Ellipse Packing

by

Hannah Cairns

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF MATHEMATICS

in

The Faculty of Graduate Studies

(Mathematics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

October 2013

Table of Contents

Appendices

List of Figures

Acknowledgements

I will take this opportunity to express my gratitude to the University of British Columbia for giving me the opportunity to consider this question, and to Asaf for suggesting it to me.

Dedication

I would like to thank everyone in sight.

Statement of Co-Authorship

The question in this paper was asked me by my advisor, Asaf Nachmias, and I had some fruitful discussions with him about it.

Chapter 1

Introduction

Circle packings are a subject of recent interest. They are a window into a discrete version of complex analysis. For example, given a circle packing on a space, the graph distance can be thought of as a discrete version of a conformal metric.

The starting point of the theory is the Koebe-Andreev-Thurston theorem, which says that, if G is a planar triangulation, there is a unique circle packing with tangency graph G in the unit circle, up to reflections and Möbius transformations.

This allows one to take a wide variety of graphs and put a geometric structure on them.

In 1987, Rodin and Sullivan [8] proved a result relating this 'discrete' complex analysis to the usual kind. They constructed a series of maps between two circle packings in two regions, and showed that the sequence converges to a conformal mapping as the circles in the packing become smaller. The procedure that they used was suggested by William Thurston in a talk in 1985, so we refer to it as 'Thurston's procedure.'

There is a natural temptation to generalize this procedure. One way to do this is to use ellipses instead of circles in one of the two regions. The result should no longer be conformal: it should have some special differential structure depending on how we choose the shape of the ellipses.

We prove that this is indeed the case. We are not able to prove that the maps converge, but we do show that any subsequential limit must have a certain differential structure. It must solve the *bar-Beltrami equation*

$$f_{\bar{z}} = \nu(z)\overline{f_z},$$

where ν is a complex-valued function that specifies the shape of the ellipses. This is a variation on the usual Beltrami equation, $f_{\bar{z}} = \mu f_z$.

1.1 The structure of the essay

We will explain the structure of the essay and the function of each chapter.

In Chapters 2 and 3, we explain how to modify Thurston's procedure to work on ellipses. In Chapter 4, we will talk about the image of a circle under a linear map, and make an elementary approach to the bar-Beltrami equation.

The maps in our procedure are not even approximately conformal, but they are 'K-quasiconformal' for a certain constant K. We will define a quasiconformal map, and recall some standard convergence theorems, and then we will recall and slightly extend the rigidity theorem from Rodin and Sullivan's paper.

Finally, in Chapter 7, we will tie everything together and show that, when we carry out our revision of Thurston's procedure, any subsequential limit solves the bar-Beltrami equation with the prescribed function ν.

In Chapter A, we find bounds on the Beltrami coefficient of a linear map between two triangles. We then examine in detail the piecewise linear maps that the procedure generates and make sure that there is a uniform constant K so that all maps are K-quasiconformal.

For all of this to work, we need an ellipse packing theorem, for which we go to Oded Schramm's paper [9]. Unfortunately, the theorem in that paper is not precisely suitable for this procedure. In Chapter B, we do the work to prove that Schramm's theorem implies a suitable ellipse packing theorem, namely Theorem 3.1.1.

Chapter 2

Circle packing

2.1 The definition of a packing

A *packing* is a collection of closed sets E_α in the complex plane so that the interiors, $E_\alpha \setminus \partial E_\alpha$, are pairwise disjoint. We allow the sets to touch each other at the boundary, and if two sets do touch, we say they are *tangent*.

A packing is a *circle packing* if every set is a nondegenerate disc. It is an *ellipse packing* if every set is a nondegenerate ellipse (including the interior).

The *tangency graph* of a packing is the graph with one vertex for each set, and an edge between two vertices if and only if their corresponding sets are tangent.

Let G be a planar graph with a distinguished outside face, and let X be a domain. A *packing of a graph* G is a packing whose sets are labeled by the vertices of G and whose tangency graph is G.

It is a *packing in the domain* X if every set E_α is contained in $\overline{X}$, and the sets on the outside face, and only those sets, intersect ∂X.

Here is an example of a circle packing in the unit disc:

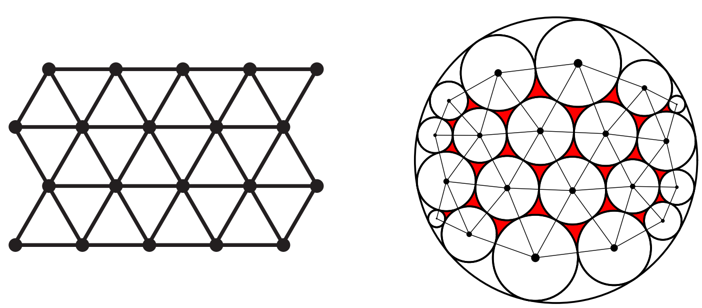

Figure 2.1: A graph and a circle packing of it in the unit disc. The circles on the boundary are tangent to the outside circle.

2.1.1 Circle and ellipse packings have planar embeddings

Every circle packing has a natural planar embedding. Vertices v are mapped to the centre of the corresponding circle E_v, and the edges are mapped to straight lines between the centres of tangent circles.

Ellipse packings also have planar embeddings, but there's no completely natural way to choose the lines. One way is to connect the centres e and f of two ellipses which are tangent at a point t with the broken line etf.

The interior faces of the graph correspond to gaps between circles that we call *interstices.* Our convention will be that interstices are open. In Figure 2.1, interstices are coloured red. The *carrier* of a circle packing is the union of every disc in the packing and every interstice in the packing.

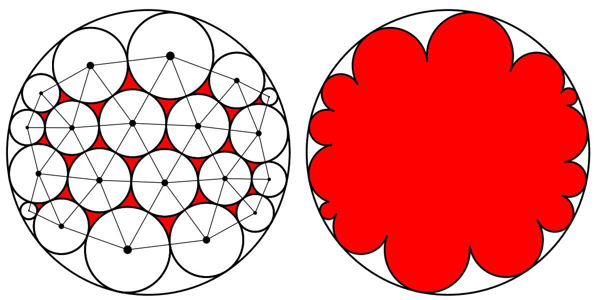

Figure 2.2: The carrier of the circle packing in Figure 2.1.

The same definitions work for ellipse packings.

Suppose G is an oriented planar triangulation, and $E_v : v \in V(G)$ is a circle or ellipse packing of it.

Let F be any interior face of the graph. Then there are three vertices on F: say a, b, c going around counterclockwise. The chain of three ellipses E_a, E_b, E_c may go clockwise or counterclockwise around the interstice corresponding to F.[1] If they go counterclockwise, we say that the packing *respects the orientation* of that face.

We say that the packing is oriented if it respects the orientation of every face. Two interstices whose faces share an edge must have compatible orientations, so it's enough for one of the interstices to be oriented.

[1]To give this a precise meaning, pick a Jordan curve $\gamma \subset E_{v_1} \cup E_{v_2} \cup E_{v_3}$ which goes through the three ellipses $E_{v_1}, E_{v_2}, E_{v_3}$ in that order. Since γ is Jordan, the winding number around any point in the interstice must be ± 1. All such curves are homotopic to each other through $E_{v_1} \cup E_{v_2} \cup E_{v_3}$, which means that that number is independent of the choice of γ. If the winding number is $+1$, then we say the ellipses go counterclockwise.

2.2 Circle packings: existence and uniqueness

Let G be a planar triangulation. Then there is a way to pack it in the unit disc, and that way is unique up to reflections and Möbius transformations.

Koebe-Andreev-Thurston Theorem. *Let G be a finite planar graph with a distinguished 'outside' face. There exists a circle packing of G in the unit disc. If G is oriented, the circle packing can be oriented too.*

If G is also a triangulation, then there is a unique circle packing of G in the unit disc, up to reflections and Möbius transformations that preserve the circle and send the disc to itself.

Corollary. *Let v be an interior vertex of G and w be another vertex. There is a unique oriented packing so that E_v is centred at zero and E_w is centred on the positive real line.*

For an elementary proof, see Brightwell and Schneierman [7]. Neither of the original proofs were elementary, however.

William Thurston sparked contemporary interest in this theorem with a talk at Purdue University in 1985. He gave a proof in his book [2], Chapter 13, Corollary 13.6.2, making it a corollary of Andreev's theorem [3] on hyperbolic reflection groups.

However, Reiner Kuhnau pointed out that it was originally proven by Koebe [5] in 1936 as an application of his Kriegsnormierungstheorem or circle domain theorem to 'contact domains.' See [4], page 141–142.

This history is taken from Kenneth Stephenson's article [6].

We discuss the second part of the theorem a little. It is an exercise in algebra to prove that any Möbius transformation which preserves the unit circle and sends the disc to itself must look like

$$f_{a,\theta}(z) = e^{i\theta}\frac{z-a}{1-\bar{a}z} \qquad \text{for } a \in \mathbb{C}, |a| < 1 \text{ and } \theta \in \mathbb{R}.$$

Let G be a graph. Let C_v be a circle packing of G in the unit disc. Then this packing cannot be unique. For example, $f_{1/2,0}(C_v)$ will be a circle packing with the same tangency relations.

However, if G is a triangulation, then the theorem states that any two packings C_v and D_v of the same graph are related by $D_v = f_{a,\theta}(C_v)$.

2.3 How to approximate conformal homeomorphisms with circle packing

The procedure in this section is taken from Rodin and Sullivan [8], and it was originally proposed by William Thurston in 1985 in a talk at Purdue University.

The *infinite regular hexagonal circle packing* $\mathbb{H}$ is constructed as follows. Start with a unit circle at the origin, and then add layers of unit circles around the outside. Do this forever, and then stop. See Figure 2.3.

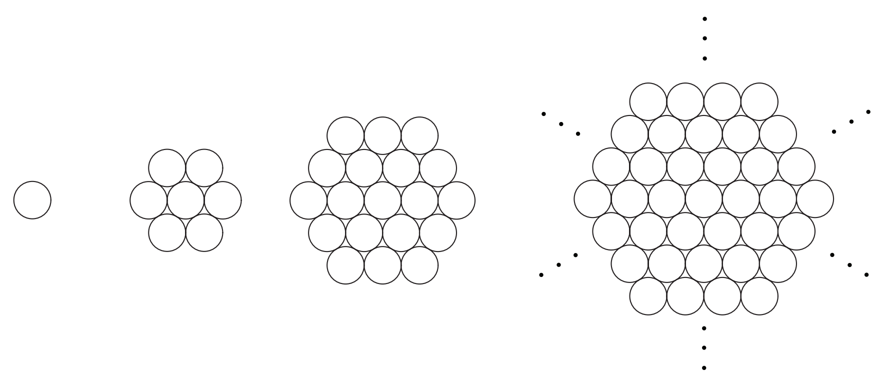

Figure 2.3: Generations of the regular hexagonal circle packing $\mathbb{H}$.

The tangency graph is the infinite triangular lattice. We will call it T. Then $\mathbb{H}$ is a circle packing of T in the plane. We will mention here a fact that becomes very important later: this packing $\mathbb{H}$ is the *only* circle packing of T in the plane, up to uniform scaling and rotation.

This fact is due to Rodin and Sullivan. We will discuss it in Chapter 6, although for a proof we will refer the reader to Oded Schramm's paper [15].

Thurston's original procedure

Let X be a simply connected, bounded domain in $\mathbb{R}^2$. Let z_0 and z_1 be points in its interior. Let δ be small enough that the circle of radius δ at z_0 is contained in X.

