Angle Doubling, the Convex Hull Boundary, and the $K = 2$ Conjecture

D.B.A. Epstein A. Marden V. Markovic

30 March 2001

1 Introduction

In this paper we will develop the theory of a certain class of simply connected regions $\Omega \subset \mathbb{C}$, $\Omega \neq \mathbb{C}$. Our study will be coupled with the study of the boundaries of the corresponding hyperbolic convex hulls $\mathcal{CH}(X)$, where $X = \mathbb{S}^2 \setminus \Omega$.

The relative boundary $\partial\mathcal{CH}(X) \subset \mathbb{H}^3$ faces Ω. It is helpful to picture a domed stadium, such as one finds in Minneapolis, with Ω its *floor* and the *dome* given by $\mathrm{Dome}(\Omega) = \partial\mathcal{CH}(X)$. We will continue to use this terminology throughout the paper.

We prove (see Theorem 23.9) that the nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$ is a continuous, 2-Lipschitz mapping with respect to the induced hyperbolic path metric on the dome and the hyperbolic metric on the floor. Canary and Bridgeman have proved 4-Lipschitz.

Each bending line in $\mathrm{Dome}(\Omega)$ has a tangent line field. The inverse image under r is an arc of circle or a crescent, also with a tangent line field. The integral curves of the tangent line field are disjoint arcs of circle in Ω, each arc having the same endpoints as the corresponding bending line. This is a canonical lamination of any such region Ω by arcs of circle. The underlying point-set of the lamination is a closed set of positive area, unless $\Omega \subset \mathbb{S}^2$ is a round circle, in which case the bending lamination is empty.

The class of regions we will study is the collection of simply connected regions Ω which have the property that the result of applying any earthquake supported on the bending lamination (see Definition 6.2) is the dome of some other simply connected region Ω'. We denote this class of regions by $\mathfrak{QR}$, and call any such region *quake-resistant* (see Definition 10.2 and Definition 16.5).

There is an isometry (see Theorem 2.1) that takes the dome to the unit disk $\mathbb{D}^2$ and the bending lamination to a measured geodesic lamination (Λ, μ) in $\mathbb{D}^2 \equiv \mathbb{H}^2$. Our class $\mathfrak{QR}$ can be characterized by saying that the measured lamination (Λ, μ) satisfies the "ortho-condition" (see Definition 8.3 and Definition 13.1). This is a purely geometric condition on (Λ, μ).

In general it seems difficult to say when a given measured lamination in $\mathbb{D}^2$ comes from the bending measure of a dome. Measured laminations satisfying the ortho-condition are just such a fairly general class. During our work we show (see Proposition 20.2) that there exists a universal $\varepsilon > 0$ with the following property. Any measured lamination in $\mathbb{D}^2$, whose measure $\mu > 0$ satisfies $\mu(I) < \varepsilon$ for every transverse geodesic segment I of unit length, is the bending measure of a dome. D. Saric obtained similar results in his Ph.D. thesis.

We will introduce in $\mathbb{D}^2$ a lamination not just by hyperbolic geodesics but more generally by mutually disjoint circular arcs. We call a lamination a "crescent limit" (see Definition 14.5). A crescent limit is associated with a measured geodesic lamination that satisfies the ortho-condition (see Theorem 15.3 and Theorem 15.4). It carries its own transverse measure which we call the "circular measure". The endpoints of the leaves are the endpoints of the associated geodesic lamination. This more general lamination is needed to characterize limits of situations satisfying the finite ortho-condition.

During our study we will introduce the class of "angle-scaling" mappings $\{A_t(z; \Lambda, \mu)\}$ defined on the union of $\mathbb{D}^2$ and its exterior $\mathcal{E}\mathbb{D}^2$ with image in $\mathbb{S}^2$. These are related to "quakebends" or "complex earthquakes", but are a more encompassing class of maps. They are automatically defined on the points of $\mathbb{S}^1$ which are endpoints of the leaves.

Each measured lamination (Λ, μ) which satisfies the ortho-condition gives an associated crescent limit on $\mathbb{D}^2$ plus its reflection into $\mathcal{E}\mathbb{D}^2$. This data spawns a family of angle-scaling mappings $\{A_t(z; \Lambda, \mu)\}$, defined on $\mathbb{D}^2 \cup \mathcal{E}\mathbb{D}^2$, with parameter $t \in \mathbb{C}$. This family has the following properties.

- For each $-1 < \Im t \leq 1$, the angle-scaling maps are $K(t)$-quasiconformal mappings maps $\mathbb{D}^2 \to \Omega_t \in \mathfrak{QR}$, where $1 \leq K(t) \leq 2$. Each such map extends continuously to $\partial\mathbb{D}^2$ (see Theorem 18.1).

- When $-1 < \Im t < 1$, the maps are the restriction of $K(t)$-quasiconformal maps of $\mathbb{S}^2$, with $1 \le K(t) < 2$. The corresponding images Ω_t are all Jordan regions if and only if (Λ, μ) satisfies the ortho-condition.

- When $\Im t = 1$ the angle-doubling mappings collapse the exterior $\mathcal{E}\mathbb{D}^2$ in a manner which depends on the geodesic lamination. When restricted to $\mathbb{D}^2$, the mapping is referred to as an "angle-doubling mapping". Each image Ω_t is a Jordan region if and only if (Λ, μ) satisfies the strong ortho-condition (see Theorem 19.1).

- When $\Im t > 1$, the angle-scaling mapping restricted to $\mathbb{D}^2$ becomes a local quasiconformal developing mapping, or, more precisely, a quasiregular function. At the same time in $\mathcal{E}\mathbb{D}^2$ is folded.

- For each z in the closure $\overline{\mathbb{D}^2}$, $A_t(z)$ is holomorphic function of t on $\{t \in \mathbb{C} : -1 < \Im t\}$.

- For each $z \in \mathbb{S}^2$, $A_t(z)$ is holomorphic function of t on $\{t \in \mathbb{C} : -1 < \Im t < +1\}$. In fact the family $\{A_t(z; \Lambda, \mu)\}$ is a holomorphic motion of $\mathbb{S}^2$ for $\{t \in \mathbb{C} : -1 < \Im t < 1\}$ (see Theorem 21.3).

We add that the support of the Beltrami differentials of the restriction of the angle-scaling maps to $\mathbb{D}^2$ is contained in the crescent limit associated with (Λ, μ). Moreover, for parameter values in $\{t \in \mathbb{C} : \Im t \le 1\}$, the properties cited above for $\mathbb{D}^2$ apply equally well to $\mathcal{E}\mathbb{D}^2$.

Every geodesic lamination (Λ, μ) satisfying the ortho-condition is the bending lamination for the dome of some $\Omega \in \mathcal{QR}$. Moreover the region Ω is uniquely determined up to Möbius equivalence (see Theorem 16.1). It is the image of $\mathbb{D}^2$ under an angle-doubling map, defined on a crescent limit associated with the measured lamination. Conversely, each image of $\mathbb{D}^2$ under an angle-doubling map is a region $\Omega \in \mathcal{QR}$ associated with the bending measure on the dome.

The regions in $\mathcal{QR}$ have another property as well. Namely there is a 2-quasiconformal map $\Psi : \mathrm{Dome}(\Omega) \to \Omega$ that pointwise fixes the common boundary $\partial\Omega$ (see Theorem 17.8). This brings us to Thurston's $K = 2$ Conjecture Our original aim was to identify a large class of regions that satisfy Thurston's $K = 2$ conjecture. This conjecture states that $K = 2$ works for all simply connected Ω (Theorem 2.3,and Conjecture 2.4). Our work began in response to a paper of C. Bishop drawing attention to the potential importance in complex analysis of the $K = 2$ conjecture

A crescent limit is invariant under a fuchsian group G representing a finite area surface R if and only if it is associated with a finite measured lamination on R. When $-1 < \Im t < 1$, the angle-scaling mappings conjugate the action of G on $\mathbb{D}^2 \cup \mathcal{E}\mathbb{D}^2$ to a quasifuchsian group. When $\Im t > 1$, the image group H, which is a homomorphic image of G, is not necessarily discrete.

Our methods allow us to prove, with the help of McMullen's Disk Theorem, that if equivariance is required, then the $K = 2$ conjecture is false (see Theorem 22.8).

Along the way, other interesting facts appear as lemmas or corollaries. For example, a K-quasiconformal mapping of the unit disk $\mathbb{D}^2 \to \mathbb{D}^2$ is automatically a (K, a)-quasiisometry with additive constant $a = K \log 2$ when $1 < K \leq 2$ and $a = 2.37(K - 1)$ otherwise (see Theorem 17.1). This has the following consequence (see Corollary 17.3): If K is the minimum maximal dilation of quasiconformal mappings in a homotopy class of maps $R \to S$ between two Riemann surfaces of finite area, then the infimum of the constants for Lipschitz mappings in the same class satisfies $L \leq K$.

We have developed a theory that is reminiscent of quasifuchsian and projective deformations. At the same time our theory can be regarded as a general theory of "grafting". It seems useful in the further exploration of these topics.

2 Background

Brennan's conjecture is that the derivative f' of a Riemann map $f : \Omega \to \mathbb{D}^2$ of *any* simply connected region $\Omega \subset \mathbb{C}$, $\Omega \neq \mathbb{C}$, onto the unit disk $\mathbb{D}^2$ satisfies, for all p such that $4/3 < p < 4$,

$$\iint_\Omega |f'(z)|^p \, dxdy < \infty.$$

More specifically, the conjecture is that the upper endpoint of the open interval of values of p for which the integral is finite, is 4; the example of the Riemann map of the complement of a half-infinite line segment shows that the upper endpoint cannot be greater than or equal to 4 if the integral is to be finite for all such regions Ω. The lower endpoint is known to be 4/3. Many leading complex analysts have worked on this problem. We refer to Bishop's paper [4] for the background and history.

In his stimulating paper, Bishop introduced a new approach to the conjecture by showing there is perhaps a close connection to a seemingly unrelated conjecture that arises in hyperbolic geometry.

Throughout our work, $\mathbb{H}^3$ will denote hyperbolic 3-space and $\partial\mathbb{H}^3 \equiv \mathbb{S}^2 \equiv \mathbb{C} \cup \infty$ its boundary or sphere at ∞. We will use either the Poincaré ball model $\mathbb{D}^3$ or the upper halfspace model $\mathbb{U}^3$ depending on the situation.

Let Λ be a closed set of $\mathbb{S}^2 \equiv \mathbb{C}\cup\infty$ containing at least two points. We will begin by reviewing the notion of the hyperbolic convex hull of Λ, using first the Poincaré ball model (though the construction is of course independent of the model used to explain it). In the ball model, a hyperbolic plane is the intersection with the unit ball of a euclidean plane through the origin, or of a sphere which is orthogonal to the boundary of the unit ball. A closed hyperbolic halfspace is the closure of one of the two components into which a hyperbolic plane divides hyperbolic 3-space, and we take the closure in the closed unit ball. To form the convex hull $\mathcal{CH}(\Lambda)$, we take the intersection of all closed hyperbolic halfspaces containing Λ. It is easy to see that we can restrict attention to those halfspaces for which the closure C of the hyperbolic plane in the boundary of the halfspace contains at least two points of Λ, and we do restrict attention in this way. $\mathcal{CH}(C \cap \Lambda) = \mathcal{CH}(\Lambda) \cap C$ is the convex hull in the two-dimensional subspace C. If $C \cap \Lambda$ consists of exactly two points, then its convex hull is a geodesic and is known as a *bending line*.

Suppose $C \cap \Lambda$ contains more than two points. Then the convex hull has non-empty interior in C. It is known as a *flat piece*. The boundary components of $\mathcal{CH}(C \cap \Lambda)$ in C are geodesics, and these are also known as bending lines. They may or may not be bending lines under the previous definition. If $C \cap \Lambda$ is a finite set, then its convex hull is a finite sided ideal polygon.

Now regard $\mathbb{C} \cup \infty \equiv \mathbb{S}^2$ as the boundary (or sphere at infinity) of the upper halfspace model $\mathbb{U}^3$. We parameterize $\mathbb{U}^3$ as $\{(z,t) : z \in \mathbb{C}, 0 < t < \infty\}$. We recall that in $\mathbb{U}^3$ the (hyperbolic) geodesics are semicircles orthogonal to $\mathbb{C}$ and euclidean halflines orthogonal to $\mathbb{C}$. The hyperbolic planes are hemispheres orthogonal to $\mathbb{C}$ and euclidean halfplanes orthogonal to $\mathbb{C}$. The metric is

$$ds = \frac{de}{t},$$

where de is the infinitesimal euclidean metric.

Given a simply connected planar region $\Omega \neq \mathbb{C}$, let Ω^c denote the complement of Ω with respect to the extended plane. Every euclidean line or circle $c \subset \Omega$ is the boundary of a hyperbolic plane C. The plane C divides $\mathbb{H}^3$ into two half spaces one of which contains Ω^c in its closure. The convex hull of Ω^c, $\mathcal{CH}(\Omega^c)$, is the intersection of all these closed halfspaces.

We are interested in its boundary $\partial\mathcal{CH}(\Omega^c) \subset \mathbb{H}^3$. This boundary is connected and faces the *floor* Ω. We refer to the bounday as the *dome* and denote by $\mathrm{Dome}(\Omega)$. The common boundary of the dome and the floor is $\partial\Omega$. More precisely, denoting by Y either the dome or the floor, we have $\overline{Y} \setminus Y$ equal to the frontier of Ω in $\mathbb{S}^2$.

The dome is the union of the bending lines and the flat pieces. The closed set which is the complement of the open flat pieces is a union, in general an uncountable union, of mutually disjoint bending lines. A bending line is a complete geodesic whose end points lie in the boundary $\partial\Omega$. If ℓ is an isolated bending line, the two flat pieces abutting on it are bent with respect to each other. The exterior dihedral angle θ with $0 < \theta \leq \pi$ is called the *bending angle*. Thus, in the extreme case $\theta = 0$, the two flat pieces are not bent with respect to each other at all. If the geodesics are not isolated it is necessary to generalize the notion of "bending angle" by a process of integration.

If Ω is the interior of an ellipse, the dome is a half-ellipsoid. There are no flat pieces and the ellipsoid is swept out by the bending lines whose end points lie on symmetric points of the ellipse. More generally the dome of any euclidean convex region is a C^1-submanifold of $\mathbb{H}^3$. The dome of a euclidean convex region may contain flat pieces—this is the case for a triangular region, for example.

On the other hand suppose Ω is the complement of a half-infinite euclidean line. In this degenerate case the convex hull of Ω^c has no interior and the boundary of the convex hull can be considered as the union of two distinct sides of the vertical wall rising from the halfline.

We refer to [7] for a complete description of convex hulls and for a proof of the following theorems.

2.1 Theorem (W.P.Thurston). *The hyperbolic metric in $\mathbb{H}^3$ induces a path metric on the dome, which we will refer to as its hyperbolic metric. There is an isometry of the dome with its hyperbolic metric onto $\mathbb{D}^2$ with its hyperbolic metric. The set of bending lines is taken to a geodesic lamination of $\mathbb{D}^2$ and the bending measure is taken to a transverse measure associated with the lamination.*

2.2 Remark. Suppose $S_1 = \mathrm{Dome}(\Omega_1)$ and $S_2 = \mathrm{Dome}(\Omega_2)$ are isometric in their respective path metrics such that the corresponding bending measures are the same. Then Ω_1 and Ω_2 are Möbius equivalent. □

2.3 Theorem (Sullivan, Epstein-Marden). *There exists K such that, for any simply connected $\Omega \neq \mathbb{C}$, there is a K-quasiconformal map $\Psi : \mathrm{Dome}(\Omega) \to \Omega$, which extends continuously to the identity map on the common boundary $\partial\Omega$.*

Query: If $\Omega \subset \mathbb{S}^2$ is not a round disk, can $\Psi : \mathrm{Dome}(\Omega) \to \Omega$ be conformal?

It is known for the universal K that $K \geq 2$ (see Section 3). In working with a kleinian group which fixes Ω setwise, and therefore $\mathrm{Dome}(\Omega)$, one wants the map Ψ to be equivariant. In [7] it is shown that the maximal dilatation for equivariant maps satisfies $K_{\mathrm{eq}} < 82.7$. Using some of the same methods, but dropping the requirement of equivariance, Bishop [4] improved this to $K = 7.82$. In addition, Bishop presented another, much shorter proof of Theorem 2.3, which however does not seem to allow a good estimate of the constant. Bishop's proof reveals what seems to lie at the heart of the matter. Namely let $\Phi : \mathbb{D}^2 \to \Omega$ be a Riemann map. Bishop showed that $\Psi \circ \Phi(\zeta)$ can be defined for all $\zeta \in \partial\mathbb{D}^2$ and is a quasisymmetric homeomorphism $\partial\mathbb{D}^2 \to \partial\mathbb{D}^2$. Initially, this seems very surprising, as Φ itself may have bad boundary behavior, for example if $\partial\Omega$ is not locally connected.

2.4 Conjecture (Thurston). *$K = 2$.*

At the time (around 1984) that Thurston first made this conjecture, the distinction between the equivariant case and the general case was not clear. At that time, it didn't occur to anyone, as far as we know, that the answers might be different.

2.5 Theorem (Bishop). *If the conjectured $K = 2$ holds for all simply connected regions $\Omega \neq \mathbb{C}$, then Brennan's conjecture is true. If $K = 2$ for a particular region Ω then Brennan's conjecture holds for that particular region.*

In §23.14, using the notion of substantial boundary dilatation, we will relax the hypothesis of Bishop's theorem, while retaining the conclusion. The example in § 8.11 shows cases for which it is immediate that our new hypothesis holds. so that Brennan's Conjecture is true. On the other hand, it seems quite difficult to decide whether $K = 2$ holds for these cases.

The original purpose of this paper was to explore the $K = 2$ conjecture by presenting new evidence in its favour. Our class $\mathfrak{QR}$, which includes the sets of discontinuity for certain quasifuchsian groups, satisfies $K = 2$, and hence Brennan's conjecture holds for them.

We can prove that $K < 2$ for euclidean convex regions in the plane, as we will show in § 23.12. In Section 3, we show that $K = 2$ is the best possible quasiconformal constant in the case of the plane slit along a half-infinite straight line.

Yet if equivariance is required, K_{eq} *must* exceed two (see Theorem 22.8). We have found invariant regions Ω of quasifuchsian groups representing once-punctured tori for which the equivariant $K = 2$ cannot hold. The provably exceptional groups correspond to small open sets near the boundary of McMullen's embedding of Teichmüller space.

Figure 1: point at angle 2beta from 0. edge shown of length tan beta.

3 Easy examples

3.1 The round disk in $\mathbb{C}$. Let Ω be the open unit disk in $\mathbb{C} = \partial\mathbb{U}^3$. The convex hull $\mathcal{CH}(\Omega^c)$ is the complement of the open three-dimensional halfdisk of radius one, with centre $0 \in \mathbb{C}$. Its boundary is the hyperbolic plane H which is the hemisphere of radius one, with centre 0.

3.2 Nearest point retraction. If $X \subset \mathbb{H}^n$ is closed, convex and non-empty, then we have a well-defined continuous map $r : \mathbb{H}^n \to X$ which maps each point p to the point $r(p) \in X$ which minimizes $d(p, r(p))$. If $p \in X$, then $r(p) = p$. If $p \notin X$, then we take a ball centred at p and expand it until it first hits X. The first point of contact is $r(p)$.

This domain of r can be extended to $\mathbb{H}^n \cup \partial\mathbb{H}^n$, provided X is a closed subset of $\mathbb{H}^n \cup \partial\mathbb{H}^n$. If $p \in \partial\mathbb{H}^n \setminus X$, then $r(p)$ is defined by expanding the horosphere centred at p until the horosphere first hits X. This will be at a point of $\mathbb{H}^n$. Under this definition, r is again continuous.

The nearest point retraction $r : \Omega \to H$ is given by euclidean straight line projection from the point $S = (0, 0, -1) \in \mathbb{R}^3$. S is the South Pole for the euclidean unit sphere containing H. A line drawn from S to a point $v \in H$ meets $\mathbb{C}$ in z and the tangent plane to H at v at the same angle, proving the standard fact that projection from S is conformal. This also shows that there is a euclidean circular arc orthogonal to H and $\mathbb{C}$, passing through v and z. This arc is the hyperbolic geodesic from v to $r(v) = z$.

In this case, r is itself the quasiconformal map between the floor Ω and the dome H and the quasiconformality constant is 1.

It will be useful to have some data about r in euclidean terms. We take the intersection with a vertical plane through 0. So r can be regarded as a map from the interval $[-1, 1]$ to the semicircle. If the point $z \in [-1, 1]$ is given the coordinate $\tan(\beta)$, as in Figure 1, then the point $r(z) = v$ on the semicircle has angular coordinate 2β. Note that these coordinates are those corresponding to euclidean distance in domain and range. With respect to these coordinates $r(t) = 2\arctan(t)$, so that $r'(t)$ varies between 2, when $t = 0$, and 1, when $t = \pm 1$.

3.3 Definition. Let $\alpha \in \mathbb{R}$ with $0 < \alpha \leq 2\pi$. By a *wedge* of angle α, we mean the open subset $\{z \in \mathbb{C} : 0 < \arg(z) < \alpha\}$. When $\alpha = 2\pi$, Ω is the complement in $\mathbb{C}$ of the closure of the positive x-axis. We also regard a wedge as lying in the extended complex plane. The image in $\mathbb{S}^2$ or $\mathbb{C}$ of a wedge under a Möbius transformation is called a *crescent* of angle α. □

Although it may at first seem that all wedges should be discussed within the same framework, our problem actually looks quite different when $\alpha > \pi$ and when $\alpha < \pi$. The reason is that the structure of the boundary of the convex hull of the complement Ω^c in $\mathbb{S}^2$ of Ω is quite different in the two cases. Ω^c is itself a (closed) wedge of angle $2\pi - \alpha$.

3.4 The fat wedge. Suppose $\pi \leq \alpha < 2\pi$. The boundary of the wedge consists of two rays starting at 0, and at an angle α to each other. There are two vertical planes with boundary containing one of these rays. Each such plane cuts upper half 3-space $\mathbb{U}^3$ into two components, one of which contains Ω^c in its boundary in $\mathbb{C} \cup \{\infty\}$. We get two components, one for each of the two planes. Their intersection is a 3-dimensional wedge in $\mathbb{U}^3$, with angle $2\pi - \alpha$ between the 2-dimensional faces. This 3-dimensional wedge is equal to the convex hull $\mathcal{CH}(\Omega^c)$. Each 2-dimensional face is half a hyperbolic plane. The two 2-dimensional faces meet along the vertical geodesic through 0. In the ball model, the 3-dimensional wedge looks like a segment of an orange. Here we have exactly one bending line, and the bending angle is $\alpha - \pi$.

Let $S = \partial\mathcal{CH}(\Omega^c)$. The isometry ι of S onto the upper halfplane $\mathbb{U}^2$ unbends S. The boundary value of ι on the positive real axis is the identity. On the ray $\arg(z) = \alpha$, we have $\iota(re^{i\alpha}) = -r = re^{i\pi}$. Note that there is no choice about ι, except for composition on the left with an isometry of the hyperbolic plane.

