MAT 627, Spring 2025, Stony Brook University

Topics in Complex Analysis: Quasiconformal Mappings

Christopher Bishop

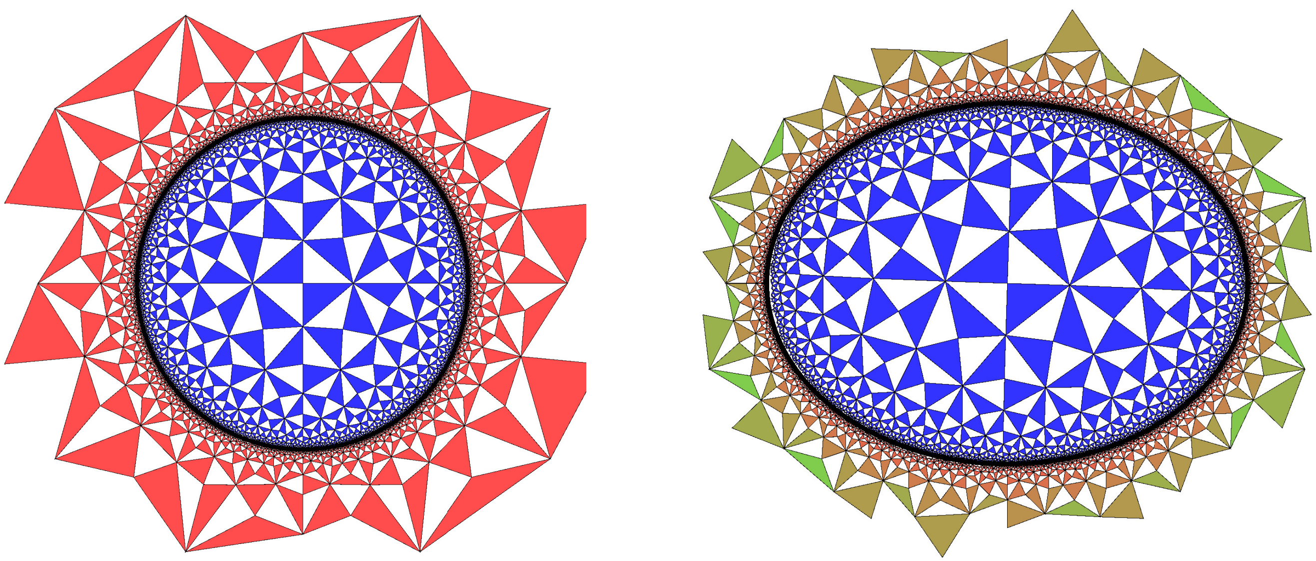

This semester I hope to cover the following topics:

- Review of complex analysis
- Extremal length and conformal modulus,
- Logarithmic capacity, harmonic measure
- Geometric definition of quasiconformal mappings, compactness
- Compactness corollaries: quasisymmetry, extension, removability, weldings
- Analytic definition and the measurable Riemann mapping theorem
- Analytic dependence on the dilatation
- Astala's theorems on area and dimension distortion
- More topics?: QC maps metric spaces, David maps, conformal dimension,...

Quasisymmetric maps

A homeomorphism $h : \mathbb{R} \to \mathbb{R}$ is called M**-quasisymmetric** if $|h(I)| \leq M|h(J)|$ whenever I and J are adjacent intervals of equal length. Equivalently,

$$\sup_{t \in \mathbb{R}, x > 0} \frac{h(x+t) - h(t)}{h(t) - h(x-t)} \leq M.$$

A homeomorphism is called **quasisymmetric** if it is M-quasisymmetric for some $M < \infty$. Later we will discuss quasisymmetric map of the unit circle to itself, but for the moment we stick to maps of $\mathbb{R}$ to $\mathbb{R}$.

The cross ratio of four points a, b, c, d on the real line

$$\frac{(a-c)(b-d)}{(b-c)(a-d)},$$

and is equal to

$$\frac{a-c}{b-c},$$

if $d = \infty$.

When f is M-quasisymmetric on $\mathbb{R}$ and a, b, c, d equal $x+t, x-t, x, \infty$, the cross ratio is -1. The cross ratio of the image points is between $-M$ and $-1/M$.

Theorem 5.1. *A homeomorphism $h : \mathbb{R} \to$ reals is quasisymmetric if and only if it extends to a quasiconformal mapping of the plane to itself.*

Proof. First we show that if f is a K-quasiconformal map of the plane that maps $\mathbb{R}$ to itself, then the restriction of f to $\mathbb{R}$ is quasisymmetric.

Without loss of generality we may assume $I = [0, 1/2]$ and $J = [1/2, 1]$ and that f fixes 0 and 1.

Consider the modulus of the topological annulus $A = \mathbb{C} \setminus ([0,1] \cup [2, \infty)$. This has a fixed finite, non-zero modulus, so its image $B = f(A) = \mathbb{C} \setminus ([0, x] \cup [1, \infty))$ also has modulus bounded between two positive real numbers that depend only on K.

If $x = f(1/2)$ is too close to 0 or 1, then B clear has modulus close to 0 or ∞ respectively, a contradiction. Thus x is bounded away from both 0 and 1 with an estimate depending only on K, and hence h is M-quasiconformal with a constant depending only on K.

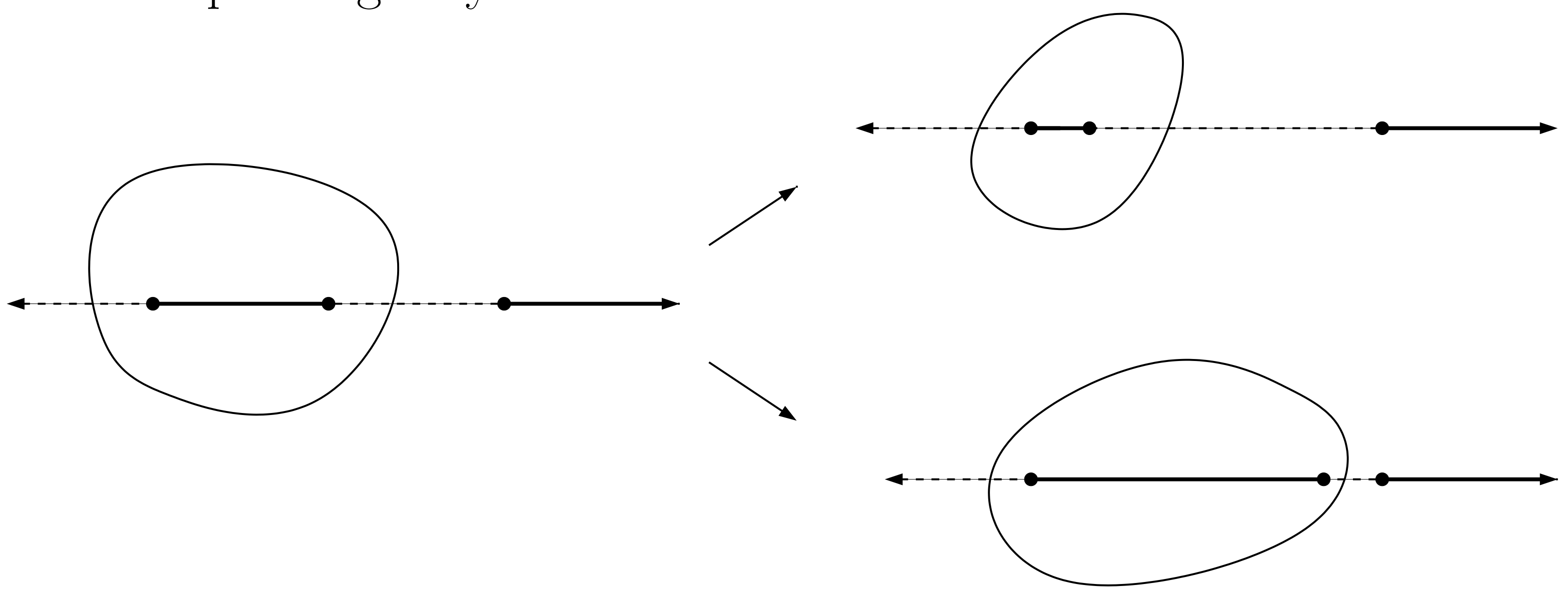

Next suppose $h : \mathbb{R} \to \mathbb{R}$ is M-quasisymmetric. We will assume h is increasing; the other case is handled by a similar argument. We will use the fact that the hyperbolic upper half-plane can be tesselated by hyperbolically identical right pentagons. The corresponding picture for the disk is shown below.

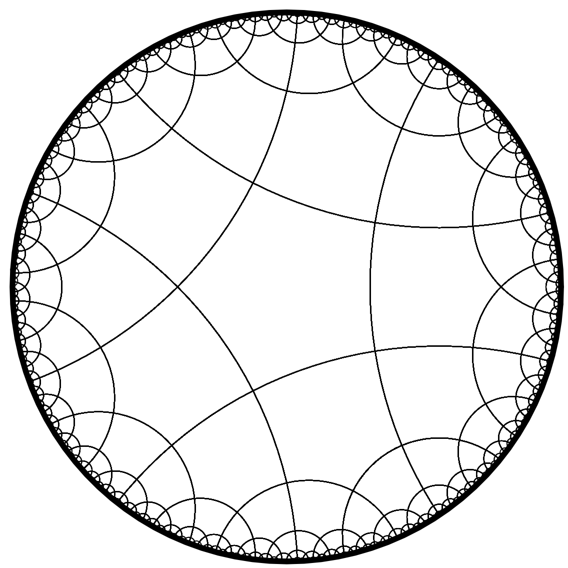

Hyperbolic space is tesselated by hyperbolically identical right pentagons. There is a corresponding picture on the upper half-plane model.

Each right pentagon in the tesselation of the upper half-plane determines five hyperbolic geodesics containing its sides, and these determine ten distinct points on the real line.

The h images of these point are also ten distinct points and the same pairs of point determine five new geodesics that define a hyperbolic pentagon (it need not be regular or right).

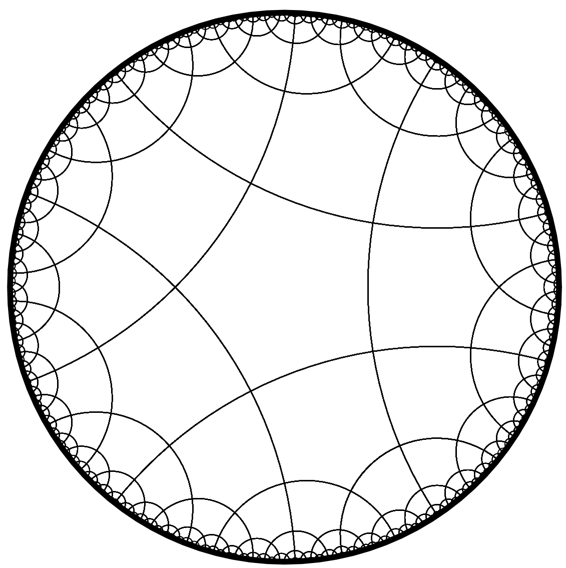

There is a diffeomorphism of the right pentagon to this new one that preserves arc-length along the edges in the sense that on each side of the pentagon length are multiplied by the ratio of the image length over the starting length.

This ensures that the diffeomorphisms defined on adjacent pentagons agree on the common sides. These diffeomorphisms come from a compact family of possibilities, thus have uniformly bounded dilatations, and hence define a quasiconformal map of the half-plane to itself that agrees with h on the boundary.

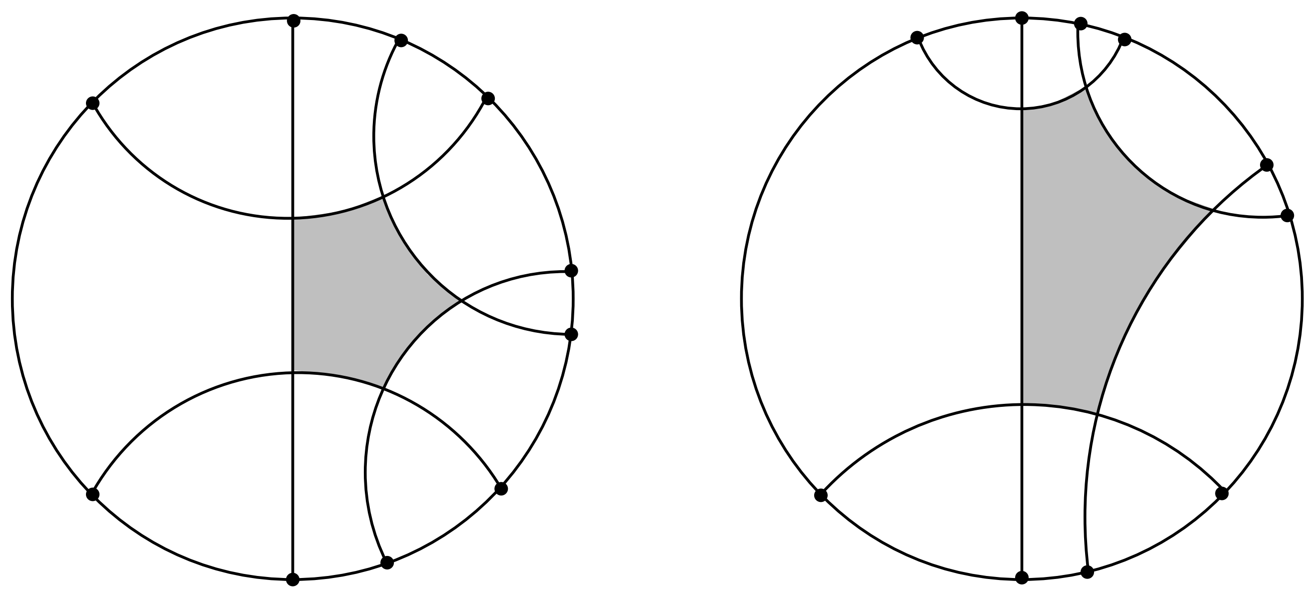

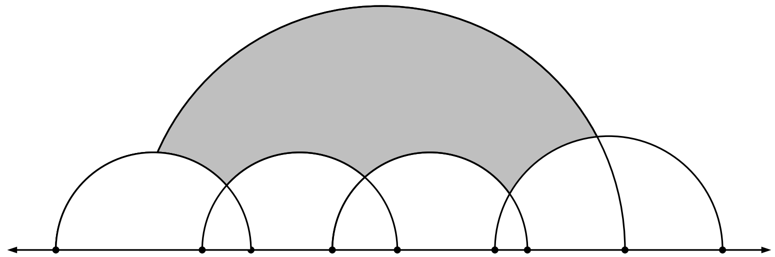

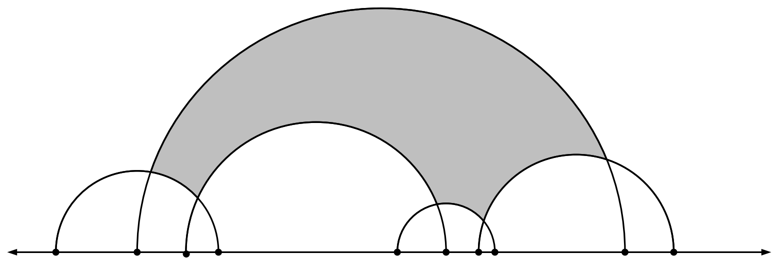

Sides of a hyperbolic right pentagon determine 5 geodesics and 10 boundary points. The images of these 10 points determine 5 geodesics, which give a hyperbolic pentagon.