For every positive integer $n > 1/\delta$, perform the following tasks:

Step 1. Let H_n be a copy of $\mathbb{H}$ that is scaled down by a factor of n and translated so that the origin moves to z_0. Each circle has diameter $1/n$ and is surrounded by six other circles, and one of the circles is centred at z_0.

Step 2. Let $A'_n = \{C \in H_n : C \subset X\}$, a finite set. Let A_n be the connected component of A_n that contains the circle at z_0 in H_n.

Step 3. Pack the tangency graph of A_n in the unit circle using the Koebe-Andreev-Thurston theorem, sending the circle centred at z_0 to the origin, and some circle adjacent to z_1 to the positive real line.

We get a circle packing. Call it B_n. It has the same tangent graph as A_n, but it's packed in the unit disc.

Step 4. Let G_n and H_n be the embedded tangency graphs of A_n and B_n. Let f_n map the vertices of G_n to the vertices of H_n. Extend it linearly onto every face of G_n.

Then f_n is defined on every vertex, edge, and interior face of A_n, and it is a homeomorphism on its domain of definition.

See Figure 2.4.

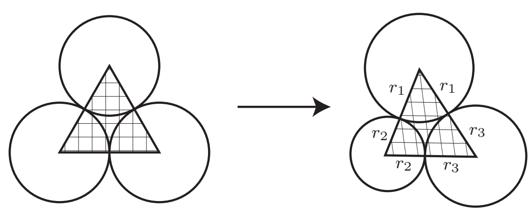

Figure 2.4: A piece of a piecewise linear map.

This will be a homeomorphism as long as the faces of B_n are indeed disjoint. One needs to check that every face contains three sectors of three circles and its own interstice, and nothing else.

Then f_n is a piecewise linear homeomorphism defined on some closed set inside the unit disc.

Lemma 2.3.1. *The interior of the domain of definition of f_n contains every circle that isn't on the boundary of the packing, and every interstice.*

Here f is defined on every triangle in A_n. All triangles connect the centres of adjacent circles.

Let c be a circle that isn't on the boundary. Then there is a ring of six other circles around it, and there are six triangles in A_n that connect the central circle with its neighbours. Their union is a large hexagon which contains c in its interior.

All interstices are in the interior of a single triangle. (Look at the left size of Figure 2.4. The triangle just grazes the side of the interstice, but recall that interstices are *open* sets.)

This proves the result. □

LEMMA 2.3.2. *Fix a compact set $K \subset X$. If n is large enough, the domain of f_n contains K in its interior.*

PROOF. Every circle in the domain packing has diameter $d = 1/n$. Let Y_n be the set of points z in X for which there is a broken line path γ in the unit disc which goes from z_0 to z, and where the distance $d(\gamma, E^c)$ is strictly greater than $2d$. Then Y_n is open and $\bigcup Y_n = K$, so K is contained in some Y_n.

We now show that f_n is defined on Y_n. Fix $z \in Y_n$. Choose a path of line segments γ connecting z to z_0 so that $d(\gamma, X^c) > 2d$. Let $c_1, \ldots, c_m \in T_n$ be the set of (closed) circles that this path enters, in order.

Each pair c_j, c_{j+1} is adjacent, because once the path has left a circle, it enters an adjacent circle or an interstice. And when it leaves an interstice, it can only enter an adjacent circle. Therefore, the list of circles is connected.

By the bound on the distance, if γ is in some circle c_j, then $d(c_j, X^c)$ is greater than d. Therefore, we can fit another layer of circles around c_j that are still contained in X, so c_j is an interior circle.

Therefore, the whole path is inside interior circles or in interstices. And z is in the path, so z is in the interior of the domain of definition of f_n, by Lemma 1.1.1. □

In view of Lemma 2.3.2, it makes sense to talk about the convergence of the functions f_n on any compact set in X. In Chapter 3 we will prove the surprising fact that f_n do in fact converge to a conformal limit.

This is old news. What we want to talk about is ellipse packing. In order to do this, we have to make some small modifications to Thurston's procedure.

Chapter 3

Ellipse packing

We want to replace the circle packing in Thurston's original procedure by an ellipse packing, where we get to choose the shape of each ellipse.

3.1 Revising Thurston's procedure to work with ellipses

We say that two ellipses E, F have the same 'shape' if one is a uniform scaling and translation of the other. That is, $E = aF + c$ for some $a > 0$ and $c \in \mathbb{C}$.

We can obtain an ellipse packing theorem similar to the Koebe-Andreev-Thurston theorem from Oded Schramm's blunt packing theorem [9].

THEOREM 3.1.1. (after Schramm) *Let T be an oriented planar triangulation with an outside edge. For every vertex, let E_v be some fixed ellipse. Let t be a vertex that is not on the boundary, and let u be some other vertex.*

There is an oriented packing of T in the unit disc so that F_v has the same shape as E_v, and the centre of E_t is zero and the centre of E_u is on the positive real line.

Schramm's theorem is much more general than this, and handles objects that are not necessarily even convex. On the other hand, the normalization is different: we are not allowed to select an interior circle and put it at the origin.

We will discuss how to get from his theorem to this one in Appendix B.

However, the construction of the piecewise linear map in Step 4 doesn't work for ellipses. The problem is that a line between the centre of two ellipses may no longer go through the tangent point.

This leads to the kind of thing on the left in Figure 3.1. The two image faces overlap each other, so the corresponding piecewise linear map is not injective.

We solve this problem by breaking every line in half and moving the midpoints to tangent points, as on the right in the figure.

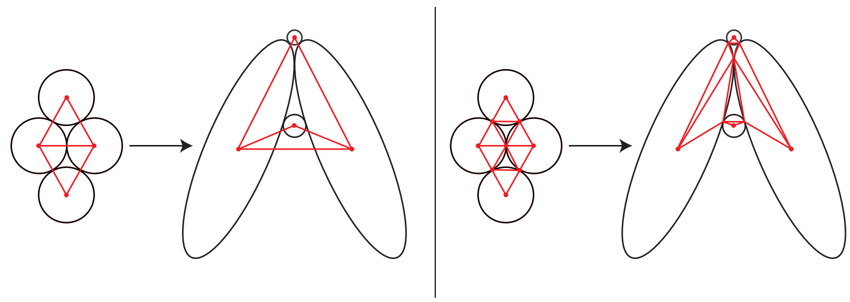

Figure 3.1: The problem with the simple way, and its solution

Having made those two changes, we get the following revised procedure.

3.1.1 Thurston's revised procedure

Let X be a simply connected, bounded domain in $\mathbb{R}^2$. Let z_0 and z_1 be points in its interior. Let $\nu : X \to \mathbb{C}$ be a continuous 'shape field.'

The shape of an ellipse is determined by the ratio M between its major and minor axes, and the angle θ of the major axis to the real line. We say that an ellipse E has shape coefficient ν if $\nu = e^{2i\theta}(M-1)/(M+1)$.

We will elaborate on the reason for this choice in Section 4.3. It turns out that ν uniquely determines the ellipse, and the linear map $z + \nu\bar{z}$ sends circles to ellipses of shape ν.

The steps in the revised procedure are as follows:

Step 1 and 2 of the revised procedure are the same. They give us a connected circle packing A_n in E.

Step 3e. Take the tangency graph G_n of A_n and pack it in the unit circle using the ellipse packing theorem on the previous page, Theorem 3.1.1.

Choose the shapes of the image ellipses depending on the centres of the corresponding circles in A_n. The shape of the ellipse at v is $\nu(z_v)$.

We get an ellipse packing, E_n.

Step 4e. Let f_n be the piecewise linear map which sends vertices of A'_n to vertices of B_n and which sends tangent points in A'_n to tangent points in B_n. It can't be linear on every face, but we can make it piecewise linear on every face, broken up as in Figure 3.2.

This will be a homeomorphism. One can check that every one of the image triangles is either part of one of the ellipses, or covers three sectors of an ellipse and an interstice. One can also check that the interiors don't overlap. (For example, check that the interstice doesn't escape the triangle that connects its cusps, and that the angle between successive tangent points around the inside of the ellipse is less than 180°.)

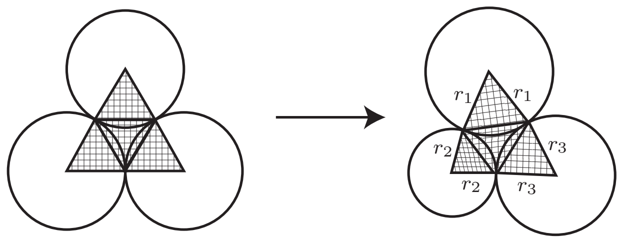

Figure 3.2: The piecewise linear maps in the revised procedure.

Whether we use the original or revised version of Thurston's procedure, the function f_n is defined on every interior face of the embedded graph, and Lemma 2.3.2 goes through as before.

3.1.2 Conclusion

We have a sequence of piecewise linear maps f_n, which are eventually well-defined on every compact set in X. The individual linear maps are determined by the sizes and positions of the image circles.

We now wish to study the limits of these maps.

Chapter 4

Oriented linear maps

4.1 A way of writing oriented linear maps

Recall that the maps f_n from Thurston's procedure are pieced together out of oriented linear maps. Our first step is a helpful lemma that characterises this type of map.

LEMMA 4.1.1. *An oriented linear map can be written as* $f(z) = \alpha(z+\mu\bar{z})$ *with* $\alpha, \mu \in \mathbb{C}$ *and* $\alpha \neq 0$ *and* $|\mu| < 1$*, and any map that can be written like that is an oriented linear map.*

Let $f : \mathbb{R}^2 \to \mathbb{R}^2, f(x, y) = (ax + by, cx + dy)$ be an arbitrary linear map. This map is oriented when $ad - bc > 0$. Write

$$\begin{aligned} f(z) &= (a+ic)\left(\frac{z+\bar{z}}{2}\right) + (b+id)\left(\frac{z-\bar{z}}{2i}\right) \\ &= \underbrace{\frac{a+ic-ib+d}{2}}_{\alpha}\left(z + \underbrace{\frac{a+ic+ib-d}{a+ic-ib+d}}_{\mu}\,\bar{z}\right). \end{aligned} \tag{4.1}$$

This is always possible if the Jacobian is positive, because if $a+ic-ib+d$ is zero then $ad - bc = -a^2 - b^2 \leq 0$.

We can therefore write any oriented linear map in the form $\alpha(z + \mu\bar{z})$. This proves one direction.

Consider a general map $f(z) = \alpha(z + \mu\bar{z})$ with $\alpha \neq 0$. Let a, b, c, d be as above. They are well-defined: for example, we have $c = \mathrm{Im}(\alpha + \mu\alpha)$. Then

$$1 - |\mu|^2 = 1 - \frac{(a-d)^2 + (c+b)^2}{(a+d)^2 + (c-b)^2} = \frac{4(ad-bc)}{(a+d)^2 + (c-b)^2},$$

The denominator is always positive. (Otherwise $\alpha = 0$.) So the sign of the Jacobian is the same as the sign of $1 - |\mu|^2$.

Therefore, a linear map $f(z) = \alpha(z + \mu\bar{z})$ with $\alpha \neq 0$ is oriented if and only if $|\mu| < 1$. $\square$

4.2 The image of a circle under a linear map

We have a convenient way of writing linear maps. We now ask what the image of a circle is under such a map.

The answer turns out to be reasonably nice. Recall that the 'distortion' of an ellipse is the ratio between the major and minor axes. Then:

THEOREM 4.2.1. *An oriented linear map* $f : z \mapsto \alpha(z + \mu\bar{z})$ *takes circles to ellipses with distortion* $(1 + |\mu|)/(1 - |\mu|)$. *The angle of the major axis to the real line is* $\arg \alpha + \frac{1}{2} \arg \mu$.