The uniquely extremal quasiconformal map $\psi : \Omega \to \mathbb{U}^2$, from Ω to the upper halfplane $\mathbb{U}^2$ with these boundary values is $re^{i\theta} \mapsto re^{i\theta\alpha/\pi}$ (see [14]). Its maximal dilatation is $\frac{\alpha}{\pi}$. The composition $\psi^{-1} \circ \iota : S \to \Omega$ is $\frac{\alpha}{\pi}$-quasiconformal and is the identity on $\partial\Omega$. Thus the $K = 2$ conjecture is verified in this case.

When $\alpha = \pi$, the map from S to Ω is actually conformal.

The limiting case $\alpha = 2\pi$ corresponds to the plane slit along a half-infinite line. It is treated much as before, except that we need the intersection of three halfspaces rather than of two to express the convex hull of Ω^c. The first two halfspaces lie over $\{z : \Im(z) \geq 0\}$ and over $\{z : \Im(z) \leq 0\}$. There are many choices for the third halfspace, for example the halfspace lying over

Figure 2: The fat wedge

Figure 3: Orange segment

$\{z : \Re(z) \geq 0\} \subset \mathbb{C}$. For this case the minimum quasiconformal constant is exactly two.

3.5 The thin wedge. The other type of wedge has $0 < \alpha \leq \pi$. Vertical halfspaces are now no good—no vertical halfspace contains the complementary wedge Ω^c. The convex hull of Ω^c is the complement of an open euclidean half-cone, with vertex at 0, angle $\alpha/2$ and axis along the ray $Y = \{z : \arg(z) = \alpha/2\}$. The half-cone is the union of all euclidean hemispheres with center on Y and which are tangent to the positive real axis and to the ray through 0 at angle α. The convex hull boundary is the surface of the half-cone. There are no flat pieces. The bending lines are the half-circles joining the two points of tangency just mentioned.

Once again, there is a uniquely extremal quasiconformal map $\Psi : \mathrm{Dome}(\Omega) \to \Omega$. To see what this is, note that the inclusion of $\mathbb{U}^3$ with the hyperbolic metric in $\mathbb{R}^3$ is conformal. In particular, $S = \partial\mathcal{CH}(\Omega^c)$ is conformally equivalent to the same half-cone with the euclidean metric. Such a euclidean cone is flat and unrolls onto a wedge with angle $\pi\sin(\alpha/2)$. Converting Ψ so that it maps this wedge to the wedge of angle α, the condition that Ψ should be the identity on the boundary translates to the condition that on the boundary the map between wedges should preserve the euclidean distance to the cone-point. Under this translation, the uniquely extremal quasiconformal map between the two wedges is given by $re^{i\theta} \mapsto re^{i\theta\alpha/\pi\sin(\alpha/2)}$. The dilatation K is equal to $\pi\sin(\alpha/2)/\alpha$, which decreases from near $\pi/2$ for α small to 1 when $\alpha = \pi$. Again the $K = 2$ conjecture holds for this case.

It is also interesting to work out the isometry $\iota : S \to \mathbb{U}^2$ in the case of a wedge with $\alpha \leq \pi$. To construct the isometry, consider the generator σ of the cone lying vertically above the ray Y. This is a straight line through 0, at an angle $\alpha/2$ to the horizontal. We have to define ι to be an isometry of σ with the vertical geodesic σ_0 over 0 in $\mathbb{U}^2$, and then extend it to be an isometry on each bending line, noting that the bending lines are orthogonal

Figure 4: orange segment in upper half 3-space

Figure 5: The bending angle for a fat wedge

Figure 6: The thin wedge showing the bisector Y, angle $\alpha/2$ and trig $r\sin(\alpha/2)$

to σ. The isometry on σ is given by mapping a point $p \in \sigma$ at a euclidean distance r from the origin to the point $ir^{1/\sin(\alpha/2)} \in \sigma_0$. One can check that this is an isometry. We now define ι on each bending line, that is, on each semicircle orthogonal to $\mathbb{C}$, intersecting σ orthogonally. We achieve an isometry by using a euclidean similarity to map it to the appropriate semicircle in $\mathbb{U}^2$, orthogonal to σ_0. The restriction of ι to the boundary has the form $r \mapsto r^{1/\sin(\alpha/2)}$ and $re^{i\alpha} \mapsto -r^{1/\sin(\alpha/2)}$.

We can also compute the bending measure explicitly by looking at the induced measure on σ. The measure is invariant under hyperbolic isometries preserving the various subsets. That is, it is invariant under euclidean multiplication by any positive real number. This means that, if we identify σ with the positive reals, and then with the reals by taking the logarithm, the measure is taken to a positive multiple of Lebesgue measure. The factor turns out to be $\cos(\alpha/2)$; we will not provide details of the computation. It is not a coincidence that this is the geodesic curvature of σ as a curve in hyperbolic 3-space.

We will also compute the nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$ for this case and show that it too is a 2-quasiconformal map.

Let Ω be a thin wedge with vertex at $z = 0$. First note that the retraction satisfies $|r(z)| = |z|$, that is, euclidean distances from the origin are preserved. Let c be a circular arc which is the intersection with Ω of a circle centered at 0. Let $c^* \subset \mathbb{U}^3$ denote the geodesic (half-circle) with the same endpoints as c. Next note that $r(c) = c^*$.

Let $\mathbb{D}^2 \subset \overline{\Omega}$ denote the disk which is tangent to the boundary of Ω at the endpoints of c. It will be no restriction to assume that it is a unit disk. Then let H denote the hyperbolic plane rising from $\partial\mathbb{D}^2$. It is tangent to $\mathrm{Dome}(\Omega)$ along c^*.

Let r^* denote the nearest point map $\mathbb{D}^2 \to H$. This is a conformal map

Figure 7: thin wedge with circles

and $r^*(z) = r(z)$ for $z \in c$, We want to compute the distortion $|r'(z)|$ along c, using euclidean coordinates. For this derivative, $|r'(z)| = |r^{*\prime}(z)|$. Moreover r^* is conformal, and so $|r^{*\prime}(z)|$ itself is independent of the direction in which it is computed. We compute $r^{*\prime}$ using § 3.1, where it is shown the derivative required is equal to the derivative of $t \to 2\arctan(t)$, where t is the euclidean distance $|z|$. The largest value of the derivative is given when $|z|$ is as small as possible, namely $|z| = \tan(\alpha/4)$, when the derivative is $2\cos^2(\alpha/4)$. Since $0 < \alpha \leq \pi$, this derivative lies between 1 and 2.

Now r sends rays in Ω from $z = 0$ isometrically in the euclidean metric to rays in the half-cone $\mathrm{Dome}(\Omega)$. Therefore the distortion $|r'| = 1$ in the radial direction. Consequently r is 2-quasiconformal as claimed. In fact its constant is equal to $2\cos^2(\alpha/4) < 2$.

3.6 The parallel strip $\Omega = \{z : 0 < \Im(z) < 2\}$. By moving the vertex of the wedge to infinity and letting the angle α tend to zero, we obtain the parallel strip as a limit of wedges. The convex hull of the strip is the complement of the solid half-cylinder of radius 1 lying on Ω. Its boundary is the half-cylinder. There are no flat pieces and the bending lines are the half-circles with endpoints at the symmetric points. $\mathrm{Dome}(\Omega)$ unrolls onto to the strip $\{z : 0 < \Im(z) < \pi\}$, This is a conformal map that pointwise fixes the lower edge and preserves symmetric points in the edges. The mapping $\Phi : \mathrm{Dome}(\Omega) \to \Omega$ that pointwise fixes the boundary is given via the infinite strips by $x \mapsto x, \;\; y \mapsto \frac{2}{\pi}y\}$. This is a $\frac{\pi}{2}$-quasiconformal map.

For later use we will also compute the nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$. To do so note that r is a euclidean isometry of each horizontal line in Ω to a generating line in $S = \mathrm{Dome}(\Omega)$. r takes each vertical segment $[a, a + 2i]$ to the half-circle in S with the same endpoints. We know that r does not distort in the horizontal direction. The distortion in the vertical direction has been calculated at the end of § 3.1, where it is shown that r' varies between 1 and 2.

Therefore $r : \Omega \to \mathrm{Dome}(\Omega)$ is 2-quasiconformal.

4 Maximal disks

Every open subset $\Omega \subset \mathbb{S}^2$ is a union of round open disks. As usual we are assuming that the complement Ω^c is not the null-set, not a single point, and not the whole of $\mathbb{S}^2$. We partially order these disks by inclusion. It is then clear that Ω is the union of maximal open disks. A corresponding closed disk meets the complement Ω^c in one, two or more points.

4.1 Example. For example, working in $\mathbb{C}$, we can take $\Omega = \{(x, y) : y > x^2\}$. The boundary $\partial\Omega$ is a parabola (together with the point at infinity). For each $x > 0$, there is a maximal disk meeting Ω^c in the two points $(\pm x, x^2)$. There is exactly one maximal disk which is not of this form, namely the disk of radius $1/2$, with centre at $(0, 1/2)$. This meets Ω^c only at $(0, 0)$. $\square$

4.2 Lemma. *Ω is the union of round open disks in $\mathbb{S}^2$ whose closures meet Ω^c in at least two points. The hyperbolic convex hull of the boundary circle of such a disk is a support plane for $\mathcal{CH}(\Omega^c)$, meeting $\mathcal{CH}(\Omega^c)$ at least along a geodesic.*

Proof: Let $p \in \mathbb{D}^3 \cup \mathbb{S}^2 \setminus \mathcal{CH}(\Omega^c)$. Let $r(p) \in \mathcal{CH}(\Omega^c)$ be the nearest point to p. Let H be the plane through $r(p)$ orthogonal to the geodesic $[p, r(p)]$. Then H is a support plane for $\mathcal{CH}(\Omega^c)$.

The closure of H meets Ω^c in at least two points. To see this, first note that the intersection with Ω^c cannot be empty, for then H could be moved further from p, so that the halfspace bounded by H contains Ω^c, but not $r(p)$.

The intersection $\bar{H} \cap \Omega^c$ cannot be a single point x. For consider the geodesic $\gamma \subset H$ through $r(p)$ orthogonal to $xr(p)$. We can rotate H slightly about γ to a new plane H' whose closure is disjoint from the closed set Ω^c. We then apply to H' a small translation along the geodesic through $pr(p)$ and in the direction $pr(p)$, so that it remains disjoint from Ω^c. The halfspace bounded by H' and containing Ω^c does not contain $r(p)$, and this is a contradiction.

From this the lemma follows immediately. $\blacksquare$

4.3 Philosophy. We assume throughout that Ω is a connected open subset of $\mathbb{S}^2$ which is not the whole of $\mathbb{S}^2$ nor the nullset, nor the complement of a single point. We also assume the the complement Ω^c of Ω is connected. These conditions imply that Ω is homeomorphic to an open disk.

It is the philosophy of this paper that a key to understanding Ω, from the point of view of complex analysis and the Riemann mapping theorem, is to understand how Ω is built up from its maximal disks. By Lemma 4.2, we may restrict attention to maximal open disks with closures which meet Ω^c in at least two points, and, from now on, this will be what we mean when we talk of a *maximal disk*.

For many purposes, a reasonable way to approximate Ω is to take the union of only a finite number of maximal open disks. This approximates Ω^c in the Hausdorff metric on compact subsets of $\mathbb{S}^2$. It is often sufficient to prove a result for a general Ω, by proving an appropriately powerful version of the desired result when Ω is a finite union of disks. Even if this philosophy fails for particular problems, it may provide a useful direction in which progress can be made.

A further aspect of this philosophy is that we can consider $\mathbb{S}^2$ to be the boundary of hyperbolic 3-space $\mathbb{H}^3$. In this context a round disk is the intersection with $\mathbb{S}^2$ of the closure of a hyperbolic half 3-space. The hyperbolic convex hull $\mathcal{CH}(\Omega^c)$ is the intersection of closed halfspaces which contain Ω^c. A maximal disk in Ω, whose closure meets Ω^c in at least two points, is the complement in $\mathbb{S}^2$ of a closed halfspace which contains Ω^c, such that the boundary of the halfspace is a support plane for $\mathcal{CH}(\Omega^c)$.

The convex hull boundary S is isometric to the hyperbolic plane (see Theorem 2.1). This hyperbolic plane has a measured lamination. The leaves of the lamination are called *bending lines*. These arise from a maximal disk in Ω or from a supporting halfplane of $\mathcal{CH}(\Omega^c)$. Such a supporting halfplane intersects $\partial\mathcal{CH}(\Omega^c)$ either along a single geodesic, or along a convex subset of the supporting halfplane H. In the latter case, the boundary of the convex subset in H is a union of geodesics. Each such geodesic is called a bending line. The union of the bending lines is a closed subset of S.

Given Ω, we can construct the corresponding measured lamination on the hyperbolic plane. The measured lamination then determines Ω up to Möbius transformation. In the case of a finite lamination, this is obvious. In general, the lamination forces the hyperbolic plane to be bent into hyperbolic 3-space in a certain way, dividing $\mathbb{H}^3$ into two components, one of which is convex. The closure of the convex component meets $\mathbb{S}^2$ in Ω^c and therefore determines Ω.

There is one exceptional situation to which we should draw the reader's attention.

4.4 Definition. If the measured lamination on $\mathbb{H}^2$ consists of a single geodesic, with bending angle π, then the surface S is called *folded*. It can be regarded as a hyperbolic plane in $\mathbb{H}^3$, by thinking of the two sides of the folded surface as distinct. The need to consider such surfaces was discussed just before Theorem 2.1. □

The corresponding region $\Omega \subset \mathbb{S}^2$ is the complement of a closed circular arc. Using a Möbius transformation, we may assume that this circular arc is a ray from 0 to ∞ in $\mathbb{C}$, so that Ω is a wedge of angle 2π.

If we start with an arbitrary measured lamination, or even an arbitrary finite measured lamination, we will not in general obtain an embedded hyperbolic plane when we bend it into $\mathbb{H}^3$. Even if the surface is embedded, this is not sufficient for it to be the boundary of a convex subspace—we also need the map of $\mathbb{H}^2$ into $\mathbb{H}^3$ to be a proper map.

It would be good to have a reasonable way of looking at a measured lamination so that we can determine whether this is a bending lamination corresponding to some Ω, without actually constructing the map of $\mathbb{H}^2$ into $\mathbb{H}^3$. Unfortunately criteria for this are unknown. One of the things we do in this paper is to solve this problem for a certain class of measured laminations, but we are still far from being able to deal with the general problem. We will introduce such criteria in Section 13, and not discuss them further at this point.

We now show how one can deduce global convexity from local convexity. We will make no attempt at generality: instead we will concentrate on what we actually need.

4.5 Definition. Let $f : \mathbb{H}^2 \to \mathbb{H}^3$ be a continuous map. We say that f is *locally convex* if, for each $x \in \mathbb{H}^2$, there is an open ball $N_x \subset \mathbb{H}^3$, centred at $f(x)$ and a neighbourhood U_x of x, such that $f|U_x$ is injective, $f(U_x) \subset N$ and $f(U_x)$ separates N into two components A_x and B_x such that A_x is a convex subset of $\mathbb{H}^3$. Furthermore, we assume that, if $y \in \mathbb{H}^2$ is near enough to x then A_x meets A_y, and $B_x \cap A_y = A_x \cap B_y = \emptyset$. □

4.6 Lemma. *Let $f : \mathbb{H}^2 \to \mathbb{H}^3$ be an injective, locally convex, proper map. Then $f(\mathbb{H}^2)$ separates $\mathbb{H}^3$ into two components, one of which contains all the subsets A_x of Definition 4.5. This component, which we call A, is a convex subset of $\mathbb{H}^3$, whose boundary in $\mathbb{H}^3$ is $f(\mathbb{H}^2)$. Moreover, a geodesic in the convex subset $\bar{A}$ either lies entirely in $f(\mathbb{H}^2)$, or meets $f(\mathbb{H}^2)$ at most at its endpoints.*

Proof: This is a standard result about locally convex submanifolds, which we repeat for the convenience of readers. Given any two points in the closure $\bar{A}$ of A in $\mathbb{H}^3$, the Ascoli theorem gives us a shortest path α in $\bar{A}$ parametrized by path-length and joining the two points. Let β be a hyperbolic geodesic defined on a maximal subinterval of the domain of α. We prove by contradiction that $\alpha = \beta$. So suppose x is an endpoint of β which is not an endpoint of α. Let N_x be a disk in $\mathbb{H}^3$ with centre x, such that $N \cap \bar{A}$ is hyperbolically convex. We obtain a contradiction to the maximality of β by looking at two points in N on either side of $x \in \alpha$. It follows that $\bar{A}$ is convex and that α is a hyperbolic geodesic.

Suppose α does not lie entirely in $f(\mathbb{H}^2)$. Let $x \in \alpha \setminus f(\mathbb{H}^2)$. Let $B \subset A$ be a small ball with centre x, disjoint from $f(\mathbb{H}^2)$. By convexity, $\bar{A}$ contains all hyperbolically geodesic segments joining points of B to points of α. So, except possibly for its endpoints, α is contained in A and is disjoint from $f(\mathbb{H}^2)$. ∎

5 Angle doubling and finite laminations

If $\alpha \leq \pi$, we can define a homeomorphism

$$A : z = re^{i\theta} \mapsto re^{2i\theta} = ze^{i\theta} = z^{\frac{3}{2}}\overline{z}^{-\frac{1}{2}} \text{ for } 0 \leq \theta \leq \alpha$$

of the wedge of angle α into $\mathbb{C}$. This map is the identity on the positive real axis, and on the ray of angle α its values are $re^{i\alpha} \mapsto re^{2i\alpha}$. On this ray, its boundary values agree with those of the elliptic transformation with fixed points $0, \infty$ that maps the ray of angle α to the ray of angle 2α. We will refer to the homeomorphism as the *angle doubling map* of the wedge of angle α.

More generally we can take any crescent of angle α formed by two circles intersecting at points p, q. Designate one of the two edges of the crescent as the initial and the other the terminal edge. There is a Möbius transformation T that maps the crescent onto the wedge of angle α, sending the initial edge to the positive real axis. With the vertex of the wedge at $0 \in \mathbb{C}$, the only Möbius transformations which preserve the wedge setwise are multiplications by positive real numbers. T is determined up to composition with such a Möbius transformation. Let S be any Möbius transformation. We will refer to $f = SAT$, which maps the crescent into $\mathbb{S}^2$, as an *angle doubling map* of the crescent.

Angle doubling maps are quasiconformal (Teichmüller maps) with complex dilatation $\kappa(z) = \frac{\partial A/\partial \bar{z}}{\partial A/\partial z} = -\frac{z}{3\bar{z}}$ and maximal dilatation $K = \frac{1+|\kappa(z)|}{1-|\kappa z|} = 2$. It is known (see [14]) that the angle doubling map of a wedge or crescent is the quasiconformal map with the least maximal dilatation $K = 2$ among all quasiconformal mappings with the same boundary values.

5.1 Definition. If $\nu \in \mathbb{R}$, we can define the *angle scaling* map from a wedge of angle α to a wedge of angle $\nu\alpha$. If $-2\pi < \nu\alpha < 2\pi$, the wedge of angle $\nu\alpha$ is an ordinary wedge in $\mathbb{C}$. The map is given by

$$A_\nu : z = re^{i\theta} \mapsto re^{i\nu\theta} = ze^{i(\nu-1)\theta} \text{ for } 0 \leq \theta \leq \alpha$$

of the wedge of angle α into $\mathbb{C}$. If $|\nu\alpha| \geq 2\pi$, the wedge is not embeddable in $\mathbb{C}$, but it has a well-defined projective structure (the geometric structure modelled on 1-dimensional complex projective space $\mathbb{PC}^1 = \mathbb{S}^2$ with group $\mathbb{P}GL(2, \mathbb{C})$, the group of Möbius transformations). It is clear how the definition of the angle scaling map extends.

The angle scaling map is an orientation-preserving transformation if $\nu > 0$, a retraction onto the positive real axis if $\nu = 0$ and an orientation-reversing map if $\nu < 0$.

If $\nu = 0$, the map is called the *angle collapse*.

The angle scaling map is the identity on the positive real axis, and on the ray of angle α its values are $re^{i\alpha} \mapsto re^{i\nu\alpha}$. On this ray, its boundary values agree with those of the elliptic transformation with fixed points $0, \infty$ that maps the ray of angle α to the ray of angle $\nu\alpha$.

Just as the concept of angle doubling applies naturally to any crescent, so does the concept of angle scaling.

We will eventually need to take $\nu \in \mathbb{C}$, in which case we call this kind of map *complex angle scaling*. For a single wedge, complex angle scaling is a map from a wedge of angle α to a wedge of angle $\Re(\nu)\alpha$. □

5.2 Definition. Let ℓ be a geodesic in the hyperbolic plane $\mathbb{D}^2$ and let β, with $0 < \beta \leq \pi$ be a bending angle along ℓ. We say that a crescent C is *associated* to (ℓ, β) if $C \subset \mathbb{D}^2$, C has the same endpoints as ℓ, and the angle of C is equal to β. □

Now suppose we have a finite measured lamination with geodesics $\{\ell_i\}$ and bending angles $\{\beta_i\}$. Suppose we have associated disjoint crescents $\{C_i\}$ embedded in $\mathbb{D}^2$. We extend the meaning of an *angle doubling map* to be a continuous map $f : \mathbb{D}^2 \to \mathbb{S}^2$ whose restriction to each C_i is an angle doubling map as previously defined, and whose restriction to each component of the complement in $\mathbb{D}^2$ of the union of the crescents is a Möbius transformation. Similarly for angle scaling maps.

Recall that we have defined crescents to be open subsets of $\mathbb{S}^2$, so that each component of $\mathbb{D}^2 \setminus \bigcup C_i$ corresponds to a component of $\mathbb{D}^2 \setminus \bigcup \ell_i$, although there is in general no relationship of inclusion between corresponding components. What a pair of corresponding components have in common is that their closures have the same intersection with $\mathbb{S}^1 = \partial\mathbb{D}^2$.

The point about embedding disjoint crescents in $\mathbb{D}^2$ is that we can then prove the $K = 2$ conjecture for a region which arises from $\mathbb{D}^2$ by applying angle doubling to each of the crescents. The result is stated in the next theorem.

5.3 Theorem. *Suppose $\{\ell_i\}$ is a finite system of mutually disjoint geodesics in $\mathbb{D}^2$, possibly with some common end points, and let $\{C_i\}$ be an associated* *set of disjoint crescents in* $\mathbb{D}^2$*. Then there is a corresponding angle doubling map* $f : \mathbb{D}^2 \to \mathbb{S}^2$ *which is unique up to composition with a Möbius transformation from* $\mathbb{S}^2$ *to itself. Furthermore* f *is a 2-quasiconformal homeomorphism.*

The boundary values of f *on* $\partial\mathbb{D}^2$ *depend only on the angles of the crescents and their endpoints on* ∂D*, that is, only on the measured lamination, not on exactly how the crescents are embedded in* $\mathbb{D}^2$*. As a set, the image* $f(\mathbb{D}^2) = \Omega$ *also depends only on the the measured lamination. On the other hand,* f *itself depends on the precise embeddings of the disjoint crescents.*

Proof: We will define a homotopy $f_t : \mathbb{S}^2 \to \mathbb{S}^2$ for $t \in [0, 1]$ and f will be the restriction of f_1 to the open unit disk $\mathbb{D}^2$.