We take any QC map between the pentagons that multiplies hyperbolic arclength on each edge by a constant (the ratio of the lengths of an edge and it image).

$\square$

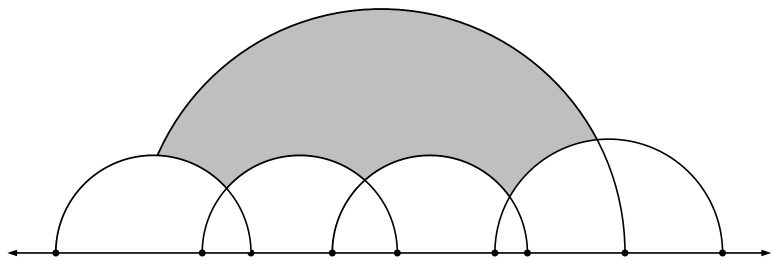

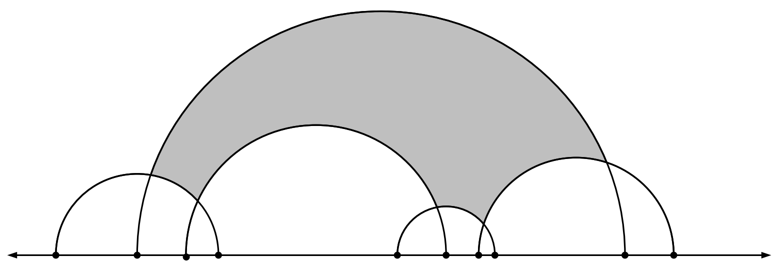

This proof is just a more hyperbolic version of a proof due to Jerison and Kenig using a tiling of the upper half-plane by rectangles (upper halves of dyadic Carleson squares).

There are several other well known extensions. We mention two without proof.

Beurling-Ahlfors extension: Given a quasisymmetric homeomorphism f on the real line define

$$u(z) = \int_0^1 f(x+ty)dt = \frac{1}{y}\int_x^{x+y} f(s)ds$$
$$v(z) = \int_0^1 f(x-ty)dt = \frac{1}{y}\int_{x-y}^{x} f(s)ds$$

and set

$$F(z) = \frac{1+i}{2}\left(u(z)+iv(z)\right).$$

Douady-Earle extension, 1986: this gives an extension E from $\mathbb{T}$ to $\mathbb{D}$ that is C^∞, biLipschitz in the hyperbolic metric (hence quasiconformal) and conformally natural, i.e.g., for any Möbius transformations ϕ and ψ, $E(\phi \circ f \circ \psi) = \phi \circ E(f) \circ \psi$. Let

$$G(z, w) = \frac{1}{2\pi} \int_{\mathbb{T}} \frac{f(\zeta) - w}{1 - \overline{w} f(\zeta)} \frac{1 - |z|^2}{|z - \zeta|^2} |d\zeta|.$$

If $z \in \mathbb{D}$ there is a unique point w so $G(z, w) = 0$. We set $F(f)(z) = w$.

If $z = 0$, we apply Möbius transformations to f until its "average" lies at the orgin.

A different equivariant extension was given by Tukia in 1985.

It was a question of Dennis Sullivan whether there was an extension operator from quasisymmetric maps on the circle to quasiconformal maps of the disk, that was a homomorphism with respect to composition.

In 2007 Epstein and Markovic proved there is no such extension.

Quasicircles

We say that a curve γ satisfies the **3-point condition**, if there is a $M < \infty$ so that given any $x, z \in \gamma$ and y on the smaller diameter arc $\gamma(x, y) \subset \gamma$ between x, y, we have

$$|x - y| \leq M|x - z|,$$

Equivalently,

$$\operatorname{diam}(\gamma(x, z)) \leq M|x - z|.$$

This is also called the **Ahlfors M-condition** or **bounded turning**.

It is immediate from Lemma 4.5 that the image of the real line under any quasiconformal mapping of the plane is bounded turning, and below we shall prove the converse is also true.

The similar looking, but stronger, condition

$$\ell(\gamma(x,z)) \leq M|x-z|$$

where we assume γ is locally rectifiable is called the chord-arc condition. Such curves are called **chord-arc curves** or **Lavrentiev curves**, and form a special, but very important, subclass of the bounded turning curves.

It turns out that chord-arc curves are exactly the images of the real line under bi-Lipschitz maps of the plane, but we will not prove this here.

Lemma 5.2. *Suppose* γ *is bounded turning with constant* M *and* $0, 1, \infty \in \gamma$.

Suppose Ω *is one of the connected components of* $\mathbb{C} \setminus \gamma$ *and suppose* x *is a point on* γ *between* 0 *and* 1.

Let $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ *denote the disjoint subarcs of* γ *from* $-\infty$ *to* 0, *from* 0 *to* x, *from* x *to* 1 *and from* 1 *to* $+\infty$ *respectively.*

Let Γ *be the path family joining the arc* $\gamma_x \subset \gamma$ *from* 0 *to* x *to the disjoint half-infinite arc* $\gamma_1 \subset \gamma$ *joining* 1 *to* ∞.

Then $M(\Gamma) \to 0$ *as* $x \to 0$ *with upper and lower bounds that depend only on* $|x|$ *and* M

Proof. The 3-point condition implies that

$$\operatorname{dist}(\gamma_2, \gamma_4) \geq \frac{1}{M} - |x|,$$

so for $|x|$ sufficiently small every path in Γ crosses the round annulus

$$\{z : M|x| < |z| < \frac{1}{2M}\} \subset \{z : \operatorname{diam}(\gamma_2) < |z| < \operatorname{dist}(\gamma_2, \gamma_4)\}.$$

For $|x|$ small, this implies $M(\Gamma)$ is small.

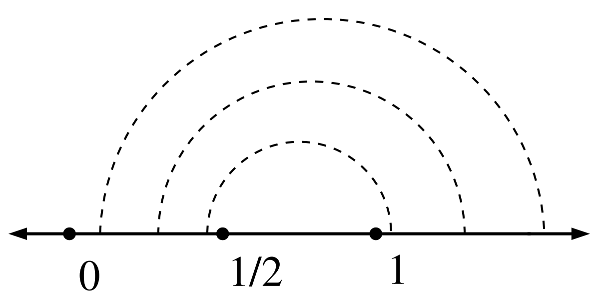

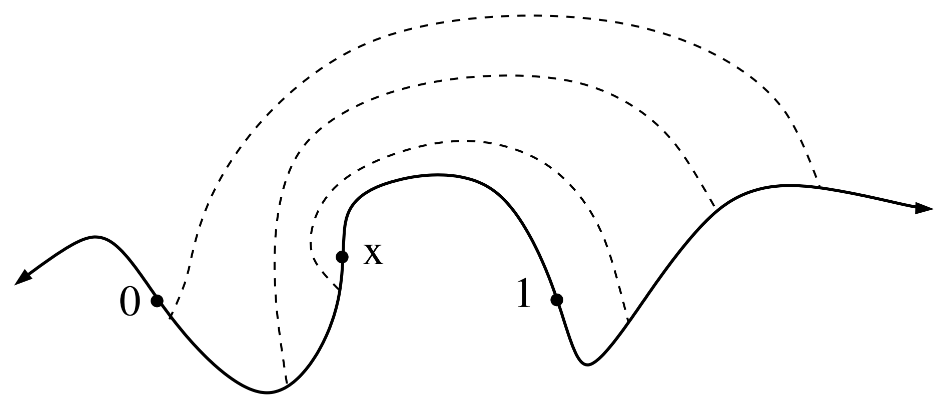

To prove a lower bound on $M(\Gamma)$ it suffices to prove an upper bound on the reciprocal modulus of the path family connecting γ_1 to γ_3.

By the 3-point condition, these arcs are at least distance $|x|/M$ apart, so the metric $\rho = M/|x|$ on the disk of radius M around the origin is admissible.

The reciprocal family has modulus at most $\pi M^4/|x|^2$, so $M(\gamma) \geq |x|^2/M^4\pi$.

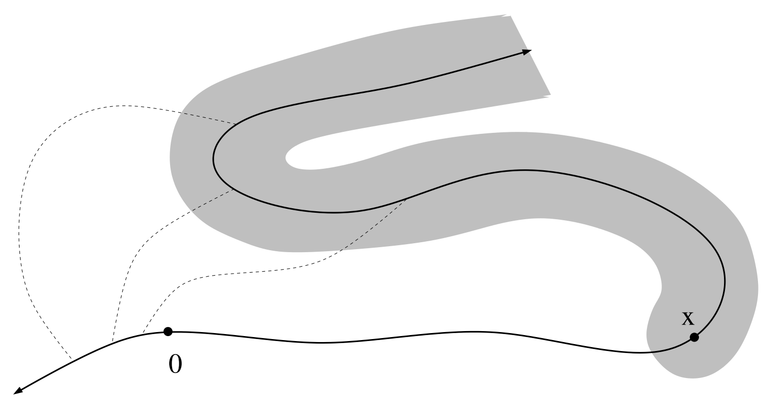

Since $\mathrm{dist}(\gamma_1, \gamma_3) \geq |x|/2M$, the metric $\rho = 1/2M$ is admissible. $\square$

Lemma 5.3. *If γ has bounded turning, and f, g are the conformal maps from the upper and lower half-planes to the two sides of γ (mapping ∞ to ∞ in both cases), then $h = g^{-1} \circ f$ is a quasisymmetric homeomorphism of the line.*

Proof. Consider two adjacent intervals of equal length on the real line.

After renormalizing by linear maps, we may assume these are $I = [0, 1/2]$ and $[[1/2, 1]$ and that h fixes both 0 and 1. By two applications of Lemma 5.2, $f(1/2)$ can't be too close to either 0 or 1, and hence $h(1/2) = g^{-1}(f(1/2))$ can't be too close to 0 or 1 either.

Thus h is quasisymmetric with a constant depending only on the 3-point constant.

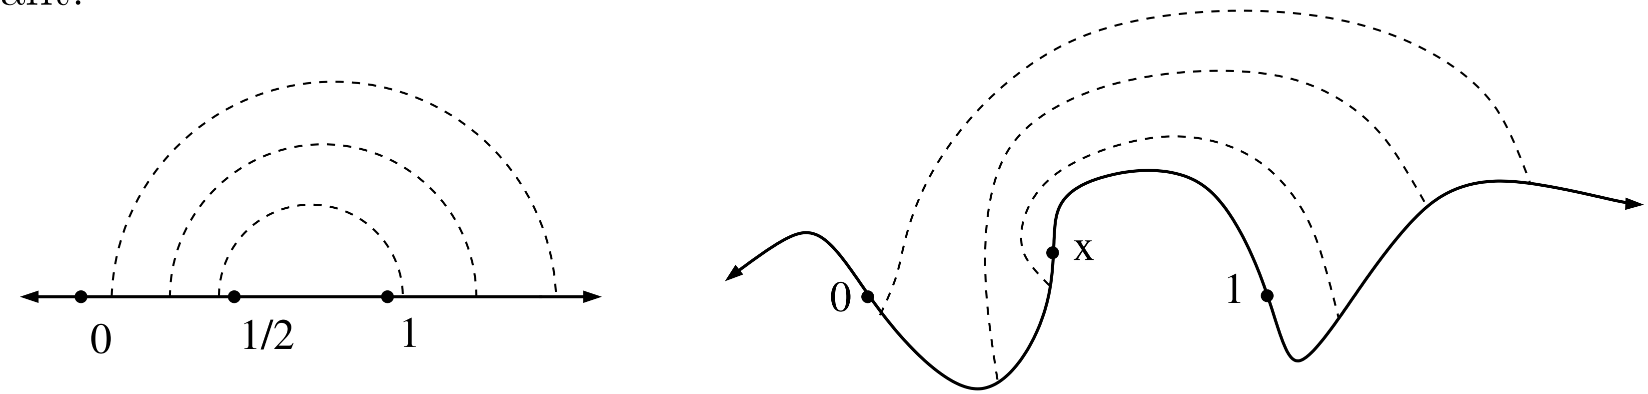

The path family in the upper half-plane connecting $[0, 1/2]$ to $[1, \infty)$ has modulus 1, so its conformal image also has modulus 1. Therefore $x = f(1/2)$ can't be too close to either 0 or 1. $\square$

Lemma 5.4. *A curve γ is a quasi-line if and only if it has bounded turning.*

Proof. If γ is the quasiconformal image of a line, then it satisfies the 3-point condition by Lemma 4.5, as mentioned earlier.

On the other hand, if γ satisfies the 3-point condition, then $h = g^{-1} \circ f$ is quasisymmetric, and hence extends to a quasiconformal map H of the whole plane.

Now set $F = f$ on the lower half-plane and $F = g \circ H$ on the upper half-plane. Clearly this is quasiconformal on each half-plane and on the real line $g \circ H = g \circ g^{-1} \circ f = f$ so the two definitions agree. Thus H is quasisymmetric on the whole plane and $F(\mathbb{R}) = f(\mathbb{R}) = \gamma$. $\square$

Actually, the previous proof has a small error.

We claimed that if a homeomorphism of the plane is quasiconformal in both the upper and lower half-planes, then it is quasiconformal in the whole plane.

This is true, but not yet proven.

It is trivial from analytic definition; a little harder from the geometric definition. We will prove a much stronger result.

For the proof of Lemma 5.4 we can assume the QC map H is piecewise smooth on a hyperbolic tesselation, since we proved the QS extension theorem using an explicit construction that did this (other extensions methods even give smooth maps on whole half-plane).

Lemma 5.5. *If F is a homeomorphism of the plane that is quasiconformal on the upper and lower half-planes, and is piecewise smooth on a countable decomposition of each of these half-planes (such as given by a hyperbolic tesselation), then F is quasiconformal on the whole plane.*

We leave the proof to reader.

It follows the proof we gave that the piecewise differentiable definition of QC implies the geometric definition, except now we use that each line hits the partition boundary countably often instead of finitely often.

However, the result is true, even without assuming any smoothness.

Removable Sets

When f is continuously differentiable, it is relatively easy to check whether it is quasiconformal; we just compute the complex dilatation $\mu = f_{\overline{z}}/f_z$ and check that $|\mu| < k < 1$ everywhere.

