^{i-1} t_j\right) \neq \emptyset$ for all $i \geq 2$.

The argument is the same for fat and skinny flows, so we shall not prove separate cases for the two types of flows. Let w be any weight function on X, and let H_w be the w-height for the flows. If H_w is infinite, then the proposition is trivial. So assume that H_w is finite. The proof is by diagonalization. We know that $H_w = \inf_p \mathrm{len}_w(p)$, where the infimum is taken over all appropriate flows p. Hence, for all positive integers n, there exists a flow p_n such that

$$\mathrm{len}_w(p_n) < H_w + \frac{1}{n}.$$

We have only finitely many bottom tiles, so there exists a subsequence $\{p_{1n}\}$ of $\{p_n\}$, all the flows of which start from the same bottom tile. Since this initial tile has only finitely many neighbors, there exists a subsequence $\{p_{2n}\}$ of $\{p_{1n}\}$, all of which share the same first two tiles. Continuing inductively, we use the condition of local finiteness to obtain a subsequence $\{p_{kn}\}$ of $\{p_{(k-1)n}\}$, all of which share the same first k tiles. Now take the diagonal sequence $\{p_{nn}\}$ and note the following two facts:

1. For all $m \geq n$, p_{mm} has the same first n tiles as p_{nn}.
2. $\mathrm{len}_w(p_{nn}) < H_w + \frac{1}{n}$ because $\{p_{nn}\}$ is a subsequence of $\{p_n\}$.

Now construct a flow p whose nth tile is the nth tile of p_{nn}. Note that by the first fact listed above, p agrees with p_{nn} for the first n tiles. Thus, p starts at the bottom and contains infinitely many distinct tiles, and it is therefore indeed a flow.

We now claim that $\mathrm{len}_w(p) \leq H_w$. To see that this claim is true, denote the weight of the ith tile of p by $p(i)$, and suppose that $\mathrm{len}_w(p) \geq H_w + 1/n$ for some n. Then there exists an N such that

$$\sum_{i=1}^{N} p(i) \geq H_w + \frac{1}{n+1}.$$

Without loss of generality, we assume $n+1 \geq N$. Then

$$\sum_{i=1}^{N} p_{(n+1)(n+1)}(i) = \sum_{i=1}^{N} p(i) \geq H_w + \frac{1}{n+1}.$$

But

$$\begin{aligned} \sum_{i=1}^{N} p_{(n+1)(n+1)}(i) &\leq \sum_{i=1}^{\infty} p_{(n+1)(n+1)}(i) \\ &= \mathrm{len}_w(p_{(n+1)(n+1)}) \\ &< H_w + \frac{1}{n+1}, \end{aligned}$$

by our second fact above; this final inequality contradicts the above result. Hence $\mathrm{len}_w(p) \leq H_w$. Since H_w is the infimum of flow lengths, $\mathrm{len}_w(p) = H_w$, and p is a w-minimal flow. □

Two facts allow us to construct this w-minimal flow—namely, the fact that each tile has finitely many neighbors and the fact that the set of bottom tiles is finite. Figure 2.1 shows

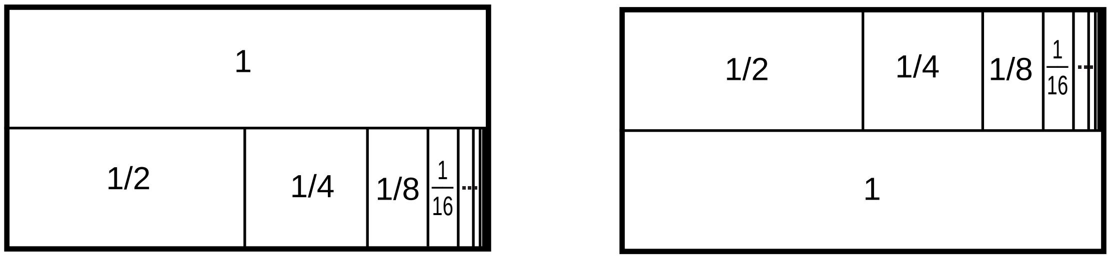

Figure 2.1: No Minimal Flows

two examples (neither of which are half-open quadrilaterals by our definition) for which there do not exist minimal flows. Note that in each case, the fat flow modulus is finite (and is in fact equal to one). The examples are designedly similar in order to emphasize that the problem is the existence of infinitely many "choices" of how to begin or extend flows.

2.2 The Vector Space Setting

As with the finite case, we can approach the situation by using a vector space. Since there are countable many tiles in X, we may well-order the tiles and consider each weight function w to be a point in the Hilbert space ℓ^2. Considering w as an element of ℓ^2, the ith component $(w)_i$ of w would be the weight assigned to the ith tile of the tiling. We would denote the set of all weight vectors by $\mathcal{W}$. Finite subsets C of X (such as cuts) would correspond to vectors in ℓ^2, with the ith component $(C)_i$ being 1 if the ith tile is in C and 0 otherwise. The w-length of C, denoted $\text{len}_w(C)$, would equal $\langle w, C\rangle$, where $\langle \cdot, \cdot \rangle$ denotes the standard inner product on ℓ^2. One could represent flows in this manner as well; however, they would not be ℓ^2 vectors since they would not have finite norm. As in the finite case, we define heights and moduli with respect to a specified set of paths P. The w-length of a path p is no longer necessarily an inner product; rather, we must define it to be

$$\text{len}_w(p) = \sum_{i=1}^{\infty} (p)_i (w)_i,$$

which of course reduces to an inner product if the path P is in the vector space.

We shall use the two viewpoints (vector space and tiling) interchangeably. Using the vector space viewpoint, we can prove analogues of two results in [5].

Proposition 2.2 *Let w be a weight vector, and let P be the set of w-minimal paths. If there exists a measure μ on P such that*

$$w = \int_P p \, d\mu,$$

then

$$A_w = H_w \int_P d\mu.$$

Proof: First note that we may view $\int_P p\, d\mu$ as a sequence, each of whose components is non-negative. The value of the ith element in this sequence is

$$\left(\int_P p\, d\mu\right)_i = \int_P (p)_i\, d\mu.$$

Since $w = \int_P p\, d\mu$, this sequence is ℓ^2 and is hence a vector in $\mathcal{W}$. So we have

$$\begin{aligned}
A_w &= \langle w, w\rangle \\
&= \left\langle w, \int_P p\, d\mu \right\rangle \\
&= \sum_i (w)_i \left(\int_P p\, d\mu\right)_i \\
&= \sum_i (w)_i \int_P (p)_i\, d\mu \\
&= \sum_i \int_P (p)_i (w)_i\, d\mu \\
&= \int_P \left(\sum_i (p)_i (w)_i\right) d\mu \quad \text{by the Monotone Convergence Theorem} \\
&= \int_P H_w\, d\mu \quad \text{by minimality of each } p \\
&= H_w \int_P d\mu.
\end{aligned}$$

□

Theorem 2.1 *Suppose $w_0 \in \mathcal{W}$, $\|w_0\| = 1$, and $w_0 = \int_P p\, d\mu$ for some measure p on the set P of w_0-minimal paths. Suppose $M < \infty$. Then w_0 is an optimal weight function.*

Proof: We know that $M = \sup_{\mathcal{W}}(H_w^2/A_w) = \sup_{\|w\|=1} H_w^2$; thus, $\sqrt{M} = \sup_{\|w\|=1} H_w$. So, let $\{w_j\}_{j=1}^{\infty}$ be weight vectors of unit norm such that $H_{w_j} > \sqrt{M} - 1/j$. Then

$$\begin{aligned}
\langle w_0, w_j\rangle &= \left\langle \int_P p\, d\mu, w_j \right\rangle \\
&= \int_P \left(\sum_i (p)_i (w_j)_i\right) d\mu \quad \text{as in the proof of Proposition 2.2} \\
&\geq \int_P H_{w_j}\, d\mu \quad \text{since } H_{w_j} = \inf_p \textstyle\sum_i (p)_i (w_j)_i \\
&> \int_P \left(\sqrt{M} - \frac{1}{j}\right) d\mu \\
&= \left(\sqrt{M} - \frac{1}{j}\right) \int_P d\mu
\end{aligned}$$

$$= \left(\sqrt{M} - \frac{1}{j}\right)\frac{A_{w_0}}{H_{w_0}} \qquad \text{by Proposition 2.2}$$
$$= \frac{\sqrt{M} - \frac{1}{j}}{H_{w_0}}$$

Now $\{w_j\}$ is a bounded sequence, so by a theorem of functional analysis (see [8]) there exists a subsequence $\{w_{j_n}\}$ converging weakly to a vector w_ℓ. Thus, for all $w \in \mathcal{W}$, $\langle w, w_{j_n}\rangle \longrightarrow \langle w, w_\ell\rangle$. In particular,

$$\frac{\sqrt{M} - 1/j}{H_{w_0}} < \langle w_0, w_{j_n}\rangle \longrightarrow \langle w_0, w_\ell\rangle.$$

Therefore,

$$\langle w_0, w_\ell\rangle \geq \frac{\sqrt{M}}{H_{w_0}} \geq 1,$$

since $\sqrt{M} = \sup_{\|w\|=1} H_w$.

Since $\|w_0\| = 1$, we have $\|w_\ell\| \geq 1$ by the Cauchy-Bunyakowski-Schwarz inequality. So $1 = \|w_{j_n}\| \leq \|w_\ell\|$ for all n. By another theorem of functional analysis (see [8]), this fact implies strong convergence $w_{j_n} \longrightarrow w_\ell$. Therefore, $\|w_\ell\| = 1$. And, since $\langle w_0, w_\ell\rangle \geq 1$ and $\|w_0\| = 1$, we know that $w_0 = w_\ell$.

However,

$$1 = \|w_0\|^2 = \langle w_0, w_\ell\rangle \geq \frac{\sqrt{M}}{H_{w_0}}$$

from above, whereas we also have

$$\frac{\sqrt{M}}{H_{w_0}} \geq 1$$

by definition of M. So, $\sqrt{M}/H_{w_0} = 1$, and $H_{w_0}^2 = M$. Hence w_0 is an optimal weight vector. $\square$

Note that in Proposition 2.5 we shall prove that $M = \infty$ implies that no weight function is the integral of its minimal paths. Thus the hypothesis that $M < \infty$ (in the statement of Theorem 2.1) is redundant.

The next proposition shows that optimal weight vectors are unique up to scaling if $M < \infty$.

Proposition 2.3 *Let P be a set of paths in the vector space setting, and suppose the combinatorial modulus with regard to P is finite. Then there exists at most one optimal weight vector of unit norm.*

Proof: Let S be the unit sphere in our vector space. Then S is convex. Suppose that w_1 and w_2 are two distinct optimal weight vectors in S. Since w_1 and w_2 are optimal, and since

$M < \infty$, we know that $H_{w_1} = H_{w_2} = \sqrt{M} < \infty$. Let $w = (1-t)w_1 + tw_2$ for some t such that $0 < t < 1$. If p is a path, then we have

$$\sum_i (p)_i (w)_i = (1-t)\sum_i (p)_i (w_1)_i + t\sum_i (p)_i (w_2)_i \geq (1-t)H_{w_1} + tH_{w_2} = \sqrt{M}.$$

Since $0 < \|w\| < 1$, since $(1/\|w\|)w$ lies in S, and since $\sum_i (p)_i((1/\|w\|)w)_i > \sum_i (p)_i (w)_i \geq \sqrt{M}$ for all paths p, then we know that

$$H_{\frac{1}{\|w\|}w}(P) > \sqrt{M},$$

which cannot be. So optimal weight vectors of unit norm are unique. $\square$

2.3 Dealing with Degeneracies

In considering the existence of optimal weight functions, we must be more careful than we were in the finite case. In particular, we must take into account the fact that moduli may be infinite. We must also be sure that the moduli are non-zero. We have the following proposition regarding figures with infinite moduli; since we prove it in the vector space setting, it applies equally well to the fat flow modulus and the skinny cut modulus.

Proposition 2.4 *Let P be a set of paths in the vector space setting. Suppose the modulus M with regard to P is infinite. Then optimal weight vectors of unit norm, if they exist, are not unique.*

Proof: Suppose w is an optimal weight vector with unit norm. Then $M = H_w^2/A_w = H_w^2$, so H_w must be infinite. Thus, for every p in P, $\sum_{i=1}^{\infty}(p)_i(w)_i = \infty$. Let w' be the vector with components

$$(w')_i = \begin{cases} (w)_i + 1 & \text{if } i = 1 \\ (w)_i & \text{otherwise.} \end{cases}$$

Clearly w' is not a scalar multiple of w, so $w'/\|w'\|$ is a unit vector different from w. Just as clearly, $H_{w'} \geq H_w = \infty$, so w' is an optimal weight vector of unit norm. $\square$

The next proposition shows that if $M = \infty$, then no ℓ^2 weight function is an integral of its minimal paths. Thus, we could remove the hypothesis of finite modulus from Theorem 2.1.

Proposition 2.5 *Let $\mathcal{W}$ and P be as described above. Suppose the modulus M is infinite. Then there is no weight vector that is an integral of its minimal paths.*

Proof: Suppose $M = \infty$, and suppose we have a weight vector w_0 that is an integral of its minimal paths. By invariance of the modulus ratio under scaling, we may assume that w_0 has unit norm. If P_0 is the set of w_0-minimal paths, then

$$w_0 = \int_{P_0} p\, d\mu$$

for some measure μ. Since $M = \infty$, we know that $\sup_{\|w\|=1} H_w = \infty$. Thus, for every positive integer j, there is a weight function w_j of unit norm such that $H_{w_j} > j$. Now we have

$$\begin{aligned}
\langle w_0, w_j \rangle &= \left\langle \int_{P_0} p\, d\mu, w_j \right\rangle \\
&= \int_{P_0} \left(\sum_i (p)_i (w_j)_i \right) d\mu \quad \text{as in the proof of Proposition 2.2} \\
&\geq \int_{P_0} H_{w_j}\, d\mu \\
&> \int_{P_0} j\, d\mu = j \int_{P_0} d\mu \\
&= j \frac{A_{w_0}}{H_{w_0}} \quad \text{by Proposition 2.2} \\
&= \frac{j}{H_{w_0}}
\end{aligned}$$

Now $\{w_j\}$ is a bounded sequence, so by the same theorem of functional analysis that we used in the proof of Theorem 2.1, we know there is a subsequence $\{w_{j_n}\}$ converging weakly to a vector w_ℓ. Thus, for all w in $\mathcal{W}$, we have $\langle w, w_{j_n} \rangle \longrightarrow \langle w, w_\ell \rangle$. Now for all $m \geq n$, we have

$$\frac{j_n}{H_{w_0}} \leq \frac{j_m}{H_{w_0}} < \langle w_0, w_{j_m} \rangle.$$

But $\langle w_0, w_{j_m} \rangle$ converges to $\langle w_0, w_\ell \rangle$ as $m \to \infty$, so we can conclude that

$$\frac{j_n}{H_{w_0}} \leq \langle w_0, w_\ell \rangle.$$

This inequality is true for all n. However, $\{j_n\}$ is a subsequence of the positive integers, and hence $\langle w_0, w_\ell \rangle = \infty$, which cannot be. Thus no weight vector is the integral of its minimal paths if $M = \infty$. $\square$

See Figure 2.2 for an example of a half-open quadrilateral with infinite fat flow modulus; the tiles are labeled with an optimal weight function that is not the sum of its minimal paths.

Thus we may run into problems when our tilings have infinite modulus. The next proposition, however, shows that we do not have to worry about half-open quadrilaterals or annuli with zero modulus.

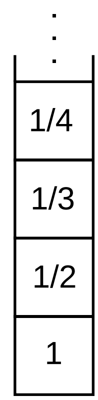

Figure 2.2: An Optimal Weight Function for a Tiling with Infinite Modulus

Proposition 2.6 *Let X be a half-open quadrilateral or annulus with a locally finite tiling. Then neither M_f nor m_s are zero.*

Proof: First consider M_f. Suppose X has n tiles whose boundaries intersect the bottom; any fat flow must contain one of these tiles. For any such tile t, let $w(t) = 1/\sqrt{n}$, and let $w(t) = 0$ for any other tile. Then $A_w = 1$, and $H_{w,f} \geq 1/\sqrt{n}$. So clearly $M_f \geq 0$.

Next consider m_s. Suppose $m_s = 0$. Then $\inf_{\|w\|=1} 1/C^2_{w,s} = 0$, so $C_{w,s}$ is unbounded on the set $\{w \mid \|w\| = 1\}$. Thus, for any skinny cut c, $\langle c, w\rangle$ is unbounded on the set $\{w \mid \|w\| = 1\}$. However, c is a finite set, and hence the supremum of $\langle c, w\rangle$ on the set of unit weight vectors must be finite. Thus we conclude that $m_s \neq 0$. □

Now let us establish some notation. Let B be the bottom boundary segment of X. If X is a topological quadrilateral, then let L and R be its left and right boundary segments, respectively. Now take finite subtilings $\{X_n\}$ such that:

1. $B \subseteq X_n$ for all n;
2. $X_n \subseteq X_{n+1}$ for all n;
3. If X is a half-open quadrilateral, then X_n is a topological quadrilateral with bottom B, left side $\partial X_n \cap L$, right side $\partial X_n \cap R$, and top T_n, where $T_n = \partial X_n \setminus \{B \cup R \cup L\}$. If X is a half-open annulus, then X_n is a topological annulus with bottom B and top T_n, where $T_n = \partial X_n \setminus B$;
4. $T_n \subseteq \text{int } X_{n+1}$ for all n;
5. $\bigcup_{n=1}^{\infty} X_n = X$.

Clearly each cut in X_n is also a cut in X. We may now establish the existence of a skinny cut optimal weight function.

Proposition 2.7 *Let X be a half-open quadrilateral or annulus such that $m_s < \infty$. Then there exists an optimal weight function for skinny cuts on X, and it is the weak (or tile-wise) limit of the normalized optimal weight functions on a subsequence of the sets X_n.*

Proof: By Proposition 2.6 we have $m_s \neq 0$. Now since

$$m_s = \inf_w \frac{A_w}{C_{w,s}^2} = \inf_{\|w\|=1} \frac{1}{C_{w,s}^2},$$

we can write

$$\frac{1}{\sqrt{m_s}} = \sup_{\|w\|=1} C_{w,s}.$$

Therefore, for every positive integer n, we can get a weight function $\bar{w}_n$ on X such that $\|\bar{w}_n\| = 1$ and $C_{\bar{w}_n,s} > \frac{1}{\sqrt{m_s}} - \frac{1}{n}$. Now consider the subtilings X_n described above. For each n, restrict $\bar{w}_n$ to X_n; denote the restriction by $w_n^\sharp$, and note that it is a weight function defined on the finite tiling X_n. Clearly

$$\|w_n^\sharp\| \leq \|\bar{w}_n\| = 1,$$

so $1/\|w_n^\sharp\| \geq 1$. Now let

$$\hat{w}_n = \frac{1}{\|w_n^\sharp\|} w_n^\sharp,$$

so that $\hat{w}_n$ has norm 1. In the rest of this proof, we shall use the notation $C_{w,s}(X_n)$ and $C_{w,s}(X)$ to distinguish between the skinny circumference (under w) of X_n and of X, respectively. If we denote a skinny cut by c, we have

$$\begin{aligned}
C_{\hat{w}_n,s}(X_n) &= \inf_{c \subseteq X_n} \langle \hat{w}_n, c \rangle = \inf_{c \subseteq X_n} \left\langle \frac{1}{\|w_n^\sharp\|} w_n^\sharp, c \right\rangle \\
&= \frac{1}{\|w_n^\sharp\|} \inf_{c \subseteq X_n} \langle w_n^\sharp, c \rangle = \frac{1}{\|w_n^\sharp\|} \inf_{c \subseteq X_n} \langle \bar{w}_n, c \rangle \\
&\geq \frac{1}{\|w_n^\sharp\|} \inf_{c \subseteq X} \langle \bar{w}_n, c \rangle \geq \inf_{c \subseteq X} \langle \bar{w}_n, c \rangle \\
&= C_{\bar{w}_n,s}(X) > \frac{1}{\sqrt{m_s}} - \frac{1}{n}.
\end{aligned}$$

By Theorem 1.4, each X_n has a skinny cut optimal weight function w_n of unit norm. Thus

$$\frac{1}{C_{w_n,s}^2(X_n)} = \frac{A_{w_n}}{C_{w_n,s}^2(X_n)} = \inf_w \frac{A_w}{C_{w,s}^2(X_n)} \leq \frac{A_{\hat{w}_n}}{C_{\hat{w}_n,s}^2(X_n)} = \frac{1}{C_{\hat{w}_n,s}^2(X_n)},$$

where the infimum is taken over all weight functions on X_n. So

$$C_{w_n,s}(X_n) \geq C_{\hat{w}_n,s}(X_n) > \frac{1}{\sqrt{m_s}} - \frac{1}{n}.$$

As before, a theorem of functional analysis guarantees the existence of a subsequence $\{w_{n_j}\}$ converging weakly to a function w_ℓ. Now for any skinny cut c in X, there exists an integer N such that $c \subseteq X_n$ for all $n \geq N$. So for all $n_j \geq N$,

$$\langle w_{n_j}, c\rangle \geq C_{w_{n_j},s}(X_{n_j}) > \frac{1}{\sqrt{m_s}} - \frac{1}{n_j}.$$

But $\langle w_{n_j}, c\rangle \longrightarrow \langle w_\ell, c\rangle$, so $\langle w_\ell, c\rangle \geq 1/\sqrt{m_s}$ for all skinny cuts c in X. Since we are assuming that m_s is finite, we know that w_ℓ is nonzero and hence is a weight function. By the Banach-Alaoglu Theorem for Hilbert spaces, $\|w_\ell\| \leq 1$. Now since $\langle w_\ell, c\rangle \geq \frac{1}{\sqrt{m_s}}$ for all skinny cuts c, we know that $C_{w_\ell,s}(X) \geq \frac{1}{\sqrt{m_s}}$. Thus we have

$$\frac{\|w_\ell\|^2}{C^2_{w_\ell,s}(X)} \leq \frac{1}{1/m_s} = m_s = \inf_w \frac{\|w\|^2}{C^2_{w,s}(X)}.$$

Therefore

$$\frac{\|w_\ell\|^2}{C^2_{w_\ell,s}(X)} = m_s,$$

so w_ℓ is a skinny cut optimal weight function for X. $\square$

The following proposition relates the finiteness of M_f to the finiteness of m_s:

Proposition 2.8 *Let X be a half-open topological quadrilateral or annulus. If M_f is finite, then m_s is finite also.*

Proof: Suppose M_f is finite. For each positive integer n, we let $\tilde{w}_n$ be a fat flow optimal weight function for X_n such that $\|\tilde{w}_n\| = 1$. Since the sequence $\{\tilde{w}_n\}$ is bounded, we know (as in the proof of Proposition 2.7) that we have a subsequence $\{\tilde{w}_{n_j}\}$ converging weakly to $\tilde{w}_\ell$, a weak limit. Pass to this subsequence. Let M_n be the fat flow modulus of X_n, and recall that M_n is also the skinny cut modulus by Theorem 1.4. Let $w_n = (\sqrt{M_n})\tilde{w}_n$. So we have

$$A_{w_n} = \|w_n\| = M_n\|\tilde{w}_n\| = M_n.$$

We know that w_n optimizes both fat flows and skinny cuts. Thus,

$$\frac{A_{w_n}}{C^2_{w_n,s}} = M_n,$$

and we conclude that $C_{w_n,s} = 1$. Now by the Finite Riemann Mapping Theorem, we have $C_{w_n,s} \cdot H_{w_n,f} = A_{w_n} = M_n$, so we know that $H_{w_n,f} = M_n$.

By the Layer Theorem, $M_{n+1} \geq M_n$ for all n; thus $\{M_n\}$ is an increasing sequence. Since the fat flow modulus of X is finite, the sequence $\{M_n\}$ is bounded and thus converges. Let $M = \lim_n M_n$. Let $w_\ell = (\sqrt{M})\tilde{w}_\ell$. Then for all t,

$$w_n(t) = \left(\sqrt{M_n}\right)\tilde{w}_n(t) \longrightarrow (\sqrt{M})\tilde{w}_\ell(t) = w_\ell(t).$$

Hence the sequence $\{w_n\}$ converges weakly to w_ℓ.

Now for all skinny cuts c, we have $\langle w_n, c\rangle \geq C_{w_n,s} = 1$ for all n large enough that $c \subseteq X_n$. Thus $\langle w_\ell, c\rangle \geq 1$ for all skinny cuts c. Hence we know that $\|w_\ell\| > 0$, so w_ℓ is a weight function. The fact that $\langle w_\ell, c\rangle \geq 1$ for all skinny cuts c also implies that $C_{w_\ell,s} \geq 1$, so

$$\frac{A_{w_\ell}}{C^2_{w_\ell,s}} \leq A_{w_\ell} < \infty.$$

Therefore

$$m_s = \inf_w \frac{A_{w_\ell}}{C^2_{w_\ell,s}} \leq A_{w_\ell} < \infty.$$

□

If we pass to the subsequence described in Proposition 2.7, we will have a sequence of fat flow/skinny cut optimal weight functions w_n on X_n and a skinny cut optimal weight function w_ℓ on X. We know that w_n converges to w_ℓ weakly. We can also establish that the w_ℓ-height of X is non-zero.

Proposition 2.9 *Let H be the w_ℓ-height of X. If $m_s < \infty$, then H does not equal zero.*

Proof: Suppose $H = 0$. Then by Proposition 2.1 we know that there exists a fat flow f such that $\text{len}_{w_\ell}(f) = 0$. Since the w_ℓ-weight of every tile in f is zero, we know that $w_n(t) \longrightarrow 0$ for all t in f. We shall say that two tiles t and s are *w-fat adjacent* if there exists a fat path p from t to s such that $\text{len}_w(p) = w(t) + w(s)$. In other words, one can connect the two tiles via tiles with zero w-weight. We know that f contains at least one bottom tile, and we may assume that f is as large as possible—namely, that it contains all tiles of zero w_ℓ-weight which are w_ℓ-fat adjacent to f.

Let T be the set of bottom tiles which are in f. We claim that T is connected; if it were not, then there would be some bottom tile t of non-zero w_ℓ-weight that is flanked by tiles of T (and hence by tiles of f). Now any w_ℓ-minimal skinny cuts can bypass t via f since all tiles of f have zero w_ℓ-weight. Thus, reducing the weight of t to zero will not affect the circumference. So we can reduce the area of X without affecting the circumference, thus

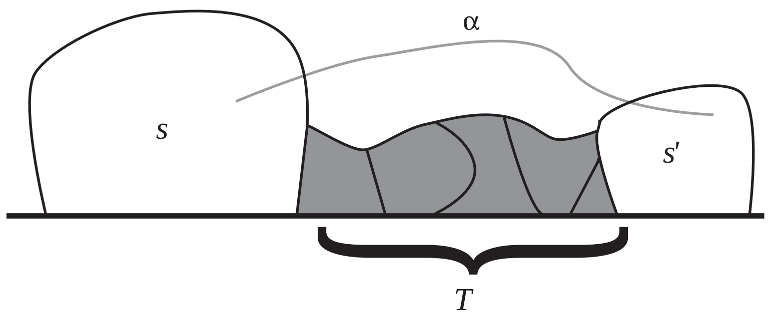

Figure 2.3: The Length of T Is Non-Zero

lowering the ratio $A_w/C_{w,s}^2$. But this result contradicts the fact that w_ℓ is a skinny cut optimal weight function. We conclude that T is a connected set.