We reflect (invert) in the unit circle to obtain copies outside the unit circle of our given crescents inside the unit circle. Recall that these crescents are open subspaces. Now take the complement X in $\mathbb{S}^2$ of the union of these crescents with the two endpoints of each crescent. Each component of X is either symmetric under inversion, if it contains an arc of the unit circle, or it has a companion component obtained by inversion, if its closure meets the unit circle in only a finite number of points. A component C may have empty interior in $\mathbb{S}^2$. In this case C is equal to an open arc of the unit circle, and one of the given crescents C_i abuts the unit circle along C.

On the crescents inside the unit disk, f_t is angle scaling by $1 + t$. On the crescents outside the unit disk it is angle scaling by $1-t$. On each component of X, it is a Möbius transformation, and the two Möbius transformations for two companion components are equal.

If f_t has these properties, then so does Af_t, where A is any Möbius transformation. Therefore f_t is completely determined by specifying which Möbius transformation to apply to one component of X, or, alternatively, to one of the arcs foliating one of the crescents.

More explicitly, we proceed as follows. The map is defined on an increasing connected union of components of X and crescents. In fact we make the definitions on the closures of the crescents and components of X. We deal simulataneously with companion crescents and with companion components.

We start with S_0, the union of a component of X with its mirror image. For all t we may as well take f_t to be the identity map. Suppose f_t has been defined on S_r, the union of some component of X with its companion. Suppose that $f_t|S_r$ is equal to a certain Möbius transformation $m_{r,t}$. We are about to define f_t on the union of a crescent A inside the unit circle and its companion crescent B outside the unit circle. Suppose that the closure of A meets S_r in a circular arc A_0 and that the closure of B meets C in a circular arc B_0. Then the remainder of the boundary of A is an arc A_1 and the remainder of the boundary of B is an arc B_1. Let α be the angle of the crescents A and B. We extend f_t continuously to $A \cup B$ so that it is angle scaling by $1+t$ on A and angle scaling by $1-t$ on B. Let E_t be the elliptic transformation whose fixed points are the two common endpoints of hA_0 and hB_0, and which rotates by an angle $t\alpha$ in the circular direction hA_0, hA_1, hB_1, hB_0. The map f_t on $A_1 \cup B_1$ is then equal to $E_t \circ m_{r,t}$, and we use this Möbius transformation to define f_t on a component of X whose closure contains A_1 or B_1. We check that the same Möbius transformation is used on the component and on its companion.

So we have defined a homotopy $f_t : \mathbb{S}^2 \to \mathbb{S}^2$.

Now f_0 is the identity map of $\mathbb{S}^2$ and so f_1 has degree one. For any point $x \in \mathbb{S}^2$ which is not in the union of a finite number of circular arcs, namely the images of the crescents outside the unit circle under angle collapse and the image of the unit circle, the inverse image under f_1 of a small disk D, centred at x, is the union of disjoint disks, each mapped homeomorphically onto D by an orientation preserving homeomorphism. Since the degree is one, there is exactly one such disk in the inverse image. It follows that $f : \mathbb{D}^2 \to \mathbb{S}^2$ is injective, and it is therefore a 2-quasiconformal homeomorphism. ∎

5.4 Corollary. *We make the same hypotheses as in Theorem 5.3 and we use the notation introduced in its proof. The restriction of the angle doubling map f to any S_i is a Möbius transformation, which we denote by m_i. The image $\Omega = f(\mathbb{D}^2)$ is the union of the open disks $m_i(\mathbb{D}^2) \subset \mathbb{S}^2$, with the sole exception of the folded case (see Definition 4.4). Moreover, these are maximal disks contained in Ω and their boundaries meet the boundary of Ω in at least three points. Any maximal disk meeting the boundary of Ω in at least three points is equal to one of the $m_i(\mathbb{D}^2)$.*

Proof: If we have only one crescent, there are two cases: either the crescent angle is equal to π or it is less than π. If the crescent angle is equal to π, then we are in the folded case (see Definition 4.4). Then X has one component, namely an open arc of circle, but we can think of it as occurring twice. We get two Möbius transformations, namely m_1, the identity, and m_2, which is rotation by an angle π around the endpoints of the arc. Ω is equal to $m_1(\mathbb{D}^2) \cup m_2(D_2) \cup m_2(X)$.

In the case of a single crescent of angle $\alpha < \pi$, m_1 is again the identity and m_2 is rotation by α around the endpoints of the crescent. The corollary is then obvious.

In general, the proof is an easy induction on the number of crescents in the disk. We recall that we started by defining f_t to be the identity on S_0 of X. We have constructed $f = f_1$ such that its restriction to S_r is equal to some Möbius transformation, which we denote by m_r. Suppose S_k needs a maximum number of crescents to be crossed to reach it from S_0 while staying in $\mathbb{D}^2$. Suppose C_i is the last crescent which has to be crossed, and that S_r is adjacent to C_i on the other side from S_k. So S_r is "nearer" to S_0 than S_k in that one crescent less needs to be crossed in order to reach it.

Suppose the corollary is known if C_i, together with its companion crescent outside the unit circle, and S_k, are eliminated, and absorbed into S_r. We obtain an angle doubling map f'. Inductively, we know that the image of f' is the union of the disks $m_j(\mathbb{D}^2)$, with $m_k(\mathbb{D}^2)$ omitted.

We may assume without loss of generality that $S_r = S_0$, so that $f'|S_r$ is the identity map. When we now reinsert C_i, together with its companion crescent, and S_k, we find, by application of the case where there is only one crescent, namely C_i, that $f(\mathbb{D}^2) = \Omega$ is the union of $f'(\mathbb{D}^2)$ with $m_k(\mathbb{D}^2)$. This completes the proof of the corollary.

The disks $m_i(\mathbb{D}^2)$ have boundaries $m_i(\mathbb{S}^1)$ which are easily seen to meet the boundary of Ω in at least three points. Therefore these disks are maximal. It is easy to see, by the same type of induction used above, that any disk in Ω whose boundary meets the boundary of Ω in at least three points is one of our disks $m_i(\mathbb{D}^2)$. ∎

5.5 Remark. The angle doubling mapping is not the quasiconformal mapping with the least maximal dilatation for the same boundary values, except in the case that Ω is $\mathbb{C}$ slit along a half-infinite euclidean line. Except in this special case, the extremal map will have maximal dilatation strictly less than two. See Theorem 23.6. □

6 Finite complex earthquakes

In Section 4 we saw how the boundary of the convex hull of a closed subspace of $\mathbb{S}^2$ has a natural measured lamination. If the closed subset is the complement of a finite union of round open disks, then we get a finite measured lamination.

We now change our point of view a little. We have been regarding the bending angles as positive numbers. But we now multiply each bending angle by $i = \sqrt{-1}$. To understand why this is a reasonable convention, note that bending through an angle θ is associated with a rotation of $\mathbb{H}^3$ through an angle θ about an axis. If the axis is the vertical axis through 0 in the upper halfspace model, than such an isometry is provided by the Möbius transformation $z \mapsto ze^{i\theta}$. We also want to consider earthquakes. An earthquake along the same axis will correspond to a Möbius transformation $z \mapsto ze^d$, where d is a hyperbolic distance. This all fits in with the natural requirements that θ should normally be taken modulo 2π and that d is an arbitrary real number, the sign giving the direction of the earthquake.

The above Möbius transformations will be used frequently and we therefore provide names for them, and for associated conjugates. Let $E(u)(z) = e^u z$. We regard $E(u)$ both as a Möbius transformation and as an orientation-preserving isometry of $\mathbb{H}^3 = \mathbb{U}^3$. It fixes 0 and ∞ and is elliptic if and only if u is purely imaginary. In matrix terms we have

$$E(u) = \begin{bmatrix} e^{\frac{u}{2}} & 0 \\ 0 & e^{-\frac{u}{2}} \end{bmatrix}.$$

Given an oriented geodesic $\ell \subset \mathbb{U}^3$, and $u \in \mathbb{C}$ we will use the notation $E(\ell, u) = AE(u)A^{-1}$, where A is a Möbius transformation sending the vertical geodesic above 0 in upper half 3-space, oriented from 0 to ∞ to ℓ.

6.1 Definition. We fix an isometric embedding $f_0 : \mathbb{U}^2 \to \mathbb{U}^3$. Let $\Lambda = (\ell_1, \ldots, \ell_N)$ be a finite lamination of the hyperbolic plane (specified here as the upper halfplane) and let $z = (z_1, \ldots, z_N) \in \mathbb{C}^N$. If we cut the hyperbolic plane along the disjoint geodesics of Λ, we get $N + 1$ open flat pieces $(P_0, \ldots, P_N)$. Each P_i has boundary, which is disjoint from P_i, consisting of the union of some of the ℓ_j. We fix one flat piece, say P_0, and use the given isometric embedding $f_0 : P_0 \to \mathbb{U}^3$.

If we change f_0 by composing with an orientation-preserving isometry of $\mathbb{U}^3$, then all our constructions change by composition with the same isometry.

Moreover, it will be straightforward for the reader to check as the discussion proceeds, that, if we start with some isometric embedding $f_r : P_r \to \mathbb{U}^3$ instead of with f_0, the only eventual difference in the objects constructed is again the effect of applying a further orientation-preserving isometry of $\mathbb{U}^3$ with itself.

For each j, there is one side of ℓ_j that is nearer to P_0 and one side that is further from P_0. Using the standard orientation of $\mathbb{U}^2$, this enables us to fix an orientation for ℓ_j as follows. Fix a point $x \in \ell_j$, and let v_1 be the unit normal to ℓ_j pointing away from P_0. Let v_2 be the unit vector pointing along ℓ_j, so that (v_1, v_2) gives the usual orientation of U^2. Then v_2 orients ℓ_j.

For each r, we will define an isometric embedding $f_r : P_r \to \mathbb{U}^3$. This will be done inductively, using induction on the number of $\ell_j \in \Lambda$ separating P_r from P_0. So suppose P_k is the next flat piece for which $f_k : P_k \to \mathbb{H}^3$ has to be defined. Let $\ell_j = \overline{P_r} \cap \overline{P_k}$ where $f_r : P_r \to \mathbb{H}^3$ is already defined. Then ℓ_j lies between P_k and $P_0 \cup P_r$.

We define $f_k : P_k \to \mathbb{H}^3$ as the restriction to P_k of $E(f_r(\ell_j), z_j) \circ f_r$, where ℓ_j is oriented as described above.

We would also like to define isometric embeddings $\ell_j \to \mathbb{U}^3$—so far, isometric embeddings have been defined only on the complement in $\mathbb{H}^2$ of these geodesics. If the two flat pieces whose frontiers contain ℓ_j are P_r and P_k, then f_r and f_k coincide on ℓ_j if and only if $\Re(z_j) = 0$. In general, we use the isometric embedding $E(f_r(\ell_j), z_j/2) \circ f_r : \mathbb{U}^2 \to \mathbb{U}^3$. The restriction to ℓ_j depends only on the real part of z_j.

We denote the (usually) discontinuous earthquake map by $E_{\Lambda,z} : \mathbb{H}^2 \to \mathbb{H}^3$. □

6.2 Definition. If all the z_j are real, the complex earthquake is called an *earthquake*. This definition is due to Thurston. If the z_j are all positive, we obtain a *left earthquake*. □

For a left earthquake, if we stand on one side of a geodesic ℓ_j, then the other side moves to the left. Here we are using the fixed orientations of $\mathbb{U}^2$ and $\mathbb{U}^3$ inherited from standard vector space bases of $\mathbb{R}^2$ and $\mathbb{R}^3$.

Note that the isometric embeddings restricted to the P_r and to the ℓ_j do not fit together to gives us a continuous map $\mathbb{U}^2 \to \mathbb{U}^3$ unless all the z_j are purely imaginary. On the other hand, we always obtain a continuous map $\mathbb{S}^1 = \partial\mathbb{U}^2 \to \mathbb{S}^2 = \partial\mathbb{U}^3$ which is well-defined up to composition with a Möbius transformation acting on $\mathbb{S}^2$. We denote this map also by $E_{\Lambda,z} : \mathbb{S}^1 \to \mathbb{S}^2$.

Figure 8: show an earthquake by a geodesic in the disk, with crossing geodesics, before and after

Figure 9: earthquake: in upper half plane show a picture on each side of a vertical geodesic, and then how it changes after an earthquake

We can construct a continuous map $\Phi_{\Lambda,z} : \mathbb{U}^2 \to \mathbb{U}^3$, sending any rectifiable path in the domain to a rectifiable path of the same length, with $\Phi(P_r) = f_r(P_r)$ for each r. To define Φ, we write $z_j = x_j + iy_j$, and take the earthquake corresponding to $(x_1, \ldots, x_N)$. We regard the initial embedding $P_0 \to \mathbb{U}^3$ as an inclusion. Let the inductively defined isometries of the flat pieces be denoted by $f_r' : P_r \to \mathbb{U}^3$. It is immediate by induction that, for each r, The earthquake moves the finite lamination Λ to a new lamination Λ', by defining $\ell_j' = f_r'(\ell_j)$, where $\ell_j \subset P_r$. This definition of ℓ_j' is independent of the two possible choices of r. We then apply the purely imaginary earthquake (that is, a pure bending) to $(\mathbb{U}^2, \Lambda', iy_1, \ldots, iy_N)$ to obtain our map Φ. $\Phi_{\Lambda,z} : \mathbb{U}^2 \to \mathbb{U}^3$ is determined up to composition with an isometry of $\mathbb{U}^3$. Once the map is given on one of the P_r, it is then completely determined.

We have the following lemma by the same induction that is used in the construction.

6.3 Lemma. *If Λ is fixed, we have a continuous map $\Phi : \mathbb{C}^N \times \mathbb{U}^2 \to \mathbb{U}^3$.*

Figure 10: how a pattern of circles changes under an earthquake

7 The $K = 2$ theorem for finite laminations

Suppose we have a finite measured lamination with geodesics $\{\ell_i\}$ and bending angles $\{\beta_i\}$. Suppose we have associated disjoint crescents $\{C_i\}$ embedded in $\mathbb{D}^2$. We have seen in Theorem 5.3 that the angle doubling map gives us a 2-quasiconformal homeomorphism $f : \mathbb{D}^2 \to f(\mathbb{D}^2) = \Omega \subset \mathbb{S}^2$. Let $\iota : \mathrm{Dome}(\Omega) \to \mathbb{D}^2$ be an isometry (as given by Theorem 2.1).

7.1 Theorem. *The isometry $\iota^{-1} : \mathbb{D}^2 \to \mathrm{Dome}(\Omega)$ extends to a continuous map $\mathbb{S}^1 \to \mathbb{S}^2$ which is equal on $\mathbb{S}^1$ to the extension of the angle doubling map f, after composition with a suitable Möbius transformation of $\mathbb{S}^2$.*

Proof: We will start by briefly recalling the inductive construction of ι and of f. However, instead of ι, we will use its inverse $\Phi : \mathbb{H}^2 \to \mathbb{H}^3$, defined just before Lemma 6.3. In fact, Φ was defined using the upper halfplane $\mathbb{U}^2$ and upper halfspace models $\mathbb{U}^3$. We can transfer the definition to $\mathbb{H}^2$ and $\mathbb{H}^3$ by using fixed isometries $\mathbb{U}^n \to \mathbb{H}^n$. (In fact, all that is really needed is to fix orientations of $\mathbb{H}^2$ and $\mathbb{H}^3$. Our fixed isometries accomplish this.)

Since we want to relate angle doubling, which was defined using the unit disk, with $\Phi : \mathbb{H}^2 \to \mathbb{H}^3$, we will take $\mathbb{H}^3 = \mathbb{D}^3$ as before, and we will take $\mathbb{H}^2 \subset \mathbb{H}^3$ to be the equatorial disk, denoted by $\mathbb{D}^2$. We orient $\mathbb{D}^2$, $\mathbb{D}^3$, $\mathbb{H}^2$ and $\mathbb{H}^3$ using the natural orientations which come from $\mathbb{R}^2$ and $\mathbb{R}^3$ by fixing taking the standard bases of the vector spaces in their standard order.

The lamination $\Lambda = (\ell_1, \dots, \ell_N)$ lies in $\mathbb{D}^2$ and cuts $\mathbb{D}^2$ up into flat pieces $\{S_0, \dots, S_N\}$. (The two sets of indices are related only by their cardinality). We have N crescents $\{C_1, \dots, C_N\}$ in the lower hemisphere L and their companion crescents in the upper hemisphere U. (We can identify L with $\mathbb{D}^2$ and U with the complement of $\mathbb{D}^2$ by using stereographic projection from the North Pole. But, to avoid confusion, we won't do this.) C_i and ℓ_i have the same endpoints on $\mathbb{S}^1$. Take the union of the open crescents with the endpoints of the crescents in the closed lower hemisphere $\bar{L}$. Let Y be the complement in $\bar{L}$ of this union. Each component S'_i of Y corresponds to a unique flat piece S_i—the closures of S'_i and S_i have the same intersection with $\mathbb{S}^1$.

By Corollary 5.4, Ω is a finite union of open disks in $\mathbb{S}^2$, each of the form $m_r(\mathbb{D}^2)$, where m_r is a Möbius transformation constructed inductively.

The map $\Phi : \mathbb{H}^2 \to \mathbb{H}^3$ preserves the lengths of rectifiable paths in the domain, and is therefore equal to the inverse of ι, provided we choose suitable normalizations of these maps by composing with Möbius transformations.

Both Φ and f are constructed by inductive processes. We first fix a flat piece $S_0 \subset \mathbb{D}^2$. We map this into $\mathbb{D}^3$ by the identity. Correspondingly, we map S_0' into $\mathbb{S}^2$ by the identity. It is then straightforward to see inductively that precisely the same Möbius transformations are used to place each S_i and S_i' in final position. (The only question that arises is whether the orientations are consistently chosen, and we take care to do so.) ■

The following corollary is then immediate.

7.2 Corollary. *Using the hypotheses and notation of Theorem 7.1, there is a choice of isometry* $\iota : \mathrm{Dome}(\Omega) \to \mathbb{D}^2$, *such that the 2-quasiconformal map* $f \circ \iota : \mathrm{Dome}(\Omega) \to \Omega$ *extends to the identity on the intersection of the closures of domain and range.*

Figure 11: A ladder with the names of the regions and geodesics

Figure 12: three geodesics with points in common, not the same point

8 The ortho-condition

Recall that, in the hypotheses of Theorem 5.3, we were given a finite positive measured lamination satisifying a certain condition, namely that one could find a disjoint collection of crescents in the unit disk associated to this lamination. In this section, we show how to transform this condition into another equivalent condition, which we call the ortho-condition. We will show in § 8.11 how it is possible to check the ortho-condition using practical computations.

We need to concentrate our attention on certain special kinds of finite lamination.

8.1 Definition. We consider a measured lamination in the hyperbolic plane, $\Lambda = \{(\ell_i, \alpha_i)\}$ a finite measured lamination consisting of disjoint geodesics $\{\ell_i\}$ (except that they might meet at infinity) We will suppose that the leaves lie in a particular type of configuration, where they are "lined up". More precisely we assume that there is a geodesic σ which crosses each leaf of the lamination transversely. We call such a measured lamination a *finite ladder*. □

Let $(\ell_1, \alpha_1), \dots, (\ell_N, \alpha_N)$ be a finite ladder with bending angles α_i and with a transverse geodesic σ. We fix an orientation of σ and choose the indices for the lamination so that the leaves of the lamination are denoted $\ell_1, \dots, \ell_N$ in order along σ.

Let d denote the distance in the hyperbolic plane. We denote $\sigma \cap \ell_i$ by p_i, and define $\rho(p_i, p_{i+1}) = d_i = d(\ell_i, \ell_{i+1})$. The geodesics ℓ_i and ℓ_{i+1} meet at infinity if and only if $\rho(p_i, p_{i+1}) = 0$. If $i < j$, we define $\rho(\ell_i, \ell_j) = d_i + \dots + d_{j-1}$. This obviously defines a pseudometric on the finite set of leaves. We define $x_1, \dots, x_r$ to be the points of the quotient metric space.

Note that the leaves forming x_s do not necessarily all have a common endpoint. For example, ℓ_1 and ℓ_2 could have an endpoint in common, and ℓ_2 and ℓ_3 could have an endpoint in common, while ℓ_1 need have no endpoint in common with ℓ_3. The equivalence class x_1 will contain all three of these leaves.

8.2 Definition. We define the ortho-curve associated with a given finite ladder as follows. If $x_s = \left\{\ell_{i(s)}, \ldots, \ell_{i(s+1)-1}\right\}$, we define the new bending angle at x_s to be $\beta_s = \alpha_{i(s)} + \ldots + \alpha_{i(s+1)-1}$. This curve is obtained by drawing geodesic segments $\sigma_0, \ldots, \sigma_r$, where σ_0 is infinite to the left and finite to the right, σ_r is finite to the left and infinite to the right, and, for $0 < s < r$, σ_i has length $\rho(\ell_{i(s)}, \ell_{i(s+1)})$. The bending angle as we go from σ_{s-1} to σ_s is β_s. □

The segments σ_s are not subintervals of the transverse geodesic σ. However, an orientation of σ does specify whether we are to use the order $(\sigma_0, \ldots, \sigma_r)$ or the reverse order. We fix an orientation of $\mathbb{H}^2$. As we proceed along σ we always bend in the counter-clockwise direction. Once the orientations of $\mathbb{H}^2$ and σ are fixed, the ortho-curve is determined up to orientation preserving isometry. Changing either orientation changes the ortho-curve by an orientation reversing isometry.

8.3 Definition. We say that a given finite ladder and the associated ortho-curve satisfy the *ortho-condition* if the following conditions hold.

8.3.1. Each bending angle for the ortho-curve satisfies $0 < \alpha_i \leq \pi$. If $\alpha_i = \pi$, then $i = k = 1$, and the ortho-curve is a geodesic ray.

8.3.2. The ortho-curve gives an embedding of $\mathbb{R}$ in $\mathbb{H}^2$

In the special case when $i = k = 1$ and $\alpha_1 = \pi$, one traverses the ortho-curve by going from infinity to the finite endpoint along a geodesic ray, and then retraces the path in the reverse direction. This special case corresponds to a folded surface (see Definition 4.4).

By, for example, identifying to a point the boundary circle of $\mathbb{D}^2$ and applying the Jordan curve theorem, we see that the ortho-curve separates $\mathbb{H}^2$ into two components. By Lemma 4.6 the above conditions imply that one of these two components is convex, and its pointset boundary in $\mathbb{H}^2$ is the ortho-curve.

It is possible that σ_0 and σ_r, defined above, have a common point at infinity.

8.3.3. There is a stronger condition, which we call the *strong ortho-condition*. In this case σ_0 and σ_r do not meet at infinity and the ortho-curve has closure in $\overline{\mathbb{H}^2} = \mathbb{H}^2 \cup \partial\mathbb{H}^2$ which is a non-trivial arc in $\partial\mathbb{H}^2$. □

Figure 13: The ortho-condition

For a finite measured lamination, it is easy to compute whether the ortho-condition is satisfied or not.