For some applications in dynamics, functions arise that that are homeomorphisms f on $\mathbb{C}$, but which are only C^1 on an open set $\Omega = \mathbb{C} \setminus K$. If we know the dilatation is bounded on just Ω, can we still deduce that f is quasiconformal? If we can, then we say K is removable for quasiconformal mappings.

Removability depends on the “size” and “shape” of K.

We have already (implicitly) seen that K is removable if it a countable union of analytic arcs.

If K has interior, then it is easy to construct counterexamples; choose a disk $D \subset K$ and any non-quasiconformal homeomorphism of the disk to itself that is the identity on the boundary and extend it to be the identity off D.

If K has positive area, there are also counterexamples corresponding to applications of the measurable Riemann mapping theorem to a dilatation that is a non-zero constant on K and zero off K.

Even if K is quite small, there can be counterexamples. For example, given any guage function h such that $h(t) = o(t)$ as $t \searrow 0$, there is a closed Jordan curve γ and a homeomorphism of the sphere that is conformal on both components of $\mathbb{C} \setminus \gamma$ but which is not Möbius.

On the other hand, if K has finite or sigma-finite 1-measure then it is removable. These examples show that it is the "shape" rather than the "size" of K that is crucial in most cases of interest.

Recall that we proved this earlier:

Lemma 5.6. *Suppose Q is a square, $\lambda > 1$ and f is K-quasiconformal on λQ. Then*

$$\text{area}(f(Q)) \geq \epsilon \text{diam}(f(Q))^2,$$

where $\epsilon > 0$ depends only on λ and K.

A **Whitney decomposition** of an open set Ω consists of a collection of dyadic squares $\{Q_j\}$ contained in Ω so that

(1) the interiors are disjoint,

(2) the union of the closures is all of Ω,

(3) for each Q_j, $\text{diam}(Q_j) \simeq \text{dist}(Q_j, \partial\Omega)$.

The existence of such a collection is easy to verify be taking the set of dyadic squares Q so that

$$\text{diam}(Q) \leq \frac{1}{4}\text{dist}(Q, \partial\Omega),$$

and that are maximal with respect to this property (i.e., the parent square fails this condition).

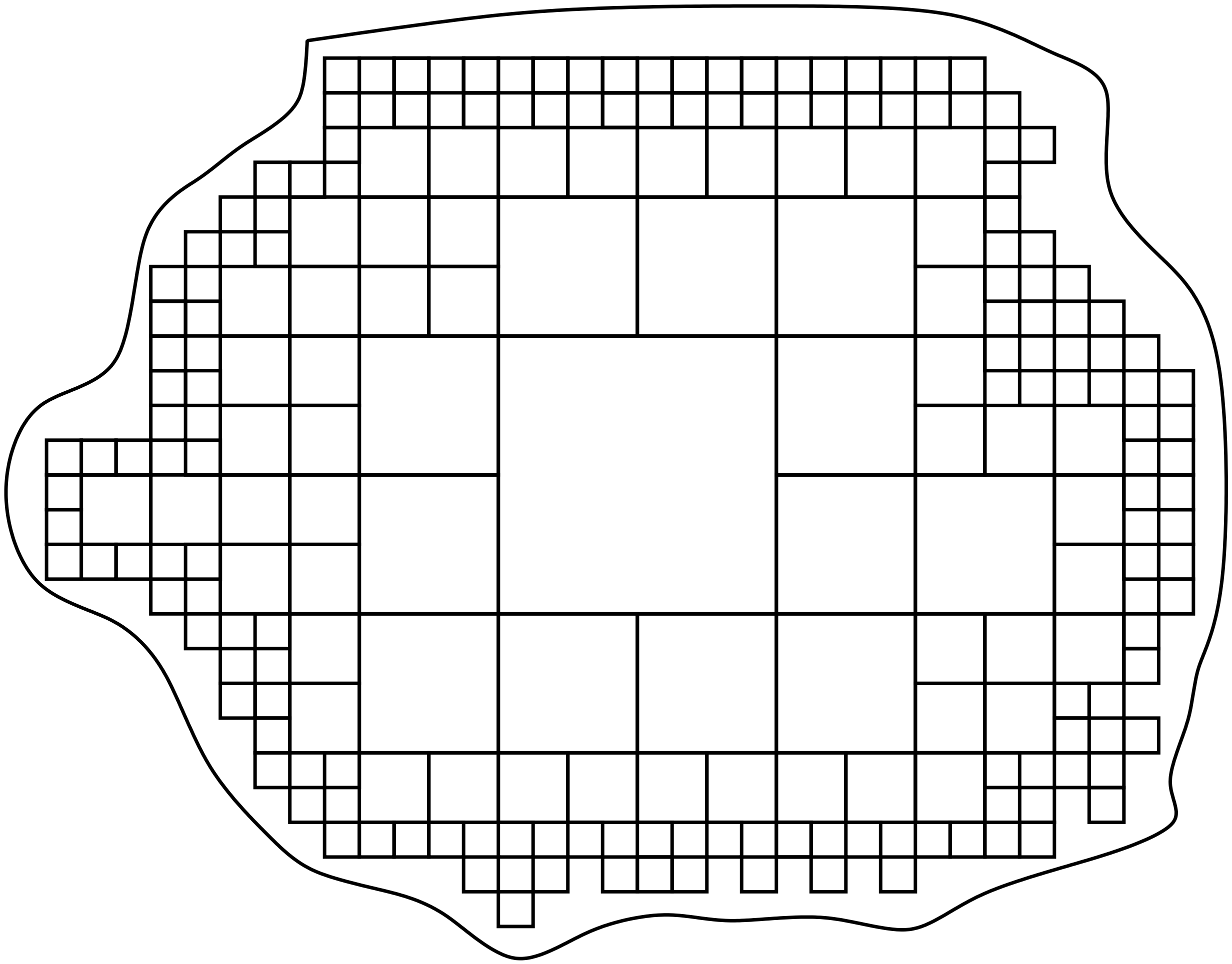

Suppose K is compact, $\delta > 0$ and for each $x \in K$ let γ_x be a Jordan arc in $\Omega = \mathbb{C} \setminus K$ that connects x to $\Omega_\delta = \{z \in \Omega : \text{dist}(z, K) \geq \delta\}$. For a single x, γ_x may consist of several arcs that connect x to Ω_δ.

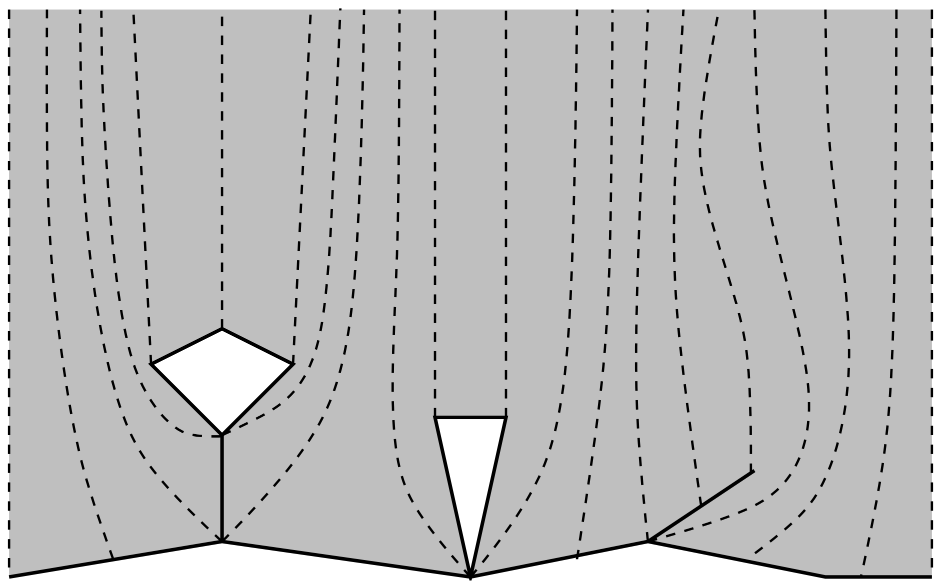

Each boundary point is connected to a point distance δ from $\partial\Omega$. Some points may be connected by more than one curve.

For each Whitney square $Q \subset \Omega$, let $S(Q) = \{x \in K : \gamma_x \cap Q \neq \emptyset\}$.

This is called the "shadow" of Q on K; the name comes from the special case when K is connected and does not separate the plane and γ_x is a hyperbolic geodesic connecting x to ∞.

If we think of ∞ as the "sun" and the geodesics as light rays, then $S(Q)$ is the part of K that blocked from ∞ by Q, i.e., it is Q's shadow.

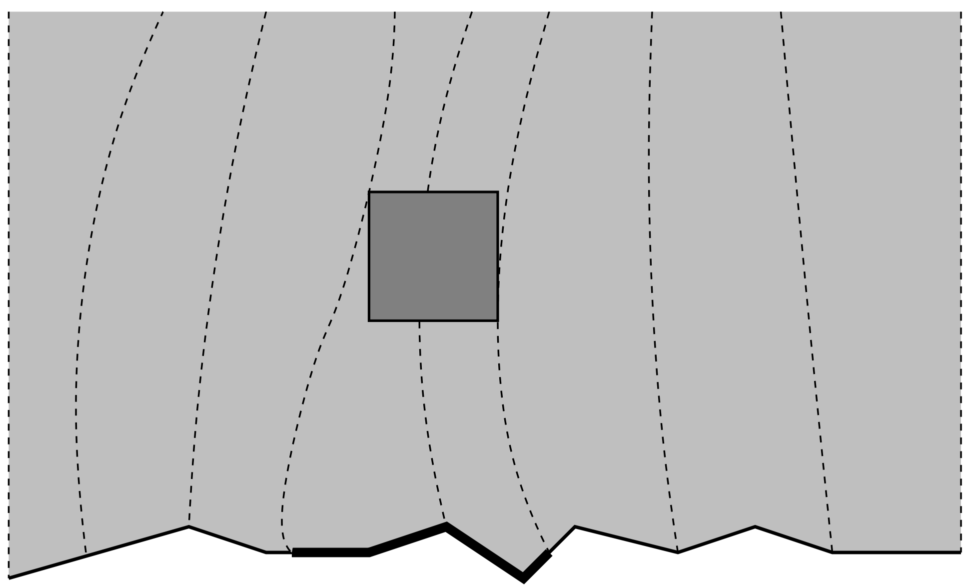

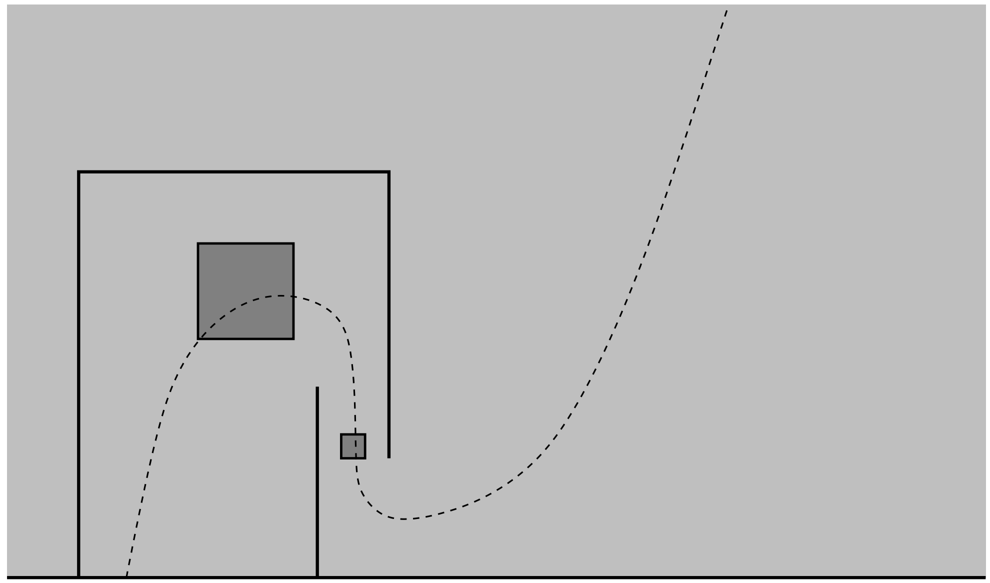

The paths connecting a Whitney square to its shadow can sometimes hit larger Whitney squares after hitting smaller ones.

However the size of the hit squares tends to zero as the path approaches the boundary. Hence there is a“largest” square hit.

Let $C(Q)$ be the union of all Whitney squares hit by the arc γ connecting Q to some point of its shadow; this is the "filled shadow" and corresponds to a Carleson square in the unit disk.

We will assume three things about the Whitney squares and their shadows:

(S1) $S(Q)$ is closed.

(S2) $\mathrm{diam}(S(Q)) \to 0$ as $\mathrm{diam}(Q) \to 0$,

(S3) $\lim_{n\to\infty} \sum_{Q:\ell(Q)\le 2^{-n}} \mathrm{diam}(S(Q))^2 = 0$, where the sum is over all Whitney squares for Ω of side length 2^{-n}.

Theorem 5.7 (Jones-Smirnov)**.** *Suppose Ω has a Whitney decomposition so that the corresponding shadow sets satisfy conditions (S1)-(S3) above. Then $K = \partial\Omega$ is removable for quasiconformal maps.*

The Jones-Smirnov paper

If the map f is conformal off $\partial\Omega$ (i.e., $K = 1$), then we will show that the extension is conformal everywhere.

If the map f is K-quasiconformal off $\partial\Omega$ then we only prove that it is C-quasiconformal for some $C < \infty$.

However, it then follows from the analytic definition that f is actually K-quasiconformal on the whole plane.

The weaker version is sufficient for many applications.

Proof of Theorem 5.7. Recall that Ω is the complement of K. Suppose F is a homeomorphism of the plane that is quasiconformal on Ω.

Suppose that W is any bounded quadrilateral in the plane, say of modulus m and that $W' = F(W)$ has modulus m'. We want to show that $m' \leq Cm$ where $C < \infty$ depends only on K as in the statement of the theorem.

We will do this by mimicking the proof of Theorem 4.1, that showed that any piecewise differentiable map with bounded dilatation was quasiconformal (in the geometric sense).

Let $\varphi : W \to R = [0, m] \times [0, 1]$ and $\psi : W' \to [0, m'] \times [0, 1]$ be conformal maps of these quadrilaterals to rectangles R, R' (vertices mapping to vertices).

Define $X = \varphi(\partial\Omega \cap W) \subset R$.

The proof is somewhat involved because we are going to to consider three different Whitney decompositions. Let

- $\{W_j\}$ denote a Whitney decomposition for W,
- $\{Q_j\}$ a Whitney decomposition for Ω, and
- $\{U_j\}$ a Whitney decomposition for $U = R \setminus X$.

Fix some $\epsilon > 0$.

Fix a Whitney cube W_j for W.

We assume the decomposition is chosen so that $2W_j \subset W$.

Suppose $\delta > 0$ is so small (depending on our choice of W_j) that the following two conditions all hold.