Now $m_s < \infty$, so $C_{w_\ell,s} \neq 0$. Thus we know that the set of bottom tiles does not have zero w_ℓ-weight as a whole. So first we assume that there exist bottom tiles s and s' immediately adjacent (on either side) to the set of tiles T. Since $w_\ell(s)$ and $w_\ell(s')$ are greater than zero, there exists a positive integer N such that for all $n \geq N$, we have $w_n(s) > 0$ and $w_n(s') > 0$. Now we claim that $\mathrm{len}_{w_n}(T) > 0$ for all $n \geq N$. Suppose this were not true. Let $n \geq N$ be such that $\mathrm{len}_{w_n}(T) = 0$. Then $w_n(s)$ and $w_n(s')$ are both greater than zero. By Theorem 1.4, there is a w_n-minimal skinny cut α containing both s and s', such that none of the tiles in α have zero w_n-weight. (See Figure 2.3.) Since α is w_n-minimal, rerouting it through T would not shorten its w_n-length, and hence α must pass directly from s to s', going through no intervening tiles. So s and s' must touch each other above T, and it follows that f, being a fat flow, would have to contain either s or s'. This conclusion contradicts the fact that neither s nor s' have zero w_ℓ-weight.

Now since T is finite, we know that $\mathrm{len}_{w_n}(T) \longrightarrow 0$, whereas $w_n(s)$ and $w_n(s')$ both approach positive values. Let K be the maximum of the valence of s and the valence of s'. Now choose N' such that for all $n \geq N'$, we have $0 < \mathrm{len}_{w_n}(T) < 1/(2K)\min\{w_n(s), w_n(s')\}$. If any two tiles share part of a side in the squared rectangle diagram for X_n (for any n), then there is a w_n-minimal skinny cut or fat flow passing directly from one to the other. Thus two adjacent tiles in the squared rectangle diagram are also adjacent in the original tiling. By the same reasoning, if a tile has an edge on the bottom in the squared rectangle diagram, it is part of a minimal fat flow, all of whose tiles have non-zero w_n-weight. It follows that any tile with an edge on the bottom in the squared rectangle diagram is also a bottom tile in the original tiling.

Consider the squared rectangle diagram for some $n \geq N'$, shown in Figure 2.4. Since all tiles which touch the bottom in the squared rectangle diagram are bottom tiles in the tiling, the distance d between the squares representing s and s' is equal to $\mathrm{len}_{w_n}(T) < 1/(2K)\min\{w_n(s), w_n(s')\}$. Therefore, it will take at least $2K$ tiles to fill up the gap between the two squares. However, any two tiles adjacent in the diagram are adjacent in the original

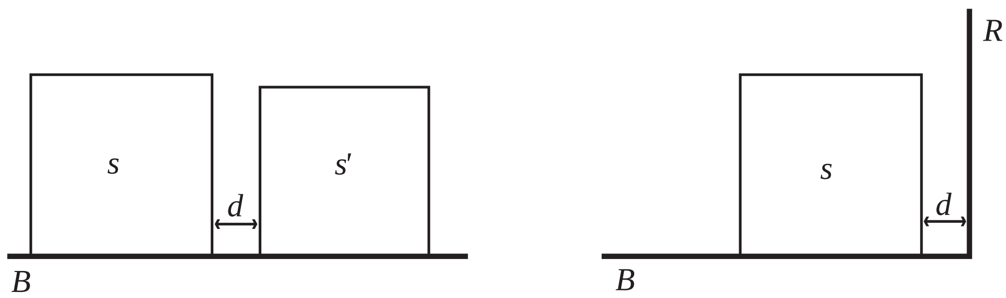

Figure 2.4: Contradiction If a Fat Flow Has Zero Length

tiling, and s is only adjacent to at most K tiles. This is a contradiction.

We supposed that there existed two bottom tiles adjacent to T. This could fail to be true if X is a half-open quadrilateral and if T includes a bottom tile adjacent to one of the two sides (L or R). Supposing this to be the case, we can use the same analysis, except that now there is just one bottom tile s of non-zero w_ℓ-weight that is adjacent to T. We still have positive $\mathrm{len}_{w_n}(T)$ for all n sufficiently large, and we can deduce that $\mathrm{len}_{w_n}(T) \longrightarrow 0$ and that $w_n(s)$ approaches a positive value. If K is the valence of s, then we may choose N such that for all $n \geq N$, we have $0 < \mathrm{len}_{w_n}(T) < 1/(2K)w_n(s)$. In the squared rectangle diagram for $n \geq N$, the distance d between the square representing s and the side (L or R) adjacent to T is equal to $\mathrm{len}_{w_n}(T) < 1/(2K)w_n(s)$. (See Figure 2.4 again.) Therefore, it will take at least $2K$ tiles to fill up the gap between the square and the side. However, any two tiles adjacent in the diagram are adjacent in the original tiling, and s is only adjacent to at most K tiles. Once again, we have a contradiction. $\square$

2.4 Working Backwards

In Chapter 1, we saw that we can convert a finite tiling of a quadrilateral to a squared rectangle. We can also try to work backwards. Suppose we are given a tiling that is itself a squared rectangle. We will allow this squared rectangle to be either a finite tiling of a geometric rectangle by squares or a locally finite tiling of a half-open geometric rectangle by squares. Suppose we define a weight function w on each tile by setting $w(s)$ equal to the side length of the square s. According to the following proposition, this function w is an optimal weight function.

Proposition 2.10 *Suppose one is given a tiling that is a squared rectangle, possibly half-open. Then the weight function determined by the lengths of the sides of the squares is an optimal weight function for fat flows, skinny flows, fat cuts, and skinny cuts. An analogous result holds for square tilings of right circular cylinders, possibly half-open.*

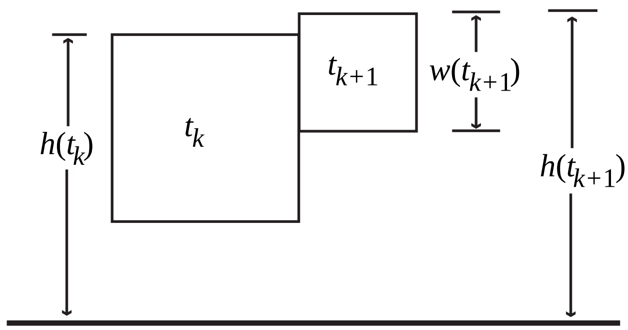

Figure 2.5: Adding Heights in a Squared Rectangle

Proof: Let w be the weight function determined by side lengths, so that $w(s)$ is the length of the side of the square s. We shall first show that w optimizes both fat and skinny flows. The argument is the same for either case, so for the rest of the proof we shall refer to "flows," instead of proving separate cases for fat and skinny flows. Let P be the set of all such flows which are w-minimal. Now if v is a vertical line on the squared rectangle which does not contain a vertical edge of some square, then the set of all tiles intersecting v is a flow. This flow is both fat and skinny. Call such flows *vertical flows*, and let $\hat{P}$ be the set of all vertical flows. If H is the (geometric) height of the rectangle, then each vertical flow has a length of H. Thus the length of a w-minimal flow is less than or equal to H.

We claim that $\hat{P} \subseteq P$. Let $f = \{t_1, t_2, \ldots\}$ be a flow in P, where t_1 is a bottom tile. For any tile t, let $h(t)$ be the height of the top of t above the bottom of the rectangle. Then, since t_1 is a bottom tile, we have $w(t_1) = h(t_1)$. We also know that if $\sum_{i=1}^{k} w(t_i) \geq h(t_k)$, we have

$$\sum_{i=1}^{k+1} w(t_i) \geq h(t_k) + w(t_{k+1}) \geq h(t_{k+1}).$$

(See Figure 2.5.) This fact allows us to see (by induction) that $\mathrm{len}_w(t_1, \ldots, t_k) \geq h(t_k)$ for all positive integers k. Now since f is a flow, we know that $\lim_k h(t_k) = H$. Thus, $\mathrm{len}_w(f) = \lim_k \mathrm{len}_w(t_1, \ldots, t_k) \geq H$. But since f is a w-minimal flow, $\mathrm{len}_w(f) \leq H$. It follows that the w-height of the rectangle is H. Since all vertical flows have height H, it follows that they are all w-minimal, so $\hat{P} \subseteq P$.

Let W be the width of the rectangle. For each r in $(0, W)$, take the vertical line at a distance of r from the left side of the rectangle. As long as this line does not contain an edge of a square, it defines a vertical flow f_r. Since there are only finitely many squares, we have f_r defined for almost every r. Clearly $\hat{P}$ is simply the set of all f_r. So define a measure μ on P as follows. We shall let $\mu(P \setminus \hat{P}) = 0$. If $S \subseteq \hat{P}$ such that the set $\{r \mid f_r \in S\}$ is Lebesgue measurable, then we let $\mu(S) = m\,\{r \mid f_r \in S\}$, where m is Lebesgue measure.

We now show that $w = \int_P f\, d\mu$. Take any tile t, and let $P|_t$ be the set of paths in P which contain t. Likewise, we shall let $\hat{P}|_t$ be the set of paths in $\hat{P}$ which contain t. Let x and y be the left and right endpoints, respectively, of the bottom edge of t. Let a and b be the distances of x and y, respectively, from the left side of the squared rectangle. Since $b - a$ is the side length of the square t, it follows that $b - a = w(t)$. A vertical flow f_r is in $\hat{P}|_t$ if and only if r is in (a, b). We shall use the convention, from the vector space setting, that $f(t) = 1$ if t is in f, and that $f(t) = 0$ otherwise. We may thus say,

$$\int_P f(t)\, d\mu = \mu P|_t = \mu\hat{P}|_t = m(a, b) = b - a = w(t).$$

So $w = \int_P f\, d\mu$, and w is an optimal weight function for flows by Theorem 2.1.

An analogous procedure shows that w is an optimal weight function for both fat and skinny cuts.

If we are given a right circular cylinder tiled by squares, we have a tiling of a topological annulus. Adopt the same procedure, except let r go from 0 to 2π. The same result follows. □

2.5 Kirchhoff Inequalities for Infinite Tilings

In this section we shall prove versions of the Kirchhoff Inequalities for locally finite tilings of half-open quadrilaterals and annuli.

Let G be a locally finite, possibly infinite, graph. Suppose that a finite contiguous number of boundary vertices are denoted as *bottom* vertices. We shall add a new vertex t_0 and add a single edge from t_0 to each bottom vertex. Call these edges *bottom edges* and denote the set of bottom edges by B. If we begin with a locally finite tiling of a half-open quadrilateral or annulus, then we can realize G as the dual graph of the tiling. (In the finite case, we added a vertex t_1 to the top. Since half-open quadrilaterals and annuli do not have a set of top vertices, we cannot do so in the case under consideration.)

Now, if s and t are two adjacent vertices in G, then denote the directed edge from s to t by (s, t). Denote the same edge in the opposite direction by (t, s). Define a *current function* i on the set of directed edges such that $i(s, t) = -i(t, s)$ for all directed edges (s, t). If $i(s, t)$ is positive, then $i(s, t)$ represents current flowing from s into t; if $i(s, t)$ is negative, then $i(t, s)$ is positive, and the current flows from t into s instead. The *weight* of a vertex will be the sum of the currents flowing *into* that vertex; we denote the weight of a vertex t by $w(t)$. Note that $w(t) \geq 0$ for all vertices t.

The Infinite Kirchhoff Inequalities are as follows:

1. **The Current Equations:** For every vertex s other than t_0, the sum of the currents flowing into s equals the sum of the currents flowing out of s.

2. The Loop Equations: Let $s_0, \ldots, s_k$ denote a sequence of vertices forming a loop in G. Assume in addition that each edge $(s_0, s_1), (s_1, s_2), \ldots, (s_k, s_0)$ carries nonzero current. We say that the loop is *rising* at vertex s_i if the currents (s_{i-1}, s_i) and (s_i, s_{i+1}) are both positive, and we say that the loop is *falling* at s_i if these two currents are both negative. Then if we add the weights of all vertices in the loop where the current is rising, and then subtract the weights of all vertices in the loop where the current is falling, the sum will be zero. In addition, let $\ldots, s_{-1}, s_0, s_1, \ldots$ be a sequence of vertices forming a path with no beginning, no end, and no repeated vertices. Let R_n be the set of vertices in $\{s_{-n}, \ldots, s_n\}$ at which the path is rising; let F_n be the set of vertices in $\{s_{-n}, \ldots, s_n\}$ at which the path is falling. Then we require that the sums $\sum_{R_n} w(s) - \sum_{F_n} w(s)$ approach 0 as n approaches infinity.

3. The Loop Inequalities: We define the length of a path in G to be the sum of the weights of its vertices. We require that there exist at least one strictly rising path p, starting at t_0 and containing infinitely many vertices, such that p has minimal w-weight. (We take the minimum over *all* paths which start at t_0 and have infinitely many distinct vertices, not just over rising paths.) This condition essentially says that p, being minimal, cannot be shortened by shunting the path through edges that carry no current.

We shall show that any current function which satisfies the Infinite Kirchhoff Inequalities defines a fat flow optimal weight function. The proof of this theorem will require Theorem 2.1 and will involve integration against a measure determined by the current function. First, however, let us consider strictly rising paths. Let p be such a path; if p is not properly contained in another strictly rising path, then we say that p is a *current path* or simply a *path*. The next proposition describes what current paths look like for functions which satisfy the Infinite Kirchhoff Inequalities.

Proposition 2.11 *Suppose i is a current function which satisfies the Infinite Kirchhoff Inequalities. Then each current path begins at t_0, contains no loops, and contains infinitely many distinct vertices.*

Proof: Let t be some vertex other than t_0. By the Current Equations, if there exists an edge carrying current into t, then there must exist an edge carrying current out of t. Thus one can extend strictly rising paths through any tile besides t_0. We conclude that the only vertex at which a current path can start or end is t_0. Furthermore, the Loop Equations guarantee that there are no closed strictly rising paths. If there were such paths, then the path would be rising at each vertex, and the sum of the weights around the loop would be positive, instead of being zero as required. So current paths cannot contain loops; thus each current path is an infinitely long arc with at most one endpoint, which would be t_0. Furthermore, we know that each current path must have at least one endpoint; otherwise, the Loop Equations would once again require that the sum of the weights of the vertices in the path be zero, whereas in reality it would be positive. So each current path is an infinitely long arc with one endpoint at t_0.

By the Loop Inequalities, there is at least one strictly rising path p starting at t_0. If there were a strictly rising path p' which ended at t_0 (instead of starting at t_0), then we could concatenate the two paths. The result would be a strictly rising path with no endpoints, contradicting the Loop Equations once again.

Thus all current paths start at t_0, do not contain any loops, and contain infinitely many vertices. $\square$

Let us now define the measure against which we shall integrate the minimal paths. Suppose we have a current function which satisfies the Infinite Kirchhoff Inequalities. Let P be the set of all current paths. Let $\mathcal{E}$ be the set of all current-carrying edges, and direct the edges such that $i(e) > 0$ for all edges e in $\mathcal{E}$. For all edges e in $\mathcal{E}$, let $Q(e)$ be the set of all paths in P which contain e. With this notation, we have the following theorem:

Theorem 2.2 *Suppose i is a current function which satisfies the Infinite Kirchhoff Inequalities. Then there exists a finite measure μ on P such that*

$$\mu Q(e) = i(e)$$

for all edges e in $\mathcal{E}$.

We shall call this measure the *current measure*. To prove Theorem 2.2, we shall define μ on a semialgebra $\mathcal{Q}$ of subsets of P and then extend it to a σ-algebra of measurable subsets of P. This σ-algebra will include all paths and edges. First, however, we need a definition. We shall define a *path trunk* to be a finite connected subset of a current path such that the path trunk begins at a bottom edge. Note that by Proposition 2.11, all current paths contain path trunks of arbitrarily many edges. We will use a caret to denote a path trunk ($\hat{p}$) as opposed to a path (p) which contains it. As with individual edges, let $Q(\hat{p})$ be the set of paths in P which contain $\hat{p}$. Note that

$$\hat{p}_1 \subseteq \hat{p}_2 \qquad \Longrightarrow \qquad Q(\hat{p}_1) \supseteq Q(\hat{p}_2).$$

For convenience, we shall define $Q(\emptyset)$ to be the empty set.

Furthermore, let us establish some notation that will be useful later. For any $e \in \mathcal{E}$, let H_e denote the set of edges adjacent to e which carry current *away from* e. Thus, current flows through e into elements of H_e. (H stands for "head," since these edges emanate from the "front end," so to speak, of e.) Likewise, T_e shall denote the adjacent edges which carry current *toward* e, so that current flows from elements of T_e into e. (T stands for "tail"; it could also stand for "toward.")

Now let us define a collection $\mathcal{Q}$ of subsets of P:

$$\mathcal{Q} = \{Q(\hat{p}) \,|\, \hat{p} \text{ is a path trunk}\}.$$

We will consider the empty set to be in $\mathcal{Q}$.

Proposition 2.12 *$\mathcal{Q}$ is a semialgebra.*

Proof: To show that $\mathcal{Q}$ is a semialgebra, we must prove two facts. First, $\mathcal{Q}$ must be closed under finite intersection; second, the complement of any set in $\mathcal{Q}$ must itself be a finite disjoint union of sets in $\mathcal{Q}$. As we shall see, the first condition follows from the definition of path trunks. Let $Q(\hat{p}_1)$ and $Q(\hat{p}_2)$ be nonempty sets in $\mathcal{Q}$. If $Q(\hat{p}_1) \cap Q(\hat{p}_2) = \emptyset$, then the intersection is in $\mathcal{Q}$ by default. If the intersection is non-empty, then either $\hat{p}_1 \subseteq \hat{p}_2$ or $\hat{p}_2 \subseteq \hat{p}_1$. Thus, either $Q(\hat{p}_2) \subseteq Q(\hat{p}_1)$ or $Q(\hat{p}_1) \subseteq Q(\hat{p}_2)$. Hence, the intersection $Q(\hat{p}_1) \cap Q(\hat{p}_2)$ equals either $Q(\hat{p}_2)$ or $Q(\hat{p}_1)$ and is thus in $\mathcal{Q}$.

To prove that the second condition holds, consider $Q(\hat{p})$ in $\mathcal{Q}$. Since $\hat{p}$ is a path trunk, $\hat{p} = \langle e_1, \ldots, e_n \rangle$. (Note that we use angle brackets to emphasize the ordered quality of these current paths.) We shall use induction on n (the length of the path trunk) to prove that $Q(\hat{p})^\sim$ is the finite disjoint union of sets in $\mathcal{Q}$. If $n = 1$, then $\hat{p} = \langle e_1 \rangle$. Hence,

$$Q(\hat{p})^\sim = \bigcup_{\substack{e \in B \\ e \neq e_1}} Q(\langle e \rangle),$$

where B is the set of bottom edges. Clearly this union is finite (because we have a finite bottom) and disjoint.

Suppose now that the second condition is true for all positive integers less than n. If $\hat{p} = \langle e_1, \ldots, e_n \rangle$, then let $\hat{p}_0 = \langle e_1, \ldots, e_{n-1} \rangle$. If $H_{e_{n-1}} = \{e_n, f_1, \ldots, f_k\}$, then let $\hat{p}_j = \langle e_1, \ldots, e_{n-1}, f_j \rangle$, for $j = 1, \ldots, k$. So $\{\hat{p}_j\}_{j=1}^k$ consists of all path trunks that match $\hat{p}$ except for the last edge. Now $Q(\hat{p}) = Q(\hat{p}_0) \cap Q(e_n)$. Furthermore, we know that the sets $Q(\hat{p}_j)$ are disjoint. Therefore,

$$\begin{aligned} Q(\hat{p})^\sim &= Q(\hat{p}_0)^\sim \cup Q(e_n)^\sim \\ &= Q(\hat{p}_0)^\sim \cup \{\text{paths including } \hat{p}_0 \text{ but not } e_n\} \\ &= Q(\hat{p}_0)^\sim \cup \left(\bigcup_{j=1}^{k} Q(\hat{p}_j) \right). \end{aligned}$$

By our inductive hypothesis, $Q(\hat{p}_0)^\sim$ is a finite disjoint union of sets in $\mathcal{Q}$, and thus $Q(\hat{p})^\sim$ is also. Hence, the second condition is true for all sets $Q(\hat{p})$ in $\mathcal{Q}$, and so $\mathcal{Q}$ is a semialgebra. $\square$

We shall now define a set function μ on $\mathcal{Q}$. The function μ will serve as a kind of "protomeasure," which we shall extend to an actual measure on a σ-algebra of subsets of P. Before we can do so, we need yet another definition:

Let $\langle e_1, e_2, \ldots, e_n\rangle$ be an ordered set of contiguous edges forming part of a current path. Then the *flowthrough* $f(\langle e_1, \ldots, e_n\rangle)$ is defined inductively as follows:

$$\begin{aligned} f(\langle e\rangle) &= i(e) \\ f(\langle e_1, e_2, \ldots, e_n\rangle) &= f(\langle e_1, e_2, \ldots, e_{n-1}\rangle)\frac{i(e_n)}{c}, \end{aligned}$$

where c is the total current flowing into the vertex at the terminal end of e_{n-1} (which is also the vertex at the initial end of e_n). If $\langle e_1, e_2, \ldots, e_n\rangle$ is a path trunk $\hat{p}$, then we write $f(\hat{p})$ for the flowthrough of $\hat{p}$.

We may interpret flowthrough as follows. Since we are assuming that the function i satisfies the Infinite Kirchhoff Inequalities, we know that the amount of current flowing into a vertex is equal to the amount flowing out of the vertex, by the Current Equations. If we suppose that the current flowing into a vertex evenly distributes itself through the edges leaving the vertex, then the flowthrough is the total amount of current that flows through *all* of the specified edges. (I.e., in the case of electric current, the electrons that go through all of the specified edges would constitute the flowthrough.) Thus, at each vertex, we multiply the amount of current which has made it through all of the previous edges by the ratio of the current through the next edge to the total current out of the vertex.

We define the "protomeasure" μ in terms of flowthrough. Suppose $Q \in \mathcal{Q}$. If $Q = \emptyset$, then we define $\mu Q = 0$. If Q is nonempty, then $Q = Q(\hat{p})$ for some path trunk $\hat{p}$, and we define $\mu Q = f(\hat{p})$.

Proposition 2.13 *The set function μ on $\mathcal{Q}$ is well-defined.*

Proof: Suppose $Q(\hat{p}_1) = Q(\hat{p}_2)$. Then either $\hat{p}_1 \subseteq \hat{p}_2$ or $\hat{p}_2 \subseteq \hat{p}_1$. Without loss of generality, suppose $\hat{p}_1 \subseteq \hat{p}_2$. Then $\hat{p}_1 = \langle e_1, \ldots, e_k\rangle$ and $\hat{p}_2 = \langle e_1, \ldots, e_k, e_{k+1}, \ldots, e_n\rangle$. The fact that $Q(\hat{p}_1) = Q(\hat{p}_2)$ implies that every path which travels through $\langle e_1, \ldots, e_k\rangle$ also travels through $\langle e_{k+1}, \ldots, e_n\rangle$. Hence, for $k \le m \le n-1$, the vertex terminating e_m has exactly one outgoing edge, namely, e_{m+1}. So if c is the total current flowing into the vertex terminating e_m, we have $i(e_m)/c = 1$. Thus $f(\langle e_1, \ldots, e_{m+1}\rangle) = f(\langle e_1, \ldots, e_m\rangle)$ for $m = k, \ldots, n-1$. Hence

$$\begin{aligned} f(\hat{p}_1) &= f(\langle e_1, \ldots, e_k\rangle) \\ &= f(\langle e_1, \ldots, e_n\rangle) \\ &= f(\hat{p}_2), \end{aligned}$$

and thus the function μ is well-defined on $Q(\hat{p}_1) = Q(\hat{p}_2)$. $\square$

Now that we have μ defined on $\mathcal{Q}$, we want to be able to extend it to the algebra generated by $\mathcal{Q}$ and, eventually, to a σ-algebra containing $\mathcal{Q}$. The next proposition will allow us to do so. We precede it with a lemma:

Lemma 2.1 *Let $\hat{p}$ be a path trunk ending at e. Suppose $H_e = \{e_1, \ldots, e_n\}$. For $k = 1, \ldots, n$, let $\hat{p}_k$ be the path trunk containing $\hat{p}$ and ending with e_k. Then $f(\hat{p}) = \sum_{k=1}^{n} f(\hat{p}_k)$.*

Proof: By definition,

$$f(\hat{p}_k) = \frac{f(\hat{p})i(e_k)}{c},$$

where $c = \sum_{k=1}^{n} i(e_k)$. Hence,

$$\begin{aligned} \sum_{k=1}^{n} f(\hat{p}_k) &= \sum_{k=1}^{n} \frac{f(\hat{p})i(e_k)}{c} \\ &= \frac{f(\hat{p})}{c} \sum_{k=1}^{n} i(e_k) \\ &= f(\hat{p}). \end{aligned}$$

□

Proposition 2.14 *Let $\{Q(\hat{p}_k)\}_{k=1}^{\infty}$ be disjoint sets in $\mathcal{Q}$. Suppose $Q = \bigcup_{k=1}^{\infty} Q(\hat{p}_k)$ is also in $\mathcal{Q}$. Then*

$$\mu Q = \sum_{k=1}^{\infty} \mu Q(\hat{p}_k).$$

Proof: First let us establish some notation. For any path trunk $\hat{p}$, let $\ell(\hat{p})$ denote the length of, or the number of edges in, $\hat{p}$. For any path trunk $\hat{p}$ and positive integer n, let $\hat{p}^n$ denote the path trunk consisting of the first n edges of $\hat{p}$, or $\hat{p}$ itself if $n \geq \ell(\hat{p})$.

Now since Q is in $\mathcal{Q}$, we know that $Q = Q(\hat{p})$ for some path trunk $\hat{p}$. Since $Q = \bigcup_{k=1}^{\infty} Q(\hat{p}_k)$, we know that $\hat{p}_k$ contains $\hat{p}$ for all k.

We claim that $\sup_k \ell(\hat{p}_k) \neq \infty$. The proof is a diagonalization argument. Suppose to the contrary that, for all n, there exists a k_n with $\ell(\hat{p}_{k_n}) > n$. Reordering and passing to a subsequence, we may say that $\ell(\hat{p}_k) > k$ for all positive integers k. Let $\tilde{e}$ be the terminal edge of $\hat{p}$. Since $H_{\tilde{e}}$ is finite, there exists a subsequence $\{\hat{p}_{1k}\}$ of $\{\hat{p}_k\}$ sharing the same first edge after $\hat{p}$. Call this first edge e_1. Since H_{e_1} is finite, there exists a subsequence $\{\hat{p}_{2k}\}$ of $\{\hat{p}_{1k}\}$ sharing the same first two edges after $\hat{p}$. Continuing inductively, we use local finiteness to obtain a subsequence $\{\hat{p}_{nk}\}$ of $\{\hat{p}_{(n-1)k}\}$ sharing the same first n edges after $\hat{p}$. Take the diagonal subsequence $\{\hat{p}_{kk}\}_{k=1}^{\infty}$ and note the following two facts about it:

1. For all $h \geq k$, $\hat{p}_{hh}$ has the same first k edges after $\hat{p}$ as does $\hat{p}_{kk}$.
2. Second, $\ell(\hat{p}_{kk}) > k$ because $\{\hat{p}_{kk}\}$ is a subsequence of $\{\hat{p}_k\}$.

Thus we may form a path p such that, for any positive integer k, p shares its first k edges with $\hat{p}_{kk}$. So if $p|_k$ denotes the path p truncated at the kth edge, we may write $p|_k = \hat{p}^k_{kk}$.