8.4 Definition. Conversely, given a piecewise geodesic with a finite number of bends, which bounds a convex region of $\mathbb{H}^2$, we can define a finite ladder with this piecewise geodesic as its ortho-curve. We embed the hyperbolic plane as some totally geodesic plane in $\mathbb{H}^3$. We then project $\mathbb{H}^3$ orthogonally onto the hyperbolic plane containing the ortho-curve. The inverse image of the piecewise geodesic under this projection separates $\mathbb{H}^3$ into two components, one of which, denoted by C, is convex. The boundary of C is isometric to $\mathbb{H}^2$ and the associated bending lamination gives us a finite measured lamination of $\mathbb{H}^2$, called an *ortho-ladder*. □

8.5 Remark. The ortho-condition is satisfied for a ladder if the angle sum satisfies $\sum_i \alpha_i \leq \pi$, irrespective of the distances between the sucessive leaves. □

We will prove the following theorem.

8.6 Theorem. *Let S be a finite ladder (see Definition 8.1). Then the ortho-condition is true if and only if disjoint crescents in $\mathbb{D}^2$ can be associated to the leaves of the measured lamination, with crescent angles equal to the transverse measure of the corresponding leaf. The strong ortho-condition is false if and only if the embedding of crescents is maximal, in the sense that, if we enlarge any of the crescents then it can no longer be embedded. This is the same as saying that as we proceed along the ladder, (the closures of) adjacent crescents always touch. Another equivalent condition is that the strong ortho-condition fails if and only if the above embedding of disjoint crescents is unique.*

We will build up to this result slowly.

8.7 Inserting crescents for small numbers of bending lines. First we deal with the case of a single bending line ℓ. The ways of associating a crescent of angle α with ℓ can be parametrized by $\pi - \alpha$, because the crescent can be rotated around the endpoints of ℓ, using elliptic transformations with

Figure 14: ℓ, c, xyp and angles and σ

Figure 15: the ortho-curve when the strong ortho-condition fails but the ortho-condition is true

these fixed points. By the definition of "crescent associated", we must not rotate the crescent out of $\mathbb{D}^2$. If $\alpha = \pi$, then there is exactly one associated crescent. Clearly Theorem 8.6 is true if there is only one bending line.

8.8 A geodesic and a constant distance curve. Let ℓ be a geodesic in $\mathbb{D}^2$ with a fixed normal direction. Let c be the constant distance curve, a distance d from ℓ in the normal direction (d may be positive or negative). If ℓ and c are extended to the boundary of $\mathbb{D}^2$, then they have the same endpoints p_1 and p_2 at infinity, and meet at an angle β ($|\beta| < \pi/2$). We have $\beta < 0$ if and only if $d < 0$, and $\sin(\beta) = \tanh(d)$. This is a standard formula in hyperbolic geometry. Let σ be a geodesic orthogonal to ℓ, meeting ℓ at x and c at y. Then xyp_1 is a right-angled hyperbolic triangle, with one ideal vertex. The side opposite the ideal vertex has length $|d|$ and the vertex y has interior angle $\pi/2 - |\beta|$. To see this, note that the angle between c and the geodesic yp_1 is the same as the angle between c and ℓ, since yp and ℓ are tangent at infinity. (See Figure 14.)

8.9 Lemma. *Theorem 8.6 is true if there are exactly two bending lines, and they have no common endpoints.*

Proof: Suppose first that the ortho-condition is satisfied, but not the strong ortho-condition, and let the bending angles be α'_1 and α'_2 at ℓ_1 and ℓ_2 respectively. Then, for $i = 1, 2$, $0 < \alpha'_i < \pi$ and $\alpha'_1 + \alpha'_2 > \pi$. We assume without loss of generality that $\alpha'_1 > \pi/2$. The ortho-curve has vertices x_+, x_1, x_2, x_+, where $d(x_1, x_2) = d(\ell_1, \ell_2) = d$.

Let C_1 be the crescent associated with ℓ_1 of angle $\alpha_1 = \alpha'_1$, pushed over as far from ℓ_2 as possible. That is, one of the circular arcs bounding C_1 is contained in $\mathbb{S}^1 = \partial\mathbb{D}^2 = \partial\mathbb{D}^2$. We denote by c_1 the other boundary arc of C_1, hyperbolically a constant distance curve from ℓ_1.

Let C_2 be the maximal crescent associated to ℓ_2 and disjoint from C_1. (Our crescents are always open, unless we specifically say they are closed.) Maximality means that one boundary arc of C_2 is contained in $\mathbb{S}^1 = \partial\mathbb{D}^2 =$

$\partial\mathbb{D}^2$, and the other, denoted by c_2, is tangent to c_1. Then c_2 is, hyperbolically, a constant distance curve from ℓ_2. Let the angle of C_2 be α_2. We must show that $\alpha_2' = \alpha_2$.

The curves c_1 and c_2 meet at a unique point $y \in \mathbb{D}^2$. Let d_1 be the signed distance from ℓ_1 to c_1 (with respect to the normal direction towards ℓ_2). Let d_2 be the signed distance from ℓ_2 to c_2 (with respect to the normal direction towards ℓ_1). Let σ be the geodesic which is perpendicular to the ℓ_i, meeting ℓ_i at x_i. Then σ meets c_1 and c_2 perpendicularly at y.

Let p_i be an endpoint at infinity of ℓ_i. By § 8.8, the triangles $x_i y p_i$ have interior angles $\min\{\alpha_i, \pi - \alpha_i\}$, The first alternative corresponds to $d_i \leq 0$ and $\alpha_i \geq \pi/2$ and the second to $d_i \geq 0$ and $\alpha_i \leq \pi/2$. Under our assumptions, $d_1 > 0$.

We take copies of the two right-angled triangles $\Delta_1 = x_1 y_1 p_1 = x_1 y p_1$ and $\Delta_2 = x_2 y_2 p_2 = x_2 y p_2$ and identify $x_1 p_1$ with $x_2 p_2$, calling the result xp. If $d_2 > 0$, we regard Δ_1 and Δ_2 as being on opposite sides of xp and take their union. If $d_2 \leq 0$, we regard Δ_1 and Δ_2 as being on the same side sides of xp and take their difference.

We obtain a hyperbolic triangle $y_1 y_2 p$ with one ideal vertex p, and $d(y_1, y_2) = d_1 + d_2 = d$. The bending angles are α_1, and α_2, no matter what the sign is of d_2. Since the ortho-curve has the same initial point and endpoint, the triangle bounded by the ortho-curve is isometric to the triangle $y_1 y_2 p$. It follows that $\alpha_2 = \alpha_2'$ as required.

Now suppose that the strong ortho-condition is satisfied. To see that we can embed disjoint crescents C_1 and C_2 associated with ℓ_1 and ℓ_2 with angles α_1 and α_2, note that we can increase α_2 without changing α_1, until the first point at which the strong ortho-condition fails. It follows that the embedding of the crescents is not unique.

Finally suppose that we can embed disjoint crescents C_1 and C_2 associated with ℓ_1 and ℓ_2, with angles α_1 and α_2. We want to prove that the ortho-condition is satisfied by the corresponding measured lamination. We increase the angles α_1 and α_2 until they crescents are maximal. Then the same argument as that just given shows that the ortho-condition is satisfied, but not the strong ortho-condition. It follows that the ortho-condition was satisfied with the original α_i, and that the strong ortho-condition was satisfied unless the crescents were already maximal. ∎

8.10 Lemma. *Theorem 8.6 is true if there are exactly two bending lines, and they have a common endpoint.*

Figure 16: Arcs at fixed distance from two geodesics

Figure 17: The associated ortho-curve

Proof: We suppose that the ortho-condition is satisfied, but not the strong ortho-condition, and let the bending angles be α_1 and α_2 at ℓ_1 and ℓ_2 respectively. We work in the upper halfplane, with ℓ_1 and ℓ_2 vertical. Then, for $i = 1, 2$, $0 < \alpha_i < \pi$ and $\alpha_1 + \alpha_2 = \pi$. We may suppose without loss of generality that $\alpha_1 \geq \pi/2$.

The associated crescents are wedges in $\mathbb{C}$. Let C_1 be the wedge associated with ℓ_1 of angle α_1, pushed over as far from ℓ_2 as possible. That is, one of the edges bounding C_1 is contained in $\mathbb{R}$. We may take this edge to be $(-\infty, 0]$ We denote by c_1 the other boundary edge of C_1, hyperbolically a constant distance curve from ℓ_1. Then c_1 is a ray in the upper halfplane, starting at 0.

Let C_2 be the maximal wedge associated to ℓ_2, which is disjoint from C_1. Maximality means that one boundary edge of C_2 is contained in $(0, \infty)$, and the other, denoted by c_2, is a euclidean ray parallel to c_1. Then c_2 is a constant (euclidean) distance curve from ℓ_2. Then the angle of C_2 is obviously $\alpha_2 = \pi - \alpha_1$. It is clear that the embeddings are unique.

Now suppose that the strong ortho-condition is satisfied. Then $\alpha_1 + \alpha_2 < \pi$. It follows that we can embed disjoint crescents C_1 and C_2 associated with ℓ_1 and ℓ_2 with angles α_1 and α_2, and that the embeddings are not unique.

Finally suppose that we can embed disjoint wedges C_1 and C_2 in the upper halfplane associated with ℓ_1 and ℓ_2 respectively, and with angles α_1 and α_2. Then $\alpha_1 + \alpha_2 \leq \pi$. It follows that the ortho-condition is satisfied by the corresponding measured lamination. ∎

Lemma 8.10 and Lemma 8.9 give the first step in the proof by induction of Theorem 8.6.

Proof (Theorem 8.6): Let $k > 2$ and let $\{\ell_1, \ldots, \ell_N\}$ be the sequence of geodesics forming the ladder, numbered consecutively along an oriented common transversal geodesic. Let the bending angle at ℓ_i be α_i. Suppose that the ortho-condition is satisfied, but not the strong ortho-condition, so that, for each i, $0 < \alpha_i < \pi$. Let the corresponding ortho-curve have vertices $x_+, x_1, \ldots, x_r, x_+$, where $x_+ \in \mathbb{S}^1 = \partial\mathbb{D}^2$.

Figure 18: The region Ω_r

We proceed by induction; we have already proved the statement for two geodesics. Assume the statement is true for $k-1$ geodesics.

If we erase ℓ_1 from the lamination, then α_2 can be increased by a maximum angle $\beta > 0$, while retaining the ortho-condition. By induction, we have unique disjoint embeddings of crescents $C'_2, C_3, \dots, C_N$ associated to $\ell_2, \ell_3, \dots, \ell_N$, with angles $\alpha_2 + \beta, \alpha_3, \dots, \alpha_N$, and this embedding is maximal. Maximality implies that one of the boundary arcs b of C'_2 lies in $\mathbb{S}^1$. Let c_2 be the arc of circle with the same endpoints as ℓ_2, lying inside C'_2 and at an angle β to b. Then c_2 cuts C'_2 into two crescents associated to ℓ_2, one of angle β and the other, which we call C_2, of angle α_2.

Let C''_2 be the crescent associated to ℓ_2 bounded by c_2 and by $\mathbb{S}^1 \setminus b$. The angle of C''_2 is $\pi - \beta$. It follows from the case $k = 2$ that we can insert a crescent C_1 associated to ℓ_1 of angle α_1 which is disjoint from C''_2. Then C_1 touches C''_2 and therefore c_2 and C_2 at the same single point (which will be at infinity if and only if ℓ_1 and ℓ_2 have a common endpoint). The crescents $\{C_1, \dots, C_N\}$ form a maximal set and the embedding is unique.

If the strong ortho-condition is satisfied, then we can increase α_1, while keeping α_i fixed for $i > 1$, until the strong ortho-condition first fails. So we see that, with the original value of α_1, we can insert crescents associated to the ℓ_i with angles α_i.

Finally, suppose we can insert disjoint crescents associated to the ℓ_i, with angles α_i. We want to prove the ortho-condition. First $0 < \alpha_i \leq \pi$, which is one of the conditions we need to verify. If, for some i, $\alpha_i = \pi$, then there is room for the one crescent C_i, so that $k = 1$. So we suppose this doesn't happen. If adjacent crescents touch, then a similar argument to that used in the reverse direction shows that the ortho-condition is satisfied, but not the strong ortho-condition. If they don't all touch, then the angles can be increased until they do. Since the ortho-condition is then satisfied, the strong ortho-condition must have been satisfied originally. ∎

Figure 19: The lamination corresponding to dome(Ω)

8.11 Example illustrating limitations of angle doubling. Following a challenge of Bishop [4], we will analyze the situation that Ω_r is the union of the nine round disks shown in Figure 18. We are unable to meet Bishop's challenge in precisely the form in which it was presented, but our computations do reveal interesting information about the situation.

The central disk will have radius one and the others radius $r < 1$. The disks are orthogonal to each other at their points of intersection.

We ask, for what values of r is Ω_r the result of angle doubling induced by a measured lamination in $\mathbb{D}^2$? Here the lamination has eight leaves with the angle $\pi/2$ assigned to each. Figure 19 illustrates the symmetric arrangement the leaves satisfy. Take one pair of leaves at hyperbolic distance d apart. The other three pairs come from $\frac{\pi}{2}$ rotations of $\mathbb{D}^2$. The hyperbolic distance between an opposite pair of inner leaves is d^* as indicated in Figure 19 and the distance between an adjacent pair of inner leaves is c.

By taking a plane in $\mathbb{H}^3$ orthogonal to the dome over Ω_r and to a set of four bending lines we are able to compute d, d^* as functions of r as displayed below. The distance c can be computed by the trigonometry of a right angled pentagon.

$$d = \log \frac{\sqrt{2}+1}{\sqrt{1+r^2}-r}, \quad d^* = 2\log \frac{1+\sqrt{1+r^2}}{r}, \quad \cosh c = \frac{1}{r^2}.$$

We have to find out for what values of r we can insert crescents of angle $\frac{\pi}{2}$ for each leaf. This can clearly be done for the four outer leaves which face an arc of $\partial\mathbb{D}^2$. Examine one of the four inner leaves ℓ. The outer edge of the maximal crescent that can be inserted for ℓ will be tangent to the adjacent inner leaf, while by symmetry the inner edge will be of distance $\frac{c}{2}$ from ℓ. Let ϕ_1, ϕ_2 denote the angles the two edges of the crescent make with ℓ at an a vertex. We find that

$$\phi_1 + \phi_2 = \arctan[\sqrt{1+r^2} + r\sqrt{2}] + \arctan[\frac{\sqrt{1-r^2}}{r\sqrt{2}}].$$

Figure 20: Mathematica graph of $\phi_1 + \phi_2$ as a function of r

The graph of $\phi_1 + \phi_2$ as a function of r is shown in Figure 20. It takes the value $\frac{\pi}{2}$ when $r = r_0 = .875972\ldots$. Therefore when $r_0 < r < 1$, Ω_r cannot be obtained from $\mathbb{D}^2$ by an angle doubling map. For the values of $0 < r \leq r_0$, there is a 2-quasiconformal mapping $\mathrm{dome}(\Omega_r) \to \Omega_r$.

For all values $0 < r < 1$, the map $\Psi : \mathrm{Dome}(\Omega_r) \to \Omega_r$ has boundary dilatation $H(\Psi) < 2$.

9 The crescent theorem

Suppose $\{\ell_i, \alpha_i\}$ is a finite measured lamination with geodesics in $\mathbb{H}^2$, which are disjoint, except possibly at the endpoints, with positive bending angles α_i. For each geodesic σ which is not equal to any of the ℓ_i, we define a ladder Λ_σ, using only those ℓ_i which meet σ inside $\mathbb{H}^2$.

9.1 Definition. We say that this measured lamination satisfies the ortho-condition if each Λ_σ satisfies the conditions of Definition 8.3. □

9.2 Theorem. *Suppose (Λ, μ) is a finite measured lamination of $\mathbb{H}^2$ satisfying Definition 9.1. Then for each geodesic ℓ_i of Λ a crescent of angle α_i associated with ℓ_i can be inserted in $\mathbb{D}^2$, so that the resulting crescents (which are open subsets) are mutually disjoint. Conversely if such disjoint crescents can be so inserted for a configuration of finite measured lamination, then it satisfies the ortho-condition.*

It follows from Corollary 7.2 that the ortho-condition implies the truth of the $K = 2$ conjecture.

Proof: The converse statement follows immediately from Theorem 8.6. So we need only show that the crescents can be inserted if the ortho-condition is satisfied.

Given a finite measured lamination with geodesics $\{\ell_i\}$ and bending angles $\{\alpha_i\}$ and a point p_0 which does not lie on any of the geodesics, we will describe a systematic procedure to insert the crescents. We have a partial order on the geodesics, where $\ell_i < \ell_j$ means that every geodesic from p_0 to a point of $\mathbb{H}^2 \cap \ell_i$ crosses ℓ_j. An equivalent condition is that for some point $x \in \ell_i$, the geodesic $p_0 x$ crosses ℓ_j. We will also write $i < j$.

For each ℓ_i, we have a normal direction, defined as the direction pointing towards p_0. This normal direction defines a signed distance from ℓ_i, with p_0 having positive distance. The signed distance has a constant value on each of the two edges of a crescent associated to ℓ_i. We can therefore speak of the *greater* and *lesser* edges of the crescent. (But note that the value of the signed distance may be positive on each edge, or negative on each edge, or positive on the greater edge and negative on the lesser edge. Moreover, the lesser edge may be nearer to p_0 than the greater edge.)

We will insert disjoint crescents one by one, always working on a crescent associated to a geodesic which has not yet been dealt with and is minimal with respect to the partial order. Each crescent will be placed in position so that its lesser edge touches the greater edge of at least one crescent which has already been placed in position. The point of contact will be at infinity if the two geodesics concerned have a common endpoint; otherwise it will be at a point of $\mathbb{H}^2$.

We claim that there is always room to place each new crescent, and in fact its position is uniquely determined by p_0. Indeed, we need only specify in which flat piece p_0 lies in order to make the positioning of the crescents unique under the above rules. To see this, we need to analyze what kind of trouble we might run into, in placing the crescent of angle α_i associated with the geodesic ℓ_i. The crescents C_j associated with any $\ell_j < \ell_i$ are already in position. Let the signed distance from ℓ_i be denoted by h_i. We choose $r < i$ and $M \in (-\infty, \infty]$ such that

$$M = \sup_{j<i} h_i | C_j = \sup h_i | C_r.$$

There may be more than one such r.

After changing the names of the indices, we have a finite sequence

$$\ell_1 < \ell_2 < \ldots < \ell_{s-1} < \ell_s = \ell_r < \ell_i,$$

such that the lesser edge of each ℓ_j touches the greater edge of ℓ_{j-1} and such that the lesser edge of ℓ_1 lies in $\mathbb{S}^1 = \partial\mathbb{H}^2$. This sequence forms a ladder to which we can apply Theorem 8.6. It follows that there is a unique way to place the crescent C_i in position in $\mathbb{H}^2$ so that its lesser edge touches the greater edge of C_r.

It is now conceivable that C_i intersects some C_j which is already in position. (Recall that our crescents are open subsets of $\mathbb{D}^2$.) We will show that such an intersection is not possible. First note that i cannot be comparable with j in the partial order. This means that the greater edge of C_i must cross the greater edge of C_j. Moreover ℓ_j is on the positive side of ℓ_i and ℓ_i is on the positive side of ℓ_j. As in the preceding paragraph, we form a ladder starting from ℓ_i and going away from p_0. We similarly form a ladder starting from ℓ_j and going away from p_0. In each of these two cases, there is a finite sequence of crescents, starting with a crescent having one edge in $\mathbb{S}^1$, and with the greater edge of each crescent touching the lesser crescent of the next. The union of these two ladders is again a ladder. We apply Theorem 8.6 to this combined ladder, in order to show that C_i will not meet C_j. This contradiction proves the required result. ∎

10 Some equivalent conditions

We have already shown (Theorem 9.2) that a necessary and sufficient condition to be able to embed disjoint crescents in $\mathbb{D}^2$ is that the associated measured lamination should satisfy the ortho-condition (Definition 9.1). In this section we will show that these two conditions are equivalent to another condition, which we call *quake-resistance*.

According to Corollary 7.2, the image of the angle doubling map is a simply connected region $\Omega \subset \mathbb{S}^2$.

10.1 Lemma. *If we have a finite set of disjoint embedded crescents, there are two ways to get a measured lamination in $\mathbb{H}^2$. The first way is to use the geodesics which join the two ends of a crescent, and assign a measure equal to the angle of the crescent. The second way is to take the hyperbolic convex hull of the complement of Ω in $\mathbb{S}^2$ and then take the bending lamination of its boundary. These two measured laminations coincide.*

Proof: This result follows from Corollary 5.4. A maximal disk in Ω whose boundary meets the boundary of Ω in at least three points gives rise to a support plane for the convex hull of the complement of Ω. This support plane meets the convex hull boundary in a flat piece. The angle between adjacent flat pieces is equal to the angle between the corresponding maximal circles. This is clearly the same as the angle of the crescent which separates the corresponding regions of the unit disk. ∎

Let $\Omega \subset \mathbb{S}^2$ be an open simply connected subset whose complement contains more than one point, and let S be the boundary of the convex hull of its complement. We know from Theorem 2.1 that S is isometric (via the path metric in S) to the hyperbolic plane. In Definition 6.2 an earthquake is defined on the hyperbolic plane. Using the isometry, we obtain an earthquake defined on S. Since S is bent, the earthquake will move S to a new surface in $\mathbb{H}^3$, whereas the original earthquake, defined on the hyperbolic plane itself, did not change the underlying pointset.

An earthquake does not change the isometry class of S, only the way flat pieces of S is put together, provided that the earthquake does not cause S to intersect itself inside $\mathbb{H}^3$. We can understand this by restricting our attention to an earthquake along a single geodesic. The statement then amounts to the fact that if we take a geodesic in the hyperbolic plane, divide the hyperbolic plane into two halfplanes along the geodesic, slide one of the

Figure 21: An example showing that many embedded surfaces are not quake-resistant.

Figure 22: Is it possible to make a reasonable drawing in $\mathbb{U}^3$ of a quake-resistant surface?

halfplanes along the geodesic and reglue, we once again obtain an isometric copy of the hyperbolic plane.

10.2 Definition. We consider surfaces S as above, with the property that all earthquakes along its finite set of bending lines for S continue to give a surface which is either embedded in $\mathbb{H}^3$ or is a folded surface, as in Definition 4.4. We call such an S *quake-resistant*, because its basic structure is unchanged by an earthquake (along the bending lines). A folded surface is sent to the same folded surface by an earthquake, but by a discontinuous map. It may at first seem unnatural to allow folded surfaces as well as embedded surfaces. However, we would like the class of quake-resistant surfaces to be closed under certain operations, and these operations do sometimes produce a folded surface from an initial embedded surface. □

We will see that, while there are many quake-resistant surfaces, there are also many convex hull boundary surfaces which are not quake-resistant (see Figure 21).

Let Ω be an open simply connected subset of $\mathbb{S}^2$. Let $S = \mathrm{Dome}(\Omega)$. The following lemma is straightforward and the proof is left to the reader.