(1) If Q_k is a Whitney square for Ω with diameter less than δ and the shadow $S(Q_k)$ hits W_j, then $S(Q_k) \subset 2W_j$ and the entire Whitney chain connecting any point $x \in S(Q_k)$ to Q_k is contained in $2W_j$.

This is possible by condition (S2) on shadow sets (small squares have small shadows).

Note that two points $x, y \in S(Q_k)$ can be connected by a chain of adjacent Whitney squares for Ω, all in the "shadow" of Q_k.

(2) Let $\mathcal{S}(W_j)$ denote the collections of all Whitney squares Q_k for Ω so that $\mathrm{diam}(Q_k) \le \delta$ and $S(Q_k)) \cap W_j \neq \emptyset$. Then

$$\sum_{Q_k \in \mathcal{S}(W_j)} \mathrm{diam}(S(Q_k))^2 \le \epsilon \cdot \mathrm{area}(W_j).$$

This holds for small enough δ, because by condition (S3) on shadows, this sum over all Whitney squares for Ω is finite, so removing all the squares bigger than δ gives a sum that tends to 0 as δ tends to zero.

The sum is less than $\epsilon \cdot \mathrm{area}(W_j)$ if δ is small enough (depending on W_j).

Let $\mathcal{S} = \cup_{W_j} \mathcal{S}(W_j)$.

This is the collection of all shadow sets of all Whitney squares Q_k for Ω so that (1) $\mathrm{diam}(Q_k) < \delta$ and (2) $S(Q_k)$ is contained in $\mathcal{S}(W_j)$ for some Whitney square W_j of W.

Note that each point $x \in \partial\Omega \cap W_j$ is associated to a Whitney chain that contains a square with diameter comparable to δ. There are only finitely many such squares, so their shadows form a finite collection that covers $\partial\Omega \cap W_j$.

Suppose $L = [a + iy, b + iy]$ is a horizontal segment, compactly contained in the interior of R at height y.

Let $g : R \to R'$ be the composition $\psi \circ f \circ \varphi^{-1}$. We wish to show that

$$\int_0^1 |g(b + iy) - g(a + iy)|dy \leq C\sqrt{mm'}, \tag{5.11}$$

where C depends only on K.

If we can do this, then by letting $a \to 0$ and $b \to m$ we get

$$m' \le \lim_{a\to 0, b\to m} |g(b+iy) - g(a+iy)|,$$

and hence

$$m' \le \lim_{a\to 0, b\to m} \int_0^1 |g(b+iy) - g(a+iy)| dy \le C\sqrt{mm'},$$

which gives the desired inequality $m' = O(m)$.

The reversed inequality, $m = O(m')$, can be deduced from the same argument applied to the other pair of opposite sides of Q, since the corresponding path families have the reciprocal moduli. Thus it suffices to prove (5.11).

Since L is compactly contained in the interior of R and X is relatively closed in the interior of R, $L \cap X$ is compact. Thus $\varphi^{-1}(L \cap X)$ is a compact set of W, hence covered by finitely many Whitney squares for W and hence is covered by finitely many shadows sets in $\mathcal{S}$.

Let $\mathcal{X}$ be the image of the elements of $\mathcal{S}$ under φ. Then $L \cap X$ is covered by finitely many elements of $\mathcal{X}$, say $X_1, \ldots X_n$.

For $k = 1, \ldots, n$, let $Y_k = [a_k, b_k]$ be the smallest closed interval in L that contains $X_k \cap L$ (this is the convex hull of $X_k \cap L$, i.e., the interval with the same leftmost and rightmost point as $X_k \cap L$).

Then $Y_1, \ldots, Y_n$ also cover $L \cap X$ and we can extract a subcover with the property that $Y_j \cap Y_k \neq \emptyset$ implies $|j - k| \leq 1$.

Since the points a_k, b_k are both in the same set X_k, we can deduce that the preimage points $\varphi^{-1}(a_k), \varphi^{-1}(b_k)$ are both in the same element of $\mathcal{S}$.

Thus they are both in the shadow set of some Whitney square for Ω and are associated to a two sided chain of distinct Whitney squares $\{Q_m\}_{-\infty}^{\infty}$ of Whitney squares for Ω.

If two chains arising in this way, say from Y_k and Y_m with $m > k$, have a Whitney square in common, then we can combine the chains to form a chain connecting a_k to b_m consisting of distinct squares. We can replace $Y_k, Y_{k+1}, \ldots Y_m$ by the single interval $Z = [a_k, b_m]$ which covers the same part of $L \cap X$..

Doing this for all intersections, we obtain a finite collection of closed intervals Z_k in L which covers the same set as the union of the Y_k's.

Furthermore, the two endpoints of each Z_k correspond to a two-sided Whitney chain in Ω and that different intervals use different Whitney squares (no overlapping chains).

Moreover, if Z_k has endpoints c_k, d_k and the corresponding chain is $\{Q_n\}$, then

$$|g(c_k) - g(d_k)| \leq \sum_n \operatorname{diam}(\psi(f(Q_n))).$$

The set $V = L \setminus \cup_k Z_k$ consists of finitely many open intervals in $U = R \setminus X$ with their endpoints in X.

We break V into countable many sub-intervals by intersecting it with the Whitney squares for $U = R \setminus X$ (without loss of generality, we can assume the endpoints of L occur on the boundary of a Whitney square for U).

On each Whitney square U_k for U we define the constant function

$$Dg = \frac{\mathrm{diam}(g(U_k))}{\mathrm{diam}(U_k)}.$$

Then if $L_j = L \cap U_j$,

$$\int_{L_j} Dg dx = \mathrm{diam}(g(U_j))/\sqrt{2}.$$

Thus if Z_L is the union of all the $Z_k \cap L$, we get

$$\int_{L \setminus Z_L} Dg dx \simeq \sum_j \operatorname{diam}(g(U_j)),$$

where the sum is over Whitney squares for U that hit L.

Thus

$$|g(b+iy) - g(a+iy)| \lesssim \int_{L \cap U} Dg dx + \sum_n \operatorname{diam}(\psi(f(Q_n))).$$

The sum is over Whitney squares Q_j for Ω that have diameter $\leq \delta$.

Now integrate in y to get

$$\int_0^1 |g(b+iy) - g(a+iy)|dy \lesssim \iint_U Dg dx + \sum_n \operatorname{diam}(\psi(f(Q_n)))\mu_n,$$

where μ_n is the Lebesgue measure in $[0,1]$ of the set of lines L_y that use the Whitney square Q_n in at least one of the two-sided chains associated to a interval $Z \subset L_y$.

The Lebesgue measure of this set is no more than its diameter, which is no more than the diameter of $X_n = \varphi(S(Q_n))$. Thus

$$\int_0^1 |g(b+iy) - g(a+iy)|dy \lesssim \iint_U Dg dx dy + \sum_n \mathrm{diam}(\psi(f(Q_n)))\mathrm{diam}(X_n),$$

We want this to be $= O(\sqrt{m \cdot m'})$.

We now estimate the second term using the Cauchy-Schwarz inequality.

$$
\begin{aligned}
&\sum_n \operatorname{diam}(\psi(f(Q_n)))\operatorname{diam}(X_n) \\
&\quad \leq \left(\sum_n \operatorname{diam}(\psi(f(Q_n)))^2\right)^{1/2} \left(\sum_n \operatorname{diam}(X_n)^2\right)^{1/2} \\
&\quad \leq A \left(\sum_n \operatorname{area}(\psi(f(Q_n)))\right)^{1/2} \times \\
&\qquad\qquad \left(\sum_{W_k} \left[\frac{\operatorname{diam}(\varphi(W_k))}{\operatorname{diam}(W_k)}\right]^2 \sum_{Q_n \in \mathcal{S}(W_k)} \operatorname{diam}(S(Q_n))^2\right)^{1/2}.
\end{aligned}
$$

We have used Koebe's theorem to estimate the size of the images.

Now use Lemma 5.6,

$$
\begin{aligned}
&\leq A\left(\sum_n \operatorname{area}(\psi(f(Q_n)))\right)^{1/2}\left(\sum_{W_k}\left[\frac{\operatorname{diam}(\varphi(W_k))}{\operatorname{diam}(W_k)}\right]^2\cdot\epsilon\cdot\operatorname{area}(W_k))\right)^{1/2}\\
&\leq A\left[\operatorname{area}(R')^{1/2}\cdot\epsilon\cdot\operatorname{area}(R)\right]^{1/2}\\
&\leq A\sqrt{\epsilon\cdot m\cdot m'}.
\end{aligned}
$$

where A just depends on the distortion estimate for conformal maps and ϵ is as small as we wish (this was Condition 2 in our choice of δ).

The other term is also bounded by Cauchy-Schwarz

$$\begin{aligned}
\iint_U Dg dx &= \sum_k \iint_{U_k} Dg dx dy \\
&\leq \left(\sum_k \iint_{U_k} (Dg)^2 dx dy \right)^{1/2} \left(\sum_k \iint_{U_k} dx dy \right)^{1/2} \\
&\leq \left(\sum_k (\mathrm{diam}(g(U_k))^2 \right)^{1/2} \left(\sum_k \mathrm{area}(U_k) \right)^{1/2} \\
&\leq C \left(\sum_k (\mathrm{area}(g(U_k)) \right)^{1/2} \mathrm{area}(R)^{1/2} \\
&\leq C \mathrm{area}(R')^{1/2} \cdot \mathrm{area}(R)^{1/2} \leq C\sqrt{m'm}.
\end{aligned}$$

Thus

$$\int_0^1 |g(b+iy) - g(a+iy)| dy \lesssim \sqrt{m'm} + O(\epsilon),$$

Taking $\epsilon \to$ gives the desired inequality. $\square$

Corollary 5.8. *If K satisfies the conditions of Theorem 5.7, then K is removable for conformal homeomorphisms, i.e., any homeomorphism of the plane that is conformal off K is conformal everywhere.*

Proof. Theorem 5.7 implies that f is quasiconformal on the plane, so the point is to show that we can take the quasiconformal constant to be 1.

If we redo the proof assuming f is conformal off $\partial\Omega$, then the piecewise constant function Dg can be replaced by the usual derivative $|g'|$.

This leads to the inequality $m' \leq \sqrt{m'm}$, or $m' \leq m$.

The reverse inequality follows by considering the reciprocal path family in each quadrilateral. Together, these imply f is 1-quasiconformal, and hence conformal.

$\square$

Corollary 5.9. *If f, g are quasiconformal maps of the upper and lower half-planes that agree on the real line, then they define a quasiconformal map on the whole plane.*

Proof. This is immediate since a line clearly satisfies the Jones-Smirnov criteria: just consider $\mathbb{R}$ as the boundary of the upper half-plane and for $x \in \mathbb{R}$, let γ_x be a vertical line ray.

Then the shadow of any square is its vertical projection, and the square of the shadows length is comparable to the area of the square.

Thus any compact segment of $\mathbb{R}$ is removable, and since quasiconformality is a local property (Theorem 4.15), the whole line is removable. □

Corollary 5.10. *If f is a quasiconformal map of the upper half-plane to itself, mapping the real line to itself, then the extension of f to the whole plane by $f(\bar{z}) = \overline{f(z)}$ is quasiconformal in the whole plane.*

Proof. Immediate from the previous result since composing a quasiconformal map with reflections gives another quasiconformal map. $\square$

Corollary 5.11. *Quasicircles are removable.*

Proof. If $\Gamma = g(\mathbb{R})$ is a quasiconformal image of the reals and f is a homeomorphism that is quasiconformal on each side of Γ, then $h = f \circ g$ is a homeomorphism that is quasiconformal on each side of $\mathbb{R}$, then quasiconformal on the whole plane.

Thus $f = h \circ g^{-1}$ is a composition of quasiconformal maps and hence is QC. $\square$

An open, connected set Ω in $\mathbb{R}^2$ is called a **John domain** if any two points $a, b \in \Omega$ can be connected by a path γ in Ω with the property that $\mathrm{dist}(z, \partial\Omega) \gtrsim \min(|z-a|, |z-b|)$.

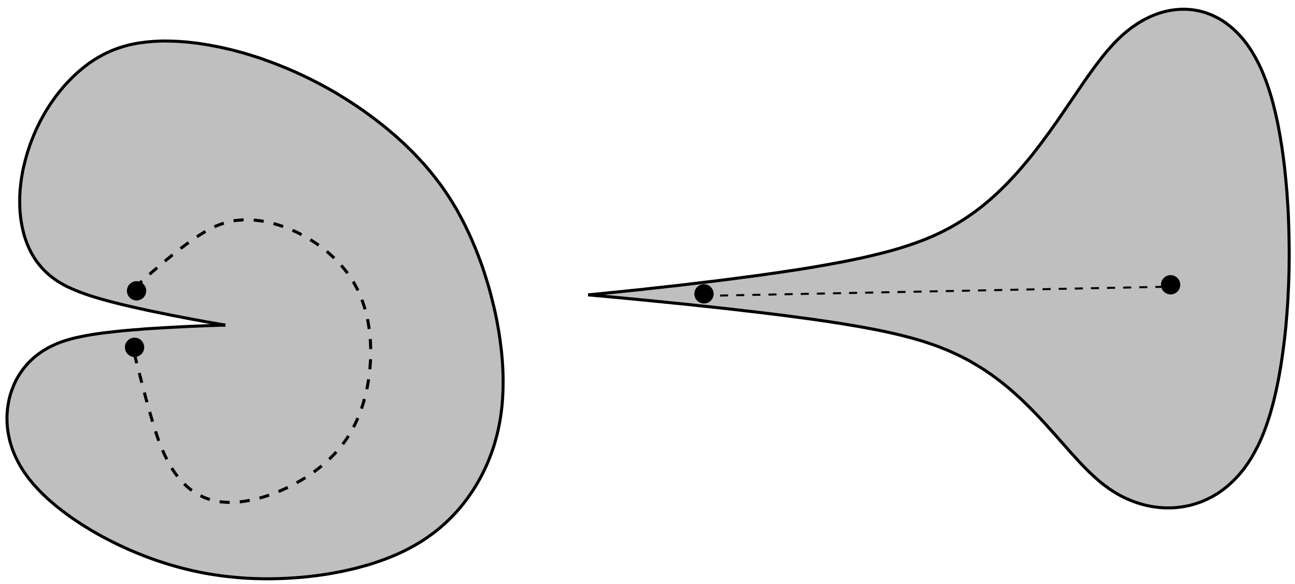

The domain on the left is a John domain, but the one on the left is not; inward pointing cusps are OK, but outward pointing cusps are not.

Lemma 5.12. *The Riemann map φ from the unit disk to a bounded John domain satisfies*

$$\operatorname{diam}(\varphi(I(Q))) \leq C \operatorname{diam}(\varphi(Q)),$$
$$\operatorname{dist}(\varphi(Q), \varphi(I(Q))) \leq C \operatorname{diam}(\varphi(Q)),$$

for some constant $C < \infty$ and any Whitney square Q and is shadow $I(Q)$.