At this point we return to our original sequence $\{Q(\hat{p}_k)\}_{k=1}^{\infty}$ as opposed to the subsequence mentioned above. Now since p contains $\hat{p}$, we know that $p \in Q(\hat{p}) = Q$, and hence $p \in Q(\hat{p}_k)$ for some k. Let $n > \ell(\hat{p}_k)$. Then $\ell(\hat{p}_{nn}) > n > \ell(\hat{p}_k)$. By disjointness of the sets $Q(\hat{p}_k)$, we must conclude that $\hat{p}_{nn}$ does not contain this particular $\hat{p}_k$. In fact, since $n > \ell(\hat{p}_k)$, we can say that $\hat{p}^n_{nn}$ does not contain $\hat{p}_k$. And since $p|_n = \hat{p}^n_{nn}$, we know that $p|_n$ does not contain $\hat{p}_k$. Once again we observe that $n > \ell(\hat{p}_k)$, so p does not contain $\hat{p}_k$. Hence p cannot be in $Q(\hat{p}_k)$. But this contradicts our choice of k. Thus we conclude that $\sup_k \ell(\hat{p}_k) \neq \infty$, since it was our assumption to the contrary that allowed us to construct the path p. So there exists an integer N such that $\ell(\hat{p}_k) < N$ for all k.

By finiteness of the bottom and local finiteness of the graph, there are only finitely many path trunks with N or fewer edges. Since any two non-empty sets $Q(\hat{p}_i)$ and $Q(\hat{p}_j)$ are disjoint if $i \neq j$, we conclude that only finitely many of the sets $Q(\hat{p}_i)$ can be non-empty. Reordering, we may say that there exists a positive integer M such that $Q(\hat{p}_i)$ is non-empty if and only if $i \leq M$.

Once again we remind ourselves that, for all k, $\hat{p}_k$ contains $\hat{p}$ for all k. So for all $n > \ell(\hat{p})$, we can use an induction argument based on Lemma 2.1 to obtain

$$f(\hat{p}) = \sum_{\text{distinct } \hat{p}^n_k} f(\hat{p}^n_k)$$

Taking $n > N$, we have

$$f(\hat{p}) = \sum_{k=1}^{M} f(\hat{p}_k).$$

So for all $n > N$, we have

$$\mu Q = \sum_{k=1}^{M} \mu Q(\hat{p}_k) = \sum_{k=1}^{\infty} \mu Q(\hat{p}_k).$$

□

So the set function μ on the semialgebra $\mathcal{Q}$ is countably additive. By a theorem of real analysis (see [12]), μ extends uniquely to a measure (which we will still call μ) on the algebra $\mathcal{A}$ generated by finite unions of sets in $\mathcal{Q}$. The measure extends in such a way that if A is the disjoint union of sets $\{Q(\hat{p}_k)\}_{k=1}^{n}$ in $\mathcal{Q}$, then

$$\mu A = \sum_{k=1}^{n} \mu Q(\hat{p}_k).$$

The next lemma shows that there are only finitely many path trunks terminating at any current-carrying edge.

Lemma 2.2 *Let e be in $\mathcal{E}$. Then there exist only finitely many path trunks terminating at e.*

Proof: Suppose there exist infinitely many distinct path trunks $\{\hat{p}_k\}_{k=1}^{\infty}$ terminating at e. Once again we use a diagonalization argument, this time starting at e and working backwards. Since T_e is finite, there exists a subsequence $\{\hat{p}_{1k}\}$ of $\{\hat{p}_k\}$ sharing the same first edge e_1 immediately preceding e. So all path trunks $\hat{p}_{1k}$ share the same two final edges, namely, e_1 and e. In turn, since T_{e_1} is finite, there exists a subsequence $\{\hat{p}_{2k}\}$ of $\{\hat{p}_{1k}\}$ sharing the same last two edges before e. So each path trunk $\hat{p}_{2k}$ ends with the sequence $\langle e_2, e_1, e\rangle$. Continuing inductively, we use local finiteness to obtain a subsequence $\{\hat{p}_{nk}\}$ of $\{\hat{p}_{(n-1)k}\}$ sharing the same last n edges before e. So each path trunk $\hat{p}_{nk}$ ends with the sequence $\langle e_n, e_{n-1}, \ldots, e_1, e\rangle$.

Now take any current path through e. Let $\langle f_1, f_2, \ldots\rangle$ be all edges of this path after e. We may construct a path $p = \langle \ldots e_2, e_1, e, f_1, f_2, \ldots, \rangle$. For any n, the edges e_n and e_{n-1} are contained in $\hat{p}_{nk}$ for any k; thus, we know that the path is rising at the vertex shared by e_n and e_{n-1}. Hence, the path p is rising at each of its vertices, but it has no beginning and no end. However, this construction contradicts Proposition 2.11, which states that all current paths begin at t_0. Thus we cannot continue our diagonalization indefinitely, and hence we could not have had infinitely many distinct path trunks terminating at e. $\square$

Since, by Lemma 2.2, there are only finitely many path trunks terminating at e, we know that $\mathcal{A}$ contains the set $Q(e)$ for each edge $e \in \mathcal{E}$. We may now prove the following fact:

Proposition 2.15 *For any edge $e \in \mathcal{E}$, $\mu Q(e) = i(e)$.*

Proof: Lemma 2.2 states that there are only finitely many path trunks terminating in e. Let m_e be the maximum length of such path trunks. We shall argue by induction on m_e. If e is a bottom edge, then $Q(e) = Q(\langle e\rangle)$, $\langle e\rangle$ being a path trunk one edge long; thus $\mu Q(e) = f(\langle e\rangle) = i(e)$.

Now suppose the proposition is true for all edges e such that $m_e < n$. Suppose e_0 is an edge such that $m_{e_0} = n$, and suppose that $T_{e_0} = \{e_1, \ldots, e_k\}$. Let P_0 be the set of path trunks terminating in e_0; let P_j be the set of all path trunks terminating in e_0 but containing e_j. It is clear that $P_0 = \bigcup_{j=1}^{k} P_j$ and that this union is disjoint. For all $\hat{p} \in P_0$, let $\hat{p}'$ be $\hat{p}$ with the final edge (e_0) removed. Then we have

$$\mu Q(e_0) \;=\; \sum_{\hat{p}\in P_0} f(\hat{p})$$

$$
\begin{aligned}
&= \sum_{j=1}^{k} \sum_{\hat{p} \in P_j} f(\hat{p}) \\
&= \sum_{j=1}^{k} \sum_{\hat{p} \in P_j} f(\hat{p}') \frac{i(e_0)}{c}, \quad \text{where } c = \textstyle\sum_{j=1}^{k} i(e_j) \\
&= \frac{i(e_0)}{c} \sum_{j=1}^{k} \sum_{\hat{p} \in P_j} f(\hat{p}') \\
&= \frac{i(e_0)}{c} \sum_{j=1}^{k} \mu Q(e_j) \\
&= \frac{i(e_0)}{c} \sum_{j=1}^{k} i(e_j) \quad \text{by our inductive hypothesis} \\
&= i(e_0) \quad \text{by definition of } c.
\end{aligned}
$$

□

Note that if we denote the bottom edges by $\{e_1, \ldots, e_k\}$, then we have $P = \bigcup_{i=1}^{k} Q(e_i)$. Thus, P is the finite union of sets with finite measure. Hence μ is itself a finite measure.

A *σ-algebra* is a collection of sets that is closed under complements and countable unions. By De Morgan's Laws, it will therefore also be closed under countable intersection. Thus, since $\mathcal{A}$ contains $Q(e)$ for all edges e, and since a path is defined by a sequence of contiguous edges, it follows that any σ-algebra containing $\mathcal{A}$ also contains each path p. If $\mathcal{S}$ is the smallest σ-algebra containing $\mathcal{A}$, then $\mathcal{S}$ contains $\{p\}$ for all paths $p \in P$, $Q(e)$ for all edges $e \in \mathcal{E}$, and countable intersections or unions of the same.

By a theorem of Carathéodory (see [12]), μ extends uniquely to a finite measure on $\mathcal{S}$. Continuing to call this extended measure μ, we now have the current measure defined on a σ-algebra of subsets of P, with $\mu Q(e) = i(e)$ for all edges e in $\mathcal{E}$. This statement is Theorem 2.2. □

We may now prove a result analogous to Theorem 1.8.

Theorem 2.3 *Let G be the dual graph of a locally finite tiling of a half-open quadrilateral or annulus. Suppose that there exist currents $i(s,t)$, not all 0, defined on the edges of G, and suppose that this current function i satisfies the Kirchhoff inequalities. Then the weights defined by this current function give us a fat flow optimal weight function.*

Proof: By Proposition 2.11, the set P of current paths contains only paths which start at t_0, contain no loops, and contain infinitely many vertices. By Theorem 2.2, there exists a measure μ on P such that

$$w(t) = \int_P p(t)\, d\mu.$$

By the Loop Inequalities, we have at least one w-minimal current path p. Suppose there is a current path p' such that $\text{len}_w(p') > \text{len}_w(p)$. By Proposition 2.11, both p and p' have t_0 as an endpoint. Concatenate p and p', and apply the Loop Equations. This infinite sequence of vertices would be falling along p (coming in the reverse direction) and rising along p'. The sum of the weights along p would be less than the sum of the weights along p', and the total sums would have a positive limit instead of going to zero. The Loop Equations forbid this type of behavior, so we conclude that all current paths are w-minimal paths. Hence, each current path represents a w-minimal fat flow in the associated tiling. We conclude that w is the integral of w-minimal fat flows; it is therefore a fat flow optimal weight function by Theorem 2.1. □

2.6 Extensions of Two Results

We can easily extend the Bounded Overlap Theorem and the Layer Theorem from the finite case to the infinite. As noted in Section 1.5.2, neither the statement nor the proof of the Bounded Overlap Theorem depend on the finiteness of the tiling. Therefore, now that we have defined combinatorial moduli for half-open quadrilaterals and annuli, we may state this theorem for these figures as well.

Theorem 2.4 (Bounded Overlap Theorem) *Suppose that a half-open quadrilateral or annulus X has two locally finite tilings T and T', such that no element of T intersects more than K elements of T' and such that no element of T' intersects more than K elements of T. Let $M(X,T)$ be the fat flow modulus of X with the tiling T, and let $M(X,T')$ be the fat flow modulus of X with the tiling T'. Then*

$$M(Q,T) \leq K^3 \cdot M(Q,T').$$

It is also simple to extend the Layer Theorem to deal with an infinite number of layers:

Theorem 2.5 (Layer Theorem) *Suppose a half-open quadrilateral or annulus X with a locally finite tiling is divided into a family $\{X_j\}_{j=1}^{\infty}$ of tiled quadrilaterals or annuli, respectively, such that two distinct members of the family are disjoint except possibly at their boundaries. Then the fat flow modulus of X is greater than or equal to the sum of the fat flow moduli of the sets X_j.*

Proof: Assume that we have numbered the layers $\{X_n\}$ such that each X_n separates X_k from the bottom for all $k > n$. (I.e., we have numbered them from the bottom up.) For all n, let M_n be the fat flow modulus of X_n. Let $\bar{X}_n$ equal the union $\bigcup_{i=1}^{n} X_n$. Let w_n be an optimal weight function on $\bar{X}_n$. Let $\bar{M}_n$ be the fat flow modulus of $\bar{X}_n$. We may extend each function w_n to X by letting w_n be zero on all tiles not included in $\bar{X}_n$. Then if M is the fat flow modulus of M, we have

$$\bar{M_n} = \frac{H^2_{w_n,f}(\bar{X}_n)}{A_{w_n}(\bar{X}_n)} = \frac{H^2_{w_n,f}(X)}{A_{w_n}X} \leq M$$

for all n. Now by the finite Layer Theorem (applied to $\bar{X}_n$), we also know that

$$\bar{M_n} \geq \sum_{i=1}^{n} M_n$$

for all n.

So for all n, we have

$$M \geq \sum_{i=1}^{n} M_n.$$

Letting n go to infinity, we have

$$M \geq \sum_{i=1}^{\infty} M_n.$$

$\square$

Chapter 3

An Implication for Parabolicity of Tilings

In this chapter, we shall see that whether or not a tiling has finite fat flow modulus depends on the degrees of the vertices of the tiling. In essence, the modulus is finite if the degrees are large, causing the number of tiles to grow quickly. Otherwise it is infinite.

3.1 A Conjecture by He and Schramm

Following He and Schramm in [10], we shall define a *disk triangulation graph* to be the 1-skeleton of a triangulation of an open topological disk. If G is a disk triangulation graph, and v is a vertex of G, then consider the tiling T that is dual to G. If t_v is the interior of the tile corresponding to v, then let $T_v = T \setminus \{t_v\}$. T_v is a half-open annulus; its bottom consists of the edges that used to be adjacent to v. We shall say that G is *parabolic* if the fat flow modulus of T_v is infinite for some vertex v in G; otherwise, we say that G is *hyperbolic*. As noted in Section 1.6.2, He and Schramm use the term "vertex extremal length" to refer to combinatorial moduli; their formulation is in terms of the dual tiling.

This difference in formulation, however, need not affect the way we look at the graph, as long as we are concerned only with determining parabolicity or hyperbolicity. Suppose G is a locally finite disk triangulation graph with bounded valence, and suppose we remove a vertex from G (together with all edges incident to that vertex). Then the resulting graph $\bar{G}$ is the 1-skeleton of a locally finite triangulation T of a half-open annulus. Let M_v be the modulus of $\bar{G}$ in terms of the He/Schramm formulation (with weights concentrated at the vertices and fat paths given by edges); and let M_t be the fat flow modulus of the triangulation T (viewed as a tiling, with weights assigned to the faces). Let D be the dual tiling of the graph $\bar{G}$. Now M_v is the fat flow modulus of D. Since G has bounded valence, the tiling T spanned by $\bar{G}$ has bounded overlap with the tiling D. By the Bounded Overlap Theorem, then, either M_v

and M_t are both infinite, or they are both finite. Hence, whether or not we assign weights to the vertices or to the tiles does not affect the type (parabolic or hyperbolic) of the graph.

One can also define parabolicity and hyperbolicity in terms of circle packings, or in terms of "edge extremal length" (in which weights are assigned to the edges instead of to the vertices). He and Schramm show in [10] that all three formulations are equivalent.

In the same paper, He and Schramm proceed to discuss the relationship between the vertex valences of a disk triangulation graph and the type (parabolic or hyperbolic) of the graph. They obtain the following results:

Theorem 3.1 *Let G be the 1-skeleton of a disk triangulation, and suppose that at most finitely many vertices in G have valence greater than 6. Then G is parabolic.*

Theorem 3.2 *Let G be the 1-skeleton of a locally finite disk triangulation. Let* $\mathrm{val}(v)$ *denote the valence of the vertex v. Suppose that*

$$\sup_{W_0} \inf_{W \supseteq W_0} \left(\frac{1}{|W|} \sum_{v \in W} \mathrm{val}(v) \right) > 6,$$

where W and W_0 are nonempty finite connected sets of vertices. Then G is hyperbolic.

They note the wide gap between these two theorems and speculate on the possibility of obtaining type criteria in terms of the sequence $\{k_n\}$, where

$$k_n = \sum_v (\mathrm{val}(v) - 6),$$

the sum being taken over all vertices within n edges of a base vertex v_0. In particular, they speculate that boundedness of $\{k_n\}$ implies parabolicity of the corresponding disk triangulation graph. It is this conjecture which we shall prove.

3.2 Discussion of the Conjecture

To get a feel for what might be going on, let us consider some possible behaviors of the sequence $\{k_n\}$. It turns out that this sequence can be quite ill-behaved. For instance, the sequence can be unbounded and yet have a subsequence that is constant. For an example of such a triangulation, see Figure 3.1. Identify the left and right sides of this triangulation in order to get a triangulation of the open unit disk. The vertex at the bottom will be our base vertex. The sequence of numbers to the side is the sequence $\{k_n\}$. Note that $\limsup k_n = \infty$, whereas $\liminf k_n = -6$.

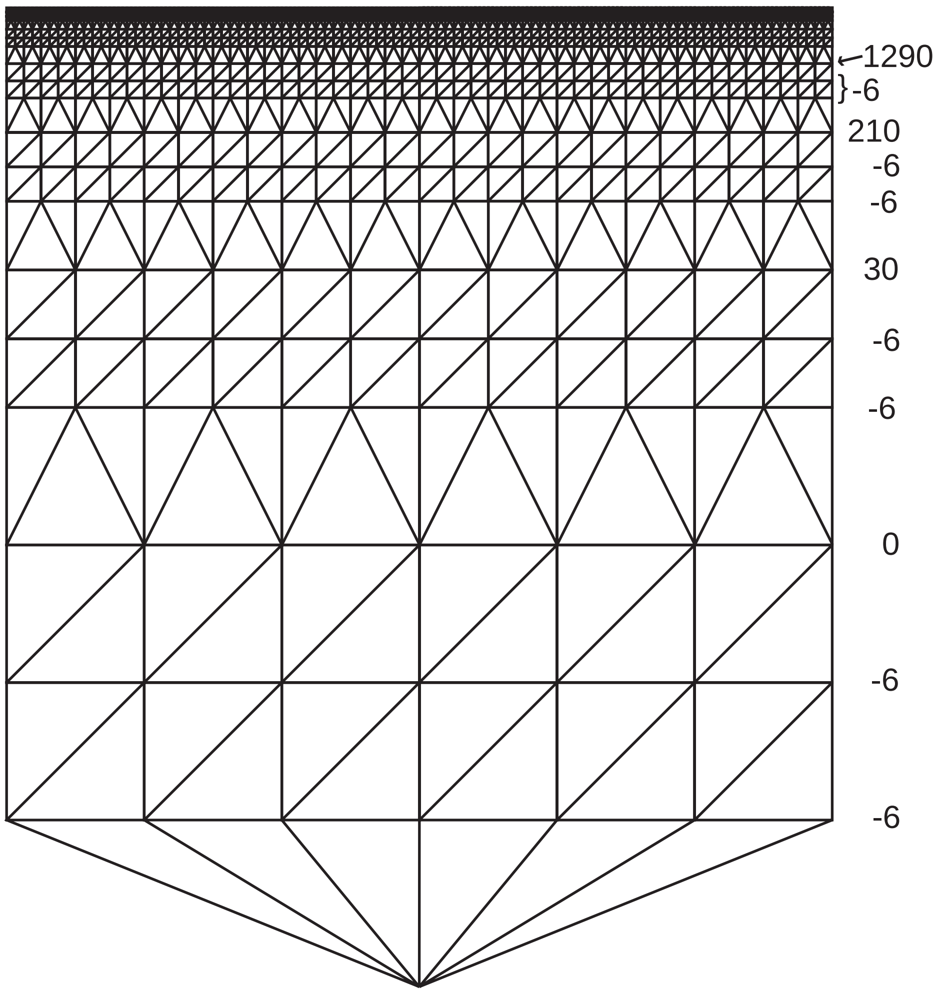

Figure 3.1: A Hyperbolic Disk Triangulation Graph

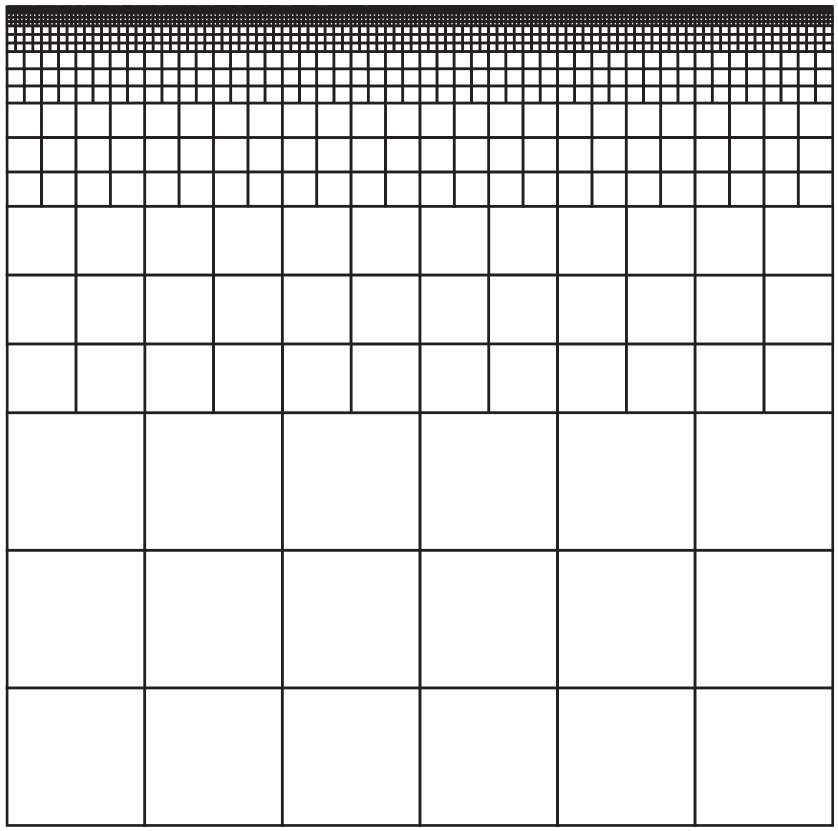

Figure 3.2: A Squared Rectangle R Corresponding to Figure 3.1

So, although $\{k_n\}$ is not bounded, it has a subsequence that is bounded. However, we claim that the graph has finite modulus and is therefore hyperbolic. To see that the fat flow modulus is finite, consider the squared rectangle R shown in Figure 3.2. Note that R has bounded overlap with the triangulation shown in Figure 3.1. By Proposition 2.10, the weight function determined by the sizes of the squares in the figure is the fat flow optimal weight function on R. If the largest square in R has unit side length, then the height of R is 6 and the area is 36. The fat flow modulus is 1, and therefore the original triangulation must be hyperbolic. Thus, even if $\{k_n\}$ has a bounded subsequence we are not guaranteed parabolicity.

However, consider R more closely. Notice that it consists of concentric layers (which become annuli after we identify the sides); the sum of the heights of these annuli gives the height of the entire squared rectangle. Furthermore, notice that the values of $\{k_n\}$ increase only when the number of tiles in a layer increases. When the number of tiles in a layer remains the same, $\{k_n\}$ stays at -6.

Suppose that we modify R. Instead of doubling the number of tiles every third layer, let the number of layers between each doubling be described by a sequence r_n. In particular, we shall say that the number of layers with $6 \cdot 2^{n-1}$ tiles is r_n. Now, if we have a layer in our modified rectangle with $6 \cdot 2^{n-1}$ tiles, the height of that layer is 2^{-n+1}. So the height of the total modified rectangle is

$$\sum_{n=1}^{\infty} r_n \, 2^{-n+1}.$$

If r_n increases as a polynomial, for instance, the height of the rectangle (and hence its modulus) will be finite. Recall that the only times $\{k_n\}$ was greater than -6 was when we doubled a layer; if r_n grows like a polynomial, then, we will have arbitrarily long strings of values of -6 in $\{k_n\}$; nevertheless, the tiling will still be hyperbolic.

Looking at these examples gives us an idea of how to approach the conjecture. If we can decompose the disk triangulation graph (and the tiling spanned by it) into concentric annuli, we can estimate the fat flow modulus of each of these annuli. The Layer Theorem will allow us to sum the moduli of these annuli, obtaining an estimate for the modulus of the entire tiling. How are we to estimate the moduli of the annuli? In the above example, the number of tiles in the annulus controlled the modulus (height); at the same time, the number of vertices on the borders of the annuli controlled what happened to $\{k_n\}$. Thus, if we assume that the sequence $\{k_n\}$ is bounded, we can estimate the number of vertices on the borders of these annuli. These estimates, in turn, will give us information about the disposition of the tiles in each annulus. The number of tiles in an annulus will help us estimate its modulus. Finally, we can add up the moduli of the annuli and see if the modulus of the entire tiling is infinite. When we do so, we obtain the following theorem.

Theorem 3.3 *Let G be a disk triangulation graph of bounded valence. Let v_0 be a vertex of G. For all positive integers n, let V_n be the set of all vertices v such that $d(v_0, v) = n$, where $d(v_0, v)$ is the minimum number of edges in a path connecting v_0 to v. Let*

$$a_n = \sum_{\substack{v \in V_k \\ k \leq n}} (\mathrm{val}(v) - 6).$$

If the sequence $\{a_n\}$ is bounded, then G is parabolic.

As explained before, the conjecture as stated by He and Schramm assigns weights to the vertices of the graph when determining the modulus; however, by the Bounded Overlap Theorem, we will obtain the same results by assigning weights to the faces of the complex.

We shall prove this theorem in the following sections, employing the approach outlined above. First we shall consider the combinatorics of certain triangulated annuli. We shall then use our knowledge of the structure of these annuli to estimate their moduli. Finally, we shall show that one can decompose the disk triangulation graph into concentric annuli and use the Layer Theorem to show that the modulus of the entire graph is infinite.

Figure 3.3: Up, Down, and Mid Triangles

3.3 Triangulated Annuli

Let A be a closed topological annulus with two (disjoint) boundary components homeomorphic to circles. (This definition allows us to exclude degenerate annuli.) Impose a triangulation on A—that is, construct a homeomorphism from A to a simplicial 2-complex. Then we say that A, with this triangulation, is a *simply triangulated annulus* if all vertices of the triangulation lie on the boundary of A and if any edges joining a top vertex to another top vertex are themselves part of the top.

Let us consider the types of triangles that a simply triangulated annulus can contain. First, we know that at least one vertex of each triangle must lie on the bottom of A by definition. Otherwise, all three vertices would lie on the top, and all three edges would have to be part of the top boundary component; this is impossible since the triangulation is a simplicial 2-complex. So if we suppose that exactly one vertex lies on the bottom, then we know that two vertices lie on the top. The edges connecting those two top vertices must be part of the top boundary component. So this type of triangle, which we shall call a *down triangle*, has one vertex on the bottom and two on the top; it also has one edge lying on the top. (See Figure 3.3.)

Next, suppose a triangle has exactly two vertices on the bottom and, consequentially, one vertex on the top. We have two possible cases: either the edge connecting the two bottom vertices lies in the bottom boundary component, or it does not. We call the the first sort of triangle an *up triangle*; it has two vertices on the bottom and one on the top, and it has one edge lying on the bottom. We shall call the other case a *mid triangle*; it has two vertices on the bottom and one on the top, but none of its edges lie on the boundary.

Finally, suppose all three vertices lie on the bottom. Then either two, one, or no edges lie on the bottom boundary component. We call such triangles *bent triangles*. (See Figure 3.4.)

These are the only types of triangles that can occur in a simply triangulated annulus. Note that this analysis has shown us that all vertices in a simply triangulated annulus are either on the bottom or are connected to the bottom by an edge.

We define a *well-triangulated annulus* to be a closed topological annulus with two (disjoint) boundary components and an imposed triangulation such that all vertices that do not lie

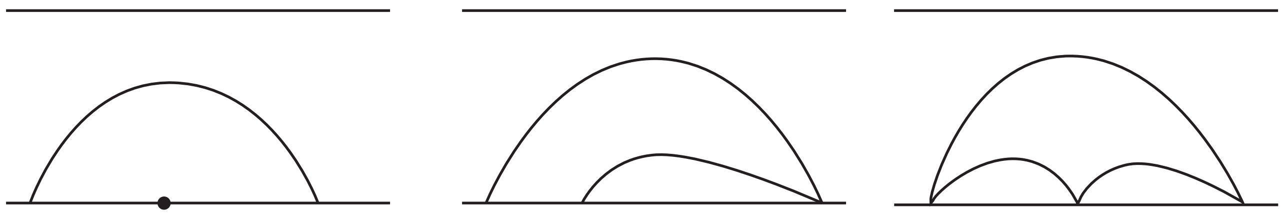

Figure 3.4: Bent Triangles

on the bottom are connected to the bottom by an edge; and such that any edges joining a top vertex to another top vertex are themselves part of the top. Note that the only difference between a simply triangulated annulus and a well-triangulated annulus is that a well-triangulated annulus may have vertices in its interior. All simply triangulated annuli are well-triangulated annuli.