10.3 Lemma. *The following conditions are equivalent:*

10.3.1. Ω is a finite union of disks.

10.3.2. Ω is a finite union of maximal disks.

10.3.3. S is a finitely bent surface.

10.4 Defining Φ_λ. In Lemma 6.3, we considered the continuous map $\Phi : \mathbb{C}^N \times \mathbb{U}^2 \to \mathbb{U}^3$, which depends on the fixed geodesic lamination Λ of the hyperbolic plane with exactly N leaves. If we also fix the bending angles, that is, the imaginary parts of the components of $\mathbb{C}^N$ and let $\lambda = (\lambda_1, \ldots, \lambda_N) \in \mathbb{R}^N$ be the real parts, we obtain a continuous mapping $\Phi_\lambda : \mathbb{H}^2 \to \mathbb{H}^3$. This map depends continuously on $\lambda \in \mathbb{R}^N$. The map can be defined by applying the earthquake determined by λ to transform the lamination Λ to a new lamination Λ'. We define $\Phi_\lambda : \mathbb{H}^2 \to \mathbb{H}^3$ to be the result of bending $\mathbb{H}^2$ into $\mathbb{H}^3$ along the geodesics of Λ'.

Each complementary component of the lamination is mapped isometrically to a flat piece in $\mathbb{H}^3$. The angle between adjacent flat pieces in $\mathbb{H}^3$ is given by the corresponding fixed bending angle, We will refer to the image S_λ of Φ_λ as the *pleated surface* associated with the given data. It is uniquely determined up to Möbius equivalence.

We also denote by $\lambda \in \mathbb{R}^N$ the earthquake which, on the i-th bending line quakes a distance λ_i to the left. (If $\lambda_i < 0$, this means that the quake is actually to the right.) Here the idea of "right" or "left" depends on fixing an orientation of the hyperbolic plane, rather than an orientation of the geodesic. Then λ determines the earthquake, except for composition with a Möbius transformation. If we fix attention on a particular flat piece, and insist that this piece be fixed by the earthquake, then the image $S_\lambda = \Phi_\lambda(\mathbb{H}^2)$ of S is completely determined by S and λ.

10.5 Lemma. *Suppose that Φ_λ is injective. Then S_λ separates $\mathbb{H}^3$ into two components and is locally convex on one side. We call this side Y_λ. Let $X_\lambda \subset \mathbb{S}^2$ be the set of limit points of Y_λ. Then $\Omega_\lambda = \mathbb{S}^2 \setminus Y_\lambda$ is a simply connected open subset of $\mathbb{S}^2$. We have $\mathcal{CH}(X_\lambda) = Y_\lambda$ and $\partial\mathcal{CH}(X_\lambda) = S_\lambda$.*

Proof: We claim that Φ_λ is a proper map. To see this, note that if the hyperbolic plane is divided into two closed halfplanes by a geodesic, then a continuous map which is proper on each closed halfplane is proper. Induction on N, the number of bending lines, then proves the claim. This shows that, topologically, S_λ is a plane separating $\mathbb{H}^3$ into two components. We can now apply Lemma 4.6 to conclude that Y_λ is convex.

It follows that $Y_\lambda \supset \mathcal{CH}(X_\lambda)$. To prove equality, let $y \in S_\lambda$. Then y is contained in some support plane P for S_λ. (In this situation we can take P to be the plane containing some flat piece F_i of S_λ. But we will later on want to generalize the statement to situations where this may not be true, and we would like the proof to continue to be valid in the more general situation.) Now ∂P contains at least two points of $\overline{S_\lambda} \cap \mathbb{S}^2 \subset X_\lambda$ and $y \in \mathcal{CH}(\partial P \cap X_\lambda)$.

It follows that there is a a maximal disk $D_P \subset \Omega_\lambda$ with $\partial D_P = \partial P$. Therefore P is a support plane for $\mathcal{CH}(X_\lambda)$. It follows that $y \in \partial\mathcal{CH}(X_\lambda)$. This proves that $Y_\lambda = \mathcal{CH}(X_\lambda)$.

We obtain a homeomorphism from S_λ to Ω_λ by projecting from a point in the interior of Y_λ. ■

Let $(\ell_1, \ldots, \ell_N)$ be the bending lines of S and let $(P_0, \ldots, P_N)$ be the corresponding flat pieces. Note that P_i is not necessarily related in any way to ℓ_i. Only the numbers of geodesics and flat pieces is related. By Theorem 2.1, we have an isometry $\iota : S \to \mathbb{H}^2 = \mathbb{D}^2$. Let (Λ, μ) be the finite bending lamination of $\mathbb{H}^2$ with leaves $(\iota(\ell_1), \ldots, \iota(\ell_N))$.

10.6 Definition. Suppose we have Ω, $S = \mathrm{Dome}(\Omega)$, and Λ as above. Let ℓ be one of the bending lines, and let R be a hyperbolic plane containing ℓ. We say that R is a *separating plane* for ℓ if the conditions below are fulfilled. We denote by S_λ the image of $\Phi_\lambda : \mathbb{H}^2 \to \mathbb{H}^3$, whether Φ_λ is an embedding or not. The first condition is that, for each $\lambda \in \mathbb{R}^n$, $S_\lambda \setminus \ell$ consists of two components, as in the embedded situation. The second condition is that, for each λ, the closure of each component of $\mathbb{H}^3 \setminus R$ contains one of the two components of $S_\lambda \setminus \ell$. The folded case—see Definition 4.4—is an exception, in which the plane containing S is regarded as the separating plane. We will say that S, or Ω, has the *separating planes property* if, for each bending line ℓ, there is a separating plane R containing ℓ.

It may be possible to vary R continuously by rotation about ℓ, while retaining its property of separation. The range of variation is a closed arc or a point. So there are either one or two extreme cases. We call the extreme cases *enveloping planes*. □

10.7 Theorem. *We use the notation introduced above. The following conditions are equivalent:*

10.7.1. The ortho-condition holds for (Λ, μ).

10.7.2. One can fit disjoint crescents into $\mathbb{D}^2 = \mathbb{H}^2$*, associated to the* $\iota(\ell_i)$ *with angles equal to* $\mu(\ell_i)$*. The corresponding* Ω *is the image of the associated angle doubling map.*

10.7.3. $S = \mathrm{Dome}(\Omega)$ *is quake-resistant.*

10.7.4. We have the separating planes property.

It follows from Corollary 7.2 that any of the above conditions implies the truth of the $K = 2$ conjecture for Ω.

We will prove the equivalences in a series of lemmas. By Theorem 9.2, we know that Condition 10.7.1 is equivalent to Condition 10.7.2.

10.8 Lemma. *Condition 10.7.2 implies Condition 10.7.3.*

Proof: We write $g_i = \iota(\ell_i) \subset \mathbb{D}^2 = \mathbb{H}^2$. Consider a crescent $C_i \subset \mathbb{D}^2$ of angle $\alpha_i = \mu(\ell_i)$ with endpoints equal to those of g_i. The two edges of C_i are curves at constant distance from g_i. It follows that an earthquake supported on g_i sends C_i to itself (possibly discontinuously). Crescents on one side of C_i are kept fixed, and crescents on the other side are moved to new crescents. It follows that the image crescents continue to be disjoint subsets of $\mathbb{D}^2$. But the whole earthquake is just the composition of earthquakes supported on the different geodesics g_i. (We need to take into account that the g_i move during the earthquakes.) So the disjointness of the crescents is maintained.

By Corollary 7.2, the the angle doubling map f' resulting from the new crescents is a homeomorphism $f' : \mathbb{D}^2 \to \Omega'$. Recall that f and f' are constructed by starting with the identity on a certain subset of $\mathbb{D}^2$ (see Theorem 5.3). The earthquake can be defined by keeping the corresponding piece of S fixed. The endpoints of the geodesics cut $\mathbb{S}^1$ up into a finite number of arcs of circle, each of which is sent by a Möbius transformation to its image. We can see inductively, working away from the fixed set, that on each such arc, the Möbius transformation induced by f, followed by the Möbius transformation induced by the earthquake, is equal to the Möbius transformation induced by f'. Let S' be the result of applying the earthquake to S. Then $S' = \mathrm{Dome}(\Omega')$, which is an embedded copy of $\mathbb{H}^2$ by Theorem 2.1. ∎

10.9 Bending lines with common endpoints. If you take a cone-shaped piece of wet soap, and squeeze it, the cone will move sideways. Given two geodesics which meet at infinity in the hyperbolic plane, the region between them has width varying from infinity down to zero, and squeezing it will move it away from the cusp end towards the thick end. In the next lemma, we show that a properly designed earthquake along two such geodesics will have the effect of a squeeze. Together with the picture given in the proof of the previous lemma, this will enable us to understand better the relationship of a quake-resistant surface to its ortho-curve.

10.10 Lemma (The squeeze lemma). *Suppose finitely bent convex surface S is quake-resistant. Suppose ℓ_1 and ℓ_2 are two geodesics which are* *boundary components of the same flat piece P. We denote the components of $S \setminus P$ by C_1 with boundary ℓ_1 and C_2 with boundary ℓ_2. Let $\lambda(x)$ be the earthquake supported on ℓ_1 and on ℓ_2, which, when normalized by keeping C_1 pointwise fixed, moves P a distance x (measured along ℓ_1) away from the common endpoint, and, when normalized by keeping C_2 fixed, moves P a distance x (measured along ℓ_2) towards the common endpoint.*

We normalize the earthquake by fixing C_1 pointwise. As x tends to infinity, P disappears and ℓ_2 moves towards ℓ_1. $S_{\lambda(x)}$ converges to a quake-resistant surface T with one less bending line. The bending angles of ℓ_1 and ℓ_2 are added to give the bending angle at the corresponding bending line of T. The ortho-curves of S and T are the same. If we think of Φ_λ (described in § 10.4) as a map $\mathbb{R}^N \to Maps(\mathbb{H}^2, \mathbb{H}^3)$, then the convergence above can be considered as taking place in the compact-open topology on the space of continuous mappings of $\mathbb{H}^2$ into $\mathbb{H}^3$.

The fact that we have to add bending angles in the limit justifies our definition of the ortho-curve. We also see that the folded case (see Definition 4.4) is a limit of cases where the convex hull boundary is an piecewise geodesic embedded copy of $\mathbb{H}^2$. For example we can have exactly two bending lines, each with bending angle $\pi/2$ and meeting at infinity. As the bending lines come together, we approach the folded case.

Proof: In the upper halfspace model, P is a vertical strip between the vertical lines ℓ_1 and ℓ_2. The boundary of P in $\mathbb{C}$ is a closed interval. The earthquake along ℓ_1 moves C_2 towards the finite endpoint of ℓ_1 by a euclidean change of scale e^{-x}. The earthquake along ℓ_2 moves $C_2 \cup \ldots \cup P_k$ away from the finite endpoint of ℓ_2 by a change of scale e^x. The combination of the two earthquakes is a parabolic transformation which is a euclidean horizontal translation, moving the finite endpoint of ℓ_2 towards the finite endpoint of ℓ_1, shortening the finite interval in the boundary of P by a factor e^x. ∎

10.11 Lemma. *Condition 10.7.3 implies Condition 10.7.1*

Proof: Let σ be any geodesic in the hyperbolic plane. We discard all leaves of the lamination Λ which σ does not meet, thus obtaining a ladder (see Definition 8.1). We need to prove the ortho-condition for this ladder. So we may assume that our lamination is a ladder. We orient σ and number $\ell_1, \ldots, \ell_N$ consecutively. We also number the flat pieces $P_1, \ldots, P_N$ consecutively.

First we consider the case where no two of the ℓ_i meet at infinity. Then, for each i, there is a unique interval $\sigma_i \subset P_i$ orthgonal to ℓ_i and to ℓ_{i+1}. Using an earthquake, we align all these intervals so that they form a single geodesic in $\mathbb{H}^2$, orthogonal to all the ℓ_i. The earthquake does not change the ortho-curve, and quake-resistance gives us a bent embedding of $\mathbb{H}^2$ in $\mathbb{H}^3$. By Lemma 4.6, we are still in the situation where the surface is the convex hull boundary of the complement of a simply connected region in $\mathbb{S}^2$. There is a hyperbolic plane $H \subset \mathbb{H}^3$ which is orthogonal to all the ℓ_i and which contains all the σ_i. Therefore the intersection of H with our surface is the ortho-curve, which clearly satisfies the ortho-condition.

The proof continues by induction on N. If $N = 1$, we are in the case already considered. Now take $N > 1$ and assume we know the lemma is true whenever we have fewer than N bending lines. By the argument above, we may assume that ℓ_i and ℓ_{i+1} meet at infinity. Let the bending angle along ℓ_j be α_j. Since S is the boundary of a convex subspace, we know that, for each j, $\alpha_j < \pi$. Also, using the upper halfspace model with the common endpoint of ℓ_i and ℓ_{i+1} at infinity, we see by looking at a horizontal horosphere that $\alpha_i + \alpha_{i+1} \leq \pi$.

In fact, equality $\alpha_i + \alpha_{i+1} = \pi$. is possible only if $N = 2$. To see this, note that the boundary of P_i at infinity is a finite interval, and the boundaries of P_{i-1} and P_{i+1} at infinity lie in two infinite rays. An earthquake along ℓ_i scales half the boundary at infinity of S, while keeping the other half fixed. Similarly for ℓ_{i+1}. A collision is inevitable if $N > 2$, so the surface S would not be quake-resistant.

If $N = 2$, then the ortho-curve satisfies the ortho-condition, as required. So we assume that $N > 2$ and $\alpha_i + \alpha_{i+1} < \pi$. We use Lemma 10.10 to bring ℓ_i and ℓ_{i+1} together, using earthquakes supported on ℓ_i and ℓ_{i+1}. We normalize all these earthquakes, fixing one particular flat piece (other than P_i, which disappears in the course of this construction). Consider the surface T given by squeezing P_i out of the picture, as in Lemma 10.10.

Let x be the parameter of Lemma 10.10, $S_x = S_{\lambda(x)}$. By assumption, S_x is the image of an embedding $\Phi_{\lambda(x)} : \mathbb{H}^2 \to \mathbb{H}^3$. By Lemma 4.6, S_x separates $\mathbb{H}^3$ into two components, one of which is convex. Let A_x be the closure of the convex component. As x tends to infinity S_x tends to the surface T and the closure of A_x tends to some closed convex subset A. Since T has $N - 1 > 1$ bending lines, A has non-empty interior, and $T = \partial A$ is an embedded convex surface.

T has the same ortho-curve as S, and each bending angle for T lies strictly between 0 and π. By induction on the number of bending lines, if we can show that T is quake-resistant, the ortho-condition will follow.

Let T' be obtained from T by an earthquake $\lambda' \in \mathbb{R}^{N-1}$ along the bending lines of T. We can approximate T' by applying an earthquake to S: we first apply the earthquake $\lambda(x)$ to S for some very large value of x, and then apply the earthquake λ', using all the bending lines of (the transformed) S except for (the transformed ℓ_{i+1}). Since S is quake-resistant, the same limiting argument for convex subspaces applied above when considering T applies again now to show that T' is an embedded bent copy of $\mathbb{H}^2$. It follows that T is quake-resistant, and the lemma is proved. ∎

10.12 The action at infinity of a complex earthquake. Before stating and proving the next lemma, we need to extend the definition of a complex earthquake to the region at infinity. Later, we will introduce some variants of this construction. Let $\rho : \Omega \to S$ be the nearest point retraction (see § 3.2). For each bending line ℓ_i, the inverse image $\rho^{-1}(\ell_i)$ is a crescent of angle α_i, which we denote by C_i. The angle α_i is the bending angle of S at ℓ_i. The crescent is foliated by arcs of circle, which we denote by $\ell_{i,t}$ with endpoints at the endpoints of ℓ_i. Here $\ell_{i,0}$ and $\ell_{i,1}$ are the boundary arcs of C_i and the angle between $\ell_{i,s}$ and $\ell_{i,t}$ is $|s-t|\alpha_i$. For each open flat piece F_i, the inverse image F'_i is an open region in Ω contained in a maximal disk, and with at least three points of its boundary in $X = \mathbb{S}^2 \setminus \Omega$. These points are also in the boundary of the maximal disk.

Under a complex earthquake supported on the bending lamination Λ of S, we move F'_i using the same Möbius transformation as is used on F_i. In order to extend the complex earthquake continuously to each C_i, note that the definition already given fixes the map on the two boundary arcs of C_i. Moreover, the effect of the complex earthquake on the two flat pieces on either side of ℓ_i is the same on the two endpoints x_i, y_i of ℓ_i. Let the earthquake map on one of these two flat pieces be g. Then the effect on the other flat piece is $f \circ g$, where f fixes $g(x_i)$ and $g(y_i)$. In order to extend the definition to C_i, we interpolate as follows. Let $(f_t)_{0 \le t \le 1}$ be the one-parameter group of Möbius transformations keeping $g(x_i)$ and $g(y_i)$ fixed, going from the identity to f, with the rotational part of f_t always less than π. We use the Möbius transformation $f_t \circ g$ on $\ell_{i,t}$.

10.13 Lemma. *Condition 10.7.3 implies Condition 10.7.4.*

Proof: Let S be a quake-resistant convex hull boundary. Let ℓ be one of the bending lines. We must find the plane R containing ℓ specified in Definition 10.6. If S is folded, then we take R to be the unique plane containing S. If S has only one bending line, we take R to be the bisector of the angle enclosing Ω. So from now on we assume that S has at least two bending lines.

In the upper halfspace model, take ℓ to be the vertical geodesic over 0, and let P' and P" be the adjacent flat pieces. Let H' and H" be the hyperbolic planes containing these flat pieces. These are two vertical halfplanes, above lines u' and u" in $\mathbb{C}$. The boundaries at infinity of P' and P" lie in rays r' and r", which have one endpoint at 0. The rays r' and r'' from 0 define a closed two-dimensional wedge $W \subset \mathbb{C}$ with an acute angle.

Let $W_0 \subset \Omega$ be the acute-angled open wedge lying between the orthogonal to r' and the orthogonal to r''. W_0 is the crescent associated to ℓ by the nearest point retraction $\Omega \to S$ (see § 10.12). Let W_0 be the wedge obtained by reflecting W through 0. Then $W_0 \subset \Omega$ and W_0 separates Ω into two components, which we call U' and U''.

Let λ be an arbitrary earthquake supported on the finite bending lamination Λ of S. We normalize the earthquake so that it keeps ℓ setwise fixed. The extension of λ to Ω—see § 10.12—preserves each ray through 0 setwise, and therefore preserves W, W_0, U' and U''.

As the earthquake varies (while preserving ℓ setwise), U' sweeps out an open subset V'. The action of λ can be factorized as the composition of two earthquakes, namely λ' which acts as the identity on U'' and λ'' which acts as the identity on U'.

We allow arbitrary earthquakes corresponding to a non-zero value of $\lambda \in \mathbb{R}^N$ on ℓ and zero on all other bending lines. In this case λ' is a change of scale map on U' and the identity on U'', and *vice versa* for λ''.

10.14 Lemma. *$V' \cap V'' = \emptyset$. There is an infinite ray $r \subset W$ from 0, such that r is disjoint from $V' \cup V''$.*

Proof: We will prove a contradiction from the assumption that there is some point $p \in V' \cap V''$. Then $p \in \lambda'(p') \cap \lambda''(p'')$ for some $p' \in U'$ and $p'' \in U''$, where λ' and λ'' are earthquakes which preserve ℓ setwise. We may assume that λ' is supported on one side of ℓ and that λ'' is supported on the other. In each case, the support is allowed to include ℓ. We now change λ', λ'' and p by composing with an earthquake supported entirely on ℓ. This allows us to assume that λ' is the identity on ℓ and on its two neighbouring pieces. Also it does not change V' or V''.

We can now combine λ' and λ'' into a single earthquake λ. Recall that we are assuming that S is quake-resistant. This means that Ω_λ is a topological disk for all earthquakes λ. Now Ω_λ is separated into the two components $\lambda(U')$ and $\lambda(U'')$ by r. This contradicts the definition of p.

Since V' and V'' are open wedges whose closures contain r' and r'' respectively, the existence of r follows. ∎

Returning to the proof of Lemma 10.13, we define R to be the vertical hyperbolic plane whose boundary contains the ray r. The remainder of the boundary of R is a ray r_0 whose interior is contained in Ω.

Let P' be the flat piece of S_λ adjacent to ℓ, and whose boundary at infinity lies in r'. Let A' be the component of $S_\lambda \setminus \ell$ containing P'. We have to prove that A' is contained in the component of $\mathbb{H}^3 \setminus R$ whose boundary at infinity contains V'.

Let $a' \in A'$. We must prove that a' lies on the correct side of R. Let H be a support plane for Y_λ, the convex subset of $\mathbb{H}^3$ with boundary S_λ, with H chosen so that $a' \in H$. Then the boundary of H at infinity is a circle or straight line. At infinity, this boundary bounds a maximal disk (possibly a halfplane) contained in Ω.

An earthquake can be considered as acting on H using the same Möbius transformation as is used on a'. It follows that the maximal disk is contained in V'. Therefore H and a' are both on the correct side of R. ∎

10.15 Lemma. *Condition 10.7.4 implies Condition 10.7.3.*

Proof: Suppose we are given, for each i, a hyperbolic plane R_i containing ℓ_i, with the properties specified in Definition 10.6. Given an earthquake of S, we move R_i with ℓ_i, maintaining the angles with the adjacent flat pieces. This clearly maintains the properties specified in Definition 10.6. It follows that the planes R_i provide barriers which prevent the different flat pieces from colliding. Quake-resistance follows. ∎

10.16 Corollary. *Suppose we have a finite measured lamination which satisfies one of the equivalent conditions of Theorem 10.7, and we then omit some of the geodesics and their associated measure, then the new finite measured lamination continues to satisfy these equivalent conditions.*

Figure 23: different drawings of enveloping planes, showing X, H(t) etc

Figure 24: enveloping plane when two geodesics meet at infinity

Proof: This follows instantly from Condition 10.7.2. ∎

This completes the proof of Theorem 10.7.

Figure 25: parabolic transformation when two geodesics meet at infinity

Figure 26: ray through two circles (orthogonal lining up)

11 Jordan curves: finite lamination case

In this section we will assume that we have the usual data: an open simply connected region $\Omega \subset \mathbb{S}^2$, the bent copy of a hyperbolic plane $S = \mathrm{Dome}(\Omega)$, and the corresponding bending lamination (Λ, μ) of the hyperbolic plane. We assume that Λ is finite and that the equivalent conditions of Theorem 10.7 are true.

11.1 Theorem. *Suppose the boundary C of S in $\mathbb{S}^2$ is not a Jordan curve. Then, there is a transverse geodesic σ, such that every leaf of Λ which it meets is crossed orthogonally. Moreover, the ladder determined by σ satisfies the ortho-condition, but not the strong ortho-condition.*

It follows that such a ladder cannot be included in a ladder with more bending lines without violating the ortho-condition.

Proof: Since C is not a Jordan curve, there is a geodesic $\sigma \subset \mathbb{H}^2$, such that the two endpoints of $\iota^{-1}(\sigma) \subset S$ are equal in $\mathbb{S}^2$. We take the ladder consisting of all geodesics intersecting σ. By Corollary 10.16, we do not lose generality in assuming that our finite lamination is equal to this ladder.