Proof. The second inequality follows directly from Lemma 2.23 by considering the path family of radial lines connecting Q to I.

To prove the first inequality, consider the Whitney-Carleson boxes Q_1 and Q_2 that are adjacent to Q and of the same size. By Lemma 2.23 each is connected to its shadow by a radial segment whose image under f has length comparable to $\mathrm{diam}(f(Q))$.

Thus there is a geodesic crosscut γ of the disk that passes through Q and whose image has length comparable to $\mathrm{diam}(f(Q))$. Now suppose x is in the shadow of Q.

Any curve connecting 0 to x crosses γ, so any curve Γ connecting $f(0)$ and $f(x)$ crosses $f(\gamma)$ and hence contains a point $z \in f(\gamma) \cap \Gamma$ that is at most distance $O(\mathrm{diam}(f(Q)))$ from $\partial\Omega$. Thus by the definition of John domain, either

$$\mathrm{dist}(f(0), z) = O(\mathrm{diam}(f(Q))) \quad \text{or} \quad \mathrm{dist}(f(x), z) = O(\mathrm{diam}(f(Q))).$$

In a bounded domain, the first can only happen for finitely many Q's; for the remainder, the second must hold and hence $f(I(Q))$ is contained in a $O(\mathrm{diam}(f(Q)))$ neighborhood of $f(Q)$. $\square$

Corollary 5.13. *Boundaries of simply connected John domains are removable.*

Proof. Let Ω be a simply connected John domain and suppose $f : \mathbb{D} \to \Omega$ is conformal.

Each Whitney square Q' for Ω is covered by a uniformly bounded number images $f(Q)$ of Whitney squares for $\mathbb{D}$ and its shadows is contained in the union of corresponding shadows.

This and Lemma 5.12 imply $\mathrm{diam}(S(Q')) = O(\mathrm{diam}(Q')$.

The three conditions (1)-(3) in Theorem 5.7 follow easily.. $\square$

A simply connected plane domain Ω is called a *Hölder domain* if the Riemann map $\mathbb{D} \to \Omega$ is Hölder.

Lemma 5.14. *Boundaries of Hölder domains are removable.*

Sketch of proof. Fix a base point in Ω. The Hölder condition implies that $\{Q_j^k\}$ lists the Whitney squares of Ω approximately hyperbolic distance k from the base point then $\text{diam}(S(Q_j^k)) \leq Ce^{-ak}$.

We also need an estimate of Jones and Makarov that for Hölder domains,

$$\sum_k \text{diam}(S(Q_j^k))^{2-\epsilon} < M < \infty$$

for some $\epsilon > 0$ and $M < \infty$ independent of k. depending on the Hölder constant. Then

$$\sum_k \sum_j \text{diam}(S(Q_j^k))^2 \leq \sum_k \sum_j \text{diam}(S(Q_k))^{2-\epsilon} \text{diam}(S(Q_j^k))^{\epsilon} \leq \sum_k C^{\epsilon} e^{-\epsilon a k} M < \infty$$

Corollary 5.15. *Julia sets of Collet-Eckmann polynomials are removable.*

The Jones-Smirnov result (Theorem 5.7) places restrictions on the set E, but none on the mapping (besides being a homeomorphism). An earlier result of Rickman makes an assumption on the mapping, but none on the set K:

Lemma 5.16 (Rickman's lemma)**.** *Suppose Ω is a planar domain and $K \subset \Omega$ is compact. Suppose f is homeomorphism of Ω that is quasiconformal on $\Omega \setminus K$ and F is quasiconformal on all of Ω. If $f = F$ on K, then f is quasiconformal on all of Ω.*

Proof. Isolated points of K are clearly removable and there are only countable many such points, so we may assume that K has only accumulation points.

The idea proof is the same as the proof of Theorem 5.7: we consider a quadrilateral W and its image $W' = f(W)$ and conformally map each to rectangles of modulus m and m' respectively.

Let $G = \psi \circ F \circ \varphi^{-1}$ and $g = \psi \circ f \circ \varphi^{-1}$.

Our assumption implies $g = G$ on X.

As before, we want to prove the estimate (5.11):

$$\int_0^1 |g(b+iy) - g(a+iy)| dy \leq C\sqrt{mm'},$$

However, this time we cover X by dyadic squares that are so small that G is quasiconformal on $6Q \subset R$ for each square Q used, and the image $G(Q)$ lies in R'.

The union of these squares plays the role of the set Z in the earlier proof.

Given a compact horizontal line segment L in R, we let $\{Y_k\}\{[c_k, d_k]\}$ enumerate the convex hulls of sets of the form $L \cap Q$ for Q in our cover of X.

Then defining Dg exactly as before on $R \setminus X$, and using $g = G$ on X, we get

$$\begin{aligned} |g(b+iy) - g(a+iy)| &\leq \int_{L\cap U} Dg dx + \sum_k |g(c_k) - g(d_k)| \\ &\leq \int_{L\cap U} Dg dx + \sum_k |G(c_k) - G(d_k)| \\ &\leq \int_{L\cap U} Dg dx + \sum_{Q:Q\cap L\neq\emptyset} \mathrm{diam}(G(Q)). \end{aligned}$$

Integrating over y then gives

$$\int_0^1 |g(b+iy) - g(a+iy)|dy \ \leq \ \int_U Dg dx + \sum_Q \operatorname{diam}(G(Q))\ell(Q).$$

The first term is bounded exactly as before and the second is bounded by

$$\begin{aligned}
\sum_Q \operatorname{diam}(G(Q))\ell(Q) \ &\leq \ [\sum_Q \operatorname{diam}(G(Q))^2]^{1/2} \cdot [\sum_Q \ell(Q)^2]^{1/2} \\
&\leq \ C[\sum_Q \operatorname{area}(G(Q))]^{1/2} \cdot [\sum_Q \operatorname{area}(Q)]^{1/2} \\
&\leq \ C[\operatorname{area}(R')]^{1/2} \cdot [\operatorname{area}(R)]^{1/2} \\
&\leq \ C\sqrt{m'm}.
\end{aligned}$$

The rest of the proof is them completed just as before. $\square$

BiLipschitz Reflections

Lemma 5.17. *A quasisymmetric map $f : \mathbb{R} \to \mathbb{R}$ can be extended to a quasiconformal map of the upper half-plane that is also biLipschitz for the hyperbolic metric.*

Proof. Go back and check the proof of the extension theorem. $\square$

Lemma 5.18. *If f is a hyperbolic biLipschitz map of the upper half-plane to itself, then f is quasiconformal.*

Proof. Easy to check that length and area change by at most a bounded factor, so modulus of any quadrilateral changes by a bounded factor (just transfer ρ without change). $\square$

Theorem 5.19. *An unbounded Jordan curve* Γ *is a quasiline iff it has a biLipschitz reflection, i.e., there is a bi-Lipschitz map of the plane that fixes* Γ *pointwise and swaps the two complements.*

Quasiline implies biLipschitz reflection. Let f and g be the conformal maps from the upper and lower half-planes to the two sides of Γ, each fixing ∞.

Since Γ is a quasiline, $h = g^{-1} \circ f$ is quasisymmetric and has a quasiconformal extension H to the lower half-plane that is biLipschitz for the hyperbolic metric.

Let $r(z) = \overline{z}$ be reflection across the real line and define $R(z) = g \circ H \circ r \circ f^{-1}$. this is a quasiconformal map from one side of Γ to the other and it fixes Γ pointwise.

If H is defined by our hyperbolic pentagon map, then each pentagon is associated to several subintervals of $\mathbb{R}$ that all have comparable harmonic measures for any point in the pentagon.

Thus the R maps the region $f^{-1}(P)$ to $g(r(P))$ and these regions have comparable diameters since the associated subintervals on Γ are the same.

Since R is a hyperbolic biLipschitz map between two domains of bounded hyperbolic diameter and comparable Euclidean size, it is a Euclidean biLipschitz map on these regions.

From this it is easy to check R is Euclidean Lipschitz everywhere. Since $R = R^{-1}$, it is automatically biLipschitz. $\square$

biLipschitz reflection implies quasiline. As above, let f and g denote conformal maps of the upper and lower half-plane to the two sides of Γ that fix ∞.

Suppose R is a biLipschitz reflection across Γ. Then $r \circ g^{-1} \circ R \circ f$ is a hyperbolic biLipschitz map of the upper half-plane to itself that extends the welding map $h = g^{-1} \circ f$.

Hyperbolic Lipschitz implies quasiconformal, so h must be quasisymmetric, which in turn implies Γ is a quasiline. $\square$

Remark: A set E is K-biLipschitz homogeneous if for any $x, y \in E$ there is a K-biLipschitz map $f : E \to E$ so that $f(x) = y$.

It is known that a biLipschitz homogeneous closed curve must be a quasicircle.

Question: is a biLipschitz homogeneous continuum a closed curve?

There are homogeneous continua for (non-biLipschitz) homeomorphisms that are not curves (e.g., the pseudo-arc and the circle of pseudo-arcs). These examples are not locally connected. Does requiring biLipschitz maps eliminate these?

A complete classification of homogeneous plane continua by L.C. Hoehn and L.G. Oversteegen, 2016.

Remark: A hyperbolic quasi-isometry $f : \mathbb{D} \to \mathbb{D}$ is a map so that

$$\frac{1}{A}\rho(z,w) - B \leq \rho(f(z), f(w)) \leq A\rho(z,w) + B.$$

Informally, these are biLipschitz at large scales.

Every quasiconformal map $f : \mathbb{D} \to \mathbb{D}$ is a hyperbolic isometry. See 2004 Annals paper by Epstein, Marden and Markovic.

Conversely, every hyperbolic quasi-isometry has boundary values that are quasisymmetric. Thus there is a quasiconformal map with the same boundary values.

Conformal Welding

Suppose Γ is a closed Jordan curve and f, g are conformal maps from $\mathbb{D}$ and $\mathbb{D}^* = \{|z| > 1\}$ to the inside and outside complementary domains of Γ.

By Carathéodorty's theorem, both these maps extend to be homeomorphisms of $\mathbb{T} \to \Gamma$, so $h = g^{-1} \circ f$ is a homeomorphism of the unit circle to itself (for brevity, we call this a circle homeomorphism).

Such a circle homeomorphism is called a **conformal welding** or **welding**. Sometimes called a conformal sewing or gluing.

There is an analogous definition for unbounded Γ and homeomorphisms of $\mathbb{R}$.

Not every homeomorphism is a welding.

Oikawa showed that if $h(x) = -|x|^\alpha$ for $x \leq 0$ and $h(x) = x^\beta$ for $x > 0$, and $\alpha \neq \beta$, then h is not a conformal welding.

Let Γ be the union of the graph of $\sin(1/x)$ and the segment $[i, -i]$. This set divides the plane into two simply connected regions, so there are an associated conformal maps f, g that define a circle homeomorphism h. One can prove that h is not a conformal welding.

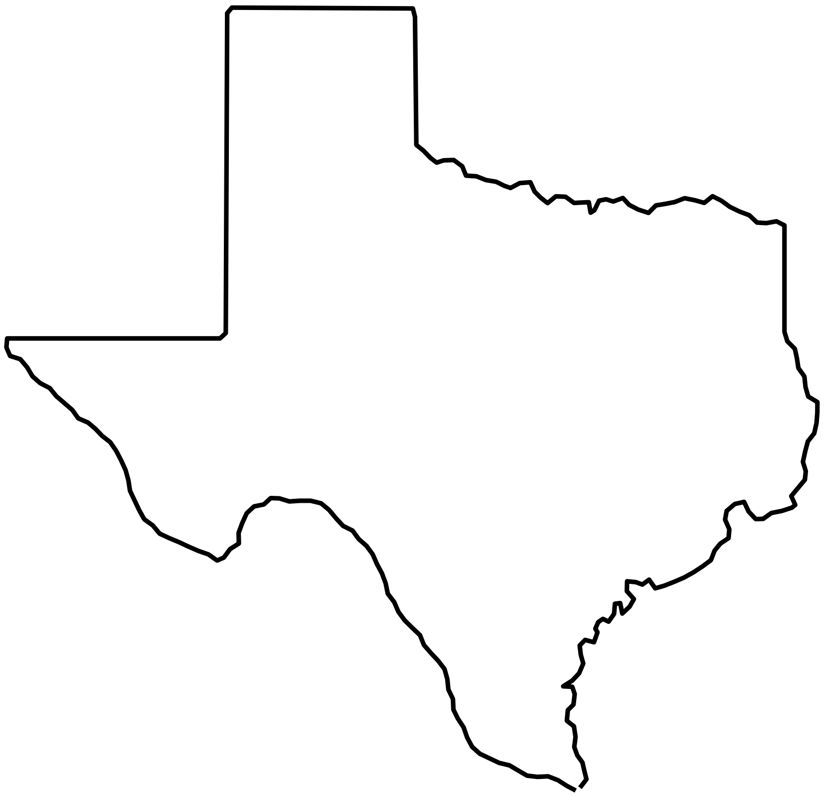

A polygon (looks like Texas)