Here $\frac{1}{2} \arg \mu$ only makes sense up to multiples of π, but that is all right, because the angle of the major axis is only determined modulo π.

PROOF. It is enough to look at what happens to the unit circle. The map is invertible, so it takes the circle to some ellipse. All we need to do is find the major and minor axes.

Suppose that the circle is centred at zero. The centre of the image ellipse will also be zero.

The distance from the centre to a point on the edge of the image is

$$|f(e^{i\theta})| = |\alpha||e^{i\theta} + \mu e^{-i\theta}| = |\alpha||1 + \mu e^{-2i\theta}|.$$

The major axis must be twice the maximum of this, namely $2|\alpha||1 + \mu|$. The minor axis is twice the minimum, $2|\alpha||1 - \mu|$. The distortion of the image ellipse is the ratio, which is $(1 + |\mu|)/(1 - |\mu|)$.

The maximum occurs at $\theta = \frac{1}{2} \arg \mu$. The angle of the image point for this value of θ is

$$\arg f(e^{i\theta}) = \arg\left(\alpha(e^{i\theta} + \mu e^{-i\theta})\right) = \arg \alpha + \arg \theta = \arg \alpha + \tfrac{1}{2} \arg \mu.$$

This is the angle of a point on the image ellipse which is as far as possible from the origin. That is, it is the angle of the major axis. □

4.3 The bar-Beltrami coefficient

The theorem in the previous section is not terribly ugly, but it would be nice if there were a single number that characterized the shape of the image ellipse. And there is!

Let $\nu = f_{\bar{z}}/\overline{f_z} = \mu\alpha/\bar{\alpha}$. This is called the bar-Beltrami coefficient, by analogy with the usual Beltrami coefficient $f_{\bar{z}}/f_z = \mu$.

We have the following corollary:

Corollary 4.3.1. *Let f be an oriented linear map. Suppose $\nu = f_{\bar{z}}/\overline{f_z}$. Then f takes circles to ellipses with distortion $(1+|\nu|)/(1-|\nu|)$ and major axis angle $\frac{1}{2}\arg\nu$.*

Proof. Let $f(z) = \alpha(z + \mu\bar{z})$ as in Lemma 4.1.1. Then

$$\nu = f_{\bar{z}}/\overline{f_z} = \mu\alpha/\bar{\alpha}.$$

So $|\nu| = |\mu|$. By the previous lemma, the distortion of image ellipses is

$$(1+|\mu|)/(1-|\mu|) = (1+|\nu|)/(1-|\nu|).$$

And, because $\arg\alpha/\bar{\alpha} = 2\arg\alpha$,

$$\frac{1}{2}\arg\nu = \frac{1}{2}(\arg\mu + 2\arg\alpha) = \frac{1}{2}\arg\mu + \arg\alpha,$$

which is the major axis angle from Lemma 4.2.1. This proves the result. □

Observation. If the distortion is d and the major axis angle is θ, then ν is uniquely determined: $\nu = e^{2i\theta}(d-1)/(d+1)$.

Corollary 4.3.2. *Let f be an oriented affine linear map. The following are equivalent:*

- *f takes circles to ellipses with distortion at most K.*
- $|f_{\bar{z}}/f_z| \leq (K-1)/(K+1)$.

Proof. Write $f(z) = f(0)+\alpha(z+\mu\bar{z})$. By Theorem 4.2.1, the distortion of the image ellipses is $(1+|\mu|)/(1-|\mu|)$. A little algebra shows

$$\frac{1+|\mu|}{1-|\mu|} \leq K \quad \Longleftrightarrow \quad |\mu| \leq \frac{K-1}{K+1}.$$

Here $\mu = f_{\bar{z}}/f_z$. This proves the result. □

4.4 Conclusion

Suppose f is a piecewise linear homeomorphism so that each map takes circles to circles. That is, f takes circles to ellipses with distortion at most 1.

By Corollary 4.3.2, we must have $f_{\bar{z}} = 0$ at every point where the function is linear, so every piece is analytic. By the removability of lines, the whole function f has to be conformal (and therefore linear).

What if the maps in f take circles to ellipses of bounded distortion $\leq K$? In that case, the piecewise map is in a class of maps called K-*quasiconformal*.

This is a natural class to study for our problem, and it turns out that the piecewise linear maps in Thurston's procedure are indeed all K-quasiconformal for a certain large K.

Chapter 5

Quasiconformal maps

In the last section, we showed that an oriented linear map $\alpha(z+\mu\bar{z})$ distorts a circle in a certain way depending only on. We present an analytic definition of a quasiconformal map, in terms of the distortion.

5.1 The first definition of a quasiconformal map

Definition. Let X be a domain. Let $f : X \to \mathbb{C}$ be a homeomorphism onto its image $f(\mathbb{C})$. Then f is K-*quasiconformal* on E if it satisfies the following three conditions.

1) The image $f(E)$ is open in $\mathbb{C}$.

2) At almost every point $z_0 \in E$, the function is once differentiable, and the first-order approximation takes circles to nondegenerate ellipses with distortion at most K.

3) The derivatives f_x and f_y exist almost everywhere on E.

   For almost every fixed x, if $[a, b]$ is an interval so that $\{x\} \times [a, b] \subset E$, then the integral of $f_y(x, y)$ from $y = a$ to $y = b$ is $f(x, b) - f(x, a)$.

   For almost every fixed y, if $[a, b]$ is an interval so that $[a, b] \times \{y\} \subset E$, then the integral of $f_x(x, y)$ from $x = a$ to $x = b$ is $f(b, y) - f(a, y)$.

   These requirements together are called *absolute continuity on lines*.

This definition has a few logical redundancies, but that's alright.[2] Notice that the third condition is always satisfied for a piecewise linear map.

LEMMA 5.1.1. *A K-quasiconformal function on a domain E is also K-quasiconformal on every subdomain.*

PROOF. Consider a subdomain $A \subseteq E$. It is clear that properties 2) and 3) are still true on A, and 1) is true because $f(A)$ is open in $f(E)$, which is open in $\mathbb{C}$. $\square$

5.2 Second definition of a quasiconformal map

DEFINITION A$'$. Let $f : E \to \mathbb{C}$ be a homeomorphism onto its image $f(E)$. Then f is *K-quasiconformal* if there is a measurable function $\mu : E \to \mathbb{C}$ with $||\mu||_\infty \leq (K-1)/(K+1)$ so that

1) the image of E under f is open, and
2) $f_{\bar{z}} = \mu(z) f_z$ almost everywhere in E, and
3) f is absolutely continuous on lines.

The equivalence comes out of Corollary 4.3.2. This function μ is called the *Beltrami coefficient* of f. Since $f_z \neq 0$ except on a null set, μ is uniquely defined up to a null set.

5.3 Convergence of quasiconformal maps

Quasiconformal maps share many nice properties with conformal maps. For example, they are a normal family.

We quote several results.

THEOREM 5.3.1. *Let X be a domain in $\mathbb{C}$, and $a \neq b$ be complex numbers. Fix $K < \infty$. The family of all K-quasiconformal functions whose image does not contain a or b is equicontinuous around every point in X.*

[2] The well-known 'invariance of domain' theorem from topology guarantees that the image is open in $\mathbb{C}$. See for example Bredon's *Topology and Geometry*, [10], Corollary 19.9.

The Gehring-Lehto theorem says that an open map which satisfies the second condition is differentiable almost everywhere on E. So the first-order approximation is automatically valid a.e. Look at Lehto and Virtanen's Theorem III.3.1 in [11], or page 17, Lemma 1 in Ahlfors's book *Lectures on Quasiconformal Mappings* [12].

Proof. See Lehto and Virtanen [11], page 69, or Astala, Iwaniec and Martin [13], page 182. Compare to Montel's theorem.

Corollary 5.3.2. If $f_n : X \to \mathbb{C}$ is a sequence of K-quasiconformal functions, and there is a uniform bound $|f_n| \leq R$, then there is a subsequence f_{n_j} that converges uniformly on compact sets to some limit f.

Proof. We can let a, b be any two points outside the disc of radius R. Then f_n are equicontinuous, and it is well-known that equicontinuous families are normal. □

Theorem 5.3.3. *Let $f_n : X \to \mathbb{C}$ be K-quasiconformal functions that converge to f uniformly on compact sets. Then f is K-quasiconformal or constant.*

Proof. See Lehto and Virtanen [11], page 74. Ahlfors [12] covers more or less this topic on page 32–33 (the fact that K-quasiconformal mappings with a fixed normalization are sequentially compact).

Not only that, but if the Beltrami coefficients μ_n of the sequence f_n converge pointwise to a limit μ, then the limit function f will have the Beltrami coefficient μ.

Theorem 5.3.4. *Suppose the K-quasiconformal functions $f_n : X \to \mathbb{C}$ with Beltrami coefficients μ_n converge uniformly on compact sets to a K-quasiconformal limit f with Beltrami coefficient ν.*

If μ_n converges pointwise almost everywhere to a limit μ, then $\nu = \mu$.

Proof. See Lehto and Virtanen [11], page 187. Alternatively, look at the proof of Lemma 5.3.5 in [13], page 171: just drop the assumption about normalization and throw away the last three sentences of the proof.

5.3.1 Convergence of functions not defined everywhere

Let X be a domain. We will say that a sequence of subsets $X_n \subset X$ 'fills' X if every compact set in X is eventually contained in X_n.

Thurston's procedure gives us functions f_n that are defined on a series of subsets of a domain X, and Lemma 2.3.2 tells us that the domains of definition of f_n fill X.

We will later prove that:

(i) f_n is K-quasiconformal for a bounded K independent of n.

(ii) The images $f_n(X_n)$ converge to the unit disc in the sense above.

(See Theorems 6.1.1 and 6.2.1.)

The first fact tells us that we can apply Theorem 5.3.1 to a sequence of compact sets K_n which fill X, and then use diagonalization to get a subsequence which converges uniformly on every set.

Theorem 5.3.3 tells us that the limit of that subsequence is either a K-quasiconformal map into the unit disc, or a constant.

With the second fact, we can show that the image of that limit is the whole unit disc, using the following theorem:

THEOREM 5.3.5. *Let X_n be domains that converge to X. Let f_n be K-quasiconformal maps from the domain X_n to the unit disc, and suppose that $f_n(z_0) = 0$.*

If the sequence of sets $f_n(X_n)$ fills the unit disc, then the subsequential limits of f_n are homeomorphisms onto the unit disc.

PROOF. Without loss of generality, let f_n converge to some limit f uniformly on compact sets in X.

By definition, f_n is a homeomorphism onto its open image. Let g_n be the inverse of f_n. It is K-quasiconformal, because the inverse of a quasiconformal map is also quasiconformal with the same modulus.[3]

The domains of the functions f_n and g_n converge to the region X and to the unit disc. By Corollary 5.3.2, there is a subsequence g_{n_j} which converges uniformly on compact sets in the unit disc. Then $g = \lim g_n$ is defined on the unit disc.

Fix $|z| < 1$. Then $g_{n_j}(z) \to g(z)$, which is in the open set X.[4] Let F be a compact set in D containing a neighbourhood of $g(z)$. Then f converges uniformly on F, so $f(g(z)) = \lim_j f_{n_j}(g_{n_j}(z)) = z$.

Similarly, for $z \in X$, $g(f(z)) = z$.

We have proven that $f : X \to \{|z| < 1\}$ has an inverse $g : \{|z| < 1\} \to X$, and it is therefore a homeomorphism. $\square$

[3] See Chapter II of Ahlfors [12] or page 17 and 168 of Lehto and Virtanen [11].