We orient σ and number the bending lines that cross σ consecutively as $(\ell_1, \ldots, \ell_N)$. We also number the flat pieces $(P_0, \ldots, P_N)$ consecutively. We minimize $j-i$ so that $i < j$ with $j - i \geq 2$, such that the closures of P_i and P_j meet at infinity. We claim that there is no loss of generality in assuming that $i = 0$ and $j = N$. For, if we can prove the theorem for the bending lamination $(\ell_{i+1}, \ldots, \ell_j)$, then the failure of the strong ortho-condition implies that one can't fit in further bending lines.

Now $\overline{P_0}$ and $\overline{P_N}$ meet at a point $p \in \mathbb{S}^2$, and we have arranged that p is not contained in the closure of any other flat piece. These two flat pieces have a common tangent at p. There is exactly one separating plane (see Definition 10.6) for each bending line, and its boundary therefore has the tangent direction already found at p in $\mathbb{S}^2$.

We use a Möbius transformation to put p at ∞. Then the separating planes become parallel vertical planes R_i in the upper halfspace model. We cannot have $R_i = R_{i+1}$ since the subsets they separate in Condition 10.7.4 are different. The geodesic ℓ_i is a geodesic in R_i which is not vertical, because, for all i, $p \notin \ell_i$. So, the endpoints of ℓ_i are two points on the boundary $r_i \subset \mathbb{C}$ of R_i.

The boundary of S has two components A and B in $\mathbb{C}$. The intersection of the boundary of P_i with $\mathbb{S}^2$ consists of two components $A_i \subset A$ and $B_i \subset B$. Also $r_i \cap \overline{\ell_i}$ is a pair of points (a_i, b_i), where $a_i \in A_i$ and $b_i \in B_i$. A_0 is a ray parallel to the y-axis, and similarly for A_N. For $0 < i < N$, A_i is a circular arc joining a_i to a_{i+1}. We have the same statements for the B_i, which is a circular arc joining b_i to b_{i+1}. Moreover, the angle between A_j and A_{j+1} is equal to the angle between B_j and B_{j+1}, because this is also the angle between P_j and P_{j+1}.

It follows that there is a horizontal line in $\mathbb{C}$, which bisects each vertical interval $[a_i, b_i] \subset r_i$. This horizontal line is the boundary of a vertical plane H, which is orthogonal to each P_i. It follows that $\sigma = H \cap S$ and the theorem is proved. ∎

12 Finite approximations of measured laminations

Suppose we are given a measured lamination $\hat{\Lambda}$ in $\mathbb{H}^2$ with transverse measure $\hat{\mu}$, and a geodesic $\hat{\sigma}$, infinite at both ends, which is not a geodesic in $\hat{\Lambda}$. Then we can omit any line which does not cross $\hat{\sigma}$, and obtain a new measured lamination. Since laminations are closed, by definition, we may need to restore either one or two geodesics passing through the endpoints of $\hat{\sigma}$, but, if so, the transverse measure is redefined so that the restored geodesics carry zero atomic measure. If we have to restore either one or two geodesics, we can then move $\hat{\sigma}$ a little, so that it is also transverse to the restored geodesics.

So we may assume that our measured lamination is a closed subset of $\mathbb{H}^2$, as usual, and we may also assume that $\hat{\sigma}$ crosses all the leaves of the lamination. We will show how to generalize to this situation the concept of an ortho-curve, previously defined only for finite measured laminations. For a finite measured lamination, we recall that the ortho-curve was geodesic except at a finite number of corners. In the new, more general, case, the measure may have no atoms, and the process of constructing the ortho-curve is then the same as constructing an arc with prescribed curvature. If there are atoms as well, then we have in addition a discrete counterpart to the construction of an arc with prescribed curvature. To do this we have to describe a noncommutative integration process.

12.1 The pseudometric induced by a geodesic lamination. Let $\hat{\Lambda}$ denote a geodesic lamination in $\mathbb{H}^2$ with an oriented geodesic $\hat{\sigma}$ which crosses all the leaves of $\hat{\Lambda}$. The intersection $\hat{\Lambda} \cap \hat{\sigma}$ is a closed subset, typically consisting of uncountably many points. The complement in $\hat{\sigma}$ is a countable union of open intervals lying in $\mathbb{H}^2 \setminus \hat{\Lambda}$. We will call these open intervals *flat segments*.

Suppose I is a flat segment with end points $p \in \ell_p$ and $q \in \ell_q$, where ℓ_p and ℓ_q are geodesics in $\hat{\Lambda}$. Define $\rho(p,q) := d(\ell_p, \ell_q) \geq 0$, the hyperbolic distance between ℓ_p and ℓ_q.

More generally suppose $p, q \in \hat{\Lambda} \cap \hat{\sigma}$. For each finite partition $P = (p = p_1, \ldots, p_n = q)$ of points in in $\hat{\Lambda} \cap [p,q]$, we define $\rho_P(p,q) = \sum_1^n d(\ell_{i-1}, \ell_i)$ where $\ell_i \subset \hat{\Lambda}$ is the geodesic through p_i and $d(\ell_{i-1}, \ell_i)$ is the hyperbolic distance between the geodesics. Then set $\rho(p,q) := \inf \{\rho_P(p,q)\}$, where we vary over all such partitions P.

The following lemma is easy and we omit the proof.

12.2 Lemma. *$\rho(p,q) \le d(\ell_p, \ell_q)$ with equality if and only if the set of geodesics of $\hat\Lambda$ that cross the segment $[p,q] \subset \hat\sigma$ have a common perpendicular.*

It may happen that $\rho(p,q) = 0$; this will be true if the geodesics $\ell_p, \ell_q \subset \hat\Lambda$ through p, q have a common end point, or if there is a geodesic $\ell \subset \hat\Lambda$ that has a common endpoint with both ℓ_p and ℓ_q, or if there is a succession of such geodesics between ℓ_p and ℓ_q.

It is also possible to have $\rho(p,q) = 0$ even if none of the intervening geodesics have common endpoints. To describe an example, we use the euclidean metric on the closed unit disk. Let X_1 be a set of two geodesics $\{\alpha, \beta\}$, such that the path distance along $\mathbb{S}^1$ from one endpoint of α to one endpoint of β is 1, while the path distance between the other two endpoints is $1/3$. We construct inductively a set X_k of 2^k disjoint geodesics with no endpoints in common. X_k is constructed so that we can form an embedded arc in the closed unit disk by following a geodesic in X_k, then an arc of $\mathbb{S}^1$ of length 3^{-k}, then another geodesic in X_k, and so on, until all the geodesics in X_k have been used up. We form X_{k+1} from X_k by inserting two additional geodesics between any two successive geodesics of X_k. One can prove that the sum of the hyperbolic distances between successive geodesics of X_k tends to zero as k tends to infinity. (A single term can be estimated as $O(6^{-k/2})$.) The closure of the union of the X_k gives us a ladder in which the ρ-distance between any two geodesics is zero.

12.3 Definition. We now form the quotient metric space Λ of $\hat\Lambda \cap \hat\sigma$. If $p_1 \in \hat\Lambda \cap \hat\sigma$ is a point between p and q then

$$\rho(p,q) = \rho(p,p_1) + \rho(p_1,q).$$

This property implies that Λ can be isometrically embedded in a real line σ, which we imagine as an oriented hyperbolic geodesic. A complementary interval (also called "flat") of Λ corresponds to a flat piece of $\mathbb{H}^2 \setminus \hat\Lambda$ bounded by two geodesics without common endpoints. A point of Λ corresponds, in general, to a subset of $\hat\Lambda \cap \hat\sigma$ rather than to a single point. The quotient map $\hat\Lambda \cap \hat\sigma \to \Lambda$ with the hyperbolic metric on the domain and ρ on the range is distance decreasing and is therefore continuous. $\square$

12.4 Local finiteness. The measure $\hat{\mu}$ on $\hat{\Lambda} \cap \hat{\sigma}$ can be pushed forward to a measure μ on σ under the composition $\hat{\Lambda} \to \Lambda \subset \sigma$, as is the case for any continuous map, though the measure may become locally infinite. In fact, the study of σ is only interesting when we can extend the ortho-condition to be meaningful in this situation. We will see in Definition 13.1 that the natural meaning of the ortho-condition is much stronger than local finiteness of the pushed forward measure. In the meantime, we will assume that μ is locally finite.

Even if $\hat{\Lambda}$ is not bounded at either end of $\hat{\sigma}$, Λ may be bounded. In other words, $\hat{\Lambda}$ may extend to the end of the infinite geodesic $\hat{\sigma}$, while Λ stretches out only over a finite interval in σ.

12.5 Example. Let Ω be the complement of the euclidean convex hull in $\mathbb{C}$ of a parabola. Ω^c is the union of the point at infinity together with the euclidean convex hull of the parabola. A support plane for the hyperbolic convex hull of Ω^c is the vertical plane whose boundary in $\mathbb{C}$ is tangent to the parabola. All the bending lines have a point in common, namely the point at infinity, and therefore their ρ-distance apart is zero. The total bending measure is π, as we can see from looking at finite approximations. It follows that Λ is a single point, with mass π. This will be one of the examples covered by our theory, and there is therefore a 2-quasiconformal map of the convex hull boundary (the product of the parabola with an infinite vertical ray) with Ω, which is the identity on the parabola (see Theorem 17.8).

12.6 Partitions of the transverse arc. Fix points $0 \in \Lambda$ as origin and $q \in \Lambda$. Given $\varepsilon > 0$, we partition $[0, q]$ by means of a sequence of points $0 = p_0 < p_1 < \ldots < p_n = q$ in Λ. The inequalities refer to the ordering in the oriented line σ. We also have points $u_i \in [p_{i-1}, p_i]$. This data is assumed to have the following properties:

- Each $p \in [0, q]$ for which ℓ_p has atomic measure greater than or equal to ε will be a point p_i.
- Each flat interval I with $\rho(I) \geq \varepsilon$ will be an interval $[p_i, p_{i+1}]$ of the partition.
- We have $\rho(p_i, p_{i+1}) < \varepsilon$, except perhaps when (p_i, p_{i+1}) is flat, or, equivalently, disjoint from Λ.

The above data $\{0 = p_0 \le u_1 \le p_1 \ldots \le u_n \le p_n = q\}$ is called an *ε-partition* of $[0, q]$.

Given an ε-partition of the transverse segment $[0, q]$, we concentrate the measure of the open interval (p_{i-1}, p_i) on u_i. We call the result an *ε-approximation* to the original bending measure on $[0, q]$. The approximation measure is purely atomic. The measure on a point p_j may increase, because we may have $u_j = p_j$ or $u_{j+1} = p_j$ or both. To avoid language which is too convoluted, we will write as though these points are distinct, leaving the reader to make the obvious adjustments when they are not distinct. (An easy way to see how to adjust is, for example, to imagine u_j tending to one of the endpoints of $[p_{j-1}, p_j]$ and use the limit construction.)

In order to be able to use a straightforward notation, we will write $\{0 = p_0 \le u_1 \le p_1 \ldots \le u_n \le p_n = q\}$ as $R = \{0 = r_0 \le \ldots \le r_{2n} = q\}$. Let β_j be the bending angle at r_j, so that $\beta_{2j} = \mu(p_j)$ and $\beta_{2j+1} = \mu(p_j, p_{j+1})$.

12.7 Convergence of measured laminations. The process above leads to an approximation by finite measured laminations $(\hat\Lambda, \hat\mu) = \lim(\hat\Lambda_n, \hat\mu_n)$ (see [7]). In particular this can be done so that a chosen leaf ℓ of $\hat\Lambda$ is also a leaf of $\hat\Lambda_n$ for all indices n.

The notion of earthquake along the leaves of a finite lamination extends as well to general laminations. Given $u > 0$, there is a unique earthquake of $\hat\Lambda$ determined by the measure $u\hat\mu$. One way to do this is to follow along the lines given below, approximating by a finite lamination, and then showing that this gives a well-defined limit (see [7]).

12.8 The sequence of polyhedral curves. With the above notation, we construct the normalized polygonal curve $\Gamma(R)$ as follows. Start the curve at the initial point $0 \in \mathbb{D}^2$ and draw a ray at counter-clockwise angle $\beta_0/2$ to the positive x-axis. This ray lies in the first quadrant. Only half of the angle β_0 is taken, because the other half is considered as belonging to the part of σ to the left of 0. Mark z_1 along the ray at a hyperbolic distance $\rho(0, r_1)$ from $0 \in \mathbb{D}^2$. Then draw the geodesic segment $[z_1, z_2]$ of hyperbolic length $\rho(r_1, r_2)$ that has counter-clockwise bending angle β_1 with respect to $[0, z_1]$. Continue in this manner. From the endpoint z_n of the last interval $[z_{n-1}, z_n]$, draw the infinite ray forming counter-clockwise bending angle β_n with $[z_{n-1}, z_n]$. The resulting polyhedral curve $\Gamma(R)$ depends on ε and on the ε-partition R. Its length from $z_0 = 0$ to z_n is equal to $\rho(0, q)$, which is an upper bound for the hyperbolic distance $d(0, z_n)$.

We will prove that this sequence of polyhedral curves converges as $\varepsilon \to 0$, independent of the partitions used.

Each point of $\Gamma(R)$ has a uniquely determined *medial support line*, which is a hyperbolic geodesic in the hyperbolic plane. The medial support line at 0 is the x-axis in $\mathbb{D}^2$. If p lies on the interior of a geodesic segment, the support line at p is the whole geodesic. If p is a vertex, the medial support line is the geodesic through the vertex angled midway between the two geodesic arcs with common endpoint p.

12.9 Proof of convergence. We will apply the technique developed in [7, Chapter 3]. Define

$$E(0, \beta) = \begin{bmatrix} e^{i\frac{\beta}{2}} & 0 \\ 0 & e^{-i\frac{\beta}{2}} \end{bmatrix}.$$

Given a point $z \in \mathbb{D}^2$, and angle β at z we will use the notation

$$E(z, \beta) = AE(0, \beta)A^{-1},$$

where A is a Möbius transformation of $\mathbb{D}^2$ sending $0 \mapsto z$, for the counterclockwise rotation about z of angle β.

We will unfold $\Gamma(R)$ to an interval on the positive real radius of $\mathbb{D}^2$; the arc from $z_0 = 0$ to z_n will have hyperbolic length $\rho(0, q)$. This is done as follows. Let $z_0 = 0, z_1, \ldots, z_{2n}$ denote the successive vertices of $\Gamma(R)$. The segment $[0, z_1]$ is rotated into the positive radius by $E(0, -\beta_0/2)$. The map $E(z_1, -\beta_1)$ lines the segment $[z_1, z_2]$ up with $[0, z_1]$, so that $E(z_0, -\beta_0/2)E(z_1, -\beta_1)$ sends it into the positive radius. The composition $E(z_0, -\beta_0/2)E(z_1, -\beta_1)E(z_2, -\beta_2)$ sends $[z_2, z_3]$ to the positive radius, and positions it adjacent to the previous segment. And so on. The composition

$$E(z_0, -\beta_0/2)E(z_1, -\beta_1) \cdots E(z_{2n}, -\beta_{2n}) \tag{12.9.1}$$

sends the ray from z_{2n} to the remainder of the positive radius.

12.10 Proposition. *We use the above notation, and fix q, as described at the beginning of § 12.6. As $\varepsilon \to 0$ and $n \to \infty$, the composition of rotations of Formula 12.9.1 converges, independent of the ε-partitions used. Given $M_1, M_2 > 0$, the convergence is uniform over all transverse segments $\sigma_q = [0, q]$ of ρ-pseudolength at most M_1 and transverse measure $\int d\mu < M_2$.*

Proof: Any two partitions have a common subpartition. Thus it is enough to consider the situation that the partitions are nested. Suppose we take the partition used in Formula 12.9.1, and subdivide each interval of that partition. We can write the composition product for the subpartition in the form $F = F_0 F_1 \cdots F_{2k}$, where the subpartition consists of k intervals and each factor is a rotation about some point of the hyperbolic plane. A single interval $[p_j, p_{j+1}]$ of the original partition gives rise to a factorization into rotations of the form $E(z_{2j}, -\beta_{2j})$ (if $j = 0$, $E(0, -\beta_0/2)$), $E(z_{2j+1}, -\beta_{2j+1})$ and $E(z_{2j+2}, -\beta_{2j+2})$. The first and third factors survive as factors in $F = F_0 F_1 \cdots F_{2k}$. If the interval $[p_j, p_{j+1}]$ is subdivided into a subintervals, then there is a corresponding product of $2a-1$ consecutive rotations in F. The sum of the bending angles for these factors is equal to β_{2j+1}. Therefore we can factor $E(z_{2j+1}, -\beta_{2j+1})$, which appears in the first factorization, into $2a-1$ rotations about z_{2j+1}, with the angles the same in the second factorization, but not the point about which the rotation is taken.

Thus we have a factorization $E = E_0 \cdots E_{2k}$, where the E_i and F_i are rotations through the same angle, about points which are at most ε apart.

We will use the formal identity

$$E_0 E_2 \cdots E_{2k} - F_0 F_2 \cdots F_{2k} = \sum_{j=0}^{2k} E_0 \cdots E_{j-1}(E_j - F_j) F_{j+1} \cdots F_{2k}.$$

Taking matrix norms of both sides we obtain

$$\|E_0 \cdots E_{2k} - F_0 \cdots F_{2k}\| \leq C \sum_{j=0}^{2k} \|E_k - F_k\|.$$

Here we are using the following lemma.

12.11 Lemma. *As we vary over all partitions, using the above notation, the set of products of the form $F_i F_{i+1} \cdots F_j$ varies within a compact subset of the group of isometries of $\mathbb{D}^2$ (which is isomorphic to* $\mathrm{PSL}(2,\mathbb{R})$*).*

Proof: We first prove the result with $i = 0$. The full result will then follow from the continuity of product and inverse in $\mathrm{PSL}(2,\mathbb{R})$. For the case $i = 0$, note that the product sends $z_j \in \mathbb{D}^2$ to $r_j \in \gamma$. Hence it sends $0 \in \mathbb{D}^2$ into the ball with centre 0 of radius $2M_1$. (M_1 is defined in the statement of Proposition 12.10.) The compactness statement follows. ∎

Continuing with the proof of Proposition 12.10, let α_j be the common angle of the rotations E_j and F_j. These are rotations about points at distance at most ε from each other. According to [7, Lemma 3.4.4],

$$\|E_j - F_j\| \leq C_1 \varepsilon \alpha_j.$$

Here the constant C_1 depends only on the ρ-length of $[0, q]$. We conclude that, with a new constant,

$$\|E_0 E_2 \cdots E_{2k} - F_0 F_2 \cdots F_{2k}\| \leq C\varepsilon \int_{[0,q]} d\mu. \tag{12.11.2}$$

Proposition 12.10 follows. ∎

Set $\tau = \{x : 0 \leq x \leq \rho(0, q)\}$. As a consequence of Proposition 12.10 there is a continuous mapping $\Upsilon : \tau \to \Gamma([0, q])$, which is the limit of the polygonal curves $\Gamma(R)$, as the mesh ε of the partitions R approaches zero.

Similarly, we can work in the negative direction from 0 along σ. Taking the union of definitions over closed finite subintervals in the positive and negative directions at the same time, we obtain a map, from the smallest subinterval of σ containing Λ, to $\mathbb{D}^2$.

Let $a = \inf \Lambda$ and $b = \sup \Lambda$ be the leftmost and rightmost points along σ. If a and b are finite, we get a well-defined proper map $f : \sigma \to \mathbb{D}^2$, with well-defined endpoints in $\partial\mathbb{D}^2$. One way to see this is to add points of mass zero to Λ just to the left of a and just to the right of b and apply Proposition 12.10 to this situation.

12.12 Pitfalls. The above path will be rectifiable, but may not be parametrized by pathlength, nor even locally injective. If $a = -\infty$, $f(t)$ may not tend to infinity in the hyperbolic plane as $t \to -\infty$, and similarly for b. Even if it does tend to infinity, there may not be well-defined endpoints. Later we will extend the definition of the ortho-condition so that it is meaningful in this situation, and we will show that if the ortho-condition is satisfied, then the pathologies just mentioned do not occur.

In summary,

12.13 Theorem. *Suppose $\hat{\Lambda}$ is a geodesic lamination of $\mathbb{H}^2$ and $\hat{\mu}$ a non-negative transverse Borel measure. Let $\hat{\sigma}$ be a geodesic transverse to $\hat{\Lambda}$ and let Λ be the quotient of $\hat{\Lambda}$ with the metric ρ, embedded isometrically in an* *oriented geodesic σ. We think of σ as $\mathbb{R}$, with metric also denoted by ρ. Let $a = \inf\Lambda$ and $b = \sup\Lambda$. Assume the pushed forward measure μ on σ is finite on every interval of finite ρ-length. Then there corresponds a uniquely determined normalized path $\Gamma(\sigma)$ in $\mathbb{H}^2$ with generalized curvature μ.*

In particular $\Gamma(\sigma)$ is independent of the choice of its origin up to Möbius equivalence.

12.14 Definition. The above curve $\Gamma(\sigma)$ is called the *ortho-curve* associated to $\hat{\Lambda}$, μ and $\hat{\sigma}$. □

13 The general ortho-condition

We now show how to generalize the ortho-condition so that it applies to an arbitrary measured lamination in the hyperbolic plane. We fix a measured lamination $\hat{\Lambda}$, with measure $\hat{\mu}$ and a transverse geodesic $\hat{\sigma}$. Following Definition 12.3, we have a quotient metric space Λ of $\hat{\Lambda}$, embedded isometrically in an oriented geodesic σ. We denote the metric on σ by ρ and the induced measure on σ by μ. We will also speak of the ortho-condition applying to ortho-curve $\Gamma(\sigma)$ (see Definition 12.14) and to the corresponding ladder.

For the convenience of the reader, we repeat all the conditions of Definition 8.3 for the ortho-condition to be satisfied in this more general situation.

13.1 Definition13.1.1. Given a point points $p \in \sigma$, we assume that $\mu(\{p\}) \leq \pi$. We insist, as in Condition 8.3.1, that if the measure of a some point is equal to π, then Λ should be the singleton set consisting of that point only. (This can occur even if none of the geodesics in Λ have common endpoints and $\hat{\Lambda} \cap \hat{\sigma}$ is unbounded both above and below.) In this special case (see Definition 4.4), the ortho-curve is a ray to infinity; regarded as a path, it is a path along the ray from infinity to a finite point, followed by the reverse of this infinite path.

13.1.2. Except for the folded case just described, the ortho-curve is embedded in $\mathbb{D}^2$, and the two ends give two well-defined endpoints which are required to lie in $\partial\mathbb{D}^2$. The two endpoints may be equal or distinct. In particular, the ortho-curve is a proper map $\mathbb{R} \to \mathbb{H}^2$.