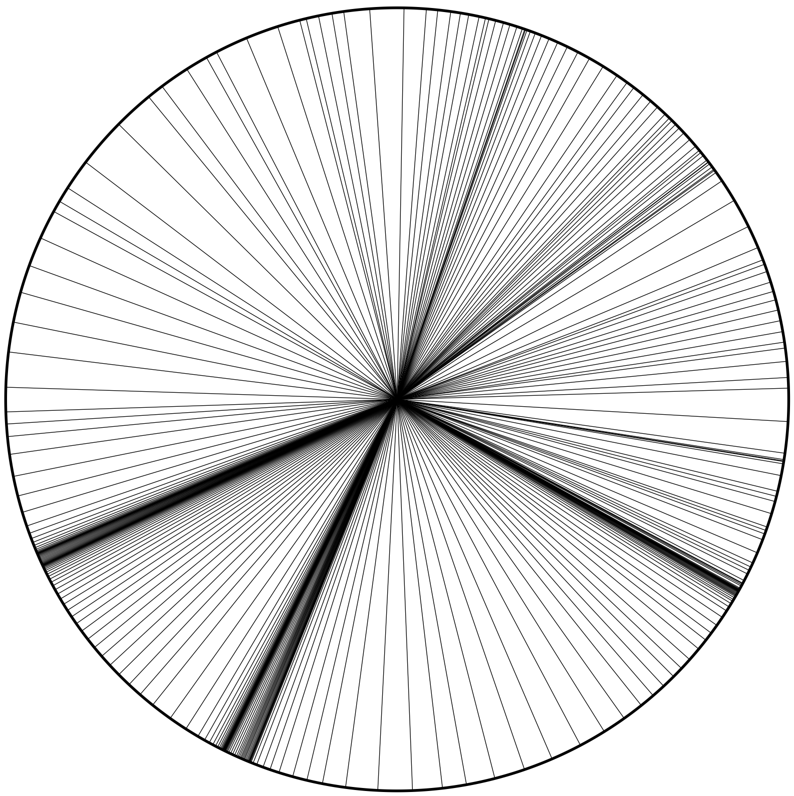

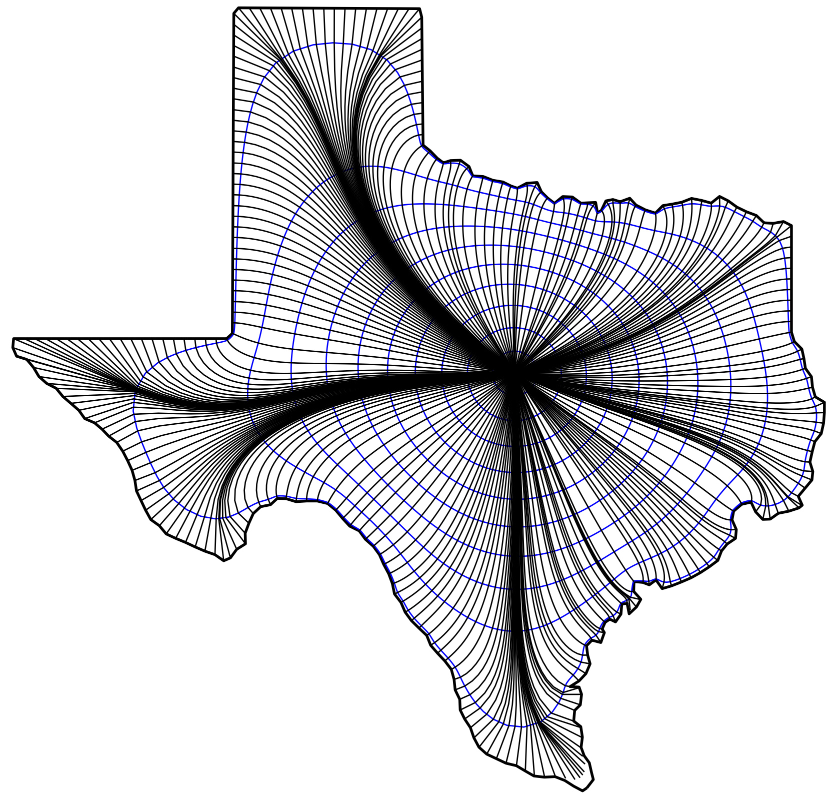

Texas reflected through a circle.

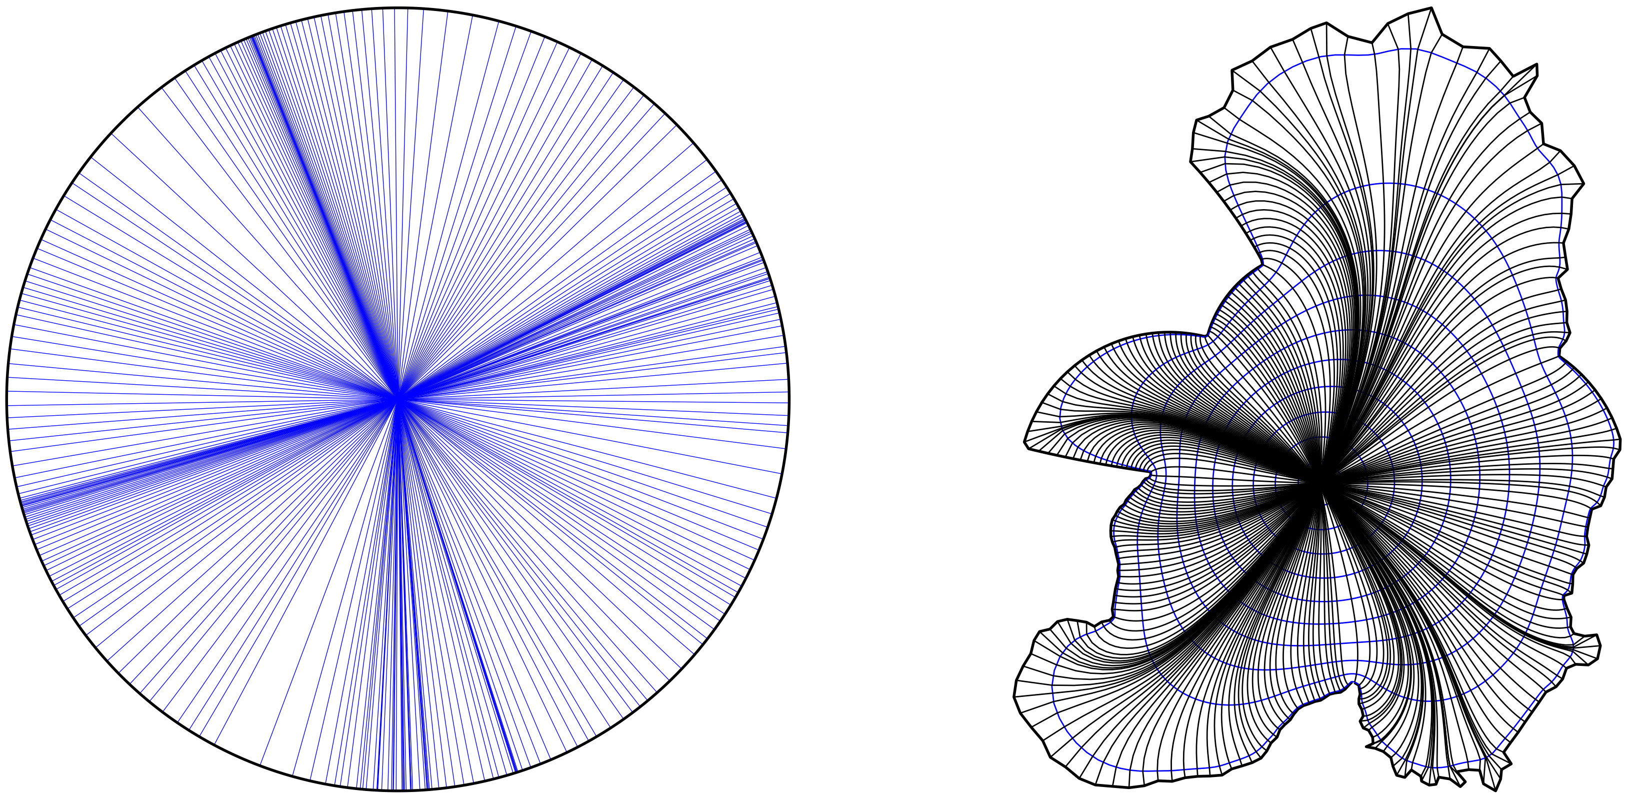

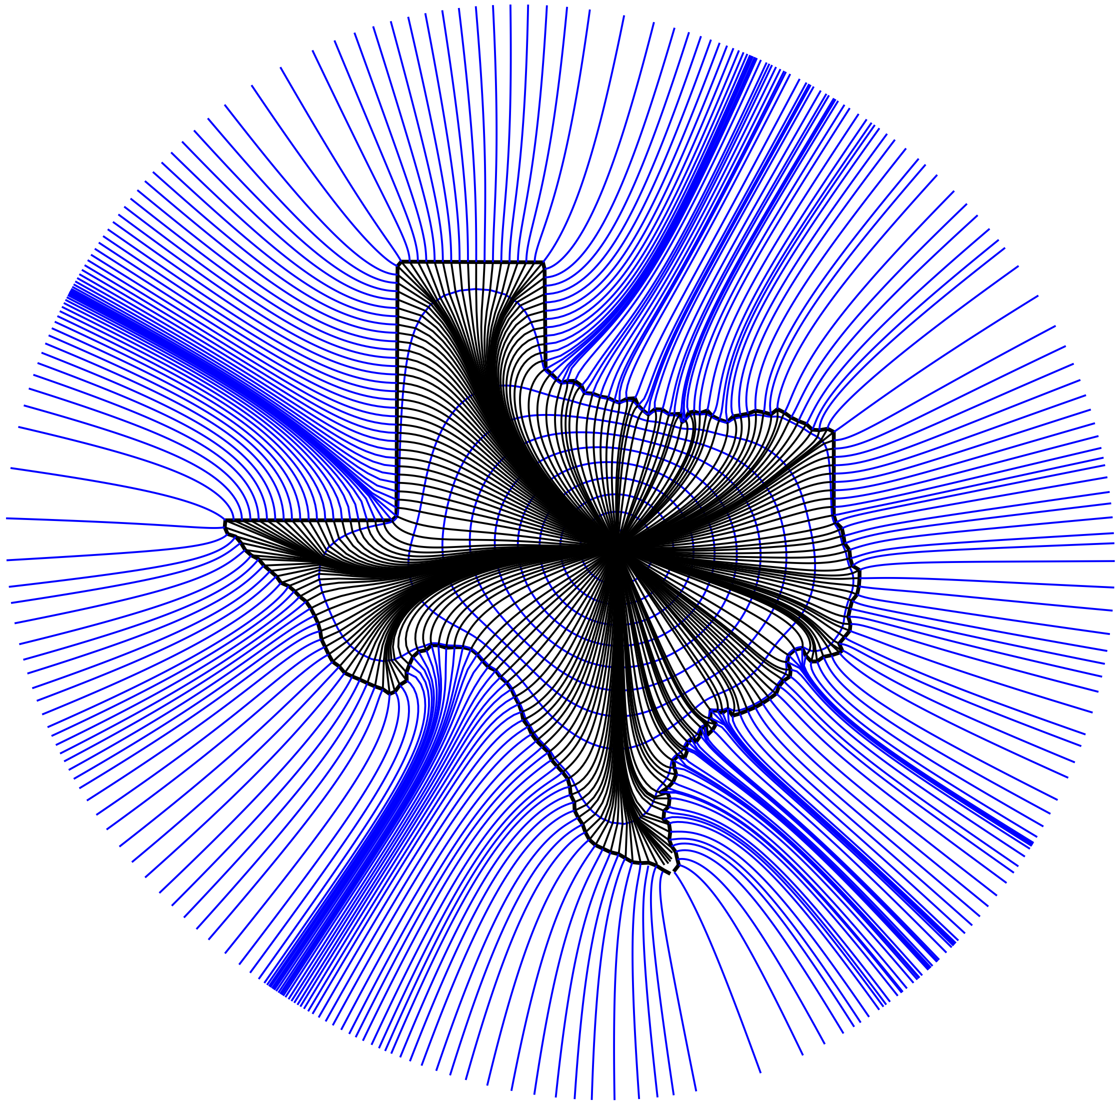

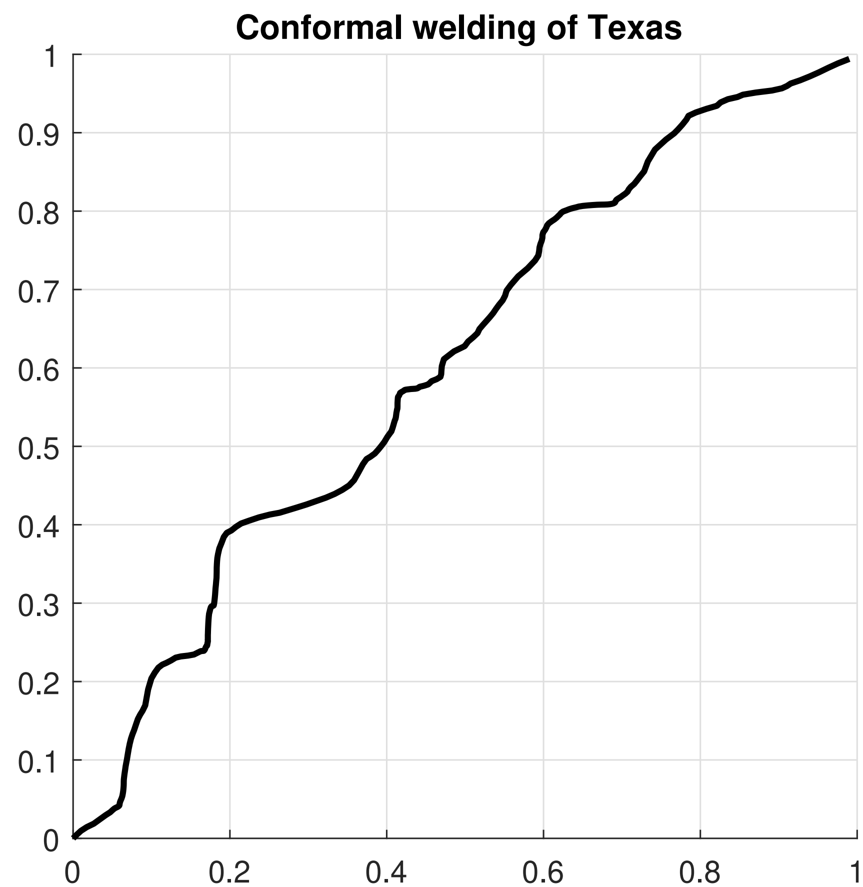

Conformal welding of Texas
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.2
0.4
0.6
0.8
1

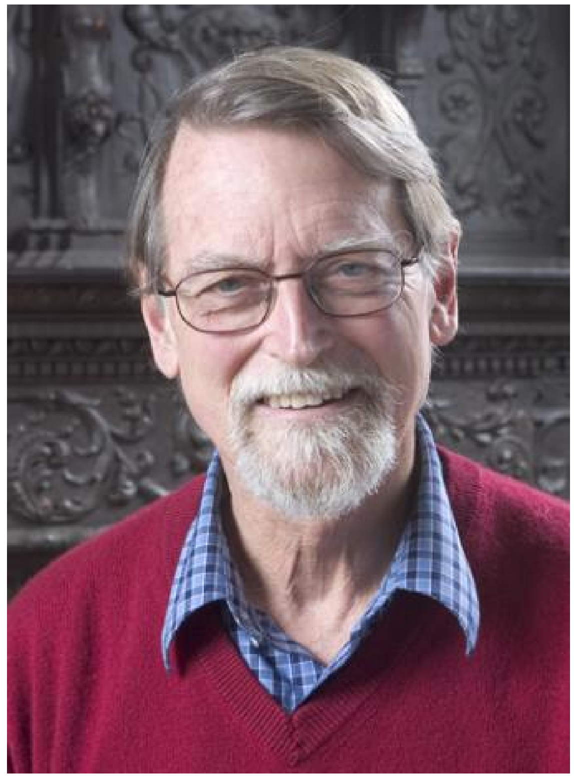

David Mumford

2D-shape analysis using conformal mappings

by E. Sharon and D. Mumford

Theorem 5.20 (Fundamental Theorem of Conformal Welding)**.** *A circle homeomorphism is quasisymmetric if and only if it is the conformal welding of a quasicircle.*

Of course, there are many weldings that are not quasisymmetric, e.g., the welding of any non-quasicircle.

Given any circle homeomorphism h and any $\epsilon > 0$ there is a welding map ϕ so that $h = \phi$ except on a set of Lebesgue measure ϵ. See My 2007 Annals paper.