Now, for each vertex v of a well-triangulated annulus, we define the *tile valence* of v (denoted $\mathrm{tval}(v)$) to be the number of tiles (triangles) in the annulus incident to v. We shall begin by considering simply triangulated annuli such that all non-boundary edges connect the bottom and the top of the annulus. In this case, A will contain only up and down triangles, and no mid or bent triangles. We have the following proposition:

Proposition 3.1 *Let A be a simply triangulated annulus such that all non-boundary edges of the triangulation connect the bottom and the top of the annulus. Suppose the bottom of the annulus contains k vertices, labeled $b_1, b_2, \ldots, b_k$, and suppose the top contains n vertices, labeled $t_1, t_2, \ldots, t_n$. Then*

$$\sum_{i=1}^{k}(\mathrm{tval}(b_i) - 3) = n - k,$$

and

$$\sum_{i=1}^{n}(\mathrm{tval}(t_i) - 3) = k - n.$$

Proof: Because we assume that all non-boundary edges of the triangulation connect the bottom and top of the annulus, we may choose an orientation for the annulus and order the non-boundary edges clockwise. Re-indexing if necessary, we assume that an edge connects b_1 and t_1, that the other vertices $b_2, \ldots, b_k$ and $t_2, \ldots, t_n$ are ordered consecutively clockwise around the annulus, and that no edge connects b_1 to t_n. (I.e., the edge from b_1 to t_1 is the "first" edge clockwise from b_1.) Now for $i = 1, \ldots, k$, we let m_i be the number of non-boundary edges incident to b_i. Note that $\mathrm{tval}(b_i) = m_i + 1$. (See Figure 3.5.)

Now for $i = 1$, the last edge from b_1 hits t_{m_1}. Proceeding by induction, suppose that for $i = s$, the last edge from b_s hits $t_{m_1+\cdots+m_s-(s-1)}$. Then the first edge from b_{s+1} must also hit $t_{m_1+\cdots+m_s-(s-1)}$, and so the last edge from b_{s+1} must hit $t_{m_1+\cdots+m_s-(s-1)+m_{s+1}-1} =$

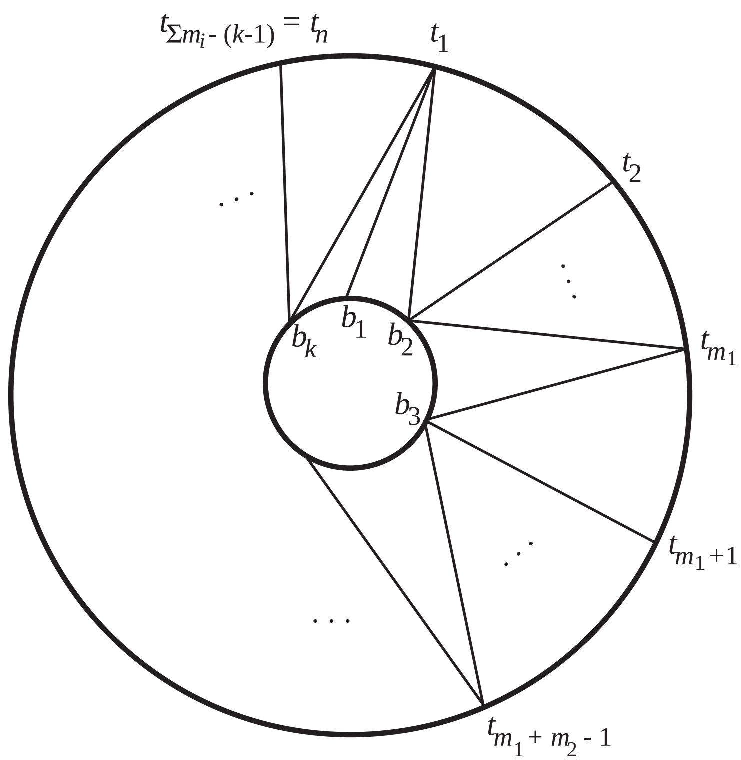

Figure 3.5: Arrangement of Edges and Vertices on A

$t_{m_1+\cdots+m_{s+1}-s}$. Hence we may conclude that the last edge of b_i hits $t_{m_1+\cdots+m_i-(i-1)}$ for $i = 1,\ldots,k$. However, the last edge from b_k must hit $t_1 = t_{n+1}$. So

$$\begin{aligned}\sum_{i=1}^{k} m_i - (k-1) &= n+1 \\ \sum_{i=1}^{k} m_i - k &= n.\end{aligned}$$

Since $\mathrm{tval}(b_i) = m_i + 1$,

$$\begin{aligned}\sum_{i=1}^{k}(\mathrm{tval}(b_i) - 3) &= \sum_{i=1}^{k}(m_i + 1 - 3) \\ &= \sum_{i=1}^{k} m_i - 2k \\ &= n - k.\end{aligned}$$

By symmetry (reversing what we call the top and bottom of the annulus), we obtain

$$\sum_{i=1}^{n}(\mathrm{tval}(t_i) - 3) = k - n.$$

□

Next, we extend this result to deal with any simply triangulated annulus. In particular, we allow some edges to connect bottom vertices to other bottom vertices.

Proposition 3.2 *Let A be a simply triangulated annulus. Suppose the bottom of the annulus contains k vertices, labeled $b_1, b_2, \ldots, b_k$, and suppose the top contains n vertices, labeled $t_1, t_2, \ldots, t_n$. Let q be the number of bottom vertices connected to the top by an edge. Then*

$$\sum_{i=1}^{k}(\mathrm{tval}(b_i) - 3) = n - q,$$

and

$$\sum_{i=1}^{n}(\mathrm{tval}(t_i) - 3) = q - n.$$

Proof: For each $i = 1,\ldots,k$, define r_i to be the number of non-boundary edges connecting b_i to the bottom, and define c_i to be the number of edges connecting b_i to the top.

Now let us consider the number of triangles incident to each bottom vertex, in terms of r_i and c_i. If $c_i = 0$, then there are $r_i + 1$ triangles incident to the vertex b_i, each of which has all

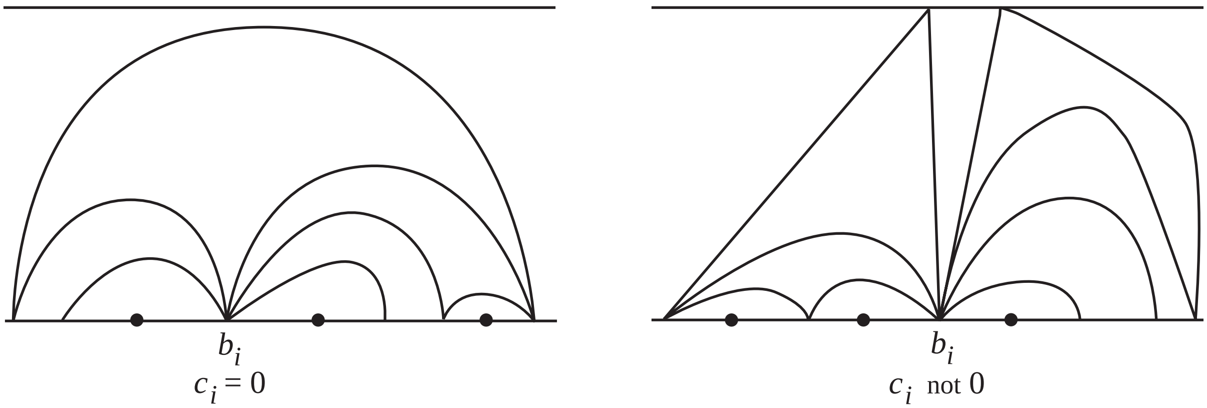

Figure 3.6: Arrangement of Triangles around a Vertex

three vertices lying on the bottom of the annulus. On the other hand, if $c_i \neq 0$, then there are r_i triangles incident to b_i which have all three vertices lying on the bottom; 2 triangles incident to b_i which have exactly two vertices on the bottom; and $c_i - 1$ triangles incident to b_i with exactly one vertex on the bottom. (See Figure 3.6.) We will need to know the location of the triangle vertices for our calculations; however, in either case the total number of triangles incident to v_i is $r_i + c_i + 1$.

So if T is the number of triangles in the triangulation, we have

$$\begin{aligned} T &= \sum_{\substack{i=1,\ldots,k \\ c_i=0}} \frac{1}{3}(r_i+1) + \sum_{\substack{i=1,\ldots,k \\ c_i\neq 0}} \left(\frac{1}{3}r_i + \frac{1}{2}\cdot 2 + (c_i - 1)\right) \\ &= \sum_{i=1}^{k} \frac{1}{3}r_i + \frac{1}{3}(k-q) + \sum_{i=1}^{k} c_i. \end{aligned}$$

Now envision the annulus girdling a sphere, and compute the Euler characteristic of the sphere. The number of vertices would be $k+n$; the number of faces would be $T+2$; and the number of edges would be

$$(k+n) + \frac{1}{2}\sum_{i=1}^{k} r_i + \sum_{i=1}^{k} c_i.$$

Since the Euler characteristic of a sphere is two, we know

$$(k+n) - \left((k+n) + \frac{1}{2}\sum_{i=1}^{k} r_i + \sum_{i=1}^{k} c_i\right) + (T+2) = 2,$$

so

$$T = \frac{1}{2}\sum_{i=1}^{k} r_i + \sum_{i=1}^{k} c_i.$$

Combining this result with our previous calculations for T, we find that

$$\begin{aligned}
\sum_{i=1}^{k} \frac{1}{3} r_i + \frac{1}{3}(k-q) + \sum_{i=1}^{k} c_i &= \frac{1}{2} \sum_{i=1}^{k} r_i + \sum_{i=1}^{k} c_i \\
2 \sum_{i=1}^{k} r_i + 2(k-q) &= 3 \sum_{i=1}^{k} r_i \\
2(k-q) &= \sum_{i=1}^{k} r_i.
\end{aligned}$$

Now let us consider the sum of the values $\mathrm{tval}(b_i) - 3$. From the previous analysis on the number of triangles incident to each vertex, we find that

$$\begin{aligned}
\sum_{i=1}^{k} (\mathrm{tval}(b_i) - 3) &= \sum_{i=1}^{k} ((r_i + c_i + 1) - 3) \\
&= \sum_{i=1}^{k} (r_i + c_i) - 2k.
\end{aligned}$$

However, since $\sum r_i = 2(k-q)$, we have

$$\begin{aligned}
\sum_{i=1}^{k} (\mathrm{tval}(b_i) - 3) &= 2(k-q) + \sum_{i=1}^{k} c_i - 2k \\
&= \sum_{i=1}^{k} c_i - 2q.
\end{aligned}$$

Now, for $i = 1, \ldots, n$, we let a_i be the number of non-boundary edges incident to the top vertex t_i. We see that

$$\begin{aligned}
\sum_{i=1}^{n} (\mathrm{tval}(t_i) - 3) &= \sum_{i=1}^{n} ((a_i + 1) - 3) \\
&= \sum_{i=1}^{n} a_i - 2n.
\end{aligned}$$

However, both $\sum a_i$ and $\sum c_i$ are the number of edges connecting the bottom to the top, and thus the two sums must be equal. Hence we may write

$$\sum_{i=1}^{n} (\mathrm{tval}(t_i) - 3) = \sum_{i=1}^{k} c_i - 2n.$$

Now consider the subannulus $\tilde{A}$ formed by all triangles of A which have a vertex on the top. (See Figure 3.7.) Note that $\tilde{A}$ will not contain any bent triangles, so all non-boundary edges of $\tilde{A}$ must connect the top and the bottom.

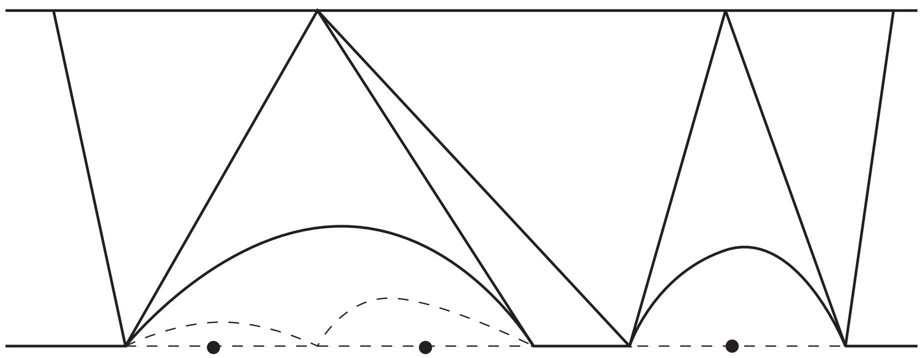

Figure 3.7: Removing Bent Triangles

The bottom of $\tilde{A}$ will consist of those bottom vertices of A such that $c_i \neq 0$, and there will be q such vertices. The top of $\tilde{A}$ is the same as the top of A, and it contains n vertices. Thus, by Proposition 3.1, $\sum_{i=1}^{n}(\text{tval}(t_i) - 3) = q - n$. This result is half of what we wanted to prove.

Since the top of A is the same as the top of $\tilde{A}$, and since we already calculated that $\sum_{i=1}^{n}(\text{tval}(t_i) - 3) = \sum_{i=1}^{k} c_i - 2n$, we may write

$$q - n = \sum_{i=1}^{k} c_i - 2n, \quad \text{and hence} \quad \sum_{i=1}^{k} c_i = q + n.$$

So

$$\begin{aligned} \sum_{i=1}^{k}(\text{tval}(b_i) - 3) &= \sum_{i=1}^{k} c_i - 2q \quad \text{from before,} \\ &= q + n - 2q \\ &= n - q, \end{aligned}$$

which is the other half of what we wanted to prove. $\square$

To conclude this section, let us consider what happens if we look at well-triangulated annuli, thus allowing for interior vertices.

Proposition 3.3 *Let A be a well-triangulated annulus. Suppose this triangulation contains n top vertices, c bottom vertices, and x interior vertices. Label the bottom vertices $v_1, \ldots, v_c$. Then*

$$\sum_{i=1}^{c}(\text{tval}(v_i) - 3) \geq n + x - c.$$

Proof: If A is a simply triangulated annulus, then there are no interior vertices, and the result follows from Proposition 3.2. So suppose A has interior vertices. Let w be an interior

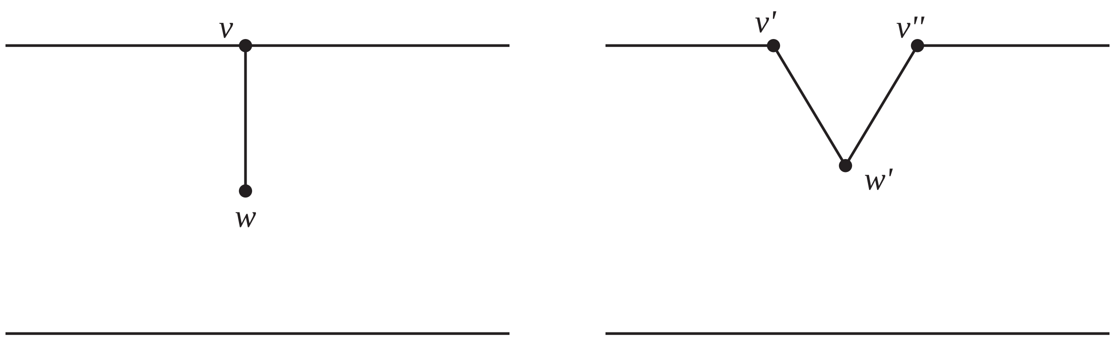

Figure 3.8: Removing an Interior Vertex

vertex of A that is connected to a top vertex v by an edge. Replace v by two top vertices v' and v'', and replace w by a top vertex w' connected to both v' and v''. (See Figure 3.8.)

We have transformed an interior vertex into two top vertices without affecting the bottom. In particular, we have not changed the sum $\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3)$. Since there are only finitely many vertices in A, we may apply this procedure repeatedly to obtain a triangulated annulus A' such that it is impossible to connect an interior vertex to the top (via a path of edges) without that path containing a bottom vertex. A' will have the same number of bottom vertices as A, and the sum $\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3)$ will be the same for A' as for A. If we eliminated x' interior vertices to obtain A' from A, then $n + 2x'$ will be the number of top vertices in A'.

If there are no remaining interior vertices, then A' is a simply triangulated annulus. By Proposition 3.2, we have

$$\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3) = n + 2x' - q > n + x' - c = n + x - c,$$

and we are done.

So suppose there are y remaining interior vertices (so that $x = x' + y$). Consider the graph consisting of these y vertices and the edges that connect them. Let G_k be a connected component of this graph. We claim that there exist two edges e and f, each of which connect the bottom to the top, such that e and f have the same endpoint on the top and different endpoints on the bottom; and we claim that G_k lies in the triangle bounded by e, f, and the included piece of the bottom.

First, observe that, since we cannot have loops (1-gons) in a triangulation, there must be at least two vertices on the bottom which are connected to the top. Therefore there are at least two edges from the bottom to the top which have distinct bottom vertices. Let e and f be the closest of these edges to the component G_k under consideration. (In other words, there are no edges that separate e from G_k in $A' \setminus f$, and vice versa.) It remains to show that e and f have the same top endpoint. Suppose they do not; call their top endpoints v_a and v_b, respectively, and their bottom endpoints w_a and w_b, respectively. (See Figure 3.9.)

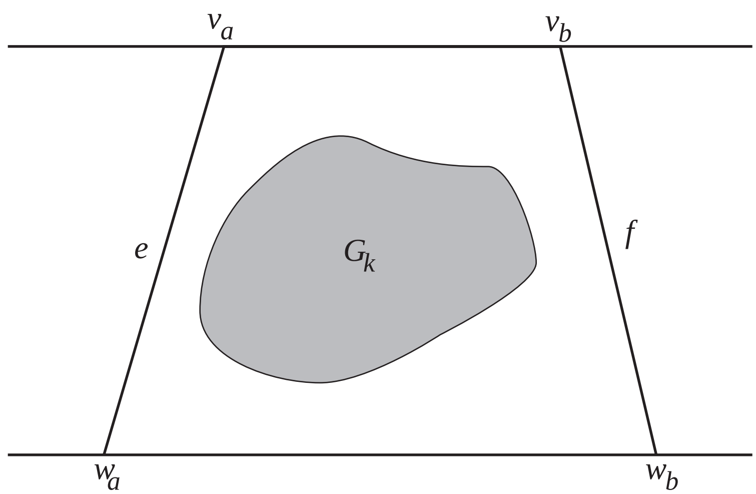

Figure 3.9: Situation of G_k

Then the edge from v_a to v_b must be part of a triangle, and the third vertex of the triangle cannot be an interior vertex. (Otherwise we would have an interior vertex connected to the top.) Therefore it must lie on the bottom. This vertex must either be w_a or w_b, or lie between them. G_k is a connected graph, so the edges from v_a and v_b to this third vertex must go around G_k. However, in each case we have violated our assumption that e and f are the closest edges to G_k that go from the bottom to the top. Therefore v_a and v_b must be the same point.

Now, let j be the number of these "triangular" subtilings which contain interior vertices. As we have just shown, these subtilings are bordered by contiguous bottom edges and by two edges which go from the bottom to a common vertex on the top, such that no edges interior to the subtiling connect the bottom to the top. We shall consider one of these subtilings, and calculate how much the interior vertices contribute to $\sum(\text{tval}(v_i) - 3)$. By Proposition 3.2, edges that go from a bottom vertex to a bottom vertex have no net effect on this sum. Thus, if there are no interior vertices, we can consider (for the purpose of counting valences) this subtiling to be a true triangle, contributing one to the tile valence sum at each of the vertices v_a and v_b. So, without interior vertices, this subtiling would contribute 2 to $\sum(\text{tval}(v_i) - 3)$. (See Figure 3.10.)

Suppose this subtiling has x_k interior vertices and c_k bottom vertices between v_a and v_b. Again, none of these vertices are connected to the top. Now by definition, the j subtilings have disjoint interiors; however, they might not cover the entire bottom of the annulus. Therefore, if q is the number of bottom vertices connected by an edge to the top, we have $\sum_{i=1}^{j} c_k \leq c - q$. We also know that $\sum_{i=1}^{j} x_k = y$. Let us imagine the kth such subtiling on the surface of a sphere and perform two Euler characteristic calculations. First, there are $3 + x_k + c_k$ vertices. Second, we must consider the edges. For each of the x_k interior vertices, let $\text{val}(w_i)$ be its valence in the subtiling, for $i = 1, \ldots, x_k$. For each of the c_k vertices on

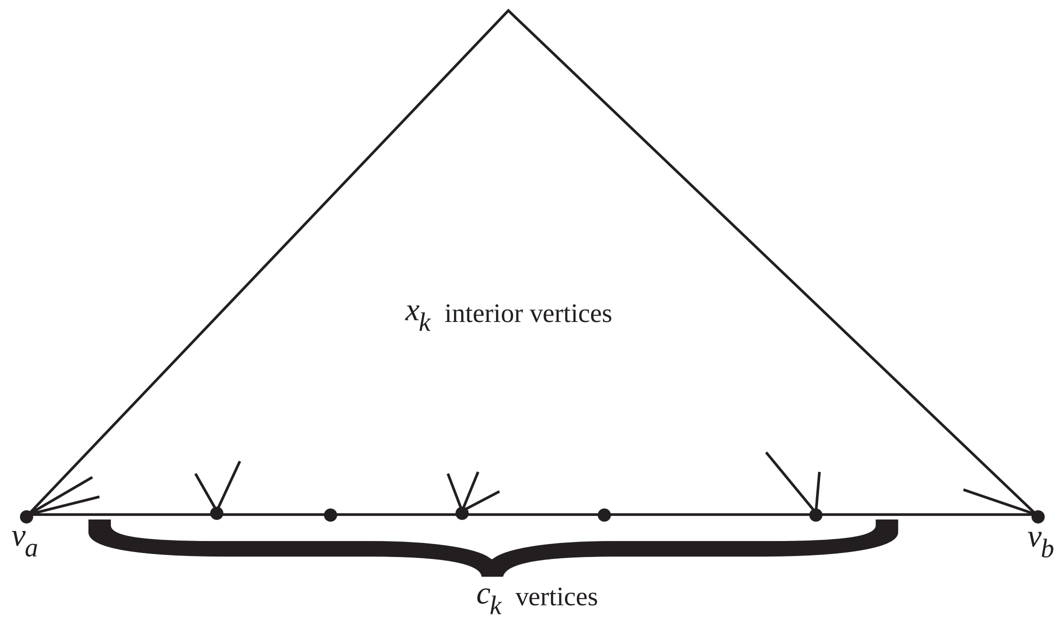

Figure 3.10: Tile Valences along the Bottom of a Subtiling

the bottom, let $\mathrm{val}(v_i)$ be its valence in the subtiling, not counting the edges that are part of the bottom, for $i = 1, \ldots, c_k$. Finally, we adopt the same convention for v_a and v_b, using $\mathrm{val}(\cdot)$ to denote their valences, not counting the edges that are part of the bottom. So, to count the non-boundary edges of the subtiling, we take

$$\frac{1}{2}\left(\sum_{i=1}^{x_k}\mathrm{val}(w_i) + \sum_{i=1}^{c_k}\mathrm{val}(v_i) + (\mathrm{val}(v_a) - 1) + (\mathrm{val}(v_b) - 1)\right).$$

We then add 2 for the edges going from bottom to top, and then we add $c_k + 1$ for the edges making up the bottom. So the total number of edges in the subtiling is

$$\frac{1}{2}\sum_{i=1}^{x_k}\mathrm{val}(w_i) + \frac{1}{2}\sum_{i=1}^{c_k}\mathrm{val}(v_i) + \frac{1}{2}\mathrm{val}(v_a) + \frac{1}{2}\mathrm{val}(v_b) + c_k + 2.$$

To compute the number of faces, we realize that each interior edge contributes $1/3$ to two faces, we realize that each exterior edge contributes $1/3$ to one face, and then we add one to account for the complement of the subtiling on the sphere. So the total number of faces is

$$\frac{2}{3}\cdot\frac{1}{2}\left(\sum_{i=1}^{x_k}\mathrm{val}(w_i) + \sum_{i=1}^{c_k}\mathrm{val}(v_i) + (\mathrm{val}(v_a) - 1) + (\mathrm{val}(v_b) - 1)\right) + \frac{1}{3}(2 + c_k + 1) + 1.$$

Simplifying, the number of faces is

$$\frac{4}{3} + \frac{1}{3}\sum_{i=1}^{x_k}\mathrm{val}(w_i) + \frac{1}{3}\sum_{i=1}^{c_k}\mathrm{val}(v_i) + \frac{1}{3}c_k + \frac{1}{3}\mathrm{val}(v_a) + \frac{1}{3}\mathrm{val}(v_b).$$

The number of vertices minus the number of edges plus the number of faces equals 2, so

$$\begin{aligned}
&(3 + x_k + c_k) \\
&\quad - \left(\frac{1}{2} \sum_{i=1}^{x_k} \mathrm{val}(w_i) + \frac{1}{2} \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \frac{1}{2} \mathrm{val}(v_a) + \frac{1}{2} \mathrm{val}(v_b) + c_k + 2 \right) \\
&\quad + \left(\frac{4}{3} + \frac{1}{3} \sum_{i=1}^{x_k} \mathrm{val}(w_i) + \frac{1}{3} \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \frac{1}{3} c_k + \frac{1}{3} \mathrm{val}(v_a) + \frac{1}{3} \mathrm{val}(v_b) \right) = 2
\end{aligned}$$

$$\begin{aligned}
\frac{7}{3} + x_k - \frac{1}{6} \sum_{i=1}^{x_k} \mathrm{val}(w_i) - \frac{1}{6} \sum_{i=1}^{c_k} \mathrm{val}(v_i) - \frac{1}{6} \mathrm{val}(v_a) - \frac{1}{6} \mathrm{val}(v_b) + \frac{1}{3} c_k &= 2 \\
14 + 6x_k - \sum_{i=1}^{x_k} \mathrm{val}(w_i) - \sum_{i=1}^{c_k} \mathrm{val}(v_i) - \mathrm{val}(v_a) - \mathrm{val}(v_b) + 2c_k &= 12
\end{aligned}$$

$$\sum_{i=1}^{c_k} \mathrm{val}(v_i) - 2c_k + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 2 = 6x_k - \sum_{i=1}^{x_k} \mathrm{val}(w_i).$$

Now we do another calculation. This time we consider only the interior vertices of this particular subtiling and the edges connecting them. The resulting graph may not be connected. Suppose it has s components, denoted $C_1, \ldots, C_s$, and suppose that component C_p contains x_{k_p} vertices, denoted $w_{i_1}, \ldots, w_{i_{(x_{k_p})}}$. Embed this component C_i in a 2-sphere. The number of vertices is x_{k_p}. To get the number of edges, we will have to subtract off the edges that attach an interior vertex to the bottom. We shall let r_p be the number of these edges. The maximum number of edges attaching an interior vertex to the bottom is the number of edges incident to the bottom. Thus, the number of edges in this component C_i is

$$\frac{1}{2} \left(\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - r_p \right),$$

where

$$\sum_{i=1}^{s} r_p \leq \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 2.$$

Now, each of these vertices has an edge connecting it to the bottom (by definition of a well-triangulated annulus), and therefore all of the vertices w_{i_l} are on the boundary of the complex spanned by this graph. Assume for the moment that this complex contains at least one 2-cell. Then the number of edges in the interior of this complex cannot be any greater than the total number of edges minus the number of vertices. Denote the number of interior edges by N, and denote the total number of edges by E.