The strong ortho-condition requires the following further condition:

13.1.3. The two endpoints of the ortho-curve are distinct.

For the boundary curve σ of any closed non-compact convex region in $\mathbb{H}^2$, we know from [7, A.6.2] that the bending measure μ satisfies the following condition: for any two points $p, q \in \sigma$, we have $\mu[p, q] \leq \pi + \rho(p, q)$, where ρ is the pathlength of the curve σ. The same condition therefore follows in our situation from Condition 13.1.2. It follows that the pushed forward measure μ is locally finite, and μ is therefore a Borel measure supported on $\Lambda \subset \sigma$. This condition implies the condition given in § 12.4.

A general measured geodesic lamination is said to satisfy the ortho-condition if it does so for each transverse geodesic γ, and similarly for the strong ortho-condition. It is easy to show that we need only check the ortho-condition for those γ such that neither endpoint is an endpoint of a geodesic of the lamination. $\square$

13.2 Properties of the ortho-curves.

13.3 Lemma. *Under the conditions of Definition 13.1, the ortho-curve (see Definition 12.14) is locally convex (always in the same normal direction from ortho-curve).*

Proof: The statement is true when Λ contains only a single point. So we assume that Λ contains more than one point. Then each point p_0 of Λ has measure strictly less than π.

The ortho-curve is formed by a process of approximation. At each finite stage, we have a map $\sigma \to \mathbb{H}^2$ where the image is finitely bent. By Proposition 12.10, it is easy to see that, at a point of σ with zero measure, the image of the unit tangent vector converges uniformly to a limit. Also, at a point with atomic measure, the images of the left and right tangents converge uniformly.

We think of p_0 as the origin of σ. Choose a small neighbourhood $(-\delta, +\delta)$, with measure only slightly more than that of $p_0 = 0$. In particular, the measure on this neighbourhood is strictly less than π.

The Gauss–Bonnet Theorem says that, for a geodesic polygon in the hyperbolic plane, the sum of the bending angles is equal to 2π plus the area. If we take a finite approximation to the measure over $(-\delta, \delta)$ as in described in § 12.8, the resulting ortho-curve has a sum of bending angles which is strictly less than π.

We now claim that this local piece of ortho-curve does not meet itself. For otherwise, two edges would meet in an unexpected way, and the bending angle at this meeting point would again be at most π, giving a total bending angle strictly less than 2π. So Gauss–Bonnet would give a negative area, which is absurd. In fact, by the same argument, the extreme edges can't even meet at infinity.

So, over $(-\delta, \delta)$, the images of these finite approximations to the ortho-curve is convex and embedded. A limiting argument then shows that the ortho-curve itself is convex over the same interval. Since p_0 was arbitrary, the ortho-curve is locally convex, always in the same direction, as claimed. $\blacksquare$

Given a subinterval $[a,b]$ of σ, we define a measure $\mu_{[a,b]} = \chi_{[a,b]}\mu$, where $\chi_{[a,b]}$ is the characteristic function. We can then construct the ortho-curve corresponding to $\mu_{[a,b]}$. If $[a,b]$ is small enough, we have seen that the ortho-curve bounds a closed convex subset $C_{[a,b]}$ of $\mathbb{H}^2$ which contains the main ortho-curve. (We normalize all the ortho-curves in the same way, so that they are consistent wherever possible.)

The next result follows immediately from Lemma 13.3 and a version of Lemma 4.6 in one dimension less.

13.4 Corollary. *An ortho-curve arising from the conditions of Definition 13.1 bounds a hyperbolically convex subset of $\mathbb{D}^2$ whose closure in the closed unit disk meets $\mathbb{S}^1 = \partial\mathbb{D}^2$ in either a point or an interval. This convex subset C_σ is the intersection of the convex sets $C_{[a,b]}$ as $[a,b]$ varies over a covering of σ. Except in the folded case (Definition 4.4), the ortho-curve separates $\mathbb{H}^2$ into two components, one of which is C_σ.*

13.5 Lemma. *Under the conditions of Definition 13.1, the ortho-curve is a rectifiable path, and can therefore be regarded as a pathlength preserving map $\mathbb{R} \to \mathbb{H}^2$. Let $K \subset L \subset \mathbb{R}$, where K and L are compact intervals with K contained in the interior of L. Take a sequence of finitely bent convex segments, converging in the Hausdorff topology to the image of L. Then the finitely bent segments can be parametrized by arclength so that they converge to the ortho-curve uniformly on K.*

Proof: We can work in a small neighbourhood, where the geometry is almost euclidean. We can even use the projective model so that convexity becomes the same as euclidean convexity. A convex curve is not necessarily continuously differentiable—there can be a countable number of points where the curve is differentiable from the left and the from the right, but the two are not equal. However, it is sufficiently close to being continuously differentiable so that arguments using differentiability can easily be made to work in this situation.

Rectifiability of a convex curve follows easily from this point of view, for example by using the mean value theorem. With rectifiability proved, we parametrize the ortho-curve by arc length. We choose a finite sequence of points in the image of L such that the sum of the successive image distances approximates L. We can then choose finitely bent convex segments which pass very near the points of this finite sequence. We can assume that derivatives are sufficiently close for ortho-curve and finitely bent approximation, so that the length of the finitely bent approximation is also nearly equal to L. Choosing consistent normalizations for the parametrizations, we see that the convergence is uniform on K. ■

13.6 Remark. In general uniform convergence does not behave well with respect to the property of preserving path distance. For example the curve $t \mapsto (t, \sin(n!t)/n)$ has length which tends to infinity with n. But its limit $t \mapsto (t, 0)$ has length one. In the case above, local convexity excludes this kind of pathology. □

13.7 How the endpoints of the ortho-curve move. We assume that we are working with a ladder which satisfies the version of the ortho-condition given in Definition 13.1. In particular this means that the endpoints of the ortho-curve are well-defined. We want to see how these endpoints move when we change the parameters defining the ortho-curve. We will keep the underlying lamination of the ladder fixed throughout the discussion.

Note that the endpoints of the ortho-curve on $\partial\mathbb{D}^2$ are determined if the pushed forward measure on σ has compact support. To see this, note that the endpoints are determined if the lamination is finite. For a general compactly supported Borel measure on σ, we have seen in Proposition 12.10 that the ortho-curves for approximating finite laminations converge uniformly to the limit ortho-curve.

Now suppose that the measure is not compactly supported on $\hat{\sigma}$. As we have pointed out above, Definition 13.1 implies that the measures on $\hat{\sigma}$ and σ are locally finite. In the discussion before the statement of Theorem 12.13, it was explained that the general ortho-curve was defined by using the definition over finite intervals and then increasing the size of the intervals to infinity.

From Corollary 13.4 we see that the convex set C_σ bounded by the ortho-curve is equal to $\bigcap_t C_{[-t,t]}$, the intersection of convex sets defined from the restricted measures $\mu_{[-t,t]}$. As we have seen above, each of these convex sets $C_{[-t,t]}$ meets $\partial\mathbb{D}^2$ in either an interval or a point. As t increases, the corresponding endpoints of the interval $C_{[-t,t]} \cap \partial\mathbb{D}^2$ Cmove towards each other, limiting at the endpoints of the interval $C_\sigma \cap \partial\mathbb{D}^2$.

As we have seen in § 12.6, we are at liberty by choice of $u_i \in [p_i, p_{i+1}]$ to shift the measure either towards the origin or away from the origin when constructing the approximating curves. Let us consistently shift the measure away from the origin. Then, as the approximation gets better, we have to shift the measure partly back towards the origin.

To see the effect on the approximating curves, we have to examine products like
$$E(z_0, -\beta_0/2)E(z_1, -\beta_1)\cdots E(z_{2n}, -\beta_{2n})$$
as in Formula 12.9.1. We have to compute the effect of changing $u_i = z_{2i-1}$, as it moves from z_{2i-2} to z_{2i}. The position of $u_i \in [z_{2i-2}, z_{2i}]$ can be recorded by means of a real parameter t, with $t = 0$ corresponding to z_{2i-2} and $t = 1$ corresponding to z_{2i}. Assuming we are working in σ to the right of 0, and that bending is upwards from a horizontal initial position, we see that the derivative with respect to t of the righthand endpoint of the approximating ortho-curve is in a clockwise direction along $\mathbb{S}^1$. In general, we will be moving all the u_i, for various values of i, at the same time. As we move away from the origin, the effect is to move the righthand endpoint of σ in a clockwise direction, and the lefthand endpoint in a counter-clockwise direction. This means that the endpoints are moved further apart from each other.

As we approximate more closely, and the measure moves back towards the origin, the righthand endpoint of σ is moved in a counter-clockwise direction along $\mathbb{S}^1$ and the lefthand endpoint is moved in a clockwise direction. Approximating over a bigger subinterval of σ surrounding the origin moves the endpoints further in these same directions, as we saw above in the discussion of the compact sets $C_{[-t,t]}$.

Therefore the approximating curves also satisfy the ortho-condition. It follows that all the approximating measured laminations satisfy the ortho-condition. Moreover, the above discussion indicates how we can deduce the strong ortho-condition for the approximations under suitable conditions.

A similar discussion to that above gives us the proof of the following lemma.

13.8 Lemma. *If the measured lamination (Λ, μ) in $\mathbb{D}^2$ satisfies the ortho-condition, then for $0 < t < 1$, $(\Lambda, t\mu)$ satisfies the strong ortho-condition.*

We next have the following result.

13.9 Lemma. *If (Λ, μ) satisfies the ortho-condition, then so does the result of applying to Λ any earthquake along Λ.*

Proof: We can suppose without loss of generality that we are dealing with a ladder. Let Λ' be the result of applying some fixed earthquake, and let μ' be the corresponding bending measure. We will adopt a similar naming convention for other items that appear in this proof, except that we use exactly the same names for corresponding geodesics in Λ and Λ'. Let ρ be the quotient distance for Λ defined in § 12.1 and let ρ' be the quotient for Λ'.

Let ℓ and m be two geodesics in Λ and let $\epsilon > 0$. Let $\ell_0 = \ell, \ldots, \ell_n = m$ be a finite sequence of geodesics in the ladder, such that

$$\rho'(\ell, m) > \sum_{i=1}^{n} d'(\ell_{i-1}, \ell_i) - \epsilon.$$

Let K be a compact subset of $\mathbb{H}^2$, whose interior meets each $\ell_i \in \Lambda$. Let $m_0 = \ell_0, m_1, \ldots, m_{N-1}, m_N = \ell_n$ be a refinement of $\ell_0, \ldots, \ell_n$, with the property that we can approximate the given earthquake by an earthquake along the m_j only. We are interested only in the effect of the approximating earthquake on the closed region bounded by ℓ and m.

We use Λ'' and μ'' for the result of applying the approximating earthquake to Λ and μ respectively, and similarly for other items, except the individual geodesics of the lamination, which retain their names. We take the approximation sufficiently close so that, for each i, $|d'(\ell_{i-1}, \ell_i) - d''(\ell_{i-1}, \ell_i)| < \epsilon/n$. This can be achieved by restricting attention to K. Since $\ell_i \cap K$ contains an interval for each i, we can ensure that the images of ℓ_i under the earthquake and under the approximating earthquake are almost the same.

We then have

$$\begin{array}{ccccc} \rho(\ell, m) & \leq & \sum_{j=1}^{N} d(m_{j-1}, m_j) & = & \sum_{j=1}^{N} d''(m_{j-1}, m_j) \quad \leq \\ \sum_{i=1}^{n} d''(\ell_{i-1}, \ell_i) & < & \sum_{i=1}^{n} d'(\ell_{i-1}, \ell_i) + \epsilon & < & \rho'(\ell, m) + 2\epsilon. \end{array}$$

This shows that $\rho \leq \rho'$, and we obtain equality by symmetry.

It now follows that the quotient metric spaces and pushed forward measures are equal for the situations before and after the earthquake. The lemma follows. ∎

13.10 Lemma. *Let Ω be a simply connected open subset of $\mathbb{S}^2$, and let $S =$ Dome$(\Omega) = \partial\mathcal{CH}(\Omega^c)$. Suppose that Ω is invariant under a reflection in a hyperbolic plane $H \subset \mathbb{D}^3$, extended to the closed unit 3-ball $\overline{\mathbb{D}^3}$. Let $\sigma = S \cap H$. Suppose that σ meets every bending line in S. Then the embedding $\sigma \subset H$ is exactly the embedding of the ortho-curve.*

Proof: The bending lamination in $\mathbb{D}^2$ consists of lines meeting σ orthogonally, since reflection in σ preserves the structure. We choose an origin in σ, which we can take to be the origin in $\mathbb{D}^3$ and $\mathbb{D}^2$. We can approximate the bending measure on σ by shifting the measure away from the origin. So, as we approximate more and more closely, the measure moves towards the origin.

The surfaces corresponding to these approximating curves are approximating convex hull boundaries in $\mathbb{D}^2$, touching S at the origin, and closing in on $\mathcal{CH}(\Omega^c)$ as the approximations become better and better. The result follows. ∎

14 Crescent limits

In this section, we discuss the generalization of Condition 10.7.2. Previously, we discussed embedding disjointly a finite collection of open crescents in the unit disk. Each crescent is foliated by arcs of circle, each arc having the same endpoints as the crescent. The edges of the crescents are also regarded as leaves. The closures of two crescents can meet at a single point of the boundary of each. We call such a point a *kissing point.* We denote by $\mathcal{A}$ the set of all leaves of the foliations of the various crescents.

14.1 Definition. A smooth arc γ, that is transverse to the arcs in $\mathcal{A}$ and that meets each arc of $\mathcal{A}$ at most once, is regarded as carrying a measure, namely the total angle of the various crescents traversed or partially traversed by γ. We call this measure the *circular measure.*

Here is another version of the circular measure on $\mathcal{A}$, in a form that makes generalization easier. Let $\gamma \subset \mathbb{D}^2$ be a closed smooth arc, transverse to $\mathcal{A}$, and parametrized by hyperbolic arclength. Let t be this parameter. Let $A(t) \in \mathcal{A}$ be the arc, if any, passing through $\gamma(t)$. If $\gamma(t)$ is a kissing point, then $A(t)$ has two possible values. However, since we are only interested in the result of integrating an associated function, and since there are only a finite number of kissing points (or a countable number—see Lemma 14.7—in the general situation, which we have not yet described), this will turn out not to matter. Let $\theta(t)$ be the angle which $A(t)$ makes $\mathbb{S}^1$. Let $\alpha(t)$ be the angle between γ and $A(t)$ at $\gamma(t)$. Let $X_\gamma = \gamma \cap \bigcup \mathcal{A}$.

$$\tau(\gamma) = \int_{X_\gamma} \sin\theta \, \sin\alpha \, dt. \qquad (14.1.3) \quad \square$$

14.2 Lemma. *Formula 14.1.3 makes $\tau(\gamma)$ equal to the total angle traversed by γ, when there are a finite number of crescents.*

Proof: We can clearly restrict to consideration of a single crescent. We use the upper halfplane model, with the crescent having endpoints at 0 and ∞. Using elementary trigonometry, we see that $\sin(\alpha)\, dt = d\theta / \sin(\theta)$. ∎

Given a finite disjoint union of crescents in $\mathbb{D}^2$, what are the properties of the measure τ? In particular, if A_0 and A_1 are two leaves in the foliations of the crescents, what can we say about the measure of a smooth arc γ going from A_0 to A_1 transversely to the foliation, and meeting any given leaf at most once? We will compute the maximum possible value of $\tau(\gamma)$, for fixed A_0 and A_1, as we vary over all configurations where these are leaves of a finite disjoint union of crescents.

Let B_0 and B_1 be the hyperbolic geodesics with the same endpoints as A_0 and A_1 respectively. Since A_0 and A_1 are disjoint, B_0 and B_1 are also disjoint (though they may meet at infinity). We specify the angle between A_0 and $\mathbb{S}^1$ as lying between 0 and $\pi/2$ positive if A_0 is on the same side (the positive side) of B_0 as B_1 and lying between $\pi/2$ and π otherwise. We specify the angle between A_1 and $\mathbb{S}^1$ similarly, interchanging the roles of B_0 and B_1.

14.3 Proposition. *Let γ be as above.*

14.3.1. Suppose both A_0 and A_1 are on the positive sides of B_0 and B_1 respectively. Then $\sup\tau(\gamma) = d(A_0, A_1)$*.*

14.3.2. Suppose both A_0 and A_1 are on the negative sides. For $i = 0, 1$, let θ_i be the angle A_i makes with $\mathbb{S}^1$, where $\pi/2 \leq \theta < \pi$. Then $\sup\tau(\gamma) = d(B_0, B_1) + \theta_0 + \theta_1 - \pi$*.*

14.3.3. Suppose A_0 is on the positive side and A_1 is on the negative side. Suppose furthermore that A_0 does not cross B_1 (they may meet tangentially). Then $\sup\tau(\gamma) = d(A_0, B_1) + \theta_1 - \pi/2$*.*

14.3.4. Suppose A_0 is on the positive side and A_1 is on the negative side. Suppose also that A_0 crosses B_1 at two points. Let A be the arc of circle with the same endpoints as B_1 which touches A_0 tangentially. Let θ be the angle between A and $\mathbb{S}^1$. Then $\sup\tau(\gamma) = \theta_1 - \theta$*.*

Proof: We consider first the case where B_0 and B_1 have a common endpoint at infinity. It is convenient to use the halfplane model, with the common endpoint equal to ∞. Then the A_i and B_i are straight lines, and the B_i are vertical. In the statement of the lemma, each term involving the hyperbolic distance between two of these lines is zero. All relevant crescents must have one endpoint at ∞, and so its two edges are also straight lines. The sum of the crescent angles is clearly bounded above by the difference between the angles of A_0 and A_1 to the x-axis.

We now assume that B_0 and B_1 do not have endpoints in common. They therefore have a common orthogonal geodesic. Only crescents whose endpoints lie on opposite sides of this common orthogonal contribute to $\tau(\gamma)$, and we may therefore discard all other crescents. We may then assume that γ is the subinterval of the common orthogonal joining A_0 to A_1.

Each crescent meets γ in a subinterval. We divide each crescent into one, two, or three subcrescents, so that the intersection of each crescent with γ lies in one of the one, two or three subintervals of γ formed by its intersection points, if any, with B_0 and B_1.

The argument now is by induction on the number of (subdivided) crescents. We think of γ as horizontal, with A_0 to the left and A_1 to the right. Along γ, from left to right, we have the different subintervals $J_1, \dots, J_N$, formed by intersecting γ with the different crescents.

First suppose that J_1 is to the left of B_0. We keep J_1 fixed, and change the leftmost crescent by moving its endpoints to those of B_0. Let $t \in J_1$. The change in the crescent changes the the angle $\alpha(t)$ of intersection between the leaves and γ to $\pi/2$ and therefore increases $\sin(\alpha)$ to 1. The change in the crescent also changes $\theta(t)$, the angle at which a leaf meets the boundary of the hyperbolic plane, by moving it towards $\pi/2$. It follows that $\int_{J_1} \sin(\alpha)\sin(\theta)\,dt$ is increased. Now we replace A_0 by the righthand edge of the new crescent. It is immediate that the claimed supremum of $\tau(\gamma)$ is decreased by at least the angle of the new crescent. By our induction, we may therefore suppose that J_1 is to the right of B_0.

We deal similarly with the case where J_N is to the right of B_1. So we may also suppose that J_N is to the left of B_1. We can then rearrange all the crescents simultaneously so that the ith crescent has endpoints equal to the endpoints of the orthogonal bisector geodesic of J_i.

In order to see that this increases the circular measure, first note that we can restrict attention to a single crescent. We may determine the circular measure by following a geodesic orthogonal to the leaves. This is already the case for the new crescent. For the old crescent, the length of the geodesic along which we integrate is reduced, as desired. In order to see the effect on the circular measure recall that, given a geodesic and a curve at constant distance x to the geodesic, the function $\arcsin\tanh(x)$ gives the angle at infinity between the geodesic and the curve. Hence the circular measure of the old crescent is at most $|\arcsin\tanh(x+d_i) - \arcsin\tanh(x)|$, where d_i is the length of J_i. This difference is in turn maximized by taking $x = -d_i/2$. But this is the value for the new crescent.

We have shown that the expressions in the statement of Proposition 14.3 are upper bounds. We now show that they are in fact suprema. Consider the subinterval of γ between B_0 and B_1. We cut this subinterval into a large number of contiguous short pieces $I_1, \ldots, I_M$. For I_j, we construct the orthogonal bisector geodesic β_j, and then form the crescent whose endpoints are the endpoints of β_j, and whose edges pass through the endpoints of I_j. The ratio of the angle of the crescent is to the length of I_j is approximately 1. So the supremum is not actually achieved in this situation, but we can get arbitrarily close to it. ■

In general, we may want to take a limit of laminations given by foliating disjoint crescents. First we have to generalize the kind of leaf we will accept. So far we have accepted curves which are at constant distance from a geodesic. However, our class of curves ought to be closed under taking limits of subsets of the hyperbolic plane. This results in the following definition.

14.4 Definition. We say that a curve in the hyperbolic plane is a *crescent leaf* if it is a curve with constant geodesic curvature between 0 and 1. If the geodesic curvature is strictly less than 1, this means it is a curve at constant distance from some geodesic, and its two endpoints are equal to those of the geodesic. If the geodesic curvature is 1, the curve is a horocycle, and when we mention its endpoint, we will mean the point on $\mathbb{S}^1$ at which the horocycle is centred. □

14.5 Definition. A set $\mathcal{A}$ of crescent leaves in $\mathbb{D}^2$ is said to be a *lamination by a crescent limit* if the following conditions are satisfied. [1]

14.5.1. $\bigcup \mathcal{A}$ is a closed subset of $\mathbb{D}^2$.

14.5.2. Given two crescent leaves, they may have a point of tangency in common, but they do not cross each other. Such a common point is called a *kissing point*.

14.5.3. Let A_0 and A_1 be two crescent leaves in $\mathcal{A}$ that are not horocycles, and let γ be any smooth path from A_0 to A_1 that is transverse to $\mathcal{A}$ and meets any crescent leaf of $\mathcal{A}$ at most once. Then the formula

$$\tau(\gamma) = \int \sin(\theta)\sin(\alpha)\, dt$$

[1] I now think that a necessary and sufficient condition for a collection of arcs and horocircles to be the limit (both topologically and in measure) of a sequence of disjoint crescent-pictures is that the measure is given by the integral of $sin(\theta)sin(\alpha)f(\theta)$, where f is a non-negative function with L_∞-norm at most one; $f = 1$ in the disjoint crescent case, but not in the limit.

defines a transverse measure on $\mathcal{A}$. The condition is that this measure is bounded above by the expressions given in Proposition 14.3 for the maximum possible value of the measure between A_0 and A_1. To be quite rigorous, we should point out that θ is not well-defined if two leaves of the lamination pass through the same point. However, we see geometrically that θ has well-defined one-sided limits along γ. We could fix, say, the positive direction along γ and then take limits in that direction in order to define θ unambiguously along γ. In Corollary 14.8, we will see that the integral is well-defined even before taking one-sided limits.