[4] In principle, $g(z)$ could be in the boundary of X. We prove not. If g is a constant map, then $g(z) = g(0) = \lim g_n(0) = z_0$ is not on the boundary.

If it is not constant, then it is a homeomorphism onto an open image. Let N be a small disc around z whose closure is contained in the unit disc. Then g_n converges uniformly to g. Let ∂N be the topological boundary of N. There is a disc around $g(z)$ which contains no point of $g(\partial N)$. If n is large enough, there is a disc that contains no point of $g_n(\partial N)$ but does contain $g(z)$ and $g_n(z)$.

If $g(z)$ is not contained in $g_n(N)$, the line between $g(z)$ and $g_n(z)$ has to have a boundary point $\partial g_n(N) = g_n(\partial N)$ in it. We know that it doesn't, so $g(z)$ is contained in $g_n(N) \subseteq X$ and is not on the boundary. This argument is taken from Lehto and Virtanen [11], page 76.

Chapter 6

The rigidity of circles and ellipses

Circle packings have a variety of rigidity properties, and they carry over readily to ellipse packings. We will present some of them.

In every theorem below, we will assume that no ellipse has distortion more than K.

6.1 The Ring Lemma

This argument is essentially taken from Rodin and Sullivan [8]. However, we have added a painfully detailed geometric justification, just to make sure that the lemma works on ellipses also. It has been hidden in small text, and we advise the reader to ignore it.

Lemma 6.1.1. (Ring Lemma) *Suppose E is an ellipse surrounded by a ring of tangent ellipses $E_1, \ldots, E_n$. Then there is a constant $c(n, K) > 0$ independent of the choice of ellipses so that*

$$\min_j \frac{diam(E_j)}{diam(E)} \geq c(n, K).$$

Proof. We suppose that the central ellipse has diameter 1. We must show that the diameter of the smallest ellipse is bounded below.

The ring of ellipses has to make a loop around the central ellipse, so the sum $\mathrm{diam}(E_1) + \cdots + \mathrm{diam}(E_n)$ is at least 1. Therefore, the *largest* ellipse in the chain must have diameter at least $1/n$.

We now proceed by induction. Suppose that there is a lower bound on the diameter of E_j. Let t be the tangent point of E_j and E. Rotate the drawing so that the tangent is horizontal, and move t to the origin. If we zoom in far enough, we will get something like this:

Figure 6.1: Closeup on the tangent point of two ellipses.

We claim that if E_{j+1} is very small, then it must be close to the tangent point t. Here is a naive argument. It is tangent to both ellipses E_j and E, so it contains points in both of them. Those points must be close together. The only place E_j and E get close is at the origin. Therefore E_{j+1} must be near the origin.

The reader may be wondering if this bound depends somehow on the precise choice of ellipses. We refer them to this tiny lemma:

OBSERVATION. The curves are approximate parabolas. If a and b are the curvatures of E_j and E at t, then the two curves are $y \sim \frac{1}{2}ax^2$ and $y \sim -\frac{1}{2}bx^2$, respectively.

Here '$\sim$' means that the two sides are asymptotic as $x, y \to 0$.

TINY LEMMA 6.1.2. *Let $d_j = diam(E_j)$. Then the distance of the centre of E_j to t is asymptotically less than $\sqrt{(a+b)/2} \times \sqrt{d_j}$ as $d_j \to 0$.*

PROOF. Let (x_1, y_1) be the point $E_j \cap E_{j-1}$, and (x_2, y_2) be the other tangent point $E_j \cap E$. They are both in the ellipse E_j, so the distance between the two points is no greater than d_j. But

$$d_j \geq |y_1 - y_2| \gtrsim \tfrac{1}{2}ax_1^2 + \tfrac{1}{2}bx_2^2 \geq \tfrac{1}{2}(a+b)\max(x_1, x_2)^2.$$

The distance from the centre of E to the tangent point t is asymptotically no greater than

$$\begin{aligned}\sqrt{x_1^2 + y_1^2} + d_j &\leq x_1 + y_1 + d_j \\ &\sim x_1 \\ &\sim \sqrt{(a+b)/2} \times \sqrt{d_j}.\end{aligned}$$

Here $y_1 \sim \frac{1}{2}ax_1^2$ and d_j are small compared to $\sqrt{d_j}$ as $d_j \to 0$. □

So E_j is close to the origin. And if one ellipse is small and close to the origin, the ellipse touching it must be small and close to the origin too, in order to fit into the dotted region.

This forces all the other ellipses $E_{j+1}, \ldots, E_n$ to be small. But this is impossible, because the ellipses have to make a ring around E! If they are all small, they cannot complete the loop.

Therefore, the diameter of E_j is bounded below. We will take a break to justify these vague remarks in excruciating detail, and then finish the proof.

TINY THEOREM 6.1.3. *Suppose that we are in the situation in Figure 6.1. Let d be small. Let E be an ellipse contained in the dotted area and in the ball $\{|x| \leq d\}$. Let F be an ellipse tangent to it, also contained in the dotted area. Then F has diameter at most Cd^2 for some constant C, where C depends only on the diameters of the large ellipses and on the ellipse distortion.*

What do we mean, "let d be small?" Let $h(\xi)$ be the distance between the top and bottom curves on the line $x = \xi$. Then $h(\xi) \sim \frac{1}{2}(a+b)\xi^2$, where a, b are the curvatures of the two large ellipses as in the tiny lemma. This is an asymptotic equality. We assume that d is small enough that the weaker exact inequality $h(\xi) \leq (a+b)\xi^2$ is actually true for $0 \leq \xi \leq 2d$.

We also assume, a little mysteriously, that $d < 1/4(a+b)$. Both these bounds are uniform in the diameters of the ellipses and in their distortion.

PROOF. Let x be the tangent point between F and E. We can inscribe a circle of diameter $\operatorname{diam}(F)/K$ inside F so that the circle also contains x, as in the footnote on page 37.

Naturally, we can inscribe a circle of smaller diameter too. Let δ be an arbitrary number less than $\operatorname{diam}(F)/K$. We also require, again a little mysteriously, that δ should be less than $\min\{d, 1/(a+b)\}$. Inscribe a circle of diameter δ in F containing x. Then there is room for the line of length δ connecting the top and bottom of the circle.

Let the centre of the circle be $z = \xi + i\eta$. Remember $\delta \leq d$. Then $\xi \leq 2d$, so our upper bound on the radius is valid: $\delta \leq h(\xi) \leq (a+b)\xi^2$.

The circle has radius $\frac{1}{2}\delta$, and x is contained in it. But x is also contained in E. So

$$\begin{aligned} d &\geq |x| \\ &\geq \sqrt{\xi^2+\eta^2} - \frac{1}{2}\delta \\ &\geq \xi - \frac{1}{2}\delta. \\ &\geq \frac{1}{\sqrt{a+b}}\sqrt{\delta} - \frac{1}{2}\delta \\ &\geq \frac{1}{2\sqrt{a+b}}\sqrt{\delta}. \end{aligned}$$

(Remember we insisted that δ be less than $1/(a+b)$.)

Solving for δ, we will get $\delta \leq 4(a+b)d^2$. But δ was an arbitrary number less than a certain upper bound. So that upper bound must be less than $4(a+b)d^2$ also:

$$\min\{\operatorname{diam}(F)/K, d, 1/(a+b)\} \leq 4(a+b)d^2,$$

but the second two terms in the minimum are guaranteed to be *larger* than $4(a+b)d^2$, because of how we chose d.

We must therefore have $\operatorname{diam}(F)/K \leq 4(a+b)d^2$, or $\operatorname{diam}(F) \leq 4K(a+b)d^2$. This proves the tiny theorem with $C = 4K(a+b)$. □

Suppose we have a lower bound on the diameter of E_{j-1}, and we use induction. We know the base case already: $\operatorname{diam}(E_1) \geq 1/n$.

Let d be 'small' as above, and also let $d \leq 1/2Cn$. Suppose the original ellipse E_j was closer than $d/2$ to the tangent point. This is certainly less than d, so, by the tiny theorem, the diameter of E_{j+1} can be at most Cd^2.

That means that the distance of E_{j+1} to the tangent point is no greater than $\frac{d}{2} + Cd^2 \leq \frac{d}{2} + \frac{d}{2n}$. This is also less than d, so we repeat the argument. The distance of the ellipse E_{j+k} to the tangent point can be at most $\frac{d}{2} + \frac{d}{2n}k$. After $n-1$ steps, the distance can be at most d.

But after $n-j+1$ steps, we've come back around to the large ellipse E_{j-1}! It can hardly be wedged in to the right of itself, so d cannot be 'small' as above.

This gives us a lower bound on the size of an ellipse, and that lower bound depends only on the sizes of the large ellipses and on the distortion. We had a lower bound on the diameter of E_{j-1}

This proves that no ellipse can be much smaller than the one before it, and that gives us a lower bound on the size of every ellipse $E_1, \ldots, E_n$. The exact lower bound doesn't matter and will be left up to the reader's imagination. □

6.2 Ellipse sizes must go to zero

This argument is identical to the one in Rodin and Sullivan.

LEMMA 6.2.1. (Length-Area Lemma) *Suppose the unit circle contains n chains of disjoint ellipses with $m_1, \ldots, m_n$ ellipses.*

Let the diameters of the ellipses in the j-th chain be called $d_{j1}, \ldots, d_{jm_j}$. Let $\ell_j = \sum_i d_{ji}$. Then we must have

$$\sum_j \frac{\ell_j^2}{n_j} \leq 4K.$$

PROOF. Let E be an ellipse of diameter d. A circle of diameter d has area $\pi(d/2)^2$, so the area of the ellipse is at least $\pi d^2/4K$.

Let the total area of the ellipses in the j-th chain be A_j. Then A_j is at least $\sum_i \pi d_{ji}^2/4K$. By the Cauchy-Schwartz inequality,

$$\ell_j^2 = \left(\sum d_{ji}\right)^2 \leq \sum d_{ji}^2 \times \sum 1 = \left(\frac{4K}{\pi} A_j\right) \times m_j.$$

Therefore, $\ell_j^2/m_j \leq (4K/\pi)A_j$. But the sum of the areas A_j can be at most π, which means

$$\sum \ell_j^2/m_j \leq (4K/\pi) \sum A_j \leq 4K.$$

This proves the result. □

COROLLARY 6.2.2. *The size of all the image ellipses in Thurston's revised procedure converge uniformly to zero.*

PROOF. We are looking at the packing B_n. This is a packing of a certain subset G_n of the regular hexagonal graph in the unit circle.

Let D be the graph distance from the vertex v_0 to the boundary of G_n. Then D goes to infinity as n does.

Suppose that the ellipse at vertex v has diameter d. There are two possibilities, and in both cases d has to be small:

- The graph distance of the vertex to the boundary is more than $D/2$. In this case, it is surrounded by $D/2$ rings containing $6, 12, 18, \ldots$ circles successively. All these rings have length at least d.

  Therefore,

  $$\frac{d^2}{6} \times \left(\sum_{j=1}^{D} \frac{1}{j} \right) = \sum_{j=1}^{D} \frac{d^2}{6j} \leq \sum_{j=1}^{D} \frac{\ell_j^2}{6j} \leq 4K.$$

  The sum in parentheses is the harmonic series, which goes to infinity as D becomes large. The other factor must become small:

  $$d \leq \sqrt{24K \left(\sum_{j=1}^{D} \frac{1}{j} \right)^{-1}}$$

  goes to zero as n goes to infinity.

- The graph distance of the vertex to z_0 is more than $D/2$. Then the vertex is surrounded by $D/2$ chains as above. Some of the chains hit the boundary this time before completing the loop, so they aren't rings, but each chain contains *at most* $6, 12, \ldots$ circles.