Any of the N edges in the interior of this graph contributes $1/3$ to two faces. Any of the $E - N$ edges on the boundary of this graph contributes $1/3$ to one face. There will also be one face for the complement of the graph on the 2-sphere. Hence, the number of faces is equal to

$$\begin{aligned}
\frac{2}{3}N + \frac{1}{3}(E-N) + 1 &= \frac{1}{3}N + \frac{1}{3}E + 1 \\
&\leq \frac{1}{3}(E - x_{k_p}) + \frac{1}{3}E + 1 \\
&= \frac{2}{3}E - \frac{1}{3}x_{k_p} + 1 \\
&= \frac{2}{3}\left(\frac{1}{2}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - \frac{1}{2}r_p\right) - \frac{1}{3}x_{k_p} + 1 \\
&= \frac{1}{3}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - \frac{1}{3}r_p - \frac{1}{3}x_{k_p} + 1.
\end{aligned}$$

Putting it all together:

$$\begin{aligned}
2 &\leq x_{k_p} - \left(\frac{1}{2}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - \frac{1}{2}r_p\right) + \frac{1}{3}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - \frac{1}{3}r_p - \frac{1}{3}x_{k_p} + 1 \\
1 &\leq \frac{2}{3}x_{k_p} - \frac{1}{6}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) + \frac{1}{6}r_p \\
6 &\leq 4x_{k_p} - \sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) + r_p \\
\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) &\leq 4x_{k_p} + r_p - 6 \\
&< 4x_{k_p} + r_p - 2.
\end{aligned}$$

On the other hand, if there are no 2-cells in the complex spanned by the graph, then the number of faces will be 1. Doing a similar calculation, we find

$$\begin{aligned}
2 &= x_{k_p} - \left(\frac{1}{2}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) - \frac{1}{2}r_p\right) + 1 \\
2 &= 2x_{k_p} - \sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) + r_p \\
\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) &= 2x_{k_p} + r_p - 2 \\
&< 4x_{k_p} + r_p - 2,
\end{aligned}$$

since $x_{k_p} > 0$. In each case, we have the same estimate. Now sum over all the components

$C_1, \ldots, C_s$:

$$\begin{aligned}
\sum_{p=1}^{s}\sum_{l=1}^{x_{k_p}} \mathrm{val}(w_{i_l}) &< 4\sum_{p=1}^{s} x_{k_p} + \sum_{p=1}^{s} r_p - \sum_{p=1}^{s} 2 \\
\sum_{i=1}^{x_k} \mathrm{val}(w_i) &< 4x_k + \sum_{p=1}^{s} r_p - 2s \\
&\leq 4x_k + \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 2 - 2 \\
&= 4x_k + \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 4.
\end{aligned}$$

Now from our previous calculations, we know that

$$\begin{aligned}
&\sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 2c_k - 2 \\
&\quad = 6x_k - \sum_{i=1}^{x_k} \mathrm{val}(w_i) \\
&\quad > 6x_k - \left(4x_k + \sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 4\right) \\
&\quad = 6x_k - 4x_k - \sum_{i=1}^{c_k} \mathrm{val}(v_i) - \mathrm{val}(v_a) - \mathrm{val}(v_b) + 4.
\end{aligned}$$

So

$$\begin{aligned}
2\sum_{i=1}^{c_k} \mathrm{val}(v_i) + 2\mathrm{val}(v_a) + 2\mathrm{val}(v_b) - 4c_k - 4 &> 2 + 2x_k - 2c_k \\
\sum_{i=1}^{c_k} \mathrm{val}(v_i) + \mathrm{val}(v_a) + \mathrm{val}(v_b) - 2c_k - 2 &> 1 + x_k - c_k \\
\sum_{i=1}^{c_k} (\mathrm{val}(v_i) - 2) + (\mathrm{val}(v_a) - 1) + (\mathrm{val}(v_b) - 1) &> 1 + x_k - c_k.
\end{aligned}$$

We noted before that, in the absence of interior vertices, the subtiling contributes 2 to the sum of the tile valences. In essence, the vertices $\{v_i\}_{i=1}^{c_k}$ would contribute nothing, whereas v_a and v_b would contribute one each. Since we are not counting boundary edges in $\mathrm{val}(v_i)$, we know that $\mathrm{tval}(v_i) - 3 = \mathrm{val}(v_i) - 2$, whereas $\mathrm{val}(v_a) = \mathrm{tval}(v_a)$ and $\mathrm{val}(v_b) = \mathrm{tval}(v_b)$. (See Figure 3.10 once again.) So now, each vertex of $\{v_i\}_{i=1}^{c_k}$ contributes $\mathrm{val}(v_i) - 2$, and v_a and v_b contribute $\mathrm{val}(v_a) - 1$ and $\mathrm{val}(v_b) - 1$, respectively, in excess of the one they would have contributed had no interior vertices been present. Thus the left hand side of the above inequality is precisely the amount contributed to tile valence sum by the presence of interior vertices. Call this amount b_k, and note that

$$b_k > 1 + x_k - c_k.$$

The total of all these contributions will be $\sum_{k=1}^{j} b_k > j+\sum_{k=1}^{j} x_k-\sum_{k=1}^{j} c_k \geq j+y-(c-q)$.

Now focus on the entire annulus again. We want to estimate $\sum_{i=1}^{c}(\text{tval}(v_i)-3)$. We have already seen that we can look at A' instead of A, and A' is a triangulated annulus with $n+2x'$ top vertices and q bottom vertices connected to the top. Hence, if there were no interior vertices in A', $\sum_{i=1}^{c}(\text{tval}(v_i)-3)$ would equal $n+2x'-q$ by Proposition 3.2. To get the actual value, we must add the amount contributed by the presence of the remaining y interior vertices, to get

$$\begin{aligned}\sum_{i=1}^{c}(\text{tval}(v_i)-3) &= n+2x'-q+\sum_{k=1}^{j} b_k \\ &> n+2x'-q+(j+y-(c-q)) \\ &= n+2x'+j+y-c \\ &\geq n+(x'+y)-c \\ &= n+x-c.\end{aligned}$$

□

Finally, we must establish one more result about well-triangulated annuli.

Proposition 3.4 *Let A be a well-triangulated annulus. Suppose A contains c bottom vertices, labeled $v_1,\ldots,v_c$; n top vertices, labeled $w_1,\ldots,w_n$; and x interior vertices, labeled $u_1,\ldots,u_x$. Then*

$$\sum_{i=1}^{c}(\text{tval}(v_i)-3)+\sum_{i=1}^{x}(\text{tval}(u_i)-6) = -\sum_{i=1}^{n}(\text{tval}(w_i)-3).$$

Proof: Let u be an interior vertex of A. Since A is a well-triangulated annulus, there must be an edge connecting u to a bottom vertex v. As shown in Figure 3.11, replace v by two bottom vertices v' and v'', and replace u by a bottom vertex u', connected on one side to v' and on the other side to v''. We have in effect "sliced" the annulus along the edge, losing one interior vertex and gaining two bottom vertices.

Notice that $\text{tval}(v)=\text{tval}(v')+\text{tval}(v'')$ and $\text{tval}(u)=\text{tval}(u')$. Therefore we can see that

$$\begin{aligned}(\text{tval}(v)-3)+(\text{tval}(u)-6) &= \text{tval}(v)+\text{tval}(u)-9 \\ &= (\text{tval}(v')+\text{tval}(v''))+\text{tval}(u')-9 \\ &= (\text{tval}(v')-3)+(\text{tval}(v'')-3)+(\text{tval}(u')-3).\end{aligned}$$

Hence, this slicing operation has not affected the sum $\sum(\text{tval}(v_i)-3)+\sum(\text{tval}(u_i)-6)$; furthermore, it has not changed the top of the annulus, so the operation has not affected the sum $\sum(\text{tval}(w_i)-3)$.

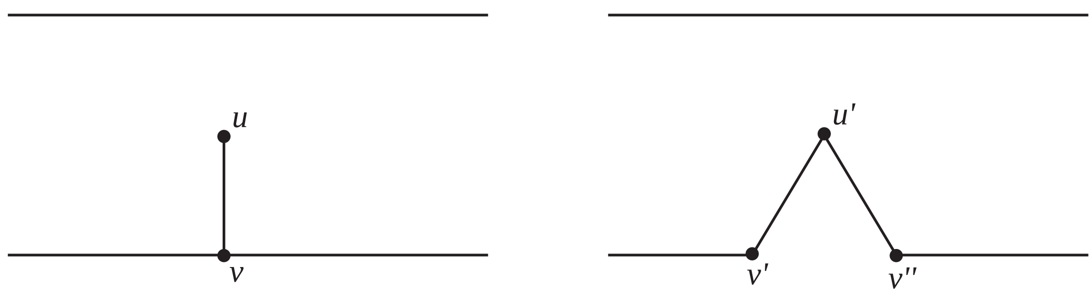

Figure 3.11: Removing an Interior Vertex

Perform this slicing operation for all interior vertices in the annulus. The result will be a new annulus A' with no interior vertices. It will have the same n top vertices $w_1, \ldots, w_n$; and it will have $c + 2x$ bottom vertices, which we will label $v'_1, \ldots, v'_{c+2x}$. Since A' has no interior vertices, it is a simply triangulated annulus, and Proposition 3.2 applies. Let q be the number of bottom vertices (in A') which are connected to the top. We see that

$$\begin{aligned}\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3) + \sum_{i=1}^{x}(\mathrm{tval}(u_i) - 6) &= \sum_{i=1}^{c+2x}(\mathrm{tval}(v'_i) - 3) \\ &= n - q.\end{aligned}$$

However, again by Proposition 3.2, we know that

$$\sum_{i=1}^{n}(\mathrm{tval}(w_i) - 3) = q - n.$$

Hence we conclude that

$$\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3) + \sum_{i=1}^{x}(\mathrm{tval}(u_i) - 6) = -\sum_{i=1}^{n}(\mathrm{tval}(w_i) - 3),$$

which was to be proven. □

3.4 The Moduli of Well-Triangulated Annuli

In this section we present a means of estimating the modulus of any well-triangulated annulus. In particular, we will obtain a lower bound on the modulus in terms of the number of vertices at the top and bottom. First we shall consider simply triangulated annuli.

We shall call a sequence of alternating up and down triangles a *basic subtiling* of a simply triangulated annulus. (We require a basic subtiling to include at least one up triangle and at least one down triangle; single triangles are not considered basic subtilings.) If the entire

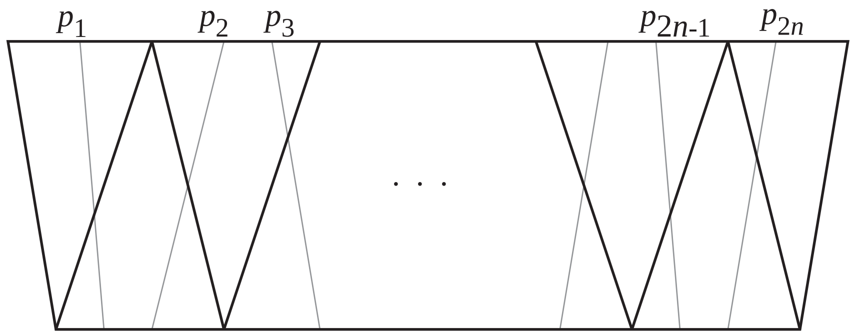

Figure 3.12: Basic Subtiling with $m = n + 1$

annulus is a basic subtiling (i.e., every up triangle is adjacent to two down triangles, and vice versa), we give it the same top and bottom as the original annulus and proceed to compute its fat flow modulus as an annulus. Otherwise, we may consider the basic subtiling to be a quadrilateral, with top and bottom inherited from the annulus. Our first proposition deals with the optimal weight function for basic subtilings:

Proposition 3.5 *Consider a basic subtiling containing $m > 0$ down triangles and $n > 0$ up triangles. Then an optimal weight function for this basic subtiling is*

$$w(t) = \begin{cases} n & \textit{if } t \textit{ is a down triangle} \\ m & \textit{if } t \textit{ is an up triangle.} \end{cases}$$

Proof: We will consider the four possible cases. In each case, we will show that the function $w(t)$ defined above is an optimal weight function by expressing it as a sum of its minimal fat flows. Note that by the fatness of these paths, all minimal fat flows must contain both an up and a down triangle. In particular, we know that minimal fat flows (from bottom to top) in a basic subtiling begin at an up triangle and end at a down triangle.

Our first case, illustrated by Figure 3.12, is a basic subtiling which is a quadrilateral with $m = n + 1$.

Note that each of the flows shown in the figure is a w-minimal fat flow. We claim that

$$w = \sum_{i=1}^{n} (n - i + 1) p_{2i-1} + \sum_{i=1}^{n} i \, p_{2i}.$$

To see the truth of this claim, first consider one of the down triangles, which we shall call t. Calculating the above sum of the paths for this triangle, we find that

$$\sum_{i=1}^{n} (n - i + 1) p_{2i-1}(t) + \sum_{i=1}^{n} i \, p_{2i}(t)$$

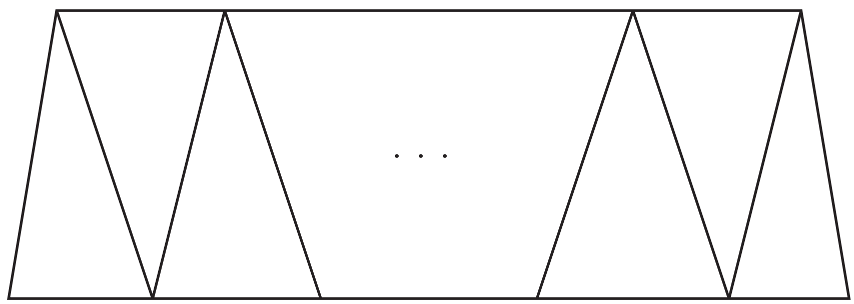

Figure 3.13: A Basic Subtiling with $m = n - 1$

$$
\begin{aligned}
&= \begin{cases} n-1+1 & \text{if } t \text{ is the leftmost down triangle} \\ n & \text{if } t \text{ is the rightmost down triangle} \\ i+(n-(i+1)+1) & \text{otherwise} \end{cases} \\
&= n = w(t).
\end{aligned}
$$

On the other hand, if t is one of the up triangles in this basic subtiling,

$$
\begin{aligned}
\sum_{i=1}^{n}(n-i+1)p_{2i-1}(t) + \sum_{i=1}^{n} i\, p_{2i}(t) &= (n-i+1)+i \\
&= n+1 \\
&= m = w(t).
\end{aligned}
$$

Thus w is the weighted sum of w-minimal flows and is therefore an optimal weight function for this basic subtiling.

Second, we consider basic subtilings which are quadrilaterals with $m = n - 1$. (See Figure 3.13.) We can see that the proposition is true for such basic subtilings by switching what we call the top and bottom, thus reducing it to the previous case.

Our third case is when the basic subtiling is a quadrilateral with $m = n$, as shown in Figure 3.14. In this case, it is clear that $w = \sum_{i=1}^{n} p_i$ gives $w(t) = n$ for all t. So w is the sum of its minimal fat flows and is thus an optimal weight function. Nor would the situation change materially if the leftmost triangle were a down triangle instead of an up triangle (which would, of course, force the rightmost triangle to be an up triangle instead of a down triangle).

Finally, we consider basic subtilings which are annuli. For an example, see Figure 3.15. Note that in this case, $m = n$. Thus it is similar to the previous case, and, as shown in the figure, $w = \sum_{i=1}^{n} p_i$. $\square$

Corollary 3.1 *If B is a basic subtiling consisting of $m > 0$ down triangles and $n > 0$ up triangles, then the modulus M of B is*

$$\frac{n+m}{nm}$$

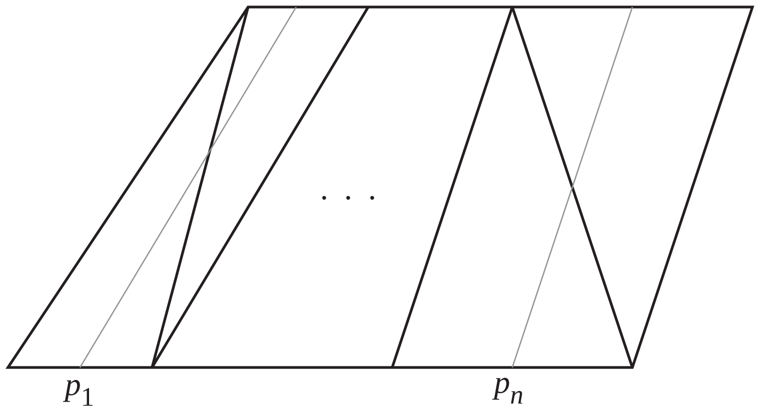

Figure 3.14: A Basic Subtiling with $m = n$

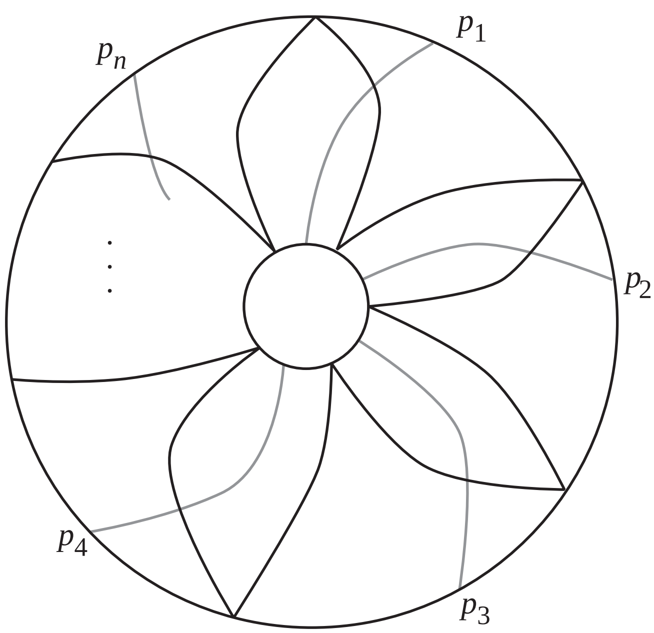

Figure 3.15: An Annular Basic Subtiling

Proof: By Proposition 3.5, the optimal weight function w on B assigns a weight of n to the down triangles and a weight of m to the up triangles. Since a fat flow will contain one up triangle and one down triangle, the length of any fat flow is $n+m$. Hence $H_{w,f}(B)$ is $n+m$. The area $A_w(B)$ will be mn^2+nm^2, since there are m down triangles (each with weight n) and n up triangles (each with weight m). The result follows. □

Next observe that one can partially order the basic subtilings by subset inclusion and obtain *maximal basic subtilings*. Clearly these maximal basic subtilings have no tiles in common, since otherwise their union would be a basic subtiling properly containing both of them. Suppose that no non-boundary edges of a simply triangulated annulus join the bottom to the bottom; then the triangulation contains only up triangles and down triangles. Also, since neither boundary component is allowed to consist of just one point, the triangulation must contain at least one up triangle and one down triangle. Thus, at some point, there must be an up triangle adjacent to a down triangle, and thus the triangulation contains a basic subtiling. It must therefore contain a maximal basic subtiling. We find that, if no non-boundary edges of a simply triangulated annulus join the bottom to the bottom, then we can optimize each maximal basic subtiling and piece them together to obtain an optimal weight function for the entire simply triangulated annulus.

Proposition 3.6 *Let A be a simply triangulated annulus such that all non-boundary edges of the triangulation connect the bottom and the top of the annulus. Let the maximal basic subtilings of A be denoted $B_1, B_2, \ldots, B_k$. For $i = 1, \ldots, k$, let n_i and m_i be the number of up and down triangles, respectively, in B_i. Then an optimal weight function $w(t)$ for A is given as follows:*

$$w(t) = \begin{cases} \frac{m_i}{n_i+m_i} & \text{if } t \text{ is an up triangle in } B_i; \\ \frac{n_i}{n_i+m_i} & \text{if } t \text{ is a down triangle in } B_i; \\ 0 & \text{otherwise.} \end{cases}$$

Proof: Observe that w is well-defined since maximal basic subtilings cannot share tiles. Note further that w is simply the concatenation of the optimal weight functions on the maximal basic subtilings, all normalized to have a height of one. Thus, in order to prove the proposition, we need to show that one cannot construct a path with w-length less than one by starting in one maximal basic subtiling and ending in another.

Suppose we have a path p which is not contained in a maximal basic subtiling. It must begin with an up triangle t. First suppose that t is part of a maximal basic subtiling. Then t must be at the edge of that maximal basic subtiling; otherwise the second tile in p would be a down triangle in the same maximal basic subtiling, and the path's length would be one. Furthermore, the path must end at the first down triangle s that it encounters, since by then it shall have reached the top of the annulus. Thus, the path must look like the one shown in Figure 3.16. The path consists of triangles $\{t, t_1, \ldots, t_n, s\}$. However, the last two tiles (t_n and s) of p are an alternating up/down pair of triangles; thus they are themselves part of a

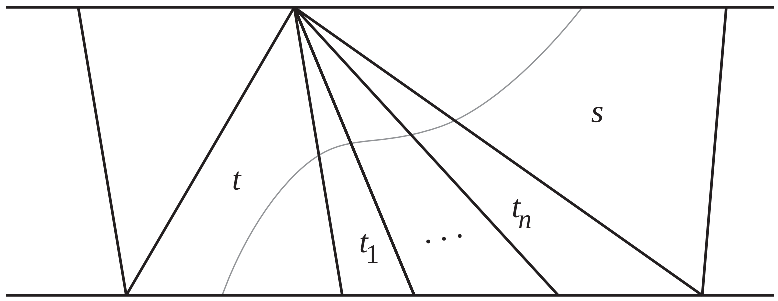

Figure 3.16: A Flow out of a Maximal Basic Subtiling

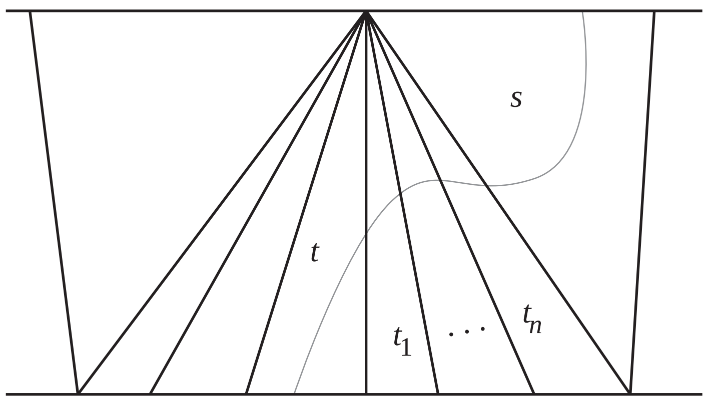

Figure 3.17: A Flow into a Maximal Basic Subtiling

maximal basic subtiling. Hence $w(t_n) + w(s)$ must equal one, and the length of p therefore cannot be less than one.

Now suppose p begins with an up triangle t that is not part of a maximal basic subtiling. In this case, both neighbors of t must be up triangles, and so p will consist of a sequence of up triangles, until the final triangle s, which will be a down triangle. In other words, $p = \{t, t_1, \ldots, t_n, s\}$, as shown in Figure 3.17. Once again, however, t_n and s must be part of the same maximal basic subtiling, and hence the length of p will not be less than one.

Therefore, we may conclude that paths of length one are w-minimal, and hence that the function w, being the sum of w-minimal fat flows, is an optimal weight function for A. $\square$

Corollary 3.2 *Let A be a simply triangulated annulus such that all non-boundary edges of the triangulation connect the bottom and the top of the annulus. Suppose A contains k maximal basic subtilings, and let M_i be the fat flow modulus of the ith maximal basic subtiling,* *for $i = 1, \ldots, k$. Then the fat flow modulus of A is*

$$\left(\sum_{i=1}^{k} \frac{1}{M_i} \right)^{-1}.$$

Proof: Proposition 3.6 states that the optimal weight function w for A is the concatenation of the optimal weight functions for the maximal basic subtilings, each being normalized to have a height of one. Note that the height of the ith maximal basic subtiling under w will be 1, and so its area will be $1/M_i$. The height of A under w will also be 1. The area of A will be the sum of the areas of the maximal basic subtilings, namely, $\sum 1/M_i$. Thus the modulus of A will be the reciprocal of $\sum 1/M_i$. $\square$

Now that we have an expression for the modulus of A, let us consider the minimum possible modulus for such a simply triangulated annulus, in terms of the number of vertices on the top and the number of vertices on the bottom.

Proposition 3.7 *Let A be a simply triangulated annulus such that all non-boundary edges of the triangulation connect the bottom and the top of the annulus. Suppose A has n vertices on the top and k vertices on the bottom. If M is the fat flow modulus of A, then*

$$M \geq \frac{3}{2m},$$

where $m = \min\{n, k\}$.

Proof: Note that n and k also denote the number of down and up triangles, respectively, in A. Denote the maximal basic subtilings of A as $B_1, B_2, \ldots, B_p$. For $i = 1, \ldots, p$, let n_i and k_i be the number of up and down triangles, respectively, in B_i. By Corollary 3.1, the fat flow modulus for B_i is

$$M_i = \frac{n_i + k_i}{n_i k_i}$$

Let $m_i = \min\{n_i, k_i\}$. We have three possible cases: either $m_i = n_i = k_i$; $m_i = n_i = k_i - 1$; or $m_i = k_i = n_i - 1$. Consider the first case. If $m_i = n_i = k_i$, then

$$M_i = \frac{2m_i}{m_i^2} = \frac{2}{m_i}.$$

The second and third cases reduce to the same situation:

$$M_i = \frac{2m_i + 1}{m_i(m_i + 1)}.$$

Note that this situation gives a smaller modulus than the former:

$$\frac{2m_i + 1}{m_i(m_i + 1)} = \frac{2}{m_i} \left(\frac{m_i + \frac{1}{2}}{m_i + 1} \right) < \frac{2}{m_i}.$$

Thus we conclude that

$$\frac{2m_i + 1}{m_i(m_i+1)} \leq M_i.$$

Now since $m_i \geq 1$, we know

$$\begin{array}{rcll} m_i^2 & \geq & m_i & \\ 4m_i^2 + 2m_i & \geq & 3m_i^2 + 3m_i & \text{by adding } 3m_i^2 + 2m_i \text{ to each side} \\ 2m_i(2m_i+1) & \geq & 3(m_i^2 + m_i) & \\ M_i \geq \dfrac{2m_i+1}{m_i^2+m_i} & \geq & \dfrac{3}{2m_i} & \\ \dfrac{1}{M_i} & \leq & \dfrac{2m_i}{3}, & \end{array}$$

for any i from 1 to p. It therefore follows that

$$\sum_{i=1}^{p} \frac{1}{M_i} \leq \sum_{i=1}^{p} \frac{2m_i}{3} = \frac{2}{3}\sum_{i=1}^{p} m_i \leq \frac{2}{3}m,$$

since m, the minimum of the total number of up triangles and the total number of down triangles in A, must be at least as large as the sum of the minima for each maximal basic subtilings (i.e., at least as large as $\sum m_i$).

However, by Corollary 3.2, we have

$$M = \left(\sum_{i=1}^{p} \frac{1}{M_i}\right)^{-1} \geq \left(\frac{2}{3}m\right)^{-1} = \frac{3}{2m},$$

which is what was to be proven. □

Note that this estimate is the best possible; one can obtain a simply triangulated annulus with this modulus by taking each n_i to be equal to 1 and letting each $k_i = 2$. Then $m_i = 1$ for all $i = 1, \ldots, m$, $M_i = 3/2$, and $\sum m_i = m$, making all the inequalities into equations.

This proposition, of course, does not deal with *all* simply triangulated annuli, only those whose non-boundary edges connect top to bottom. Such simply triangulated annuli contain only up and down triangles. What happens if we allow non-boundary edges to connect a bottom vertex to another bottom vertex, or, equivalently, allow the simply triangulated annulus to contain bent and mid triangles?

The next proposition investigates the effect of ignoring bent triangles and considering mid triangles as up triangles.