The *Annals* paper also proves that "wild" homeomorphisms are weldings.

We say a homeomorphism h is log-singular if there is set $E \subset \mathbb{T}$ of zero logarithmic capacity so that $h(\mathbb{T} \setminus E)$ also has zero capacity.

Theorem: If h is log-singular then it is a conformal welding.

The resulting curve is very far from unique: any closed curve can be approximated by a curve with welding h.

Using this, Alex Rodriguez recently proved that any circle homeomorphism is a composition of two conformal weldings. See his paper on arXiv.

The fundamental theorem is is due to Pfluger, and has several proofs, e.g, using the measurable Riemann mapping theorem.

Assuming MRMT (for smooth μ), we can argue as follows.

Suppose $h : \mathbb{R} \to \mathbb{R}$ is quasisymmetric and let H be a QC extension to $\mathbb{R}^2$.

Choose a QC map G so that $\mu_G = \mu_H$ in upper half plane amd $\mu_G = 0$ in lower half-plane. (Then $G \circ H^{-1}$ is conformal in $\mathbb{H}$ and G is conformal in $\mathbb{H}_l$).

We claim $\Gamma = G(\mathbb{R})$ has welding h. Note $f = G$ is conformal from lower half-plane to one side of Γ. Next, $g = G \circ H^{-1}$ is conformal from upper half-plane to other side of Γ. Finally $g^{-1} \circ f = H \circ G^{-1} \circ G = H = h$ on $\mathbb{R}$.

We will give a proof that is very geometric and only uses the following facts:

- K-quasiconformal maps are compact.
- Quasisymmetric maps on $\mathbb{T}$ extend to be quasiconformal on the disk.
- Circles are removable for quasisymmetric maps.
- The uniformization theorem (for finitely connected planar domains).
- Koebe's circle domain theorem.

The first four we have discussed before.

Koebe's' circle domain theorem states that every finitely connected planar domain can be conformally mapped to a domain bounded only by circles or points.

This we will accept on faith.

One proof of this theorem is given in 2005 thesis of Karyn Lundberg.

We define a circle chain $\mathcal{C}$ to be a finite union of closed disks $\{D_k\}_1^n$ in $\mathbb{R}^2$ which have pairwise disjoint interiors and such that D_k is tangent to D_{k+1} for $k = 1, \ldots, n-1$, D_n is tangent to D_1 and there are no other tangencies. We also assume the disks are numbered in counterclockwise order.

The complement, $X = S^2 \setminus \cup_k D_k$, of a circle chain consists of two disjoint Jordan domains. We denote the bounded component by Ω and the unbounded component by Ω^*.

Let $f : \mathbb{D} \to \Omega$ and $g : \mathbb{D}^* \to \Omega^*$ be Riemann maps.

We call (f, g) a normalized circle chain pair if $f(0) = 0$, $g(\infty) = \infty$ and $\operatorname{dist}(0, \partial\Omega) = 1$.

Clearly, given a circle chain, we can always obtain a normalized pair by composing with a Möbius transformation.

Lemma 5.21. *Suppose $h : \mathbb{T} \to \mathbb{T}$ is an orientation preserving homeomorphism and suppose $\{x_k\}_1^n \subset \mathbb{T}$ is a finite collection of distinct points listed in counterclockwise order. Let $I_k = (x_k, x_{k+1}), k = 1, \ldots, n$ (modulo n). Then there is a normalized circle chain pair so that for each k,*

$$f(I_k) = \partial D_k \cap \partial\Omega,$$

$$g(h(I_k)) = \partial D_k \cap \partial\Omega^*.$$

We will say that any circle chain that satisfies this conclusion corresponds to h.

Another way of stating the lemma is that given any finite positive sequences $\{a_k\}$ and $\{b_k\}$ such that $\sum_{k=1}^n a_k = \sum_{k=1}^n b_n = 1$ we can find a circle chain so that the harmonic measure of each disk in the chain satisfies

$$\omega(D_k, 0, \Omega) = a_k, \quad k = 1, \dots n,$$
$$\omega(D_k, \infty, \Omega^*) = b_k, \quad k = 1, \dots n.$$

It is a fact that this circle chain is unique up to Möbius transformations, but we will not need this here.

Proof of Lemma 5.21. We apply the Koebe circle domain theorem to a domain $\Omega = \Omega_\epsilon$ constructed as follows.

Given n points $\{x_k\}$ on the unit circle $\mathbb{T}$, let $y_k = 2h(x_k) \in 2\mathbb{T} = \{z : |z| = 2\}$. Let γ_n be disjoint smooth Jordan arcs which connect x_k to y_k in the annulus $A = \{z : 1 \le |z| \le 2\}$.

Let $\{I_k\} \subset \mathbb{T}$ be the arcs bounded by the points $\{x_k\}$ and let $\{J_k\}$ be the corresponding arcs on $2\mathbb{T}$. Thus J_k has harmonic measure $|h(I_k)|$ with respect to ∞. Let $\delta = \inf_k |h(I_k)|$ be the smallest of these harmonic measures.

Our domain Ω_ϵ is the union of $\mathbb{D}$, $2\mathbb{D}^* = \{z : |z| > 2\}$ and an ϵ-neighborhood of each γ_n, where ϵ is assumed to be so small that these neighborhoods are pairwise disjoint and $\partial\Omega$ has n components.

Let $f_\epsilon : \Omega_\epsilon \to \Omega_\epsilon^*$ be the map given by Koebe's theorem. Using a Möbius transformation we may assume $f(0) = 0$, $f(\infty) = \infty$ and $\text{dist}(0, \partial\Omega_\epsilon) = 1$.

We claim that the n circles in the complement of Ω_ϵ^*, are all contained in some disk $D(0, R)$ with R independent of ϵ (but R may depend on h and n).

To see this, suppose the union of closed disks satisfies $\cup_k D_k \subset \{1 \leq |z| \leq R\}$ and that it hits both boundary components. Let Ω_1 be the connected component of $f_\epsilon(\Omega_\epsilon \cap D(0, 3/2))$ containing 0.

Then for ϵ small enough, each interval I_k has harmonic measure $\geq 1/2n$ in Ω_1 and hence has capacity in Ω_1 which is bounded away from zero depending only on n.

Thus by Lemma 2.23, every disk must hit $\{|z| \leq M_1\}$, for some M_1 depending only on n.

Similarly for Ω_2 (the connected component of $f_\epsilon(\Omega \cap \{|z| > 3/2\})$ containing ∞), i.e., there is a M_2 depending only on δ such that every disk must hit $\{|z| \geq R/M_2\}$.

If R is so large that $R/M_2 > 3M_1$, then every disk in our chain hits both $\{|z| \leq M_1\}$ and $\{|z| \geq 3M_1\}$. Therefore For large n this contradicts the following fact:

Lemma 5.22. *At most 6 disjoint disks can hit both* $\{|z| = 1\}$ *and* $\{|z| = 3\}$.

Proof. Each such disk has a subdisk of radius 1 contained in the annulus $\{1 \leq |z| \leq 3\}$. Each of these intersects the circle $\{|z| = 2\}$ in an arc of angle measure $2\arctan(1/2) \approx .9273 > \pi/3$, and hence there can be at most 6 such disks. $\square$

Since we now know that the n disks all reamin inside a fixed annulus $\{1 \leq |z| < R\}$, every disk remains bounded.

Since each disk has a fixed harmonic measure from 0, its radius remains uniformly bounded away from zero.

Thus we can pass to the limit as $\epsilon \to 0$, and get a circle chain of n non-degenerate tangent circles, that each have the correct harmonic measure.

Consequtive circles must touch in the limit, since the extremal distance beween them is zero. Non-conequtive circles do not touch because their extremal distance is positive.

This proves Lemma 5.21. $\square$

Proof of the Fundamental Theorem. Given a homeomorphism h and n equidistributed points $\{x_k\}_1^n \subset \mathbb{T}$, let $y_k = h(x_k)$ for $k = 1, \ldots n$ and consider the corresponding circle chain $\mathcal{C}_n$ as given by Lemma 5.21.

As before, let Ω_n, Ω_n^* denote the bounded and unbounded complementary domains. By reflecting through each circle we obtain a new chain with $n(n-1)$ circles. Continuing in this way we obtain, in the limit, a Jordan curve Γ_n, with complementary components D_n (bounded) and D_n^* (unbounded).

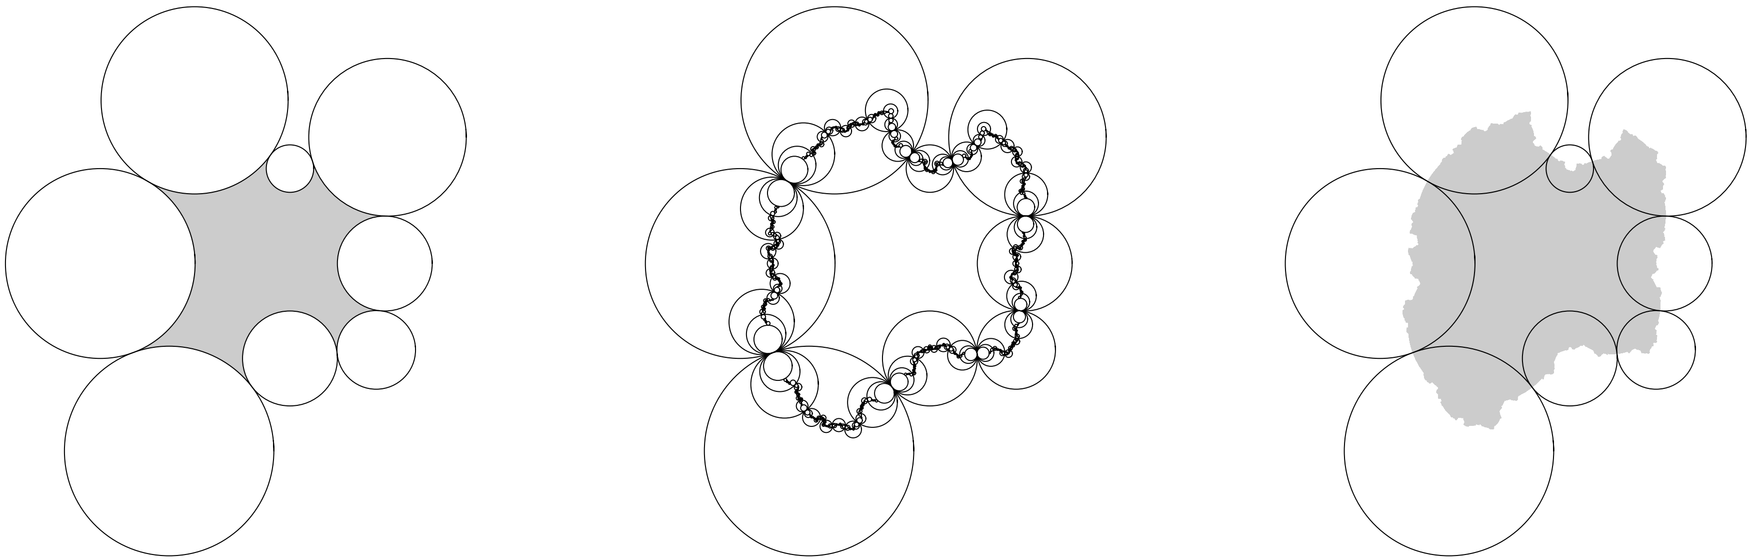

This shows the original chain and the domain Ω_n on the left, three iterations of the reflections in the center and the corresponding domain D_n on the right.

Similarly, given a circle chain $\mathcal{D}_n$ of n circles of equal size, with tangent points along the unit circle, we can reflect through the circles, getting a nested sequence of circle chains which limit on the unit circle, as shown below.

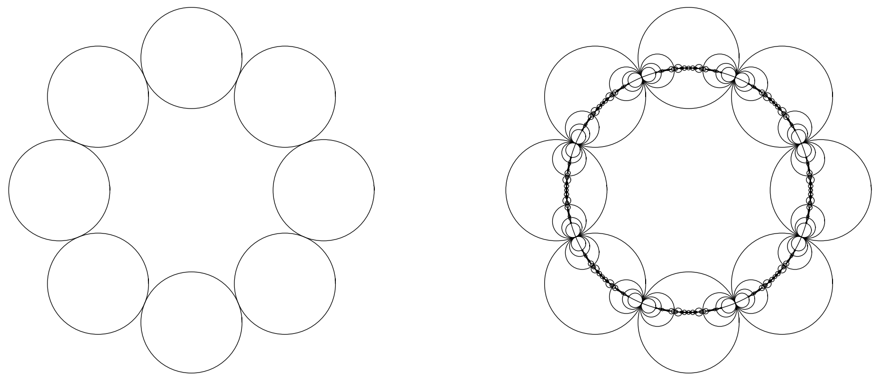

If h quasisymmetric, we know it is the boundary extension of some K-quasiconformal selfmap of the disk.

We claim there is a K-quasiconformal map of the plane sending the circles this figure to the circles in the previous figure. We will prove this by constructing the map separately inside and outside the unit circle.