  If the chains are loops, their total length is at least d. If not, they separate v and v_0, and connect two points on the boundary of the unit circle, as in Figure 6.2. One can see that the chains still have to have total length at least d. As above, this implies that d is small.

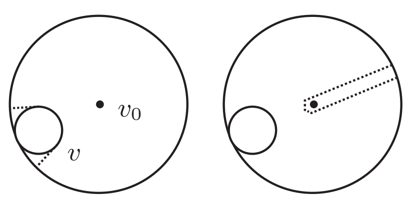

Figure 6.2: A chain around a large boundary circle. The chain has to separate v and v_0, so going around the other way doesn't help.

6.3 Packing the infinite triangular lattice T

Recall that $\mathbb{H}$ is the infinite regular hexagonal circle packing, in which every circle is surrounded by six other circles of the same size. It is a circle packing of the infinite triangular lattice, T.

Say we try to pack this infinite graph in some other way. If every circle is the same size, we will just get a rotation and scaling of the hexagonal circle packing.

Can we pack it so that two adjacent circles are different sizes? Surprisingly, this is not possible:

Rigidity Theorem. If a circle packing has graph T, all the circles are the same size.

There are several different proofs of this. The first published one seems to be in the appendix of Rodin and Sullivan's [8], relying on a difficult paper by Sullivan [14]. A later elementary proof was provided by Schramm [15]. We will not prove it here.

6.3.1 Circles are nearly the same size at the centre of a large neighbourhood that looks like T

The Rigidity Lemma has a finitary version. Again, this is taken directly from Rodin and Sullivan.

Finite Rigidity Theorem 6.3.1. (Version 1) *Let T_n be a graph neighbourhood of size n in T, centred at zero. Let E be any circle packing of T_n in the plane which takes zero to the unit disc.*

Fix m. Let $n > m$. If n is large, the diameters of the circles in the subgraph T_m are nearly 1.

Proof. Consider the circles in a neighbourhood of size m around zero. If $n > m$, they are all surrounded by a ring of circles, so the Ring Lemma applies to them. But the circle at zero is the unit circle. So by repeated application of the Ring Lemma, the diameter of every other circle in that neighbourhood is bounded above and below.

This means that the set of possible positions and diameters of the circles is compact. We claim that the diameters of the circles in every fixed neighbourhood T_m converge to one as $n \to \infty$.

Suppose not. Then we may pick a subsequence of circle packings E_m^j so that the diameters converge to something else on T_m. Use diagonalization to get a subsequence of circle packings E_n so that the diameters converge on *every* neighbourhood T_m. The limit will be a circle packing of T so that the circles are not all the same size, which is impossible. ☐

6.3.2 Even if they are not quite circles

FINITE RIGIDITY THEOREM 6.3.2. (Version 2) *Let T_n be a graph neighbourhood of size n in T, centred at zero. Let E be any ellipse packing of T_n in the plane which takes zero to the unit disc.*

Fix m and $K > 1$. Also fix an interval $a < 1 < b$.

There exists N so that, whenever $n \geq N$ and E is an ellipse packing of T_n and no ellipse in E has distortion more than K, the diameters of the ellipses in the subgraph T_m are in the interval $[a, b]$.

PROOF. We proceed by contradiction as before. Suppose that the statement of the theorem is not true. Then there is a sequence of ellipse packings E_ℓ of T_{n_ℓ} so that the maximum distortion converges to 1, but the diameters do not converge to 1 in the neighbourhood T_m.

The Ring Lemma also holds for ellipses, so the fact about compactness is still true. So there exists a subsequence of those packings which converges on every neighbourhood T_m, and some circle in the limit has a diameter different from 1. The limit will have $K = 1$, so it is a circle packing of $\mathbb{H}$ with two circles of different diameters, which is impossible. ☐

6.4 Conclusion

We have proven several rigidity results about ellipse packings, and we may apply them to our procedure.

Theorem 6.1.1 shows that if we have two adjacent ellipses in the interior the an ellipse packing, then the ratio of their diameters is bounded above and below. Theorem 6.2.1 shows that the maximum sizes of the ellipses in the image packing go to zero — even the boundary ellipses.

Finally, Theorem 6.3.2 tells us that, if the graph of a packing looks like T in a large neighbourhood of a point, and the ellipses in the packing have small distortion, then the triangles in the packing are nearly equilateral.

We will exploit all three of these results in the next section, where we finally prove that the maps from Thurston's procedure, f_n, are K-quasiconformal maps that converge to a limit.

Chapter 7

Conclusion: Thurston's procedure

7.1 The bar-Beltrami equation

We have already met the Beltrami equation in Section 5.2. It is the differential equation $f_{\bar{z}} = \mu f_z$, where μ is some function of z with $|\mu| < 1$.

Suppose f_n is a sequence of K-quasiconformal mappings that converge uniformly on compact sets to f. We know from Theorem 5.3.4 that if the Beltrami coefficients μ_n converge pointwise, the coefficient of the limit will be the limit of the coefficients (almost everywhere).

We now introduce the *bar-Beltrami equation*:

$$f_{\bar{z}} = \nu \overline{f_z}$$

where ν is some function with $\sup |\nu| < 1$. The difference between this and the Beltrami equation is that there is now a bar on f_z.

The *bar-Beltrami coefficient*, which we have already seen in Corollary 4.3.1, is that number $\nu = f_{\bar{z}}/\overline{f_z}$.

7.1.1 Bar-Beltrami and Beltrami

There is a close relation between the bar-Beltrami equation and the original equation.

LEMMA 7.1.1. *Suppose $f(z)$ has bar-Beltrami coefficient ν at z_0. Let g be the inverse of f. Then the Beltrami coefficient of g at $f(z_0)$ is $-\nu$.*

PROOF. Suppose f is linear. A linear map with bar-Beltrami coefficient ν must look like $w = \alpha z + \bar{\alpha}\nu\bar{z}$ for some constant $\alpha > 0$. If we solve for z, we will get

$$w - \nu\bar{w} = \alpha(1 - |\nu|^2)z.$$

So $z = (w - \nu\bar{w})/(\alpha(1 - |\nu|^2))$. This has Beltrami coefficient $z_{\bar{w}}/z_w = -\nu$. That proves the result for linear maps. But it is also true for nonlinear maps, because the result only depends on the derivatives of f, g at a point. $\square$

7.2 The bar-Beltrami coefficient of f_n

In this section, we prove the following theorem about the maps f_n from Thurston's procedure:

THEOREM 7.2.1. *The bar-Beltrami coefficient of the maps f_n converges to ν almost everywhere.*

If we ignore the set of measure zero where the coefficients do not exist, this convergence will be uniform on compact sets in X.

We will begin by sketching the proof.

SKETCH OF PROOF. The idea here is simple. The shape field is continuous, so neighbouring ellipses have about the same shape. We can apply a linear map, say g, to turn them all back into approximate circles.

Then, by the Finite Rigidity Theorem 6.3.2, all the circles will have about the same size, so the image triangles are all approximately equilateral. Then $g \circ f_n$ will be close to conformal, and when we apply g^{-1} to it, we get

$$f_n = g^{-1} \circ g \circ f_n$$

which will have about the same bar-Beltrami coefficient as g^{-1}. And g^{-1} turns circles into ellipses of shape ν, so it has bar-Beltrami coefficient ν.

PROOF. We recall that we have a graph G_n and two circle packings A_n and B_n, where A_n is part of the infinite regular hexagonal circle packing, and B_n is some ellipse packing.

The function f_n is differentiable except on a set of measure zero. Let x be a point in the domain X not in that set. Let x_n be the label of the closest circle to x in A_n. If n is large enough, the distance $|x - x_n|$ will be no greater than $2/n$.

For $n > m$, let $N_{m,n}$ be the graph neighbourhood of radius m around x_n in G_n. Let $A_{m,n}$ be the set of circles in A_n corresponding to vertices in N_n. Similarly, let $B_{m,n}$ be the corresponding ellipses in B_n.

A map g from ellipses of shape $\nu(x)$ to circles

By Corollary 4.3.1, the linear map $f(z) = z + \nu\bar{z}$ takes circles to ellipses of shape ν. We think of it as a map with bar-Beltrami coefficient ν.

Let g be the inverse map, $g(z) = (z - \nu\bar{z})/(1 - |\nu|^2)$. It takes ellipses of shape ν to circles, and it is not difficult to see that it will take ellipses of shape *approximately* $\nu(x)$ to approximate circles.[5]

[5] In general, g takes an ellipse of shape μ to an ellipse of shape $(\mu - \nu)/(1 - \bar{\nu}\mu)$.

The shapes of the ellipses in $B_{n,m}$ are all about the same

The circles $A_{m,n}$ are contained in the ball of radius $2(m+1)/n$ around x. Hold m fixed and make n large. Then that ball will shrink to a point. The shape field is continuous, and the shapes of the ellipses in $B_{m,n}$ are taken from points in the ball. So the shape of the ellipses $B_{m,n}$ will converge to $\nu(x)$ as $n \to \infty$ if m is held fixed.

That means that, for every fixed m, there is some N_m so that ellipses in $g(B_{m,n})$ have distortion less than $1 + 1/m$ whenever $n \geq N_m$.

Applying g to them will turn them all into approximate circles

Consider $A_{m,n}$ and $B_{m,n}$. These are ellipse packings of the same graph: a neighbourhood of size m in the infinite triangular lattice.

If $n \geq N_m$, then the ellipses in $g(B_{m,n})$ have distortion less than $1+1/m$. By the Finite Rigidity Theorem 6.3.2, as $m \to \infty$ with $n \geq N_m$, the ratio of diameters of adjacent ellipses in $g(B_{m,n})$ must converge to 1.

The map $g \circ f_n$ is close to conformal at x

Let ϕ_n be the restriction of the piecewise linear map f_n to the triangle containing x. We have seen that the ellipses in the image packing $g(B_{n,m})$ are all approximately circles and all about the same size, when m is large.

So $g \circ \phi_n$ is a linear map from an equilateral triangle to an approximate equilateral triangle. By Theorem A.1.1, the Beltrami coefficient of the map $g \circ \phi_n$ is small at x if n is large enough.

Uniformity on compact sets

To get the bound above, we had two requirements. First, ν had to be sufficiently close to $\nu(x)$ near x. Second, the point x_n had to be surrounded by a large enough neighbourhood that looks like the triangular lattice.

Suppose x is in a compact set $F \subset X$. Then ν is uniformly continuous on that set, so the first requirement is satisfied uniformly in F. The second requirement is satisfied with a uniform n, because the distance from x to the complement of X is bounded below.

Therefore, the Beltrami coefficient of the map $g \circ \phi_n$ will be uniformly small for x in that compact set.

The bar-Beltrami coefficient of f_n

Suppose that we write $g \circ \phi_n = \alpha(z + \mu\bar{z})$, as in Lemma 4.1.1. Here $|\mu|$ goes to zero as $n \to \infty$.

Then $f = z + \nu\bar{z}$ and g are inverses, so

$$\begin{aligned}\phi_n = f \circ g \circ \phi_n = f(\alpha(z + \mu\bar{z})) &= \alpha(z + \mu\bar{z}) + \nu\overline{\alpha(z + \mu\bar{z})} \\ &= (\alpha + \nu\bar{\alpha}\bar{\mu})z + (\alpha\mu + \nu\bar{\alpha})\bar{z}.\end{aligned}$$

The bar-Beltrami coefficient is the coefficient of $\bar{z}$ divided by the complex conjugate of the coefficient of z:

$$(\phi_n)_{\bar{z}}/\overline{(\phi_n)_z} = \frac{\alpha\mu + \nu\bar{\alpha}}{\bar{\alpha} + \bar{\nu}\alpha\mu} = \nu + \frac{\alpha\mu - |\nu|^2\alpha\mu}{\bar{\alpha} + \bar{\nu}\alpha\mu} = \nu + \alpha\mu\frac{1 - |\nu|^2}{\bar{\alpha} + \bar{\nu}\alpha\mu}.$$

If μ is small, this is approximately ν. So as $n \to \infty$, the bar-Beltrami coefficient of ϕ_n approaches ν. But ϕ_n is just the restriction of f_n to a neighbourhood of x.