Proposition 3.8 *Let X be a simply triangulated annulus. Let $\bar{X}$ be the simply triangulated annulus obtained by deleting all bent triangles from X. Let M and $\bar{M}$ be the fat flow moduli of X and $\bar{X}$, respectively. Then $M \geq \bar{M}$.*

Proof: Let $\bar{w}$ be a fat flow optimal weight function for $\bar{X}$. Then

$$\bar{M} = \frac{H^2_{\bar{w},f}(\bar{X})}{A_{\bar{w}}(\bar{X})}.$$

Now define a weight function w on X by

$$w(t) = \begin{cases} \bar{w}(t), & t \in \bar{X} \\ 0, & t \notin \bar{X}. \end{cases}$$

Clearly $A_{\bar{w}}(\bar{X}) = A_w(X)$. Let f be a w-minimal fat flow in X; then the restriction of f to $\bar{X}$ is a $\bar{w}$-minimal fat flow in $\bar{X}$; conversely, any $\bar{w}$-minimal fat flow $\bar{f}$ in $\bar{X}$ can be extended to a w-minimal fat flow f in X. Therefore $H_{\bar{w},f}(\bar{X}) = H_{w,f}(X)$. Now we may conclude,

$$M = \sup_v \frac{H^2_{v,f}(X)}{A_v(X)} \geq \frac{H^2_{w,f}(X)}{A_w(X)} = \frac{H^2_{\bar{w},f}(\bar{X})}{A_{\bar{w}}(\bar{X})} = \bar{M}.$$

□

This fact now allows us to estimate the modulus of an arbitrary simply triangulated annulus.

Corollary 3.3 *Let X be a simply triangulated annulus with n top vertices and with q of the bottom vertices connected to the top. If m is the minimum of n and q, and if M is the fat flow modulus of X, then*

$$M \geq \frac{3}{2m}.$$

Proof: Define $\bar{X}$ and $\bar{M}$ as in Proposition 3.8. Then $\bar{X}$ is a simply triangulated annulus, all of whose non-boundary edges connect the bottom to the top. It has n top vertices and q bottom vertices. By Propositions 3.7 and 3.8, we have

$$M \geq \bar{M} \geq \frac{3}{2m}.$$

□

Finally, let us obtain a result on the moduli of well-triangulated annuli.

Proposition 3.9 *Let X be a well-triangulated annulus. Suppose the triangulation contains n top vertices. Let M be the fat flow modulus of X. Then*

$$M \geq \frac{3}{2n}.$$

Proof: Let q be the number of bottom vertices which are connected to the top by an edge. As in the proof of Proposition 3.4, replace X by a simply triangulated annulus X', with n top vertices and $q + r$ bottom vertices connected to the top. Now if m is the minimum of n and $q + r$, then clearly $m \leq n$.

By Corollary 3.3, the fat flow modulus M' of X' is greater than or equal to $3/(2m)$.

Let w be the optimal weight function for X'. Because of the procedure we used in replacing X by X', there is a one-to-one correspondence of tiles in X onto tiles in X'. So we may apply the function w to the original tiling X as well as to the modified tiling X' and say that the w-areas of the two tilings are the same: $A_w(X) = A_w(X')$.

Next consider the relationship of flows in X to flows in X'. The top tiles of each tiling are the same, and all bottom tiles of X are bottom tiles of X'. The operation preserves all combinatorics, except that two tiles adjacent along some interior edge in X may no longer be adjacent in X'. However, both of these tiles will now be bottom tiles, so any fat flow in X containing one of these tiles has a subset that is a fat flow in X'. Thus, any fat flow in X contains a fat flow in X'. Conversely, any fat flow in X' is contained in a fat flow in X. So the infimum of the w-lengths of the fat flows may be smaller when taken over X' than when taken over X; hence we conclude that $H_w(X) \geq H_w(X')$. Thus,

$$M \geq \frac{H_w^2(X)}{A_w(X)} \geq \frac{H_w^2(X')}{A_w(X)} = M' \geq \frac{3}{2m} \geq \frac{3}{2n}.$$

□

So far we have considered well-triangulated annuli and have succeeded in expressing the sums of their vertex valences in terms of the number of vertices on top and bottom. We have in turn estimated their moduli in terms of the number of vertices on the top and bottom. Now let us consider general disk triangulation graphs for a moment and see how these results help us.

Suppose G is a disk triangulation graph. Pick a vertex v. Let the distance of a vertex w from v (denoted $d(v, w)$) be the minimum number of edges in a path connecting v to w. For any positive integer k, let $V_k = \{w \,|\, d(v, w) = k\}$. Let E_k be the set of edges e that either connect a vertex in V_k to another vertex of V_k, or that connect a vertex in V_k to a vertex in V_{k+1}. Let E'_k be the set of all edges e that connect a vertex in V_{k+1} to another vertex in V_{k+1}, such that e forms the side of a (triangular) 2-cell whose other two sides are in E_k. Finally, let $\bar{E}_k = E_k \cup E'_k$.

Let A_k be the simplicial 2-complex whose 1-skeleton is $\bar{E}_k$. We have the following proposition, which is (save for an extra hypothesis) essentially the same as Theorem 3.3:

Proposition 3.10 *Let G be a disk triangulation graph; choose a vertex v and define V_k and A_k as above. Suppose that the sets A_k are concentric well-triangulated annuli surrounding*

v_0, and suppose that the sequence

$$a_k = \sum_{\substack{v \in V_j \\ j \leq k}} (\mathrm{val}(v) - 6)$$

is bounded. Then G is parabolic.

Proof: We shall need some notation. Fix k, and consider A_k. We let n_k and c_k be the number of top and bottom vertices, respectively, and let x_k be the number of interior vertices. Label the top vertices as $v_{t1}, v_{t2}, \ldots, v_{tn_k}$; label the bottom vertices as $v_{b1}, v_{b2}, \ldots, v_{bc_k}$; label the interior vertices as $v_{i1}, v_{i2}, \ldots, v_{ix_k}$. all bottom vertices will be in V_k; all top vertices will be in V_{k+1}; and all interior vertices will be in V_{k+1}. Now let $t_k = \sum_{j=1}^{c_k}(\mathrm{tval}(v_{tj}) - 3)$, the sum of the tile valence excesses on the top; let $b_k = \sum_{j=1}^{n_k}(\mathrm{tval}(v_{bj}) - 3)$, the sum of the tile valence excesses on the bottom; and let $i_k = \sum_{j=1}^{x_k}(\mathrm{tval}(v_{ij}) - 6)$, the sum of the tile valence excesses in the interior. Thus, in order to sum the values $\mathrm{val}(v) - 6$ over V_j, say, we would consider the interior vertices of A_{j-1} and the top vertices of A_{j-1}, which would also be the bottom vertices of A_j. So we would have

$$\sum_{v \in V_j} (\mathrm{val}(v) - 6) = i_{j-1} + t_{j-1} + b_j.$$

Define a_k as in the statement of the proposition, and let $z = a_1 - b_1$. Now

$$\begin{aligned}
a_k &= \sum_{j=1}^{k} \sum_{v \in V_j} (\mathrm{val}(v) - 6) \\
&= ((a_1 - b_1) + b_1) + \sum_{j=2}^{k} (i_{j-1} + t_{j-1} + b_j) \\
&= \sum_{j=1}^{k} b_j + \sum_{j=2}^{k} t_{j-1} + \sum_{j=2}^{k} i_{j-1} + z \\
&= \sum_{j=1}^{k} b_j + \sum_{j=1}^{k-1} t_j + \sum_{j=1}^{k-1} i_j + z \\
&= b_k + \sum_{j=1}^{k-1} (b_j + t_j + i_j) + z.
\end{aligned}$$

However, by Proposition 3.4, $b_j + i_j = -t_j$, so $b_j + i_j + t_j = 0$. Thus we may say,

$$a_k = b_k + z.$$

So, since we assume the sequence $\{a_k\}$ is bounded, we also know that the sequence $\{b_k\}$ is bounded. By Proposition 3.3, $b_k \geq n_k + x_k - c_k \geq n_k - c_k$, and hence we can conclude that $n_k - c_k$ is bounded. So, for some B and all positive integers k,

$$\begin{aligned}
n_k - c_k &\leq B \\
n_k &\leq B + c_k.
\end{aligned}$$

Now let us use induction to obtain an estimate for n_k involving only B and c_1. We have $n_1 \leq B + c_1$. Next, observe that $c_{k+1} = n_k$ since the annuli $\{A_j\}$ must fit together. Suppose that $n_k \leq kB + c_1$ for some positive integer k. Then

$$\begin{aligned} n_{k+1} &\leq B + c_{k+1} = B + n_k \\ &\leq B + (kB + c_1) \\ &= (k+1)B + c_1. \end{aligned}$$

Thus, we conclude that $n_k \leq kB + c_1$ for all positive integers k.

Now let us consider the modulus M_k of A_k. The number of top vertices of A_k would be

$$n_k \leq c_1 + kB.$$

By Proposition 3.9,

$$M_k \geq \frac{3}{2n_k} \geq \frac{3}{2(c_1 + kB)}.$$

Let M be the fat flow modulus of G. Then by the Layer Theorem,

$$M \geq \sum_{k=1}^{\infty} M_k \geq \frac{3}{2} \sum_{k=1}^{\infty} \frac{1}{c_1 + kB} = \infty.$$

We conclude that G is parabolic. $\square$

3.5 Dealing with Islands

Now let us work toward getting rid of the extra hypothesis in Proposition 3.10—namely, that each A_k is a well-triangulated annulus. To do so, we will need to consider the structure of a disk triangulation graph in detail.

We shall define a *well-triangulated annulus with islands* to be a closed topological annulus with two (disjoint) boundary components homeomorphic to circles and an imposed triangulation which satisfies the following conditions: any edge joining a vertex on the top to another vertex on the top must be part of the top; and any vertex not connected to the bottom by an edge is either part of the bottom or surrounded by a cycle of vertices satisfying the following two conditions:

1. The vertices in this cycle contain no bottom vertices; and

2. The vertices in this cycle are all connected to the bottom by an edge.

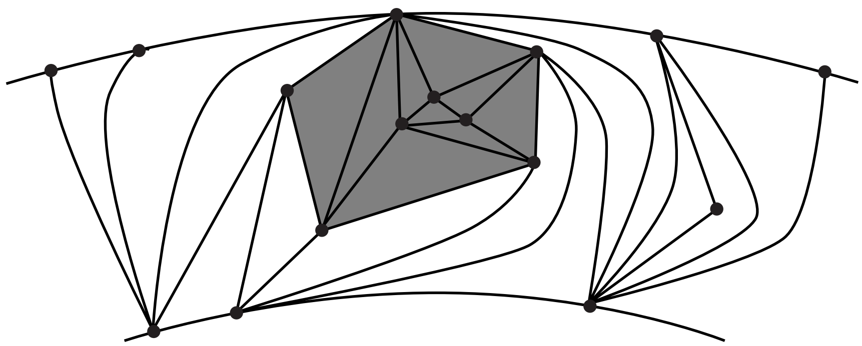

Figure 3.18: An Island

If L is the union of such a cycle with its interior, then we say L is an *island*. (See Figure 3.18.) Note that by this definition all islands must contain vertices in their interior and be simply connected. The *size* of an island is the number of vertices in its boundary cycle.

If we take a disk triangulation graph G and pick a base vertex v_0, we can define sets of vertices $\{V_i\}$ as before, letting V_i be the set of all vertices whose edge distance from v_0 is i. Suppose that A, a subset of the complex spanned by G, is a well-triangulated annulus with islands. Suppose that all vertices on the bottom of A are in V_k and that the other vertices in A are in $\bigcup_{j=k+1}^{\infty} V_j$. If L is an island in A, then we say that L has *station k* or *is stationed at k*. Note that all boundary vertices of such an island will belong to V_{k+1}. Suppose that l is the maximum value such that there is a vertex in L belonging to V_l. Then we say that the *depth* of L is $l-(k+1)$. Note that all vertices in the interior of L are in $\bigcup_{j=k+2}^{l} V_k$.

We now show that one can decompose G into the union of a disk and a series of well-triangulated annuli with islands.

Proposition 3.11 *Let G be a locally finite disk triangulation. Pick a base vertex v_0, and for each positive integer k define V_k to be the set of all vertices v such that the minimum number of edges connecting v to v_0 is k. Then there exist sets $\{\bar{A}_i\}_{i=0}^{\infty}$ such that $G = \bigcup_{i=0}^{\infty} \bar{A}_i$ and such that the following statements are true:*

1. *$\bar{A}_0$ is topologically equivalent to a closed disk.*
2. *For $i \geq 1$, $\bar{A}_i$ is a well-triangulated annulus with islands.*
3. *For $i \geq 1$, all bottom vertices of $\bar{A}_i$ are in V_i; and all top vertices of $\bar{A}_i$ are in V_{i+1}. Note that this statement implies that all islands in $\bar{A}_i$ are stationed at i.*
4. *The bottom of $\bar{A}_i$ is equal to the top of $\bar{A}_{i-1}$ if $i \geq 2$, and the bottom of $\bar{A}_1$ is equal to the boundary of $\bar{A}_0$.*

Proof: Let G be a disk triangulation with a distinguished vertex v_0. Now let A_0 be the union of all triangles (faces) incident to v_0. Note that A_0 contains all vertices in V_1. Since each of these 2-cells has v_0 as part of its boundary, we know that A_0 is connected. In fact, since A_0 consists of all triangles adjacent to v_0, its interior is also connected. Furthermore, by local finiteness, A_0 is compact. Therefore, exactly one connected component of $G \setminus A_0$ is unbounded; we shall call it K_0. By definition, K_0 is connected and open. Now, we in turn define $\bar{A}_0$ to be $G \setminus K_0$. $\bar{A}_0$ is closed, and clearly $\bar{A}_0 \supseteq A_0$. Since K_0 is connected, $\bar{A}_0$ is simply connected. Furthermore, we know that $\bar{A}_0$ is bounded (and hence compact, being closed); otherwise, $\bar{A}_0 \setminus A_0$ would be unbounded, and the complement of A_0 would have a second unbounded component besides K_0. Now A_0 is closed with connected interior, and $\bar{A}_0$ is the union of A_0 with open sets whose boundaries are subsets of ∂A_0. Therefore $\bar{A}_0$ has a connected interior. Since $\bar{A}_0$ is compact, is simply connected, and has a connected interior, it is homeomorphic to the closed disk. We know that all vertices on the boundary of $\bar{A}_0$ are on the boundary of K_0, and hence on the boundary of A_0; thus all of these vertices are in V_1. Since $\bar{A}_0 \supseteq A_0$, we know that $\bar{A}_0$ contains all vertices in V_1.

We have now proved the first conclusion of the proposition by constructing the set $\bar{A}_0$. We shall prove the rest of the conclusions inductively.

Suppose we have complexes $\{\bar{A}_i\}_{i=0}^{k}$ with a finite number of vertices. Let $D_k = \bigcup_{i=0}^{k} \bar{A}_i$, and suppose that the following statements hold:

1. D_k is a closed topological disk;
2. The vertices on ∂D_k are all in V_{k+1};
3. All vertices of $\bigcup_{i=1}^{k+1} V_i$ are contained in D_k.

We shall refer to these statements as Induction Facts in the remainder of this proof. Clearly they are satisfied for $k = 0$. Now suppose they are satisfied for some unspecified k.

Consider the edges of G that are not contained in the interior of D_k. We shall define E_{k+1} to be the set of all such edges that connect a vertex in ∂D_k to a vertex in V_{k+1} or in V_{k+2}. Note that ∂D_k is a subset of E_{k+1} since all vertices on ∂D_k are in V_{k+1}. We shall let E'_{k+1} be the set of all edges e (not contained in int D_k) from a vertex in V_{k+2} to another vertex in V_{k+2} such that e is the side of a triangle whose other two sides are in E_{k+1}. Finally, we let $\bar{E}_{k+1} = E_{k+1} \cup E'_{k+1}$. Again, $\bar{E}_{k+1}$ contains no edges in the interior of D_k. Clearly the graph spanned by $\bar{E}_{k+1}$ is a connected graph: any vertex in this graph is connected by an edge to ∂D_k, this boundary is connected, and this boundary is a subset of the graph.

Now we let A_{k+1} be the simplicial 2-complex spanned by $\bar{E}_{k+1}$. Note that any triangle that is not contained in D_k but has a vertex on ∂D_k is part of A_{k+1}. Thus, ∂D_k lies in the interior of $D_k \cup A_{k+1}$. By Induction Fact 3, every vertex in $\bigcup_{i=1}^{k+1} V_i$ lies in D_k, and by Induction Fact 2, every vertex on ∂D_k is in V_{k+1}. However, since ∂D_k lies in the interior of $D_k \cup A_{k+1}$, we must conclude that every vertex on $\partial(D_k \cup A_{k+1})$ is in V_{k+2}, being connected to ∂D_k by an edge. Since the 1-skeleton of A_{k+1} is connected, we know that A_{k+1} is connected. Since $\partial D_k = A_{k+1} \cap D_k$, we know that the set $D_k \cup A_{k+1}$ is connected.

We will follow the same procedure with A_{k+1} as we did with A_0. By local finiteness, $D_k \cup A_{k+1}$ is closed and bounded, hence compact. We thus have exactly one unbounded connected component of $G \setminus (D_k \cup A_{k+1})$; we shall call it K_{k+1}. By definition, K_{k+1} is connected and open. Note that all vertices on the boundary of K_{k+1} are in V_{k+2}. Note that $(\text{int } D_k) \cap A_{k+1} = \emptyset$, since we excluded from $\bar{E}_{k+1}$ all edges in int D_k. We shall define $\bar{A}_{k+1}$ to be $G \setminus (K_{k+1} \cup \text{int } D_k)$. Clearly $\bar{A}_{k+1}$ is closed; also, it is bounded, since K_{k+1} is the only unbounded component of the complement of $D_k \cup A_{k+1}$; also, $\bar{A}_{k+1} \supseteq A_{k+1}$. Since K_{k+1} is connected, the set D_{k+1}, which we define to be $D_k \cup \bar{A}_{k+1} = G \setminus K_{k+1}$ is simply connected. Now since A_{k+1} contains precisely the triangles adjacent to ∂D_k, it has a connected interior. However, $\bar{A}_{k+1}$ is the union of A_{k+1} with open sets whose boundaries are subsets of ∂A_{k+1}. Therefore $\bar{A}_{k+1}$ also has a connected interior. It follows that D_{k+1} also has a connected interior. So D_{k+1} is compact (being closed and bounded), is simply connected, and has a connected interior. Thus it must be homeomorphic to a closed disk. This establishes Induction Fact 1 for $k+1$. By Induction Fact 1 for k, we know that D_k is a closed disk. So $\bar{A}_{k+1}$ is what you get when you take a closed disk (D_{k+1}) and subtract the interior of another closed disk (D_k) which is properly contained in the first disk. Since ∂D_k (the boundary of the inner closed disk) is contained in the interior of D_{k+1}—and hence in the interior of D_{k+1}—we know that $\bar{A}_{k+1}$ is a closed topological annulus with two distinct boundary components. One of its boundary components is ∂D_k. All vertices on this boundary component are in V_{k+1} by Induction Fact 2 for k, and we shall call this component the *bottom* of $\bar{A}_{k+1}$. The other boundary component must be the boundary of K_{k+1}, and all the vertices on this component are in V_{k+2}. We shall call this boundary component the *top* of A_{k+1}. We have now established the third and fourth conclusions of the proposition for $k+1$. Notice that the top of $\bar{A}_{k+1}$ is the only boundary component of D_{k+1}, so all vertices on ∂D_{k+1} are in V_{k+2}. Thus we obtain Induction Fact 2 for $k+1$.

Since all vertices of V_{k+1} are contained in D_k, since all vertices on ∂D_k are in V_{k+1}, and since $\bar{A}_{k+1}$ contains all vertices which are not in D_k but which can be connected to ∂D_k by an edge, it follows that all vertices of V_{k+2} are contained in $D_k \cup \bar{A}_{k+1} = D_{k+1}$. This, along with Induction Fact 3 for k, gives us Induction Fact 3 for $k+1$.

It remains to show that $\bar{A}_{k+1}$ is a well-triangulated annulus with islands. Suppose there exists a component C of $\bar{A}_{k+1} \setminus A_{k+1}$ that is not equal to K_{k+1} or int D_k. Since K_{k+1} is the only unbounded component of the complement of $D_k \cup A_{k+1}$, and hence of A_{k+1}, we know that C is bounded. We know that each component of $\bar{A}_{k+1} \setminus A_{k+1}$ contains only vertices in $\bigcup_{i=k+2}^{\infty} V_i$ by definition combined with Induction Fact 3 for k; and we know that any vertex on the boundary of such a component is on the boundary of A_{k+1} but is not part of ∂D_k; hence any such vertices must be in V_{k+2}. Now by definition, C is connected. To see that C is also simply connected, note that

$$\begin{aligned} G \setminus C &= D_k \cup K_{k+1} \cup (\bar{A}_{k+1} \setminus C) \\ &= D_k \cup K_{k+1} \cup A_{k+1} \cup \tilde{C}, \end{aligned}$$

where $\tilde{C}$ is the union of all other connected components of $\bar{A}_{k+1} \setminus A_{k+1}$. However, A_{k+1}

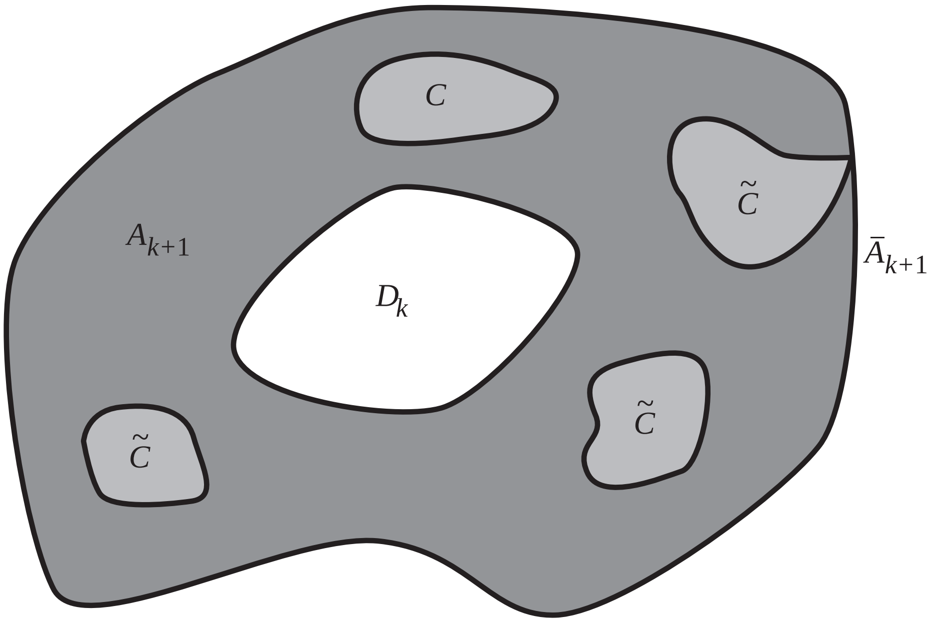

Figure 3.19: $\bar{A}_{k+1}$ Is a Well-Triangulated Annulus with Islands

shares part of its boundary with K_{k+1} and another part of its boundary with D_k; all three of them are connected sets, and A_{k+1} is closed. Therefore the set $D_k \cup K_{k+1} \cup A_{k+1}$ is connected. Furthermore, the boundary of $\tilde{C}$ is part of the boundary of A_{k+1}, and A_{k+1} is closed. Therefore $(\bigcup_{i=0}^{k} \bar{A}_i) \cup K_{k+1} \cup A_{k+1} \cup \tilde{C}$ is connected; thus $G \setminus C$ is connected, and C is simply connected.

Since C is bounded, open, and simply connected, its boundary must be a cycle of vertices. Since ∂C is part of ∂A_{k+1} but not of ∂D_k, we know that all vertices in ∂C are in V_{k+2}. Now since all triangles adjacent to the bottom of $\bar{A}_{k+1}$ are in $D_k \cup A_{k+1}$, it follows that ∂C can contain no vertices on the bottom of $\bar{A}_{k+1}$. Furthermore, since $\partial C \subseteq A_{k+1}$, all vertices in ∂C must be connected to the bottom of $\bar{A}_{k+1}$ by an edge. Hence, if C contains any vertices, it must be an island. Thus we may conclude that $\bar{A}_{k+1}$ is a well-triangulated annulus with islands, giving us the second conclusion of the proposition for $\bar{A}_{k+1}$.

With this statement, we have completed our induction. By Induction Fact 3, every vertex of G is contained in at least one set $\bar{A}_k$, and hence we have succeeded in decomposing G into the union of well-triangulated annuli with islands. $\square$

We will use the following lemma to replace an island by a set of interior vertices.

Lemma 3.1 *Let P be a graph homeomorphic to a circle. Suppose P has n vertices, labeled $v_1, \ldots, v_n$. We also assume $n \geq 3$. Then there is a triangulation T_P with the following properties:*

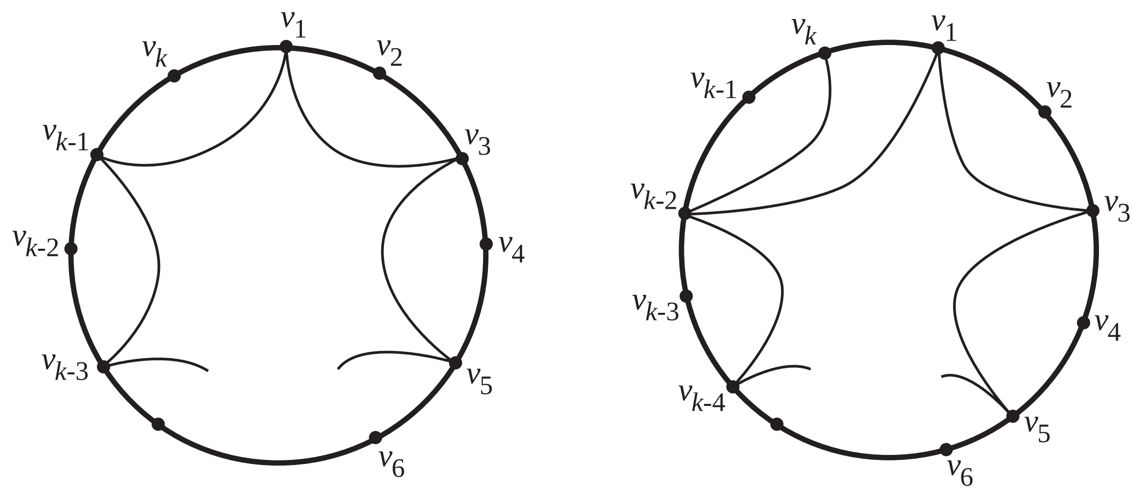

Figure 3.20: Creating T_P

1. *$|T_P|$ is a closed topological disk.*
2. *P is the boundary of T_P.*
3. *All vertices of T_P lie on its boundary.*
4. $\sum_{i=1}^{n}(\mathrm{tval}(v_i) - 3) = -6$*, where* tval$(\cdot)$ *counts only triangles inside T_p.*

Proof: In order to prove this lemma, we must triangulate the region bounded by P without introducing any new vertices. We shall do so by induction on the number of vertices n. If $n = 3$, our task is trivial—P is the boundary of a triangle. Now suppose that for n from 3 to $k - 1$ we have found such a triangulation. Consider P having k vertices. Label the vertices consecutively around the circle.

If k is even, connect v_1 to v_3, v_3 to v_5, and, in general, v_i to v_{i+2} for odd values of i. We finish by connecting v_{k-1} to v_1. (See the left side of Figure 3.20.) We now have tiled the region bounded by P with $k/2$ triangles and a $(k/2)$-gon. Since $k/2 < k$, we can triangulate the interior of the $(k/2)$-gon without introducing any new vertices, thus triangulating the entire region bounded by P.