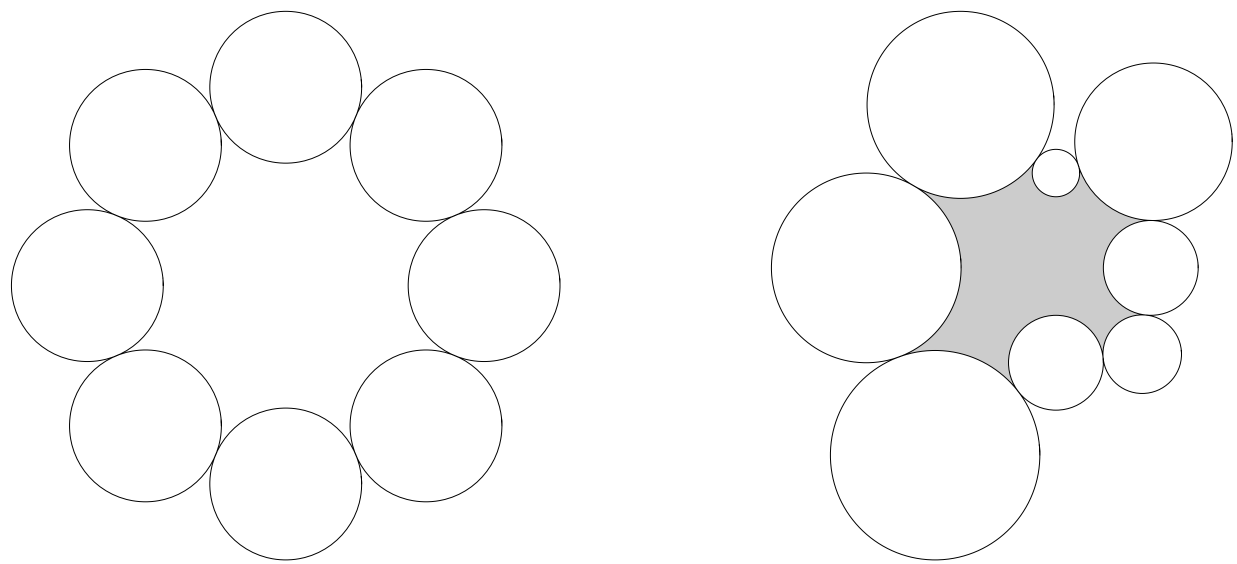

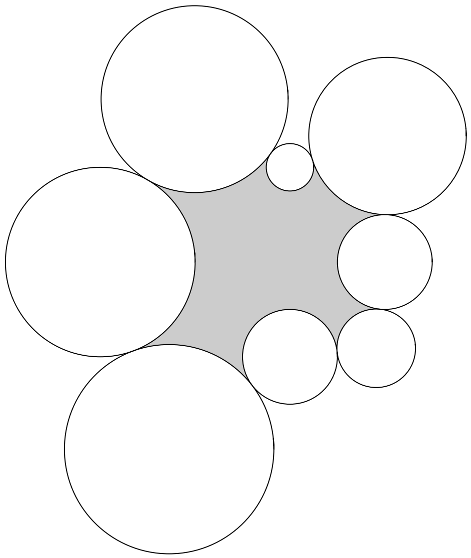
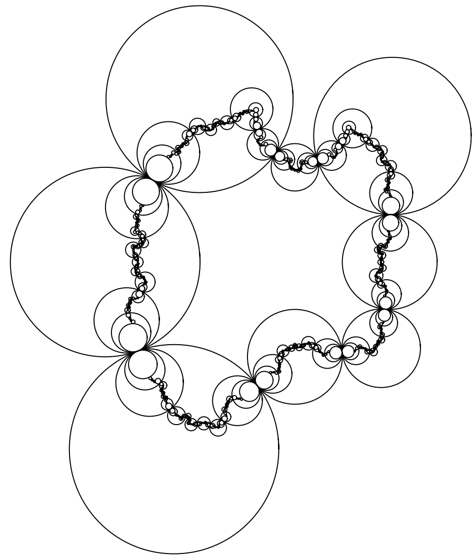
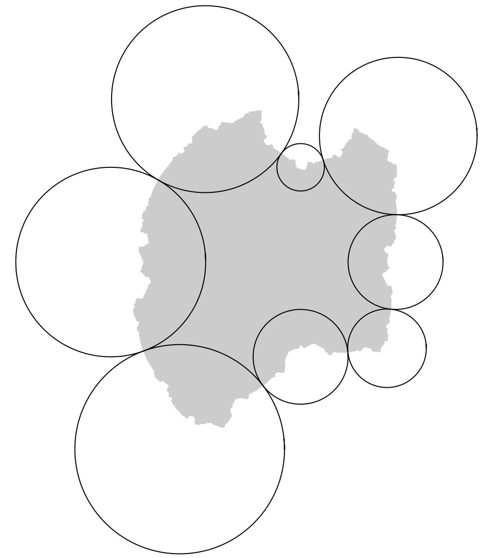
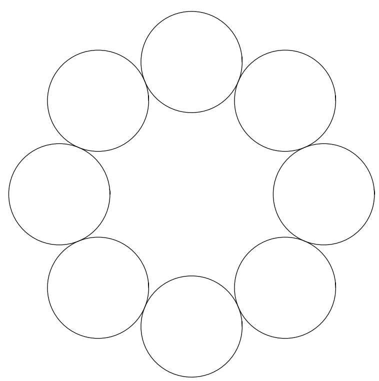
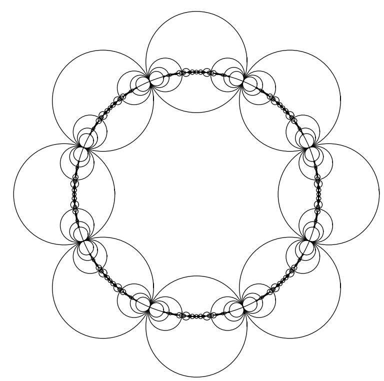

Let $W_n = S^2 \setminus \{x_1, \dots, x_n\}$. We may assume $n \geq 3$, so there is a universal covering map $\Pi : \mathbb{D} \to W_n$.

Let U_n be the component of $\Pi^{-1}(\mathbb{D})$ containing the origin, and note that by symmetry U_n may be chosen to be bounded by hyperbolic geodesics with end-points at the x_k's (the arcs $\mathbb{T} \setminus \cup\{x_k\}$ are hyperbolic geodesics in W_n; this is even clearer if we map $\mathbb{T}$ to $\mathbb{R}$ by a Möbius transformation).

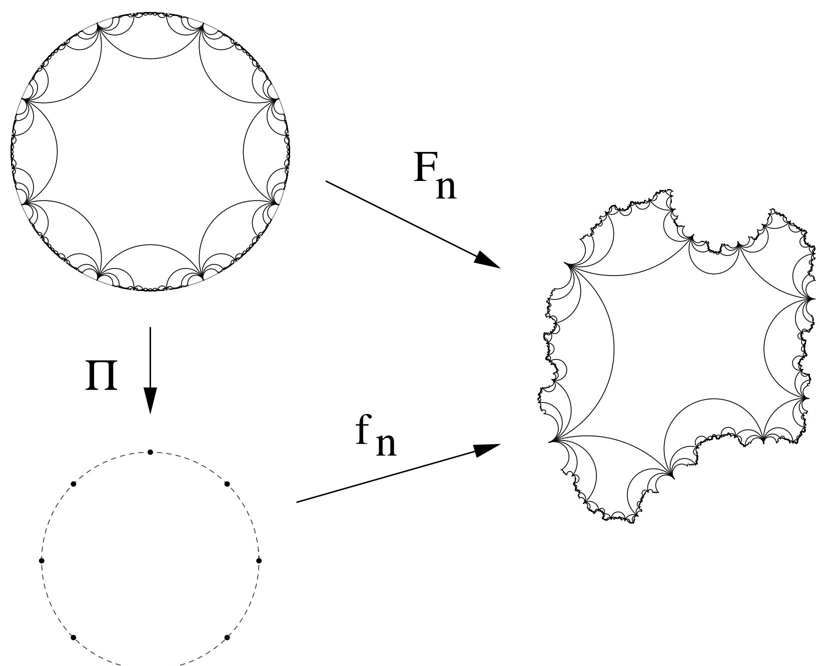

Reflecting these arcs across $\mathbb{T}$ gives the circle chain $\mathcal{D}_n$ in the figure below.

The conformal map $f_n \circ \Pi : U_n \to \Omega_n$ can be extended by repeated Schwarz reflection to a conformal map $F_n : \mathbb{D} \to D_n$.

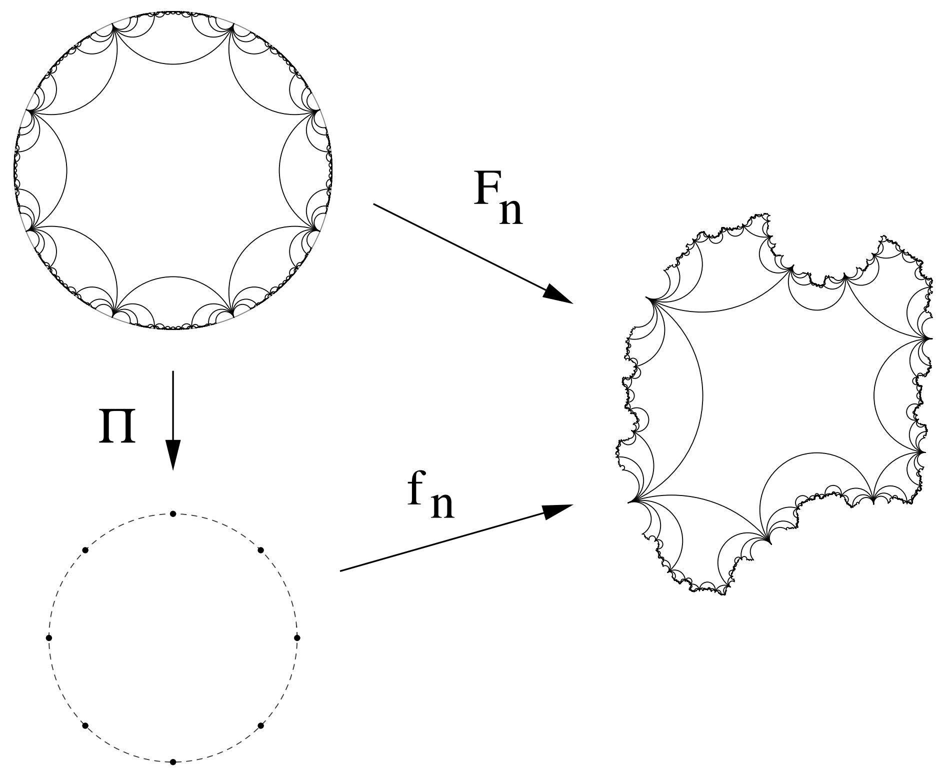

Similarly, Koebe's theorem gives a conformal map $g_n : \mathbb{D}^* \to \Omega_n^*$. Let $W_n^* = S^2 \setminus \{y_1, \ldots, y_n\}$ and consider $\Pi : \mathbb{D}^* \to W_n^*$ as the universal cover of W_n^*.

As above, we can lift g_n to map of $\Pi^{-1}(\mathbb{D}^*) \to \Omega_n^*$ and use Schwarz reflection to extend it to a map G_n from $\mathbb{D}^* \to D_n^*$. See below.

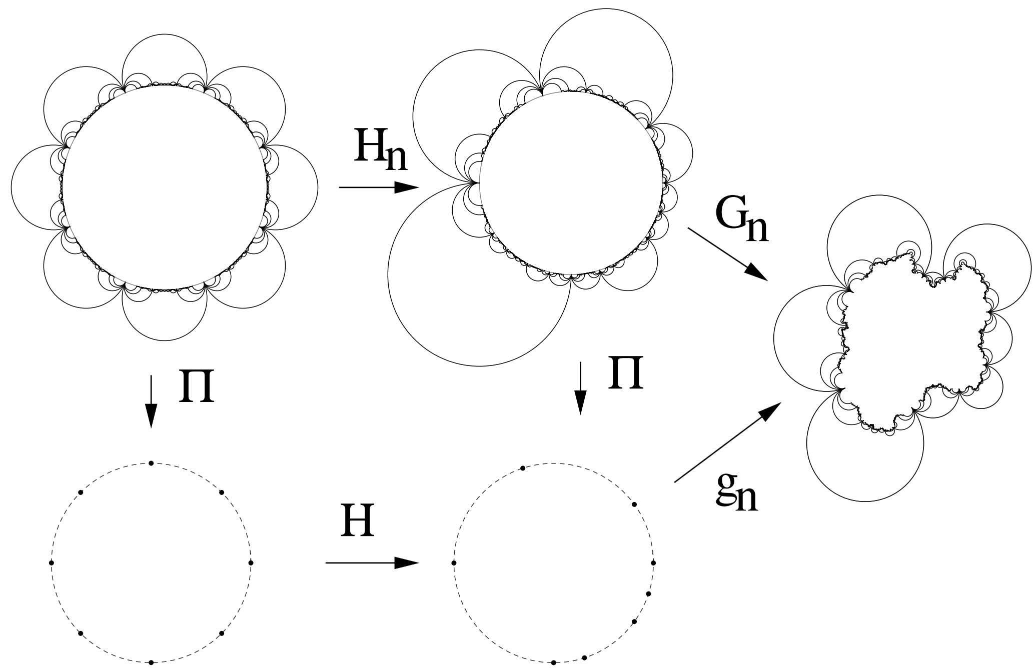

By assumption h is the boundary extension of a K-QC map of the disk to itself. By reflection we can extend this is a K-QC map H of S^2 to itself.

Then H maps W_n to W_n^* and lifts to a K-QC map of the universal covers.

We can represent these by $\mathbb{D}^*$ so we get a K-quasiconformal map $H_n : \mathbb{D}^* \to \mathbb{D}^*$ which conjugates the covering groups.

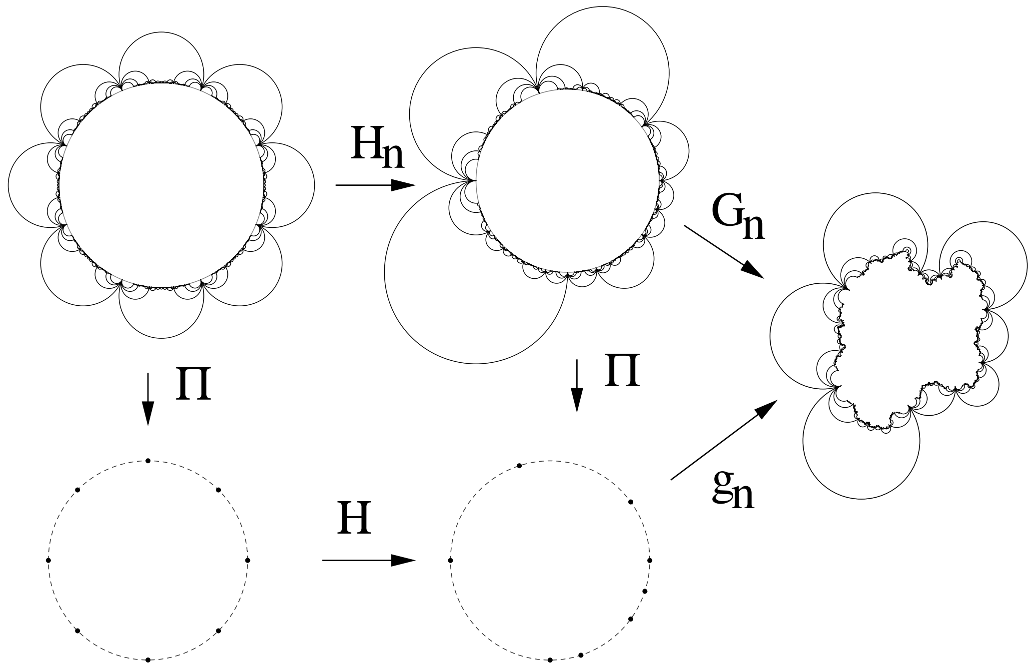

Thus $G_n \circ H_n$ is a K-quasiconformal map of $\mathbb{D}^*$ to D^* whose boundary values agree with F_n on $\mathbb{T}$, and hence these maps together define a K-quasiconformal map of S^2, since circles are removable for QC mappings.

This map takes $\mathbb{T}$ to Γ_n and the circle chain $\mathcal{D}_n$ to the chain $\mathcal{C}_n$.

Taking $n \to \infty$, using the uniform continuity of K-quasiconformal mappings and passing to a subsequence if necessary, we see that our circle chains converge uniformly to a K-quasicircle and that h is the corresponding conformal welding, as desired. $\square$