So the bar-Beltrami coefficient of the map f_n at x converges to $\nu(x)$. And this convergence is uniform on compact sets in X, if we leave out the set of measure zero of points where some f_n is not differentiable. □

7.3 Subsequences of f_n converge to a homeomorphism

The Ring Lemma 6.1.1, together with Theorem A.2.1, tells us that the maps f_n are K-quasiconformal.

The image of the map f_n is the union of every triangle in B_n. The Length-Area Lemma 6.2.1 tells us that if n is large, then the boundary ellipses in the image packing B_n are small, so the boundary of the packing must be close to the edge.

It follows that the image of f_n eventually contains every compact set in the unit disc. So we are in the situation of Theorem 5.3.5. We conclude that if f is any subsequential limit of f_n, then it is a K-quasiconformal homeomorphism of X onto the unit disc.

The convergence is uniform, so we have $f(z_0) = \lim f_n(z_0) = 0$, and similarly we have $f(z_1)$ on the real line.

7.4 The bar-Beltrami coefficient of any limit homeomorphism is $\nu(x)$

We have shown that subsequential limits of f_n all converge to homeomorphisms $X \to \{|z| < 1\}$. Replace f_n by such a subsequence, and let f be its limit.

Then we have the following theorem:

THEOREM 7.4.1. *The bar-Beltrami coefficient of f is ν.*

PROOF. The conditions of Theorem 5.3.5 are satisfied. By the proof of that theorem, the inverse maps $g_n = f_n^{-1}$ converge uniformly on compact sets to $g = f^{-1}$. The Beltrami coefficient of the function g_n is $-\nu_n(g_n(x))$, by Lemma 7.1.1.

On a set of measure zero, $\nu_n(g_n)$ does not exist. Suppose x is not in any such set for any n. Then ν_n converges uniformly to ν on compact sets in X, and $g(x) \in X$, so we will have $-\nu_n(g_n(x)) \to \nu(g(x))$ as $n \to \infty$.

But this means that the Beltrami coefficients of the inverse maps converge to $\nu(g(x))$. We can now use the standard Theorem 5.3.4, which tells us that $g(x)$ has Beltrami coefficient $\nu(g(x))$. Using Lemma 7.1.1 in reverse, we find that the bar-Beltrami coefficient of the limit f is $\nu(x)$.

7.5 Existence, but no uniqueness result

What have we proven?

FINAL THEOREM. *Let f_n be the maps from Thurston's revised procedure, with shape field ν.*

Then any subsequence of f_n has a further subsequence that converges to a quasiconformal homeomorphism $f : X \to \{|x| < 1\}$ with $f(z_0) = 0$ and $f(z_1)$ on the real line. The convergence is uniform on compact sets.

The bar-Beltrami coefficient of any such f is ν.

We have done that using the basic theory of quasiconformal functions, without using any theorems on the existence of such a map. So we have at least proven that such maps exist.

7.5.1 The problem with the normalization

Unfortunately, it is not clear to me how to prove that such a map is unique. If we take off the condition that $f(z_1)$ be on the real line, then we will get a one-parameter family of functions. This is a special case of Theorem 9.0.3 in Astala et al. [13].

That theorem uses the normalization that a point on the boundary of X should go to a specified point on the unit circle, or in general, a prime end in X should do so. Under that normalization, these functions f are unique. See [13], Section 9.2.2, page 267.

But this doesn't seem to work in our case. The functions f_n don't necessarily converge uniformly on the boundary, so it's not clear what kind of condition on f_n would force something to happen on the boundary of f.

It is reasonable to expect that some other normalization condition should be able to pin down f exactly, but I haven't been able to find one.

7.6 Future directions

7.6.1 Solving the Beltrami equation instead of the bar-Beltrami equation

We may be able to allow the shape of the ellipses to depend on *their* centre, not the centre of the corresponding circle. This doesn't work with Schramm's blunt theorem [9] in general, because we may not get a 'packable' collection, but there may be some alternative theorem that would allow it.

In that case, we could get a map f that solves the differential equation

$$f_{\bar{z}} = -\mu(f(z))\overline{f_z},$$

say, and then the inverse would solve the Beltrami equation $g_{\bar{z}} = -\mu g_z$. This would give us a geometric way to find solutions of the Beltrami equation.

Appendix A

The distortion of circle and ellipse packings

A.1 Bounds on Beltrami coefficients

In Thurston's procedure, we are piecing together oriented linear maps that take an equilateral triangle to a certain image triangle. That image triangle is defined in terms of an ellipse packing.

It is proven in Chapter 2 that tangent ellipses cannot be too different in size in an ellipse packing of a triangulation of bounded degree. It turns out that that, by itself, gives us a bound on the Beltrami coefficient of the maps in Thurston's procedure. We prove that in this appendix.

Only the first theorem in this appendix is logically necessary. We see in Chapter 2 that if an ellipse is surrounded by other ellipses packed in a large number of generations of the regular hexagonal packing, then adjacent ellipses become almost the same shape and size.

The coefficients of the map are continuous under small changes in the image triangles. The continuity is uniform if the image triangles are bounded away from degeneracy. This is clear from the proof of Theorem A.1.1.

So we know that the maps are *eventually* K-quasiconformal on every compact set, once we have enough generations surrounding every ellipse. This is all we need to use convergence.

The rest of the appendix will prove that *every* map f_n is K-quasiconformal for a uniform K, without using the Rigidity Theorem.

NOTE! Every lemma in this appendix has an unstated assumption. Whenever two circles or ellipses E_1, E_2 are described as tangent, we also assume that their diameters are *comparable.*

That is, Equation A.1 holds for some m:

$$\frac{1}{m} \leq \frac{\operatorname{diam}(E_1)}{\operatorname{diam}(E_2)} \leq m. \tag{A.1}$$

This assumption is justified by the Ring Lemma, Lemma 6.1.1.

A.1.1 Triangle maps: If the angles in the image triangle aren't too small, then μ is bounded.

THEOREM A.1.1. *Let $\theta > 0$. Let T be a triangle with vertices $a, b, c \in \mathbb{C}$. Let $f(z)$ be the oriented linear map that sends 0 to a, 1 to b, and $e^{i\pi/3}$ to c. Write $f(z) = a + \alpha(z + \mu\bar{z})$.*

Suppose no angle of T is less than θ. There exists $k < 1$ depending only on θ so that $|\mu| \le k$. If $c \to e^{i\pi/3}$, then $|\mu| \to 0$.

PROOF. The coefficient μ isn't affected by uniform scaling or rotation of the image, so we can assume that $a = 0$ and $b = 1$.

We have $f(0) = 0$, so $a = 0$. The other data that we have is

$$f(1) = \alpha(1 + \mu) = 1$$
$$f(e^{i\pi/3}) = \alpha(e^{i\pi/3} + \mu e^{-i\pi/3}) = c$$

Therefore,

$$\frac{c - e^{i\pi/3}}{c - e^{-i\pi/3}} = \frac{f(e^{i\pi/3}) - e^{i\pi/3} f(1)}{f(e^{i\pi/3}) - e^{-i\pi/3} f(1)} = \frac{\mu\alpha(e^{-i\pi/3} - e^{i\pi/3})}{\alpha(e^{i\pi/3} - e^{-i\pi/3})} = -\mu. \qquad \text{(A.2)}$$

This already shows that if $c \to e^{i\pi/3}$ then $|\mu| \to 0$.

We want to bound the absolute value of $(c - e^{i\pi/3})/(c - e^{-i\pi/3})$ by some $k < 1$. Let $\gamma = e^{i\pi/3}$. Then $|\mu| \le k$ is equivalent to the inequality

$$c\bar{c} - \bar{\gamma}c - \gamma\bar{c} + 1 \le k^2(c\bar{c} - \gamma c - \bar{\gamma}\bar{c} + 1).$$

Completing the square, this is

$$\left| c - \frac{\gamma - k^2\bar{\gamma}}{1 - k^2} \right|^2 \le \frac{k^2}{(1 - k^2)^2}.$$

This is a circle of radius $k/(1-k^2)$ and centre $\left(\frac{1}{2}, \frac{1}{2}(1+k^2)/(1-k^2)\right)$. The lowest point of the circle is $\left(\frac{1}{2}, \frac{1}{2}(1-k)/(1+k)\right)$. As $k \to 1$, that converges to $\left(\frac{1}{2}, 0\right)$. As long as the third point c is inside this circle, the distortion will be less than or equal to k. Figure A.1 shows what this looks like.

The requirement that the angles of the triangle are at least θ puts two conditions on $|c|$:

- If $|c|$ is reasonably large, then the angle $\angle acb$ will be small. It is largest when the triangle is isoceles, and in that case $|c|(2\sin\theta_{\angle acb}/2) = 1$, or $|c| \lesssim 1/\theta$.

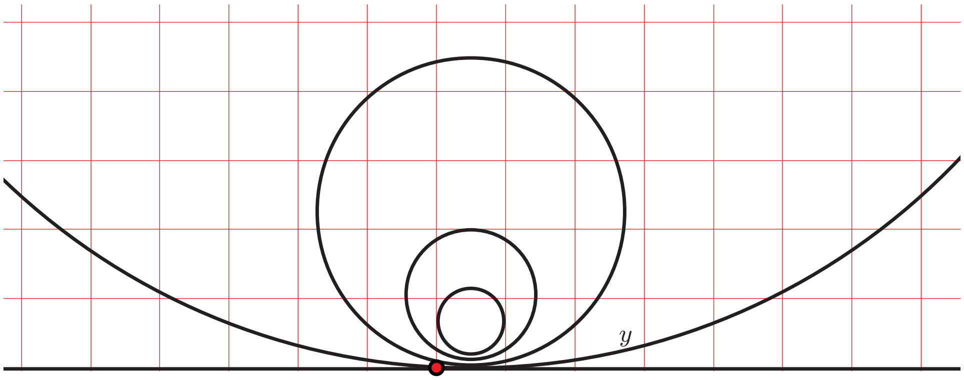

Figure A.1: The acceptable values of c, for $k = 0.4, 0.6, 0.8, 0.95$. The red circle is at the origin.

- If $\operatorname{Re} c > 1/2$, then $1/2 \sin \theta_{\angle bac} \leq |c| \sin \theta_{\angle bac} = \operatorname{Im} c$. If $\operatorname{Re} c \leq 1/2$, then we get $1/2 \sin \theta_{\angle abc} \leq \operatorname{Im} c$. In either case, we get $\operatorname{Im} c \geq 2 \sin \theta$.

Fixing θ, it is possible to pick k large enough that any point meeting both these conditions is inside the circle above. If we do that, then $|\mu| \leq k$. This proves the result. $\square$

Corollary. *Let $f(z)$ be an oriented linear map that sends an oriented equilateral triangle A, B, C to a, b, c. Write $f(z) = f(z_0) + \alpha(z + \mu\bar{z})$.*

Suppose no angle of $\triangle abc$ is less than θ. Then $|\mu| \leq k$, where k is the same bound that we got in Theorem A.1.1.