If k is odd, connect v_1 to v_3, v_3 to v_5, and, in general, v_i to v_{i+2} for odd values of i. We finish by connecting v_{k-2} to v_k. Then draw one more edge, this time from v_{k-2} to v_1. (See the right side of Figure 3.20.) We have now tiled the region bounded by P with $(k+1)/2$ triangles and a $((k-1)/2)$-gon. Since $((k-1)/2) < k$, we know we can triangulate the interior of the $((k-1)/2)$-gon, thus triangulating the entire region bounded by P.

Now we must prove the last statement in the lemma. We do so by using the Euler characteristic once again. Embed this triangulation in a sphere. It will have n vertices. Each vertex

v_i will have $\text{tval}(v_i) - 1$ non-boundary edges incident to it, and we shall count each of these edges twice if we sum over the vertices. Furthermore, there will be n boundary edges. So the total number of edges is $(1/2)\sum_{i=1}^{n}(\text{tval}(v_i) - 1) + n$, which is $(1/2)\sum_{i=1}^{n}\text{tval}(v_i) + (1/2)n$. To count the number of faces, note that each non-boundary edge will contribute 1/3 to two faces; each boundary edge will contribute 1/3 to one face; and there will be one face coming from the complement of the triangulation. Thus, the total number of faces will be

$$\frac{2}{3}\left(\frac{1}{2}\sum_{i=1}^{n}(\text{tval}(v_i) - 1)\right) + \frac{1}{3}n + 1 = \frac{1}{3}\sum_{i=1}^{n}\text{tval}(v_i) + 1.$$

Putting all of these numbers together, we obtain

$$\begin{aligned} 2 &= n - \left(\frac{1}{2}\sum_{i=1}^{n}\text{tval}(v_i) + \frac{1}{2}n\right) + \left(\frac{1}{3}\sum_{i=1}^{n}\text{tval}(v_i) + 1\right) \\ 1 &= \frac{1}{2}n - \frac{1}{6}\sum_{i=1}^{n}\text{tval}(v_i) \\ 6 &= 3n - \sum_{i=1}^{n}\text{tval}(v_i). \end{aligned}$$

Hence,

$$\sum_{i=1}^{n}(\text{tval}(v_i) - 3) = -6.$$

□

If T_P is a triangulation with n vertices as described in Lemma 3.1, we call it a *pseudoisland* of size n. We will also consider a single edge to be a pseudoisland of size 2. We will use pseudoislands to show that the presence of an island in a well-triangulated annulus with islands does not materially affect the sum of tile valences along the bottom.

Corollary 3.4 *Let X be a well-triangulated annulus with islands. Suppose X has n top vertices, c bottom vertices, y interior vertices that are part of the island boundaries, and x other interior vertices that are not contained in islands. Let the bottom vertices be labeled $v_1, \ldots, v_c$. Then*

$$\sum_{i=1}^{c}(\text{tval}(v_i) - 3) \geq n + x + y - c.$$

Proof: Replace each island by a pseudoisland of the same size, as shown in Figure 3.21. This being done, we have replaced X by a well-triangulated annulus X' without affecting the valence count along the bottom. This well-triangulated annulus will have $x + y$ interior vertices. Proposition 3.3 gives us

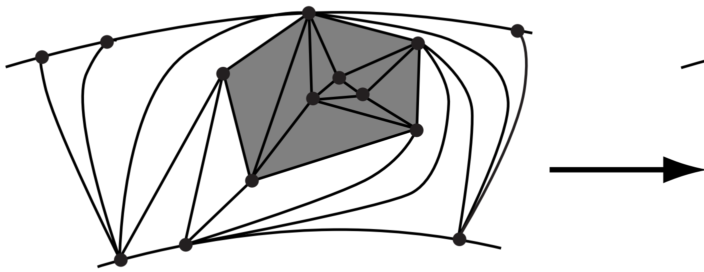

Figure 3.21: Replacing an Island by a Pseudoisland

$$\sum_{i=1}^{c}(\mathrm{tval}(v_i) - 3) \geq n + (x + y) - c.$$

□

Now we must prove some lemmas concerning the structure of islands themselves.

Lemma 3.2 *Let L be an island stationed at k_0. Suppose the island contains a vertex v which is in the set V_{k_0+n} for some $n \geq 2$. Then v is surrounded by a cycle of vertices, all of which lie in V_{k_0+n-1}. Furthermore, all vertices surrounded by this cycle will be in $\bigcup_{j=n}^{\infty} V_{k_0+j}$.*

Proof: We shall prove the lemma by induction on n. Clearly it is true for $n = 2$, since the boundary of L is a cycle of vertices in V_{k_0+1}, and L contains only vertices in $\bigcup_{j=2}^{\infty} V_{k_0+j}$.

So we assume the lemma is true for all vertices v in $L \cap V_{k_0+n}$ for some n. Suppose now that we have a vertex v in L that lies in V_{k_0+n+1}. Then there is an edge connecting v to a vertex w in V_{k_0+n}. Consider the graph $\bar{E}$ formed by all the edges connecting a vertex of L that lies in V_{k_0+n} to another vertex of L that lies in V_{k_0+n}. Let E be the component of this graph that contains w. Since $w \in V_{k_0+n}$, there is a cycle C of vertices in V_{k_0+n-1} surrounding w and hence surrounding E. We also know that all vertices enclosed by C lie in $\bigcup_{j=n}^{\infty} V_{k_0+j}$.

Consider the union of all triangles lying in the region enclosed by C and having at least one vertex on C. This union R will have at least two boundary components; one of them will be C; the vertices on all the other boundary components are enclosed by C and hence cannot lie in V_{k_0+n-1}; however, they are connected by an edge to C. Hence they must lie in V_{k_0+n}. The same argument would hold for all vertices in the interior of R. Therefore, we know that all vertices in R which are not part of C are in V_{k_0+n}.

Now let F be the region enclosed by C. Clearly $E \subseteq F$. Since all vertices in the interior of F are in $\bigcup_{j=n}^{\infty} V_{k_0+j}$, and since the only vertices in F connected to C lie in R, then we may conclude that all vertices of $F \setminus R$ lie in $\bigcup_{j=n+1}^{\infty} V_{k_0+j}$. Thus, we must conclude that all vertices of E lie in R. However, there is an edge from a vertex of E (namely, w) to a vertex of V_{k_0+n+1} (namely, v). That edge cannot be part of R. Therefore we conclude that

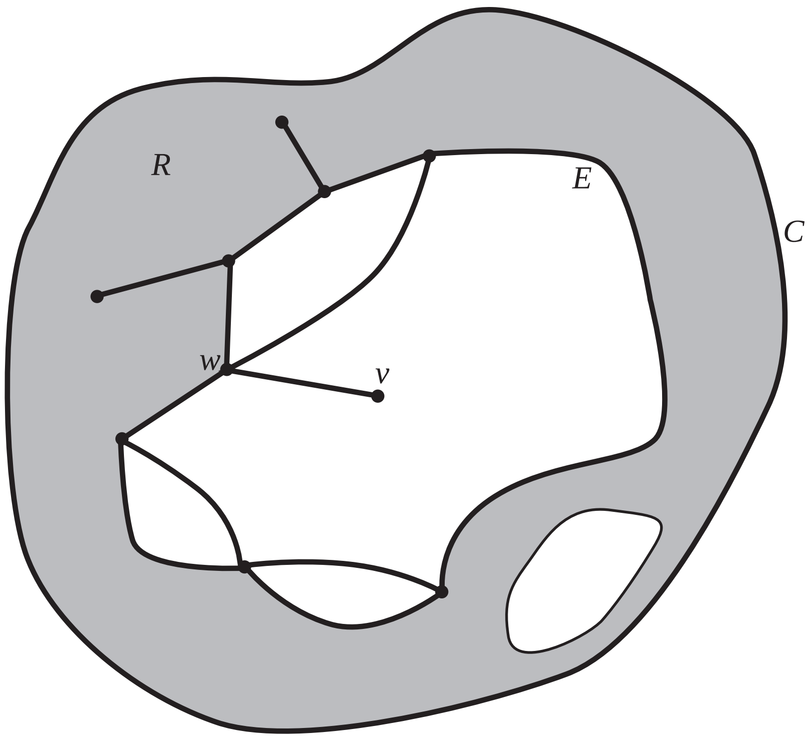

Figure 3.22: Surrounding a Vertex with a Cycle

E contains a boundary component of R. Let E' be this component. The vertex w is in E', and it is connected to v. Now all vertices in R are in V_{k_0+n-1} or V_{k_0+n}, so we know that v is not in R. Since $v \notin R$ we know that E' must enclose v. Thus there is a subset E'' of E' which is a cycle of vertices in V_{k_0+n} and surrounds v.

Consider the vertices enclosed by E''. Since they are all inside C, we know that they are in $\bigcup_{j=n}^{\infty} V_{k_0+j}$. Suppose one of these vertices v is in V_{k_0+n}. Then it would be connected to a vertex x in V_{k_0+n-1}. Now since v was enclosed by E'', either x is also enclosed by E'' or x is on E''. Either way, x is enclosed by C, and hence x cannot be in V_{k_0+n-1}. But this is a contradiction, and hence we have shown that all vertices surrounded by E'' are in $\bigcup_{j=n+1}^{\infty} V_{k_0+j}$. Since we have now proved the lemma for $n+1$, the lemma follows by induction for all values of n. $\square$

Lemma 3.3 puts a bound on the number of vertices that these cycles can share.

Lemma 3.3 *Let L be an island stationed at k_0. Pick $n > 2$ such that there are vertices of V_{k_0+n} in the interior of L. By Lemma 3.2, each of these vertices is surrounded by a cycle of vertices in V_{k_0+n-1} such that all vertices lying inside the cycle are in $\bigcup_{i=k_0+n}^{\infty} V_i$. Then one can choose these cycles such that their interiors are pairwise disjoint, and any two of these cycles can share at most two vertices.*

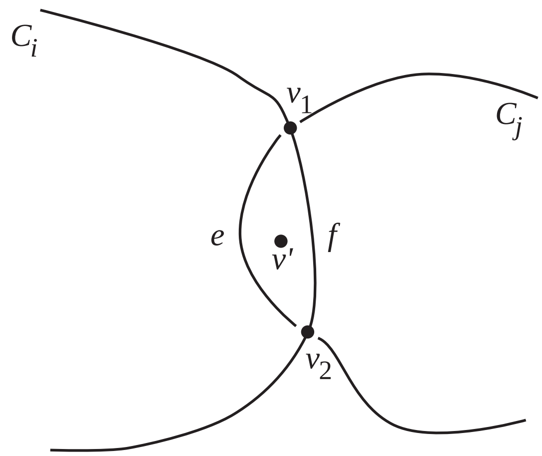

Figure 3.23: Intersecting Cycles

Proof: Let $v_1, \ldots, v_p$ be the vertices of V_{k_0+n} in the interior of L. Each such vertex v_i has a cycle C_i surrounding it, as described in Lemma 3.2. Eliminate duplications and reindex the cycles if necessary to obtain distinct cycles $C_1, \ldots, C_q$.

Consider one of these cycles C_i. The interior of C_i cannot contain any vertex on the boundary of any other cycle C_j, since then it would contain a vertex in V_{k_0+n-1}. Thus, no cycle can enclose a boundary vertex of any other cycle. Now suppose the interior of C_i contains an edge e of the boundary of C_j. (See Figure 3.23.) Then the endpoints v_1 and v_2 of e must lie on the boundary of C_i, since otherwise the interior of C_i would contain a vertex on the boundary of C_j. There cannot be any vertices of ∂C_i lying in C_j, so the segment of the boundary of C_i between v_1 and v_2 must be a single edge f. Since the region bounded by e and f is to be triangulated, there must be a vertex (or vertices) in that region; all of these vertices must be in $\bigcup_{i=k_0+n}^{\infty} V_i$, since they lie in the interior of C_i. Let v' be one such vertex. Then we can define C_i' to have the same boundary as C_i except that the edge f in the boundary will be replaced by e. Likewise, we define C_j' to be C_j with e replaced in the boundary by f. We can define a third cycle, $C_{i,j}'$ to be the cycle formed by e and f, which encloses v'.

Thus, if an edge of one cycle is in the region bounded by another, then we can redefine the cycles to eliminate the offending edge without affecting the other edges. Thus, we may choose our cycles so that none of the boundary of one is contained in the interior of another. Hence their interiors will be disjoint.

Now suppose two cycles C_i and C_j share 3 or more boundary vertices. If $n-1=1$, then C_i and C_j can contain only vertices on the boundary of the island. In this case, C_i and C_j can only intersect along one edge, and thus can share only two vertices. Hence we know $n-1$ must be greater than 1. Let the three shared vertices be v_1, v_2, and v_3; all of them will be in V_{k_0+n-1}. Only two of these vertices can be adjacent to the unbounded component of the complement of $C_i \cup C_j$, as illustrated in Figure 3.24. Assume v_2 is the one that is not.

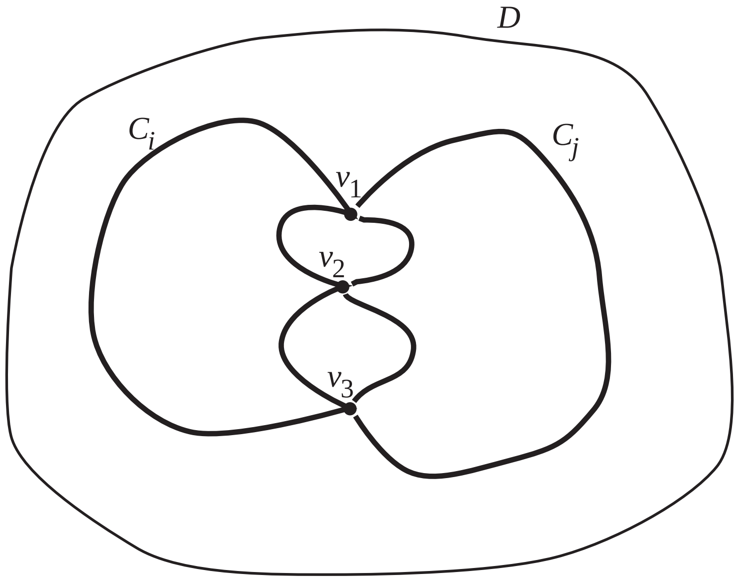

Figure 3.24: Cycles Sharing Vertices

Then by Lemma 3.2, v_2 must be enclosed by a cycle D of vertices of V_{k_0+n-2}, one of which is connected by an edge to v_2. None of these vertices can lie in C_i or C_j; hence this cycle must surround both C_i and C_j. However, it is impossible for v_2 to have an edge joining it to any vertex in the unbounded component of the complement of $C_i \cup C_j$. Therefore we have a contradiction, and two cycles C_i and C_j can share at most 2 boundary vertices.

□

Lemma 3.4 *Let L be an island stationed at k_0. Pick $n > k_0 + 1$ such that there are vertices of V_{k_0+n} in the interior of L. By Lemma 3.2, each of these vertices is surrounded by a cycle of vertices in V_{k_0+n-1} such that all vertices lying inside the cycle are in $\bigcup_{i=k_0+n}^{\infty} V_i$. By Lemma 3.3, we may choose these cycles such that the interiors of the regions enclosed by them are disjoint. Suppose there are p such distinct cycles with disjoint interiors. If we let $x^{L_1}, \ldots, x^{L_p}$ be the number of vertices in each of these cycles, and if we let x^L be the number of vertices in the union of all the cycles, then we have $\sum_{i=1}^{p} x^{L_i} < x^L + 4p$.*

Proof: All we need to do is estimate the how many times the vertices are shared among cycles. Consider a graph H, where each vertex corresponds to one of these p disjoint cycles, and where two vertices are connected by an edge if and only if the corresponding cycles share vertices. We will allow at most one edge to connect any two vertices. Since every vertex on these cycles must have a path connecting it to the boundary of L, none of the vertices in our graph H (corresponding to these cycles) can be in the interior of the complex(es) spanned by H. Let $H_1, \ldots, H_k$ be the components of H. We will also use $H_1, \ldots, H_k$ to refer to the complexes spanned in the plane by each component. We see that H_i will have p_i vertices, with $\sum_{i=1}^{k} p_i = p$. We will embed each complex H_i in a sphere and estimate the number of edges using the Euler characteristic. Suppose H_i has E_i edges of any sort and N_i edges in the interior of the complex. Since none of the vertices can be in the interior of H_i, we note that the number of edges on the boundary of H_i must be at least as great as the number of vertices in the complex minus one; hence, $E_i - N_i \geq p_i - 1$, or $N_i \leq E_i - p_i + 1$.

To compute the number of faces, note that there are no digons in this graph. Thus, each interior edge contributes at most 1/3 to two faces of the complex (because each face has at least 3 sides) and each boundary edge contributes at most 1/3 to one face of the complex. We will also need to add 1 to account for the complement of the complex on the sphere. Thus the total number of faces

$$\begin{aligned}
F &\leq 1 + \frac{1}{3}(E_i - N_i) + \frac{2}{3}N_i \\
&= 1 + \frac{1}{3}E_i + \frac{1}{3}N_i \\
&\leq 1 + \frac{1}{3}E_i + \frac{1}{3}(E_i - p_i + 1) \\
&= 1 + \frac{2}{3}E_i - \frac{1}{3}p_i + \frac{1}{3}.
\end{aligned}$$

Now, since the Euler characteristic of a sphere is 2, we have

$$\begin{aligned}
2 &\leq p_i - E_i + \left(1 + \frac{2}{3}E_i - \frac{1}{3}p_i + \frac{1}{3}\right) \\
2 &\leq \frac{2}{3}p_i - \frac{1}{3}E_i + \frac{4}{3} \\
\frac{1}{3}E_i &\leq \frac{2}{3}p_i - \frac{2}{3} \\
E_i &\leq 2p_i - 2.
\end{aligned}$$

Summing up over all the components, the total number of edges in H is

$$\sum_{i=1}^{k} E_i \leq 2\sum_{i=1}^{k} p_i - 2k = 2p - 2k.$$

However, each edge in this graph represents two of these cycles sharing vertices. Since at most two vertices can be shared by any two cycles (according to Lemma 3.3), the total number of times that vertices are shared among cycles will be less than or equal to twice the number of edges in H. This number is less than or equal to $4p - 4k$. Since k is positive, the total number of times vertices are shared among cycles will be less than $4p$, and the result follows. $\square$

The estimate is crude, but it suffices for our purposes.

Suppose a simplicial 2-complex is homeomorphic to a closed disk, suppose that it has vertices in its interior, and suppose that all interior vertices are connected to the boundary by an edge. We call such a complex an *island core*. Let us prove a fact about island cores before returning our attention to islands themselves.

Lemma 3.5 *Suppose an island core has x boundary vertices and y interior vertices. Let $v_1, \ldots, v_x$ denote the boundary vertices, and let $w_1, \ldots, w_y$ denote the interior vertices. Then*

1. $\sum_{i=1}^{x}(\mathrm{tval}(v_i) - 3) + \sum_{i=1}^{y}(\mathrm{tval}(w_i) - 6) = -6$;
2. $\sum_{i=1}^{x}(\mathrm{tval}(v_i) - 3) > -2 + y - x$.

Proof: We will use the fact that the Euler characteristic of a sphere is 2. Embed the island core into a 2-sphere. Then the number of vertices V is $x + y$. To compute the number of edges, consider that we have x boundary edges. Each vertex v_i has $\mathrm{tval}(v_i) - 1$ incident non-boundary edges; each vertex w_i has $\mathrm{tval}(w_i)$ incident (non-boundary) edges. Thus, summing up the numbers $\mathrm{tval}(v_i) - 1$ and $\mathrm{tval}(w_i)$ will give us exactly twice the number of interior edges. Hence the number of edges E is $x + \frac{1}{2}\sum_{i=1}^{x}(\mathrm{tval}(v_i) - 1) + \frac{1}{2}\sum_{i=1}^{y}\mathrm{tval}(w_i) = \frac{1}{2}x + \frac{1}{2}\left(\sum_{i=1}^{x}\mathrm{tval}(v_i) + \sum_{i=1}^{y}\mathrm{tval}(w_i)\right)$. To compute the number of faces, note that we have one face C consisting of the complement of the island core; that each of the other faces has three edges; that each interior edge borders two faces; and that each boundary edge borders only one face besides C. Thus, the total number of faces F is

$$\begin{aligned} F &= 1 + \frac{2}{3}\left(\frac{1}{2}\sum_{i=1}^{x}(\mathrm{tval}(v_i) - 1) + \frac{1}{2}\sum_{i=1}^{y}\mathrm{tval}(w_i)\right) + \frac{1}{3}x \\ &= 1 + \frac{1}{3}\left(\sum_{i=1}^{x}\mathrm{tval}(v_i) + \sum_{i=1}^{y}\mathrm{tval}(w_i)\right) \end{aligned}$$

So, calculating the Euler characteristic,

$$\begin{aligned} 2 &= V - E + F \\ 2 &= (x+y) - \left(\frac{1}{2}x + \frac{1}{2}\left(\sum_{i=1}^{x}\mathrm{tval}(v_i) + \sum_{i=1}^{y}\mathrm{tval}(v_i)\right)\right) + \\ &\quad \left(1 + \frac{1}{3}\left(\sum_{i=1}^{x}\mathrm{tval}(v_i) + \sum_{i=1}^{y}\mathrm{tval}(w_i)\right)\right) \\ 1 &= \frac{1}{2}x + y - \frac{1}{6}\left(\sum_{i=1}^{x}\mathrm{tval}(v_i) + \sum_{i=1}^{y}\mathrm{tval}(w_i)\right) \\ 6 &= 3x + 6y - \sum_{i=1}^{x}\mathrm{tval}(v_i) - \sum_{i=1}^{y}\mathrm{tval}(w_i). \end{aligned} \tag{3.1}$$

However, when we look at the expression in the first conclusion of this lemma, we find that it is the negative of the right hand side of Equation 3.1. So we may conclude

$$\begin{aligned}\sum_{i=1}^{x}(\mathrm{tval}(v_i)-3)+\sum_{i=1}^{y}(\mathrm{tval}(w_i)-6) &= -3x-6y+\sum_{i=1}^{x}\mathrm{tval}(v_i)+\sum_{i=1}^{y}\mathrm{tval}(w_i)\\ &= -6.\end{aligned}$$

To prove the second statement, we use another Euler characteristic argument. This time, we consider only the y interior vertices and the edges connecting them. The resulting graph might be disconnected, with some of the components consisting of single points. Suppose it has components $C_1,\ldots,C_q$, with the components having $y_1,\ldots,y_q$ vertices each, respectively. Clearly $\sum_{i=1}^{q} y_i = y$. Consider a component C_k embedded in a sphere. The number of vertices $V = y_k$. To express E, we will let r_k be the number of edges connecting a boundary vertex of the island core to a vertex in C_k. Note that $\sum_{i=1}^{q} r_k \leq \sum_{i=1}^{x}(\mathrm{tval}(v_i)-1)$. We will need to subtract r_k from the sum of the edges emanating from the interior vertices in order to include only the edges we want. So we have

$$E=\frac{1}{2}\left(\sum_{w_i\in C_k}\mathrm{tval}(w_i)-r_k\right).$$

Now, each of these vertices is connected to the boundary of the island core, and therefore all of the vertices in C_k are on the boundary of the complex spanned by C_k. Assume for the moment that the complex spanned by C_k contains a 2-cell. Then the number of edges in the interior of this complex cannot be any greater than the total number of edges minus the number of vertices. Denote the number of interior edges by N, and denote the total number of edges by E. Then each edge in in the interior will contribute $1/3$ to two faces, and each edge not in the interior will contribute $1/3$ to one face. Finally, one face will come from the complement of the triangulation on the sphere. Hence,

$$\begin{aligned}F &= \frac{2}{3}N+\frac{1}{3}(E-N)+1\\ &= \frac{1}{3}N+\frac{1}{3}E+1\\ &\leq \frac{1}{3}(E-y_k)+\frac{1}{3}E+1\\ &= \frac{2}{3}E-\frac{1}{3}y_k+1\\ &= \frac{2}{3}\left(\frac{1}{2}\left(\sum_{w_i\in C_k}\mathrm{tval}(w_i)-r_k\right)\right)-\frac{1}{3}y_k+1\\ &= \frac{1}{3}\sum_{w_i\in C_k}\mathrm{tval}(w_i)-\frac{1}{3}r_k-\frac{1}{3}y_k+1.\end{aligned}$$

Once again,

$$\begin{aligned}
2 &= V - E + F \\
2 &\leq y_k - \frac{1}{2}\left(\sum_{w_i \in C_k} \mathrm{tval}(w_i) - r_k\right) + \left(\frac{1}{3}\sum_{w_i \in C_k} \mathrm{tval}(w_i) - \frac{1}{3}r_k - \frac{1}{3}y_k + 1\right) \\
1 &\leq \frac{2}{3}y_k + \frac{1}{6}r_k - \frac{1}{6}\sum_{w_i \in C_k} \mathrm{tval}(w_i) \\
6 &\leq 4y_k + r_k - \sum_{w_i \in C_k} \mathrm{tval}(w_i) \\
\sum_{w_i \in C_k} \mathrm{tval}(w_i) &\leq 4y_k + r_k - 6 \\
&< 4y_k + r_k - 2.
\end{aligned}$$

On the other hand, if there are no 2-cells in the complex spanned by C_k, then the number of faces will be 1. Doing a similar calculation, we find

$$\begin{aligned}
2 &= y_k - \frac{1}{2}\left(\sum_{w_i \in C_k} \mathrm{tval}(w_i) - r_k\right) + 1 \\
4 &= 2y_k - \sum_{w_i \in C_k} \mathrm{tval}(w_i) + r_k + 2 \\
\sum_{w_i \in C_k} \mathrm{tval}(w_i) &= 2y_k + r_k - 2 \\
&< 4y_k + r_k - 2,
\end{aligned}$$

since $y_k > 0$. In each case, we have the same estimate.

Summing over all the components C_k, we have

$$\begin{aligned}
\sum_{k=1}^{q} \sum_{w_i \in C_k} \mathrm{tval}(w_i) &< \sum_{k=1}^{q}(4y_k + r_k - 2) \\
\sum_{i=1}^{y} \mathrm{tval}(w_i) &< 4y + \sum_{i=1}^{x}(\mathrm{tval}(v_i) - 1) - 2q.
\end{aligned}$$

From Equation 3.1, we know that

$$\begin{aligned}
\sum_{i=1}^{x} \mathrm{tval}(v_i) - 3x &= -6 + 6y - \sum_{i=1}^{y} \mathrm{tval}(w_i) \\
&> -6 + 6y - \left(4y + \sum_{i=1}^{x}(\mathrm{tval}(v_i) - 1) - 2q\right) \\
&= -6 + 2y - \sum_{i=1}^{x} \mathrm{tval}(v_i) + x + 2q.
\end{aligned}$$

Thus,

$$\begin{aligned}
2\sum_{i=1}^{x}\text{tval}(v_i) - 6x &> -6 + 2y - 2x + 2q \\
\sum_{i=1}^{x}\text{tval}(v_i) - 3x &> -3 + y - x + q \\
\sum_{i=1}^{x}(\text{tval}(v_i) - 3) &> -3 + y - x + q \\
&\geq -2 + y - x,
\end{aligned}$$

since $q \geq 1$. $\square$

With these lemmas in hand, we can determine the total contribution of an island to the tile valence sums. If L is an island stationed at k_0, then for all $k \geq 0$, we will let x_k^L be the number of vertices contained in L that belong to the set V_{k_0+1+k}. For any k greater than the depth of the island, of course, x_k^L will be 0.