Proof. Let $h(z) = A + Bz$. This sends $0, 1, e^{\pi i/3}$ to A, B, C. Then

$$h \circ g(z) = f(z_0) + \alpha(A + \mu\bar{A}) + \alpha B\left(z + \mu\frac{\bar{B}}{B}\bar{z}\right)$$

takes $0, 1, e^{\pi i/3}$ to a, b, c.

By Theorem A.1.1, $|\mu\bar{B}/B| = |\mu|$ is bounded by k. $\square$

A.1.2 The angle is not too small in circle packings

Lemma A.1.2. *Suppose we have one circle C_1 and two others C_2, C_3 tangent to it, and that their interiors are all disjoint.*

Let the centres of the circles be called c_1, c_2, c_3. Let their radii be r_1, r_2, r_3. Let the tangent points be called t_{12}, t_{13}.

The angle $\angle t_{12}c_1t_{13}$, the angle of the arc connecting the tangent points, is bounded below by a constant depending only on m.

PROOF. The angle between the tangent points is the same as the angle between the centres of the circles $\angle c_2c_1c_3$. The rest follows immediately from the cosine law, Theorem A.1.1, and Equation A.1.

We recall the cosine law. Let $\triangle ABC$ be a triangle which has side lengths $a = \overline{BC}, b = \overline{AC}, c = \overline{AB}$, and where the angle $\angle ACB$ is θ. Then

$$c^2 = a^2 + b^2 - 2ab\cos\theta$$

or $\cos\theta = (a^2+b^2-c^2)/2ab$. The distance between the centres of C_2 and C_3 is at least $r_2 + r_3$, and the other circles are tangent, so we can write

$$\begin{aligned}\cos\theta \leq \frac{(r_1+r_2)^2+(r_1+r_3)^2-(r_2+r_3)^2}{2(r_1+r_2)(r_1+r_3)} &= \frac{2r_1r_2+2r_1r_3+2r_1^2-2r_2r_3}{2(r_1+r_2)(r_1+r_3)} \\ &= 1-\frac{2r_2r_3}{(r_1+r_2)(r_1+r_3)} \\ &\leq 1-2\left(\frac{1}{m+1}\right)^2\end{aligned}$$

because $r_2/(r_1+r_2) = 1/(\frac{r_1}{r_2}+1) \geq 1/(m+1)$

We have $1-\frac{1}{2}\theta^2 \leq \cos\theta \leq 1-2\left(\frac{1}{1+m}\right)^2$. So θ can't be any smaller than $2/(1+m)$. $\square$

A.1.3 The curvature of an ellipse

Suppose E is an ellipse with diameter d which has distortion K. Rotate it so that the major axis is horizontal, and translate it so that the centre is zero. Then the boundary curve is

$$\mathbf{x}(\theta) = \left(\frac{d}{2}\cos\theta, \frac{d}{2K}\sin\theta\right). \qquad \mathbf{x}'(\theta) = \left(-\frac{d}{2}\sin\theta, \frac{d}{2K}\cos\theta\right)$$

We can find the curvature in the sense of differential geometry, and it is no greater than $2K/d$ at any point on the curve.

$$\kappa(\theta) = \left|\frac{d}{ds}\left(\frac{d}{ds}\mathbf{x}'(\theta)\right)\right| = \left|\frac{\mathbf{x}'(\theta)\times\mathbf{x}''(\theta)}{|\mathbf{x}'(\theta)|^3}\right| = \frac{(d/2)^2/K}{\sqrt{\left((\frac{d}{2})^2\sin^2\theta+(\frac{d}{2K})^2\cos^2\theta\right)^3}} \leq \frac{2K}{d}.$$

The curvature does not depend on rotation or translation, so any ellipse as above has curvature at most $2K/d$.

A.1.4 The angle is not too small in an ellipse

LEMMA A.1.3. *Suppose we have one ellipse E_1 and two others E_2, E_3 tangent to it, so that their interiors are all disjoint and no ellipse has distortion more than K.*

Let the centre of E_1 be e_1. Let the tangent points $E_1 \cap E_2$ and $E_1 \cap E_3$ be t_{12} and t_{13}.

Then the angle between the tangent points, $\angle t_{12}\, e_1\, t_{13}$, is bounded below by a constant depending only on m and K.

PROOF. A remark. Let E be any ellipse with diameter d which has distortion K. The curvature of the boundary of that ellipse is never more than $2K/d$, so we can pick any point on the border of the ellipse and inscribe a circle of diameter $d/2K$ in the ellipse touching that point.[6]

PROOF. Suppose E_1 is a general ellipse. Apply a non-uniform scaling to make it a circle. The distortion of the other ellipses may increase, but not to more than K^2.

At this point, we have one circle and two disjoint ellipses tangent to it. Now use the remark. The tangent points stay the same, and the diameters decrease by a bounded factor (at most K^2).

Now use Lemma A.1.2 to see that the angle of the arc between the tangent points is bounded below. The ratio of diameters may have increased to mK^2,

Finally, reverse the linear map. This might change the angles by a factor of K, but not more. Therefore, the angle is still bounded below by $2mK^3/(1+mK^2)$. □

LEMMA A.1.4. *The other two angles, $\angle e_1\, t_{12}\, t_{13}$ and $\angle e_1\, t_{13}\, t_{12}$, are also bounded below by a constant depending only on m and K.*

[6] This comes from the following general fact. Suppose X is a convex shape with a twice-differentiable border, and the curvature of the border is never more than κ. Let $x \in \partial X$. Then the circle of radius $1/\kappa$ tangent to the border at x is contained in X.

Sketch of a proof: let the tangent point be at $(0,0)$, the tangent be horizontal, and X be above the curve $\mathbf{x}(t) = (x(t), y(t)) \in \mathbb{R}^2$. For positive t,

$$\mathbf{x}(t) = \int_0^t \mathbf{e}_\theta \, ds = \tfrac{1}{\kappa} \int_0^{\theta(t)} \mathbf{e}_\theta \, d\theta + \int_0^{\theta(t)} \mathbf{e}_\theta \left(\tfrac{ds}{d\theta} - \tfrac{1}{\kappa}\right) d\theta.$$

If $0 \leq \theta(t) \leq \pi$, the right-hand summand is in the cone between 0 and $\theta(t)$ degrees, so the curve is a circle plus a vector in the half-plane determined by the outer normal. At $\theta = \pi$, the y coordinate is at least $2/\kappa$, which is the height of the circle. The same thing is true for negative t. Draw a line connecting the two endpoints with $\theta(t) = \pm\pi$; the resulting set contains the circle and is inside the convex set X.

A.1.5 The angle is not too small in an interstice

LEMMA A.1.5. *Let E_1, E_2, E_3 be mutually tangent ellipses. Let t_{12}, t_{23}, t_{31} be the tangent points. Suppose no ellipse has distortion any more than K. Then the angles in the triangle $t_{12}t_{23}t_{31}$ are bounded below.*

PROOF. The boundary of an interstice is three sub-arcs of an ellipse. For an angle to be small, at least one of the ellipse arcs would have to connect the endpoints of the long side of a sharp triangle.

Then it would have to be long and narrow. There are two ways this can happen. Either the angle of the sub-arc in the ellipse is small, or the ellipse is very distorted. The first possibility is ruled out by Lemma A.1.3, and the second one is ruled out by the bound on K. □

A.2 Thurston's procedure and μ

THEOREM A.2.1. *Let B_n be a circle packing from Thurston's original procedure, or an ellipse packing from Thurston's revised procedure.*

Let $m > 0$. Suppose that none of the ellipses in B_n have distortion more than K, and that for every tangent pair of ellipses E_1, E_2 in B_n,

$$m \operatorname{diam}(E_1) \leq \operatorname{diam}(E_2) \leq \frac{\operatorname{diam}(E_1)}{m}.$$

Then there is a constant $k < 1$, depending only on m and K, so that

$$\frac{\partial f_n}{\partial \bar{z}} = \mu \frac{\partial f_n}{\partial z},$$

with $|\mu| \leq k$ almost everywhere in the domain of definition of f_n.

PROOF. First, we show that the angles in the image triangles are bounded below. There are several different cases, each of which is covered by a lemma in this chapter.

Case I: The maps come from Thurston's original procedure. If the radii of adjacent circles are comparable, then the angles in every image triangle are bounded from below by Lemma A.1.2.

In Thurston's revised procedure, there are two different kinds of image triangles: triangles connecting the centre of an ellipse to two boundary points, and triangles covering an interstice.

Case IIa: The maps come from Thurston's revised procedure, and the triangle is contained in an ellipse. Lemma A.1.3 shows that the angle at the centre is bounded below, and Lemma A.1.4 shows that the angles at the tangent points are too.

Case IIb: The map comes from Thurston's revised procedure, and the triangle contains an interstice. Lemma A.1.5 shows that the angles in interstices are bounded below.

Every image triangle has angles greater than some uniform lower bound, so the Beltrami coefficient of every linear map in the definition of f_n is less than or equal to k for some uniform $k < 1$.

But almost every point in the domain of f_n is in one of those linear maps, so f_n is K-quasiconformal for $K = (1 + k)/(1 - k)$.

Appendix B

Schramm's blunt packing theorem and Theorem 3.1.1.

B.1 A packing theorem for ellipses

In this appendix we explain how to get the ellipse packing Theorem 3.1.1 from Schramm's blunt packing theorem.

The big difficulty is that Schramm's theorem lets us fix two adjacent sets F_a and F_b on the *boundary*. We want to choose a single vertex v not on the boundary and have the centre of that ellipse be zero.

Here is the theorem we start from, quoted from page 325 of [9].

THEOREM B.1.1. (Schramm) *Let T be a triangulation of $\mathbb{S}^2$ with vertex set V, let $[a, b, c]$ be a triangle in T, and let D be a decent trilateral $D = (D_a, D_b, D_c)$ in the sphere. For every vertex $v \in V \setminus \{a, b, c\}$, let $\mathfrak{F}_v$ be a packable collection of sets in D.*

Then there exists a unique packing $P = (P_v : v \in V)$ of T in the sphere whose tangency graph is T and which satisfies $P_v \in \mathfrak{F}_v$, $v \in V \setminus \{a, b, c\}$ and $P_v = D_v$, $v = a, b, c$.

PROOF. See the paper [9]. The reader may well ask what makes a collection 'packable' and a trilateral 'decent,' but we will not repeat the definitions here; see pages 324–325. $\square$

The class of ellipses with a certain shape is just a scaling and rotation of the 'packable' collection of circles, so it is also 'packable' in the sense of Schramm. See the paper for more details.

So Theorem B.1.1 applies to ellipses. We get the following corollary:

COROLLARY B.1.2. *Let T be an oriented planar triangulation with an outside edge. For every vertex in the triangulation, let E_v be some fixed ellipse. Let a and b be two <u>adjacent</u> vertices in the triangulation that <u>are</u> on the boundary of the graph.*

Fix two tangent ellipses E_a, E_b which are both tangent to the boundary of the unit disc. There is a unique <u>oriented</u> packing of T in the unit disc so *that F_v has the same shape as E_v for every vertex v, and $F_a = E_a$, $F_b = E_b$.*

Proof. The triangulation in Theorem B.1.1 is a triangulation of the sphere, and we want to use triangulations of the unit disc. We will therefore adjoin a point at ∞ and we will demand that E_∞ be the complement of the unit disc.

Then we use Schramm's theorem with the trilateral formed by the inside boundary of E_a, E_b, E_∞. There are two choices of inside boundary, and we choose the one that respects the orientation of the packing.