Proposition 3.12 *Let L be an island stationed at k_0 with a depth of n. Define a_k^L for all positive integers k as follows:*

$$a_k^L = \begin{cases} 0 & \text{if } k \leq k_0 \\ 6 + \sum_{v \in V_{k_0+1}}(\text{tval}(v) - 3) & \text{if } k = k_0 + 1 \\ a_{k_0+1}^L + \sum_{j=2}^{k=k_0}\sum_{v \in V_{k_0+j}}(\text{tval}(v) - 6) & \text{if } k \geq k_0 + 2, \end{cases}$$

where the sums are taken only over the appropriate vertices in L. Then the following statements are true:

1. $a_{k_0+k}^L = 0$ *if* $-k_0 < k \leq 0$;
2. $a_{k_0+k}^L > x_k^L - x_{k-1}^L$ *if* $1 \leq k \leq n$;
3. $a_{k_0+k}^L = 0$ *if* $k > n$.

A word of explanation might be in order before beginning the proof of the proposition. The values a_k^L will be less than or equal to the contribution of the island to the tile valence sums. This should be clear for values of $k \leq k_0$; however, the definition for $k = k_0 + 1$ might require some explanation. Replace L by a pseudoisland of the same size. If the size of the island is greater than 2, then, by Lemma 3.1, $\sum(\text{tval}(v) - 3) = -6$, where the sum is taken over all vertices on the boundary of the pseudoisland, and where the only triangles we count in $\text{tval}(v)$ are the triangles contained in the pseudoisland. Likewise, if the size of the island is 2, then the pseudoisland is a single edge. If you sum over these two vertices and count

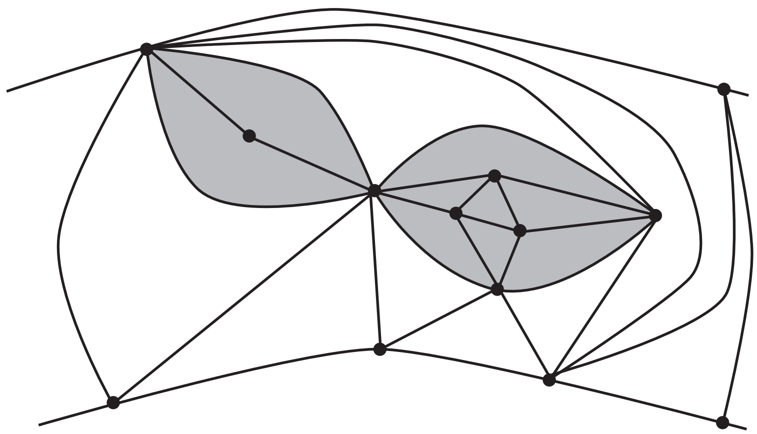

Figure 3.25: Two Islands Sharing a Vertex

only triangles that are subsets of the edge (there are none, of course), you get trivially that $\sum(\mathrm{tval}(v) - 3) = -6$. Thus, the interior of the pseudoisland contributes -6 to the tile valence sum. To determine the difference made by the presence of the island, then, one must calculate

$$\sum_{\substack{v \in V_{k_0+1} \\ v \in L}} (\mathrm{tval}(v) - 3) - (-6).$$

This fact explains our choice of definition for $a^L_{k_0+1}$.

Note also that if two distinct islands share a boundary vertex, then $a^L_{k_0+1}$ and all succeeding values of a^L_k will be 3 smaller than the actual contribution of that vertex because we subtract 3 when we count it as part of each island and also 3 when we consider the edges not contained in either island. This gives us a total of 9 subtracted from the tile valence of that vertex, instead of 6. However, all we will care about is a lower bound on the contributions of the islands, and a^L_k gives us that.

Proof: The first conclusion is simply the definition of a^L_k for $k \leq k_0$. To prove the rest, we will perform induction on the depth of the island. First suppose that the depth is 1—i.e., the only interior vertices of the island are in V_{k_0+2}. Since the boundary vertices are the only vertices of L which are in V_{k_0+1}, each of the interior vertices must be connected to a boundary vertex by an edge. Thus the island is an island core. Label the boundary vertices $v_1, \dots, v_{x^L_0}$, and label the interior vertices $w_1, \dots, w_{x^L_1}$.

$$\begin{aligned} a^L_{k_0+1} &= 6 + \sum_{i=1}^{x^L_0} (\mathrm{tval}(v_i) - 3) \\ &> 6 - 2 + x^L_1 - x^L_0 \quad \text{by Lemma 3.5} \\ &= 4 + x^L_1 - x^L_0 \\ &> x^L_1 - x^L_0. \end{aligned}$$

And also,

$$\begin{aligned}
a^L_{k_0+2} &= a^L_{k_0+1} + \sum_{i=1}^{x_1^L}(\text{tval}(v_i) - 6) \\
&= 6 + \sum_{i=1}^{x_0^L}(\text{tval}(v_i) - 3) + \sum_{i=1}^{x_1^L}(\text{tval}(v_i) - 6) \\
&= 6 - 6 \quad \text{by Lemma 3.5} \\
&= 0.
\end{aligned}$$

Clearly $a^L_{k_0+k}$ will be zero for $k > 2$.

So we have proved the proposition for islands of depth 1; now let us suppose we have proved it for islands of depth n. Let L have a depth of $n+1$. By Lemma 3.2, all vertices in the island that belong to V_{k_0+n+1} are encircled by a cycle of vertices in the island that belong to V_{k_0+n}; furthermore, by Lemma 3.3, we may choose these cycles such that the interiors of the regions bounded by them are disjoint. Each of the vertices enclosed in these cycles must be in V_{k_0+n+1} and hence have an edge connecting it to the cycle enclosing it; thus, each of these cycles is the boundary of an island core. Label these island cores $L_1, \ldots, L_p$. Replace each of them with a pseudoisland of the same size. We have now reduced the depth of the island by 1; call this reduced island $\bar{L}$. We know the following about $\bar{L}$:

$$\begin{aligned}
a^{\bar{L}}_{k_0+1} &> x_1^L - x_0^L \\
&\vdots \\
a^{\bar{L}}_{k_0+n} &> x_n^L - x_{n-1}^L \\
a^{\bar{L}}_{k_0+n+1} &= 0.
\end{aligned}$$

Replacing the island cores by the pseudoislands did not affect any vertices in $V_{k_0+1}, \ldots, V_{k_0+n}$, so we may say

$$\begin{aligned}
a^L_{k_0+1} &> x_1^L - x_0^L \\
&\vdots \\
a^L_{k_0+n} &> x_n^L - x_{n-1}^L.
\end{aligned}$$

Now the difference between $a^{\bar{L}}_{k_0+n+1}$ and $a^L_{k_0+n+1}$ will be the difference between the tile valence sums on the interior of the pseudoislands and the tile valence sums on the interior of the island cores $L_1, \ldots, L_p$. We will commit a slight abuse of terminology and notation by considering each of these island cores L_i as an island stationed at k_0+n. Then the difference between the tile valence sum on the interior of L_i and on the interior of the pseudoisland replacing it will be $a^{L_i}_{(k_0+n)+1}$. By Lemma 3.3, the interiors of the island cores L_i are disjoint; hence, $\sum_{i=1}^p x^{L_i}_{n+1} = x^L_{n+1}$. By Lemma 3.4 we can see that $\sum_{i=1}^p x^{L_i}_n < x^L_n + 4p$.

Putting this all together,

$$\begin{aligned}
a^L_{k_0+n+1} &= a^{\bar{L}}_{k_0+n+1} + \sum_{i=1}^{p} a^{L_i}_{(k_0+n)+1} \\
&> 0 + \sum_{i=1}^{p}(4 + x^{L_i}_{n+1} - x^{L_i}_{n}) \quad \text{by Lemma 3.5 and the proof above} \\
&> 4p + x^L_{n+1} - (x^L_n + 4p) \\
&= x^L_{n+1} - x^L_n.
\end{aligned}$$

Since L only has a depth of $n+1$, $a^L_{k_0+k}$ will be zero for all values of k greater than $n+1$. Thus we have proved the proposition. $\square$

We shall later have occasion to sum the numbers a^L_k. Hence we state the following corollary:

Corollary 3.5 *Let L be an island stationed at k_0 with a depth of n. Then for all positive values of k,*

$$\sum_{j=1}^{k} a^L_j > -x^L_0.$$

Proof: By Proposition 3.12, we know that

$$\begin{array}{lcll}
a^L_k &=& 0 & \text{for } 0 < k \leq k_0 \\
a^L_k &>& x^L_{k-k_0} - x^L_{k-k_0-1} & \text{for } k_0+1 \leq k \leq k_0+n \\
a^L_k &=& 0 & \text{for } k_0+n < k
\end{array}$$

If $0 < k \leq k_0$, then clearly $\sum_{j=1}^{k} a^L_j = 0 > -x^L_0$. And,

$$\sum_{j=1}^{k_0+1} a^L_j = \sum_{j=1}^{k_0} a^L_j + a^L_{k_0+1} > 0 + (x^L_1 - x^L_0).$$

We now use induction. Suppose we have shown that, for some k, such that $k_0+1 \leq k < k_0+n$, we have $\sum_{j=1}^{k} a^L_j > x^L_{k-k_0} - x^L_0$. Then for $k+1$, we have

$$\begin{aligned}
\sum_{j=1}^{k+1} a^L_j &= \sum_{j=1}^{k} a^L_j + a^L_{k+1} \\
&> (x^L_{k-k_0} - x^L_0) + (x^L_{(k+1)-k_0} - x^L_{(k+1)-k_0-1}) \\
&= x^L_{(k+1)-k_0} - x^L_0.
\end{aligned}$$

This induction shows that for all k such that $k_0+1 \leq k \leq k_0+n$, we have $\sum_{j=1}^{k} a^L_j > x^L_{k-k_0} - x^L_0$, and hence $\sum_{j=1}^{k} a^L_j > -x^L_0$.

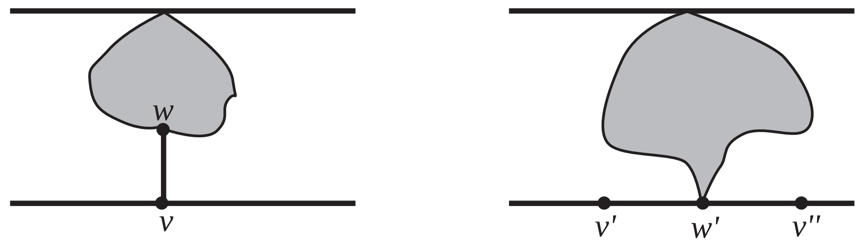

Figure 3.26: Converting an Island Vertex to a Bottom Vertex

Finally, it is clear from Proposition 3.12 that for any $k > k_0 + n$,

$$\sum_{j=1}^{k} a_j^L = \sum_{j=1}^{k_0+n} a_j^L > -x_0^L.$$

□

Now we want to show that the presence of islands in a well-triangulated annulus with islands cannot decrease its modulus.

Proposition 3.13 *Let X be a well-triangulated annulus with islands. Suppose X has n top vertices, and let M be the fat flow modulus of X. Then*

$$M \geq \frac{3}{2n}.$$

Proof: Suppose w is a vertex on the boundary of an island, and suppose w does not lie on the top of X. We know that w is connected to a bottom vertex v by an edge. As we have done before, replace v by two bottom vertices v' and v'', and replace w by a bottom vertex w', whose neighbors on the bottom to either side are v' and v''. We have succeeded in replacing an interior vertex and a bottom vertex by three bottom vertices. (See Figure 3.26.)

Perform this operation on every vertex that does not lie on the top of X but which is on the boundary of an island. The result will be a well-triangulated annulus, possibly still with islands; call it X_0'. Note that the number of top vertices of X_0' is still n. Note that any island in X has become either a region with all of its boundary vertices on the bottom, or a region with some boundary vertices on the top and the rest on the bottom. We shall consider these two cases separately.

Let M_0 be the fat flow modulus of X_0'. How does M_0 compare to M, the modulus of X? Let w_0 be an optimal fat flow weight function for X_0'. Since every tile in X_0' corresponds to a tile in X, we can apply the function w_0 to the original tiling X. By the correspondence of the tiles, $A_{w_0}(X) = A_{w_0}(X_0')$. Note that every fat flow in X_0' is a subset of a fat flow in X,

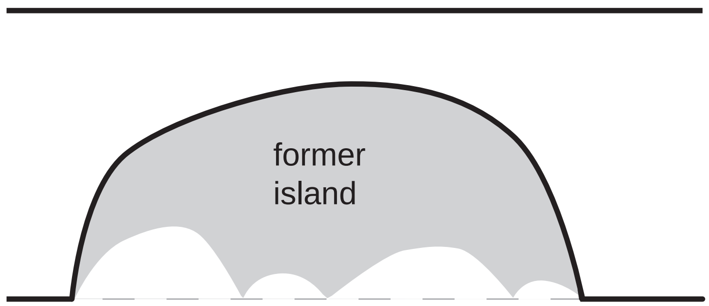

Figure 3.27: Removing an Island When No Island Vertices Were on the Top

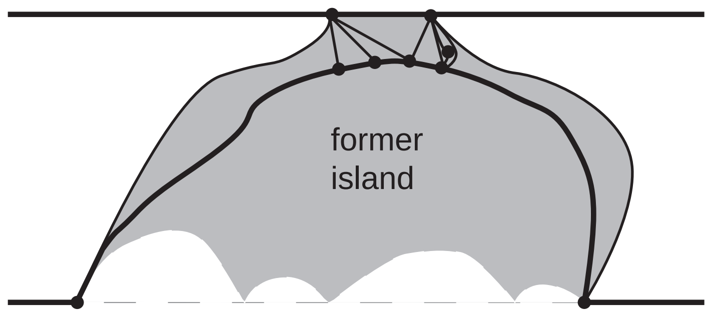

Figure 3.28: Removing an Island When Some Island Vertices Were on the Top

and every fat flow in X contains a fat flow in X_0'. Therefore, $H_{w_0,f}(X) \geq H_{w_0,f}(X_0')$. Thus, we have

$$M \geq \frac{H^2_{w_0,f}(X)}{A_{w_0}(X)} \geq \frac{H^2_{w_0,f}(X_0')}{A_{w_0}(X_0')} = M_0.$$

Now let us consider the two cases. In the the first case, all boundary vertices are on the bottom. Since the region of the former island is still bounded by a cycle of vertices, the left-most and right-most boundary vertices must be connected by a single edge lying in the interior of X_0'. We will remove all tiles lying between the bottom and this edge. (See Figure 3.27.) Note that this operation gets rid of all tiles that originally belonged to the island. In general, if the tiling was called X_n' before the operation, then we shall call the tiling resulting from the operation X_{n+1}'. Note that X_{n+1}' has the same number of top vertices as X_n'.

How does the modulus of X_{n+1}' compare to the modulus of X_n'? Let w_{n+1} be an optimal weight function on X_{n+1}'. Extend w_{n+1} to X_n' by letting $w_{n+1}(t) = 0$ for all tiles in $X_n' \setminus X_{n+1}'$. Clearly $H_{w_{n+1},f}(X_n') = H_{w_{n+1},f}(X_{n+1}')$, and $A_{w_{n+1}}(X_n') = A_{w_{n+1}}(X_{n+1}')$. Therefore, if M_{n+1} is the fat flow modulus of X_{n+1}' and M_n is the fat flow modulus of X_n', we have

$$M_{n+1} = \frac{H^2_{w_{n+1},f}(X_{n+1}')}{A_{w_{n+1}}(X_{n+1}')} = \frac{H^2_{w_{n+1},f}(X_n')}{A_{w_{n+1}}(X_n')} \leq M_n.$$

Consider the second case, in which some boundary vertices are on the top. Suppose that $t_1, \ldots, t_n$ are the tiles in the island that are adjacent to the top boundary vertices. We will remove all tiles lying between the tiles $t_1, \ldots, t_n$ and the bottom. (See Figure 3.28.) Note that islands might remain in place after this operation, coming from tiles inside the original island. As before, if the tiling before the operation was X'_n, then the tiling after the operation will be called X'_{n+1}. Note that X'_{n+1} has the same number of top vertices as A'_n. The same argument as in the previous case shows us that $M_{n+1} \leq M_n$.

Since there can only be a finite number of tiles in X, we may perform this operation until we obtain a well-triangulated annulus (without islands) X'_q ($q \geq 0$) having the same number of top vertices n as does X. By Proposition 3.9, $M_q \geq 3/(2n)$. Hence, we may conclude that

$$M \geq M_0 \geq M_1 \geq \cdots \geq M_q \geq \frac{3}{2n}.$$

□

At last we are in a position to see that islands will not affect the parabolicity of the disk triangulation graph, thus proving Theorem 3.3.

Theorem 3.3 *Let G be a disk triangulation graph of bounded valence. Let v_0 be a vertex of G. For all positive integers n, let V_n be the set of all vertices v such that $d(v_0, v) = n$, where $d(v_0, v)$ is the minimum number of edges in a path connecting v_0 to v. Let*

$$a_n = \sum_{\substack{v \in V_k \\ k \leq n}} (\mathrm{val}(v) - 6).$$

If the sequence $\{a_n\}$ is bounded, then G is parabolic.

Proof: By Proposition 3.11, we may decompose G into the union of sets $\bar{A}_0, \bar{A}_1, \ldots$. By this proposition, $\bar{A}_0$ will be a closed topological disk, and each set $\bar{A}_i$ for $i \geq 1$ will be a well-triangulated annulus with islands such that the bottom vertices are in V_i and the top vertices are in V_{i+1}. Let K be the maximum number such that there is a vertex of V_K in $\bar{A}_0$. For $k = 1, \ldots, K$, let

$$J_k = \sum_{j=1}^{k} \sum_{\substack{v \in V_j \\ v \in \bar{A}_0}} (\mathrm{val}(v) - 6).$$

For $k > K$, let $J_k = J_K$. Finally, choose a non-negative number J such that $J \geq -J_k$ for all k.

Let $\bar{G}$ be the graph G altered by replacing each island with a pseudoisland of the same size. Then each (altered) annulus $\bar{A}_j$ in $\bar{G}$ is a well-triangulated annulus such that all bottom vertices are in V_j and all top vertices are in V_{j+1}. Suppose $\bar{A}_j$ in $\bar{G}$ contains c_j bottom vertices, n_j top vertices, and x_j interior vertices. Let the set of bottom vertices be B_j, the set of top vertices be T_j, and the set of interior vertices be I_j. Then we define $b_j = \sum_{B_j}(\mathrm{tval}(v) - 3)$; we define $t_j = \sum_{T_j}(\mathrm{tval}(v) - 3)$; and we define $i_j = \sum_{I_j}(\mathrm{tval}(v) - 6)$. The tile valence counts on B_j and T_j, of course, take into account only tiles in $\bar{A}_j$. Recall that J_k eventually takes care of all vertices in $\bar{A}_0$, including all vertices on the bottom of $\bar{A}_1$. Now these sets $\bar{A}_i$ fit together, so if we sum the valences for vertices on the border between $\bar{A}_{k-1}$ and $\bar{A}_k$, we obtain

$$\sum(\mathrm{val}(v) - 6) = t_{k-1} + b_k.$$

Now in the graph $\bar{G}$, we shall define

$$\begin{aligned}
\bar{a}_l &= \sum_{\substack{v \in V_j \\ j \le l}} (\mathrm{val}(v) - 6) \\
&= J_l + \sum_{j=2}^{l} \sum_{\substack{v \in V_j \\ v \notin \bar{A}_0}} (\mathrm{val}(v) - 6) \\
&= J_l + \sum_{j=2}^{l} (b_j + i_{j-1} + t_{j-1}) \\
&= J_l + \sum_{j=2}^{l} b_j + \sum_{j=2}^{l} (i_{j-1} + t_{j-1}) \\
&= J_l - b_1 + \sum_{j=1}^{l} b_j + \sum_{j=1}^{l-1} (i_j + t_j) \\
&= J_l - b_1 + b_l \quad \text{by Proposition 3.4.}
\end{aligned}$$

Now we take into account the islands. Fix k. Suppose $L_1, \ldots, L_p$ are the islands containing vertices in V_k. For each island L_i stationed at j, let s^{L_i} be the number of vertices on the boundary of L_i that are also on the top of $\bar{A}_j$. For $l \le k$, we shall let s_l be the sum of the values s^{L_i} for which L_i has station less than l. So in calculating a_l, s_l represents the number of times islands have vertices on the top of the annulus within which they reside. For each island, let $a_l^{L_i}$ be defined for the island L_i as in Proposition 3.12. We will gather up the contributions of all the islands by letting $\mathcal{I}_l = \sum_{i=1}^{p} a_l^{L_i} + 3s_l$. The reason we add $3s_l$ to the sum is that for every top vertex on the boundary of an island stationed at j, we have already subtracted 3 from the tile valence of that vertex when considering it as part of the top of $\bar{A}_j$. We should not subtract another 3 when considering it as part of an island; however, we did subtract 3 from its tile valence when computing $a_l^{L_i}$. Hence we must add 3 back to the sum for every island containing that vertex.

Now for any $l \le k$, $a_l \ge \bar{a}_l + \mathcal{I}_l$. So for $l = 1, \ldots, k$, we have

$$a_l \ \ge \ \bar{a}_l + \mathcal{I}_l$$

$$= \quad (J_l - b_1) + b_l + \mathcal{I}_l.$$

By hypothesis, the sequence $\{a_l\}$ is bounded. Thus there exists a B such that $a_l < B$ for all l; hence, by Proposition 3.3, we know that

$$(J_l - b_1) + n_l + x_l - c_l + \mathcal{I}_l \leq (J_l - b_1) + b_l + \mathcal{I}_l \leq a_l < B,$$

and hence

$$\begin{aligned} n_l &< B + c_l - x_l - \mathcal{I}_l - J_l + b_1 \\ &\leq (B + J + b_1) + c_l - (x_l + \mathcal{I}_l). \end{aligned} \tag{3.2}$$

for $l = 1, \ldots, k$.

Let's do some more induction. We know

$$n_1 \leq (B + J + b_1) + c_1 - (x_1 + \mathcal{I}_1).$$

Suppose for some $l < k$,

$$n_l \leq l(B + J + b_1) + c_1 - \left(\sum_{j=1}^{l} x_j + \sum_{j=1}^{l} \mathcal{I}_j \right).$$

Then by Equation 3.2,

$$\begin{aligned} n_{l+1} &\leq (B + J + b_1) + c_{l+1} - x_{l+1} - \mathcal{I}_{l+1} \\ &= (B + J + b_1) + n_l - x_{l+1} - \mathcal{I}_{l+1} \quad \text{since } c_{l+1} = n_l \\ &\leq (B + J + b_1) + l(B + J + b_1) + c_1 - \left(\sum_{j=1}^{l} x_j + \sum_{j=1}^{l} \mathcal{I}_j \right) - x_{l+1} - \mathcal{I}_{l+1} \\ &\leq (l+1)(B + J + b_1) + c_1 - \left(\sum_{j=1}^{l+1} x_j + \sum_{j=1}^{l+1} \mathcal{I}_j \right). \end{aligned}$$

Thus, for k in particular,

$$n_k \leq k(B + J + b_1) + c_1 - \left(\sum_{j=1}^{k} x_j + \sum_{j=1}^{k} \mathcal{I}_j \right).$$

Now,

$$\begin{aligned} \sum_{j=1}^{k} \mathcal{I}_j &= \sum_{j=1}^{k} \left(\sum_{i=1}^{p} a_j^{L_i} + 3 s_j \right) \\ &= \sum_{i=1}^{p} \sum_{j=1}^{k} a_j^{L_i} + 3 \sum_{j=1}^{k} s_j \\ &\geq -\sum_{i=1}^{p} x_0^{L_i} + 3 \sum_{j=1}^{k} s_j \end{aligned}$$

by Corollary 3.5, with equality only if $p = 0$. Thus,

$$-\sum_{j=1}^{k} \mathcal{I}_j \leq \sum_{i=1}^{p} x_0^{L_i} - 3\sum_{j=1}^{k} s_j \leq \sum_{i=1}^{p} x_0^{L_i} - s_k.$$

However, since all of the vertices on the boundary of an island (save the ones at the top, which are counted in s_k) are interior vertices, we have

$$-\sum_{j=1}^{k} \mathcal{I}_j \leq \sum_{i=1}^{p} x_0^{L_i} - s_k \leq \sum_{i=1}^{k} x_j.$$

So

$$\sum_{j=1}^{k} x_j + \sum_{j=1}^{k} \mathcal{I}_j \geq 0,$$

and we conclude that

$$n_k \leq k(B + J + b_1) + c_1.$$

This inequality is true for all values of k. However, each $\bar{A}_k$ is a well-triangulated annulus with islands. Therefore, by Proposition 3.13, if M_k is the modulus of $\bar{A}_k$, then

$$M_k \geq \frac{3}{2n_k} \geq \frac{3}{2(k(B + J + b_1) + c_1)}.$$

Let M be the fat flow modulus of G. Then by the Layer Theorem,

$$\begin{aligned} M &\geq \sum_{k=1}^{\infty} M_k \\ &\geq \frac{3}{2}\sum_{k=1}^{\infty} \frac{1}{k(B + J + b_1) + c_1} \\ &= \infty. \end{aligned}$$

□

Bibliography

[1] Benjamini, I., and O. Schramm. “Random Walks and Harmonic Functions on Infinite Planar Graphs Using Square Tilings.” Preprint, n.d.

[2] Cannon, J. W. “The Combinatorial Riemann Mapping Theorem.” *Acta Mathematica*. 173 (1994): 155–234.

[3] Cannon, J. W., W. J. Floyd, and W. R. Parry. “Conformal Modulus: the Graph Paper Invariant, or The Conformal Shape of an Algorithm.” Preprint, n.d.

[4] ———. “On the Conformal Invariance of Tiling Systems.” Preprint, 1996.

[5] ———. “Squaring Rectangles: The Finite Riemann Mapping Theorem.” *Contemporary Mathematics*. 169 (1994): 133–212.

[6] ———. “The Length-Area Method and Discrete Riemann Mappings.” Preprint, based on a talk given by J. W. Cannon at the Ahlfors Celebration, Stanford University, Stanford, Calif., 1997.

[7] Cannon, J. W., and E. L. Swenson. “Recognizing Constant Curvature Discrete Groups in Dimension 3.” *Transactions. American Mathematical Society.* 350 (1998): 809–849.

[8] Conway, J. B. *A Course in Functional Analysis*. 2nd ed. New York: Springer, 1990.

[9] Gromov, M. “Hyperbolic Groups.” In *Essays in Group Theory*. Ed. S. Gersten. MSRI Publication 8. New York: Springer, 1987.

[10] He, Z.-X., and O. Schramm. “Hyperbolic and Parabolic Packings.” *Discrete Computational Geometry*. 14 (1995): 123–149.

[11] Rodin, B., and D. Sullivan. “The Convergence of Circle Packings to the Riemann Mapping.” *Journal of Differential Geometry*. 26 (1987): 349–360.

[12] Royden, H. L. *Real Analysis*. 3rd ed. New York: Macmillan, 1988.

[13] Schramm, O. “Square Tilings with Prescribed Combinatorics.” *Israel Journal of Mathematics*. 84 (1993): 97–118.

[14] Swenson, E. L. "Negatively Curved Groups and Related Topics." Ph.D. diss., Brigham Young University, 1993.

Vita

Andrew S. Repp was born on July 17, 1970, in Plymouth, Indiana. In 1976, his family moved to St. Louis. In 1985, they moved to Greenville, South Carolina, where Andrew began highschool. In 1989, he entered Bob Jones University, a non-denominational Christian school, where he majored in mathematics and minored in German and computer science. He graduated from Bob Jones in 1994. That same year, he began graduate work at Virginia Polytechnic Institute and State University. He completed a Master's Degree in mathematics in 1996 and in turn entered the doctoral program at Virginia Tech. Andrew's work experience includes computer programming for the Bob Jones University Data Processing Department and the teaching of college mathematics as a Graduate Teaching Assistant at Virginia Tech. His parents are Russell and Gloria Repp; he has a brother and sister-in-law, Jonathan and Krista Repp; and he has one sister, Janelle Repp.