Every point on the boundary is adjacent to ∞, so it is tangent to the unit disc, and the packing has the desired tangency graph. □

B.2 From that to Theorem 3.1.1

In Corollary B.1.2, we are only allowed to pick two adjacent vertices a, b on the boundary. On the other hand, in Theorem 3.1.1, we may pick any two vertices v, u, as long as v is not on the boundary. So we must study the behavior of the interior vertices as F_a, F_b change.

It is natural to begin by finding a nice parametrization of the ways to choose F_a, F_b.

B.2.1 Parametrization of F_a, F_b

We fix two ellipses E_a, E_b.

Suppose F_a, F_b are ellipses with the same shape as E_a, E_b respectively, and that both ellipses are tangent to each other and the unit circle, as in Corollary B.1.2.

The class of all possible choices of F_a, F_b can be written as a three-parameter family, $Z = \{|z| < 1\} \times \mathbb{S}^1$.

To see this, let $(x, \theta) \in Z$. Then:

Lemma B.2.1. *There is a unique choice of F_a and F_b so that the tangent point is at x and the tangent line at x is at angle θ to the positive real axis with F_a on the left.*

Proof. Any ellipse has exactly one point on its boundary where the angle of the tangent is θ. So x, θ determine everything but the size of each ellipse.

We are free to pick the size of the two ellipses, keeping the tangent point fixed. We must show that there is exactly one choice with both ellipses tangent to the unit circle.

If we scale up an ellipse around the tangent point while keeping the unit circle fixed, it is the same as shrinking the unit circle around the tangent point, while keeping the ellipse the same size. That is, the two pictures are 'similar' in the sense of geometry:

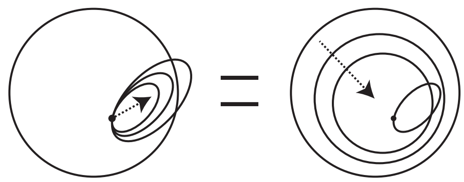

And in the second picture, the copies of the unit circle with different scaling are strictly contained in each other. So only one scaling will result in a tangent point without overlap.

Each ellipse has a unique possible scale, so the arrangement is determined by x, θ, and if we know the ellipses we can determine x, θ, so the map is bijective. □

We imagine each ellipse as an element in $\mathbb{R}^2 \times (0, \infty)$, where the first coordinate is the centre and the second one is the diameter. Then,

Lemma B.2.2. *Let Φ be the map $(x, \theta) \mapsto (F_a, F_b)$. Then Φ is continuous on Z, and the inverse is continuous on $\Phi(Z)$.*

Proof. The map is continuous because a slight change in x or θ can be corrected by a slight scaling.

The inverse takes a pair of ellipses to their tangent point and tangent angle, and that is clearly continuous on the subset $\Phi(Z)$ of $(\mathbb{R}^2 \times (0, \infty))^2$ where the ellipses are tangent.

□

B.2.2 Parametrization of the packing

Let $(x, \theta) \in Z$. There is a unique packing F given by Schramm's theorem so that the ellipses at a and b are $\Phi(x, \theta)$ respectively, and that map depends continuously on x and θ.

If the map is not continuous, we can find a subsequential limit of packings F^n so that the tangent point and angle converge to x and θ, but the other ellipses converge to something other than F_a. This would contradict uniqueness.

We have two points v, u that we want to look at. Let Ψ_v be the map that takes a point (x, θ) in Z to the centre of the ellipse F_v in that packing, and similarly for Ψ_u. Let Θ_u be the argument of the complex number $\Psi_u - \Psi_v$.

THEOREM B.2.3. *Suppose that every border vertex in T is connected to an interior vertex, and the set of interior vertices is connected.*

There is an annulus $N_r = \{r < |x| < 1\}$ depending on the graph and the ellipse shapes so that $|\Psi_v(x, \theta) - x| < 1/2$ whenever $x \in N$.

PROOF. If $|x|$ is very close to 1, then at least one of the two ellipses in $\Phi(x, \theta)$ must be small. Some interior ellipse is tangent to the small one.

The Ring Lemma states that any interior ellipse tangent to a small ellipse must be small too. For a fixed finite graph, the bound is uniform. The set of interior vertices is connected, so if $x \in N_r$ with r very close to one, we can force every interior ellipse to be uniformly small.

One of those ellipses is F_v, and we can make it arbitrarily close to x. This proves the result. $\square$

Now we can prove the existence theorem:

THEOREM B.2.4. *There exists a packing (x, θ) with F_v at the origin and F_u on the positive real line.*

REMARK. Let $r : Z \to Z$ be the map $r(x, \theta) = (-x, \theta + \pi)$. Then the pair of ellipses $\Psi(r(x, \theta))$ are the ellipses $\Psi(x, \theta)$ reflected $180°$ around the origin. Ellipses are $180°$ rotationally symmetrical, so if we reflect a packing around the origin, we get another packing so that all the ellipses have the correct shape, and their new center is the negative of their original center.

One implication of this is that, if we have a packing with F_v at the origin and F_u on the negative real line, we can reflect it around the origin to get a packing satisfying the conditions of the theorem.

PROOF. Suppose that there is no such packing. Then $(x, \theta) \mapsto (\Psi_v, \Theta)$ is a map from $\{|x| < 1\} \times \mathbb{S}^1$ into the set $\{|x| < 1\} \times \mathbb{S}^1$ with two points missing: $(0, 0)$ is missing by assumption, and, by the remark, $(0, \pi)$ is also missing.

Let M_r be the set $N_r \times \mathbb{S}^1$. Let D_θ be the disc $\{|x| \leq r\} \times \{\theta\}$, whose boundary is in M_r for every θ.

The basic idea is that D_θ is a continuous family of discs, and the edge is always in M_r. As θ goes from zero to π, the image of D_θ in the torus below travels halfway around and winds up exactly at its reflection.

But that should not be possible: the disc has to wrap around one of the missing points, so it can't wind up back at itself.

We now say the same thing in terms of relative singular homology. We start with a filled torus with two missing points, represented by dots.

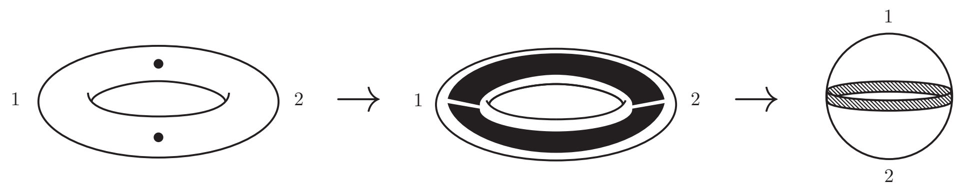

We smash the middle of the the regions 1 and 2 into two thin sheets, which form the top and bottom of the sphere. Everything else goes into a one-dimensional line around the center, represented by the shaded ring.

Call this smashed sphere S. Call the ring ρ. We have a map $\sigma : Z \to S$, the composition of the map (Ψ_v, Θ) and the smashing.

By Lemma B.2.3, $\Phi_v(M_r)$ is contained in a small layer on the outside of the torus, so $\sigma(M_r)$ is contained in ρ. Also, the map r passes through this smashing, and acts on S by reflecting the sphere through the horizontal plane and then rotating the whole thing 180 degrees around the axis.

Let D_θ be the disc $\{|x| \leq r\} \times \{\theta\}$. We think of this as a simplex in the second chain group. Its boundary is contained in M_r, so it is an element in the second relative homology group $H_2(Z, M_r)$. Then Lemma B.2.3 says that the image of ∂D_θ under σ wraps once around that shaded ring. That is, it is a generator for $H_1(\rho) \cong \mathbb{Z}$.

Expanding the ring a little and using excision, we can see that $H_2(S, \rho)$ is isomorphic to $\mathbb{Z} \times \mathbb{Z}$. The first generator, say α, is the top half of the sphere and the second one, say β, is the bottom half.

We have a boundary map $\partial : H_2(S, \rho) \to H_1(\rho) \cong \mathbb{Z}$, and we may choose orientations on α and β so that $\partial\alpha = \partial\beta = 1$. That is, the boundaries of both α and β wrap once counterclockwise around the ring ρ. Then $r : S \to S$ takes α to β, and vice versa.

The boundary $\sigma_*(\partial D_\theta) = \partial\sigma_*(D_\theta)$ is an element of homology that wraps once counterclockwise around ρ, so $\sigma_*(D_\theta)$ must be $(n+1)\alpha + n\beta$ for some integer n. Then $r_*(\sigma_*(D_\theta)) = n\alpha + (n+1)\beta$, so the reflection map r does not preserve $\sigma_*(D_\theta)$.

In particular, it doesn't preserve the class of $\sigma_*(D_0)$. But r takes D_0 to D_π, and those two discs are equal in homology.

We should have $\sigma_*(D_0) = \sigma_*(D_\pi) = \sigma_*(r(D_0)) = r_*(\sigma_*(D_0))$, but in fact they are different. This is a contradiction, so the assumption is wrong. The two points cannot be missing from the torus. $\square$

B.3 A uniqueness problem, again

This result, together with Corollary B.1.2, shows that there exists a packing of ellipses in the unit disc with the properties in Theorem 3.1.1.

It does *not* show that such a packing is unique. Although Schramm's theorem gave us uniqueness, we have lost it along the way.

We mentioned that there is not a clear choice of normalization for the bar-Beltrami equation, and there is a similar problem here. If one could find a good normalization for the finite packing in Theorem 3.1.1 so that the result was unique, it would probably lead to a full and satisfying uniqueness conclusion in the Final Theorem.

Bibliography

[1] W. P. Thurston, "The finite Riemann mapping theorem," March 1985. An invited talk given at the Symposium on the Occasion of the Proof of the Bieberbach Conjecture, at Purdue University.

[2] W. P. Thurston, "The geometry and topology of 3-manifolds," 2002. URL: `http://library.msri.org/books/gt3m/`. These were originally released as lecture notes in 1980.

[3] E. M. Andreev, "On convex polyhedra in Lobačevskiĭ spaces," *Mathematics of the USSR-Sbornik*, vol. 10, no. 3, p. 413, 1970.

[4] R. Kuhnau, *Handbook of Complex Analysis: Geometric Function Theory, Volume 2*, ch. 1, p. 141. Elsevier, 2005.

[5] P. Koebe, "Kontaktprobleme der konformen Abbildung," *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physikalische Klasse*, vol. 88, pp. 141–164, 1936.

[6] K. Stephenson, "Circle packing: a mathematical tale," *Notices of the American Mathematical Society*, vol. 50, pp. 1376–1388, 2003.

[7] G. R. Brightwell and E. R. Scheinerman, "Representations of planar graphs," *SIAM Journal of Discrete Math*, vol. 6, no. 2, pp. 214–229, 1993.

[8] B. Rodin and D. Sullivan, "The convergence of circle packings to the Riemann mapping," *Journal of Differential Geometry*, vol. 26, pp. 349–360, 1987.

[9] O. Schramm, "Existence and uniqueness of packings with specified combinatorics," *Israel Journal of Mathematics*, vol. 73, no. 3, pp. 321–341, 1991.

[10] G. E. Bredon, *Topology and Geometry*. Springer, 1993.

[11] O. Lehto and K. I. Virtanen, *Quasiconformal mappings in the plane.* Springer, 1973.

[12] L. Ahlfors, *Lectures on quasiconformal mappings.* American Mathematical Society, 1966.

[13] K. Astala, T. Iwaniec, and G. Martin, *Elliptic partial differential equations and quasiconformal mappings in the plane.* Princeton University Press, 2008.

[14] D. Sullivan, "On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions," in *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference*, pp. 465–496, 1981.

[15] O. Schramm, "Rigidity of infinite (circle) packings," *Journal of the American Mathematical Society*, vol. 4, no. 1, pp. pp. 127–149, 1991.