14.5.4. Given any smooth open arc γ, that is transverse to the subset of $\mathcal{A}$ which it meets, let $X_\gamma = \gamma \cap \bigcup \mathcal{A}$. The condition is that the support of τ in X_γ is equal to X_γ. □

14.6 Lemma. *Suppose we have a lamination $\mathcal{A}$ by a crescent limit. No point of $\mathbb{D}^2$ is contained in more than two leaves. A leaf which is a horocycle bounds a disk which contains no other leaves.*

Proof: In the case of a horocycle leaf, we must have $\sin(\theta) = 0$ for any leaves inside the corresponding horodisk. Such leaves therefore do not contribute to the measure. Then Condition 14.5.4 implies that the interior of the horodisk is free of leaves of $\mathcal{A}$.

Suppose $p \in \mathbb{D}^2$. We will deduce a contradiction by supposing that three distinct leaves of $\mathcal{A}$ pass through p. Let A be the middle leaf of the three. Taking a transverse path near p through all three leaves, we see from Condition 14.5.4 that there must be an infinite number of leaves through p, and these are all tangent to each other at p. Now, given a tangent vector in the hyperbolic plane, it is tangent to exactly two horocycles. Therefore we have three distinct non-horocyclic leaves through p, which we call L_0, L and L_1, with L between L_0 and L_1. By Condition 14.5.3 and Proposition 14.3, the circular measure between L_0 and L_1 is zero. Condition 14.5.4 then gives a contradiction, since L is not in the support of τ. ■

14.7 Lemma. *There are at most countably many kissing points.*

Proof: Since at most two crescent leaves of $\mathcal{A}$ pass through a given kissing point, no other crescent leaf of $\mathcal{A}$ can penetrate into the region near the kissing point between the two crescent leaves. It follows that the region between the two crescent leaves contains a non-empty open subset of the complement of $\bigcup \mathcal{A}$. ∎

14.8 Corollary. *The circular measure τ is well-defined for any crescent limit.*

The following proposition is now straight-forward, and the proof is left to the reader.

14.9 Proposition. *Suppose that, for each n, we have a finite union of disjoint crescents, and an associated lamination $\mathcal{A}_n$. Suppose that $(\mathcal{A}_n)_{n\in\mathbb{N}}$ converges to $\mathcal{A}$, in the sense that, firstly, for any sequence of crescent leaves $(A_n \in \mathcal{A}_n)_{n\in\mathbb{N}}$ which converges in the Hausdorff sense to a crescent leaf A, $A \in \mathcal{A}$ and, secondly, each $A \in \mathcal{A}$ is the limit of some sequence of crescent leaves $(A_n \in \mathcal{A}_n)_{n\in\mathbb{N}}$. Then $\mathcal{A}$ is a crescent limit.*

Conversely, any crescent limit is the limit of laminations $\mathcal{A}_n$ defined by a finite union of disjoint crescents.

Our next task is to define the angle doubling map $\mathbb{D}^2 \to \Omega$ determined by a crescent limit. In fact, for each $v \in \mathbb{R}$, we will define a continuous map $f_v : \mathbb{D}^2 \to \mathbb{S}^2$, so that f_v is the angle scaling map with scaling $1+v$.

The method is, as usual, to approximate the crescent limit $\mathcal{A}$ by a finite disjoint union of crescents $\mathcal{A}_n$. In more detail, we first fix a basepoint x_0 of $\mathbb{D}^2$. For each $n \in \mathbb{N}$, we pick a finite set $\mathcal{A}_n \subset \mathcal{A}$, such that each leaf of $\mathcal{A}_n$ meets disk $B(x_0, n)$ centred at x_0, with hyperbolic radius n, We assume that each point of $B(x_0,n) \cap \bigcup \mathcal{A}$ is within 2^{-n} of some leaf in $\mathcal{A}_n$. Finally, we assume that, for each n, $\mathcal{A}_{n-1} \subset \mathcal{A}_n$.

We say that two elements of $\mathcal{A}_n$ are adjacent if they are not separated by some other element of $\mathcal{A}_n$. For each adjacent pair (A_0, A_1) of elements of $\mathcal{A}_n$, we have a well-defined measure $\tau(\gamma)$ of any smooth arc γ from A_0 to A_1 which is transverse to $\mathcal{A}$ and meets each element of $\mathcal{A}$ at most once. Using Proposition 14.3 and Condition 14.5.3, we insert between each such pair a finite disjoint union of crescents, whose circular measure is approximately equal to $\tau(\gamma)$. We ensure that the ratio of the measures is between $1-2^{-n}$ and $1+2^{-n}$. This construction gives us a finite number of crescents in $\mathbb{D}^2$.

We now formulate our results about the existence of an angle doubling or angle scaling map for $\mathcal{A}$ as follows.

14.10 Theorem. *Let $\mathcal{A}$ be a crescent limit, with approximations $\mathcal{A}_n$ by a finite disjoint union of crescents. Let $v \in \mathbb{R}$ be fixed. Let $f_n : \mathbb{D}^2 \to \mathbb{S}^2$ be the continuous map given by $1+v$-scaling applied to $\mathcal{A}_n$. Then, as n tends to infinity, f_n converges to $f : \mathbb{D}^2 \to \mathbb{S}^2$, uniformly on compact subsets of $\mathbb{D}^2$.*

14.11 Definition. We call the limit f, the $1+v$*-scaling map* associated with the crescent limit $\mathcal{A}$. □

Proof: According to Theorem 5.3, we have an angle doubling map $f_n : \mathbb{D}^2 \to \Omega_n$ which is a 2-quasiconformal homeomorphism. Similarly, we have, for each $v \in \mathbb{R}$ and each n, the continuous $1+v$-scaling map $f_n = f_{n,v} : \mathbb{D}^2 \to \mathbb{S}^2$.

Let x_0 be the basepoint in $\mathbb{D}^2$. Let C_0 be either the leaf containing x_0, or, if x_0 is a kissing point, the union of the two leaves containing x_0, or, if $x_0 \notin \bigcup \mathcal{A}$, the closure of the complementary component containing x_0. There are only a finite number of leaves of C_0 meeting $B(x_0, n)$, and we include all of these as elements of $\mathcal{A}_n$. We normalize f_n by insisting that it be the identity on C_0.

We now claim that f_n converges uniformly on compact subsets. Claims like this are usually proved by showing that a certain set of isometries of $\mathbb{H}^3$ forms a compact subset. This enables us to show that various estimates are uniform as n tends to infinity. (See Lemma 12.11 and its consequences, or [7, Lemma 3.4.2].) For this reason, it is convenient to transfer the situation to $\mathbb{H}^3$. We take the upper halfspace model, and erect the unit hemisphere with boundary $\mathbb{S}^1$. Our configuration on $\mathbb{D}^2$ is transferred to this hemisphere, which is a hyperbolic plane, by stereographic projection from $(0,0,-1)$. Recall that stereographic projection is an isometry from $\mathbb{D}^2$ with the Poincaré metric to the unit hemisphere.

Let us examine more carefully the effect of the angle doubling map on the hemisphere. Consider a particular crescent C in the configuration, and let β be the hyperbolic geodesic connecting the endpoints of C. For any point $p \in C$, let $A(p)$ be the leaf containing p. Let $x = d(p,\beta) = d(A(p),\beta)$ and let θ be the angle between the leaf through p and $\mathbb{S}^1$. Then $\sin(\theta)\cosh(x) = 1$. Therefore $d\theta/dx = -1/\cosh(x)$. Let E_p be the isometry of $\mathbb{H}^3$ which acts on $A(p)$ in the angle scaling construction. (We are regarding v as fixed.)

Let q be a point near to p, with distance $\Delta x = d(A(p), A(q)) \leq d(p,q)$ between their two leaves. Then $E_q = E_p \circ \Delta E$, where ΔE is a rotation about β through an angle $|v\Delta\theta| \approx |v\Delta x/\cosh(x)|$. Then ΔE moves q a distance bounded above by $|v \sinh(x)\Delta\theta| \approx |v \tanh(x)\Delta x|$. Therefore, if $\epsilon > 0$ is small enough, and q is near enough to p,

$$d(E_p(p), E_q(q)) = d(p, (\Delta E)(q)) \leq d(p, q) + |v \tanh(x)\Delta x| + \epsilon.$$

It follows that $d(E_p(p), E_q(q)) \leq (1 + |v|)d(p, q)$.

We deduce that, without restriction on the distance from q to p, and without requiring p and q to belong to the same crescent, we have $d(E_p(p), E_q(q)) \leq (|v| + 1)d(p, q)$. It follows from this that, if p varies over a compact subset of the hemisphere, then E_p varies over a compact subset of the space of isometries of $\mathbb{H}^3$. But this is exactly what is required in order to show that the angle doubling and angle scaling maps, when restricted to any given compact subset of the hemisphere, converge uniformly. ■

If $-1 < v \leq 1$, the finite approximations are quasiconformal on $\mathbb{D}^2$. The quasiconformal constant is $1 + v$ if $v \geq 0$ and is $1/(1 + v)$ if $v < 0$. The same is therefore true of the limit (see [11, Theorem 5.1 p.72]). This proves the following theorem.

14.12 Theorem. *Let $\mathbb{D}^2$ have a lamination by a crescent limit, and let $f : \mathbb{D}^2 \to \Omega$ be the associated angle doubling map. Then f is a 2-quasiconformal homeomorphism. If $0 \leq v \leq 1$ and f is the associated angle scaling map, with scaling factor $1 + v$, then f is a $(1 + v)$-quasiconformal homeomorphism. If $-1 < v \leq 0$ and f is the associated angle scaling map, with scaling factor $1 + v$, then f is a $1/(1 + v)$-quasiconformal homeomorphism.*

14.13 Flaws. This result has a number of defects which will be attended to in the remainder of this paper. We have not yet proved, except in the case where there are a finite number of disjoint crescents, that the situation can be described in a number of distinct but equivalent ways. We do not yet have any information about the boundary values. This is a crucial point: if we don't worry about boundary values, then the Riemann mapping from $\mathbb{D}^2$ to Ω is 1-quasiconformal, and our result is much worse than the Riemann Mapping Theorem.

The analogous result to Theorem 14.12 can be proved for the more general angle scaling of Definition 5.1.

15 Equivalence: first step

In this section, we show that the ortho-condition is equivalent to the Definition 14.5. We now explain more precisely what this means. There is a map which sends a curve of geodesic curvature strictly less than 1 in the hyperbolic plane $\mathbb{D}^2$ with endpoints on $\mathbb{S}^1$ to a hyperbolic geodesic with the same endpoints. This map is continuous from the space of crescent leaves which are not horocycles to the space of geodesics.

Given a crescent limit $\mathcal{A}$ (see Definition 14.5), we have defined the corresponding circular measure by Formula 14.1.3. Consider the set of geodesics which have the same endpoints as some leaf in $\mathcal{A}$. This set of geodesics may not be closed, even though $\bigcup \mathcal{A}$ is a closed subset of $\mathbb{D}^2$, because the limit of a sequence of leaves in $\mathcal{A}$ may be an interval in $\mathbb{S}^1$. Let Λ be the set of geodesics obtained from $\mathcal{A}$, together with any limit geodesics.

We give Λ the measure μ pushed forward from the circular measure on $\mathcal{A}$.

15.1 Lemma. (Λ, μ) *is a geodesic lamination with a transverse locally finite Borel measure.*

Proof: The fact that the geodesics in Λ are disjoint is immediate. Condition 14.5.3 shows that the measure is locally finite. ■

15.2 Lemma. *There is a universal constant ϵ_0 with the following property. Let ϵ satisfy $0 < \epsilon \leq \epsilon_0$. Let A_0 and A_1 be disjoint non-horocyclic leaves of some crescent limit $\mathcal{A}$ and let their hyperbolic distance satisfy $d(A_0, A_1) < \epsilon$. Let B_0 and B_1 be the respective hyperbolic geodesics with the same endpoints, and suppose $d(B_0, B_1) < \epsilon$. Let $\tau = \tau(A_0, A_1)$ be the circular measure from $\mathcal{A}$ between A_0 and A_1. Then there is a crescent C with the same endpoints as B_0 or with the same endpoints as B_1, contained in the region between A_0 and A_1, and with angle β satisfying $\beta \geq (1-\epsilon)\tau$.*

Proof: We will use the formulas which convert between hyperbolic distance d from a given geodesic and angular distance α, namely:

$$\begin{aligned} \alpha = f(d) &= \arcsin(\tanh(d)) \text{ and} \\ d = g(\alpha) &= \operatorname{arctanh}(\sin(\alpha)). \end{aligned} \tag{15.2.4}$$

Note that $f'(0) = g'(0) = 1$ and $f''(0) = g''(0) = 0$. Therefore $|f'(d) - 1| = O(d^2)$ and $|g'(\alpha) - 1| = O(\alpha^2)$.

Let γ be a smooth arc going from A_0 to A_1 transversely to the foliation, and meeting any given leaf at most once. Leaves of $\mathcal{A}$ not meeting γ can be discarded for the purposes of this lemma. So we can think of γ as a horizontal geodesic and of A_0 on the left and A_1 on the right. Let θ_0 be the angle between A_0 and B_0, and similarly for θ_1. Note that the θ_i denote different angles from those in Proposition 14.3. We have $\left|\frac{d(A_i,B_i)}{\theta_i} - 1\right| = O(\epsilon^2)$ and $\left|\frac{\theta_i}{d(A_i,B_i)} - 1\right| = O(\epsilon^2)$. C_i, with $i = 0, 1$, will always denote a crescent with angle β_i and with the same endpoints as B_i and A_i. The left edge of C_0 will always be A_0 and the right edge of C_1 will always be A_1.

First suppose that B_0 and B_1 have an endpoint in common. In this case, we choose C_0 with β_0 equal to the angle between A_0 and A_1 at a common endpoint. Then $\beta_0 \geq \tau$.

So, from now on we assume that B_0 and B_1 do not have an endpoint in common. We also assume from now on that γ is orthogonal to A_i and B_i for $i = 0, 1$. We use the order along γ to distinguish cases.

15.2.1. Let $\gamma \cap B_0 \leq \gamma \cap A_0 \leq \gamma \cap A_1 \leq \gamma \cap B_1$. Then $\tau \leq d(A_0, A_1)$ from Proposition 14.3. Let C_0 have right edge A, which is chosen to go through $A_1 \cap \gamma$. Then, from Equation 15.2.5 and the mean value theorem,

$$\left|1 - \frac{\beta_0}{d(A_0, A_1)}\right| = O(\epsilon^2).$$

We deduce that $\beta_0 \geq \tau(1 - \epsilon)$.

15.2.2. Let $\gamma \cap A_0 \leq \gamma \cap B_0 < \gamma \cap B_1 \leq \gamma \cap A_1$. By Proposition 14.3, $\tau \leq d(B_0, B_1) + \theta_0 + \theta_1$. Let C_0 have right edge A, which is chosen to go through $A_1 \cap \gamma$. Now $\left|1 - \frac{\beta_0}{d(A_0,A_1)}\right| = O(\epsilon^2)$. Therefore

$$\begin{aligned} \beta_0 \geq d(A_0, A_1)\left(1 - \frac{\epsilon}{2}\right) &= \\ (d(A_0, B_0) + d(B_0, B_1) + d(B_1, A_1))\left(1 - \frac{\epsilon}{2}\right) &\geq \tau\left(1 - \frac{\epsilon}{2}\right)^2 \\ &\geq \tau(1 - \epsilon). \end{aligned}$$

15.2.3. Let $\gamma \cap B_0 \leq \gamma \cap A_0 \leq \gamma \cap B_1 \leq \gamma \cap A_1$. Suppose furthermore that A_0 does not cross B_1 (they may meet tangentially). Let C_1 have left edge A, which is chosen to go through $A_0 \cap \gamma$. Then $\tau \leq \theta_1 + d(A_0, B_1)$. We have

$$\beta_1 \geq \theta_1 + d(A_0, B_1)\,(1 - \epsilon) \geq \tau\,(1 - \epsilon)\,.$$

15.2.4. Let $\gamma \cap B_0 \leq \gamma \cap A_0 \leq \gamma \cap B_1 \leq \gamma \cap A_1$. Suppose further that A_0 crosses B_1 twice. Let the right edge of C_0 be A chosen so as to pass through $\gamma \cap A_1$. Then $\tau \leq \beta_0$. ∎

15.3 Theorem. *Given a crescent limit $\mathcal{A}$, the measured geodesic lamination (Λ, μ), defined above on $\mathbb{D}^2$, satisfies the ortho-condition.*

Proof: Since the ortho-condition refers only to ladders, we may assume without loss of generality that Λ is a ladder, and we do this for the remainder of the proof. Also, there is a geodesic γ which crosses each geodesic in Λ. So γ is also transverse to $\mathcal{A}$, and crosses each leaf of $\mathcal{A}$ exactly once. We will imagine γ as horizontal, so that we can talk of points on γ being to the left or the right.

For each $n \in \mathbb{N}$, we will define a finite subset $\mathcal{A}_n$. We set $\mathcal{A}_0 = \emptyset$ and fix a basepoint x_0 with $\gamma(0) = x_0$. Let the corresponding set of geodesics be denoted by Λ_n. $\mathcal{A}_n$ satisfies a number of conditions.

15.3.1. $\mathcal{A}_{n-1} \subset \mathcal{A}_n$.

15.3.2. Each geodesic of Λ meeting $\gamma[-n, n]$ is within 2^{-n} of some geodesic in Λ_n.

15.3.3. Any open interval of $\gamma[n, -n] \setminus \bigcup \Lambda$ of length at least 2^{-n} has both endpoints lying in geodesics of Λ_n.

15.3.4. Any element of Λ meeting $\gamma[-n, n]$ and having atomic measure at least 2^{-n} is in Λ_n.

15.3.5. Consider the leftmost geodesic α of Λ meeting $\gamma[-n, n]$. Then there may or may not be an element of $\mathcal{A}$ with the same endpoints as α. If there is such an element of $\mathcal{A}$, then we denote α by $\alpha_{n,0}$. Of the various elements of $\mathcal{A}$ which have the same endpoints as $\alpha_{n,0}$, we denote by $A_{n,0}$ the rightmost possibility. If there is no such element of $\mathcal{A}$, we choose $\alpha_{n,0} \in \Lambda$ such that $d(\gamma \cap \alpha_{n,0}, \gamma \cap \alpha) < 2^{-n}$ and and such that the measure $\mu[\alpha_{n,0}, \alpha] < 2^{-n}$. We then choose $A_{n,0}$ to have the same endpoints as $\alpha_{n,0}$. We insist that $A_{n,0} \in \mathcal{A}_n$ and $\alpha_{n,0} \in \Lambda_n$ We choose $A_{n,0}$ and $\alpha_{n,0}$ to be the leftmost elements of $\mathcal{A}_n$ and Λ_n respectively.

15.3.6. We define $A_{n,1} \in \mathcal{A}_n$ and $\alpha_{n,1} \in \Lambda_n$ in a similar way, but on the righthand side.

15.3.7. Any open subinterval in γ of $[A_{n,0}\cap\gamma, A_{n,1}\cap\gamma]\setminus\bigcup\mathcal{A}$ of length at least 2^{-n} has endpoints lying in leaves of $\mathcal{A}_n$.

Each component of $\mathbb{D}^2\setminus\mathcal{A}_n$ is bounded either by one or by two elements of $\mathcal{A}_n$. For each component U bounded by two elements, say $A', A'' \in \mathcal{A}_n$, we use Lemma 15.2 to find a crescent in U with the same endpoints as A' or as A'', and with angle β satisfying $\tau(A',A'')\geq\beta\geq\tau(A',A'')\,(1-2{-}n)$. We obtain a finite disjoint union of crescents in $\mathbb{D}^2$. The corresponding finite lamination is (a subset of) Λ_n with measure μ_n. By Theorem 8.6, the corresponding finitely bent ortho-curve Γ_n satisfies the ortho-condition.

We next need to check what happens when n tends to infinity.

Given two geodesics $\alpha,\beta\in\Lambda$ and $\epsilon>0$, choose geodesics $\alpha_0=\alpha,\dots,\alpha_N=\beta$ in Λ, such that $\sum_{i=1}^N d(\alpha_{i-1},\alpha_i) < \rho(\alpha,\beta)+\epsilon$. We choose n large and $\beta_i\in\Lambda_n$, such that $d(\beta_i,\alpha_i)<\epsilon/N$. Then

$$\begin{aligned}\rho(\alpha,\beta)\leq\rho_n(\beta_0,\beta_N)+\epsilon\leq\sum_{i=1}^N d(\beta_{i-1},\beta_i)+\epsilon \quad &< \\ \sum_{i=1}^N d(\alpha_{i-1},\alpha_i)+3\epsilon \quad &< \quad \rho(\alpha,\beta)+4\epsilon.\end{aligned}$$

This shows that ρ_n converges to ρ (the domains can be identified as metric copies of $\mathbb{R}$ after identifying the basepoints).

Since the measures also converge weakly on compact subsets, we see that Γ_n converges, uniformly on compact subsets, to Γ, the ortho-curve associated to (Λ,μ). Since each Γ_n bounds a non-compact closed convex subset of the hyperbolic plane, the same is true for Γ. Therefore (Λ,μ) satisfies the ortho-condition. ∎

We also need to prove the converse of Theorem 15.3.

15.4 Theorem. *Let (Λ,μ) be a measured lamination which satisfies the ortho-condition. Then there is a (non-unique) crescent limit $\mathcal{A}$ which has associated measured lamination (Λ,μ).*

Proof: We fix a basepoint $x_0\in\mathbb{D}^2$. On the ball $B(x_0,n)$ of hyperbolic radius n about x_0, we approximate $(\Lambda,(1-1/n)\mu)$ extremely closely in the weak topology by a finite lamination (Λ_n,μ_n). We include in Λ_n any geodesic γ of Λ carrying atomic measure greater than 2^{-n}, and in this case we take $\mu_n(\gamma)=(1-1/n)\,\mu(\gamma)$. We can ensure that each geodesic in Λ_n carries a non-zero amount of measure, perhaps extremely small—this is possible because the support of μ is equal to Λ. We arrange that $\Lambda_{n-1} \subset \Lambda_n \subset \Lambda$. Because $(1-1/n)\mu < \mu$ if μ is not identically zero, $(\Lambda, (1-1/n)\mu$ satisfies the strong ortho-condition. It follows that, if we approximate closely enough, (Λ_n, μ_n) will also satisfy the strong ortho-condition. (To prove the ortho-condition or the strong ortho-condition for (Λ, μ) we need to consider a possibly infinite, or even an uncountable, number of ladders. However, once we restrict to those geodesics meeting $B(x_0, n)$, we have only a finite number of ladders to consider.)

It follows from Theorem 9.2 that we can insert a finite set disjoint crescents into $\mathbb{D}^2$, with associated lamination (Λ_n, μ_n). We denote the corresponding collection of leaves by $\mathcal{A}_n$. We want to find a crescent limit $\mathcal{A}$ for this sequence of laminations. Each geodesic $\alpha \in \Lambda_n$ gives rise, for each $m \geq n$, to a non-trivial crescent $C_m(\alpha)$ with the same endpoints as α. By taking a subsequence, we can assume that $C_m(\alpha)$ converges geometrically to a limit as m tends to infinity.

By the usual Cantor trick, we can find a single subsequence so that, for each $\alpha \in \bigcup_{n\in\mathbb{N}} \Lambda_n$, the crescent $(C_m(\alpha))_{n\in\mathbb{N}}$ converges to a limit $C(\alpha)$. If $\mu(\alpha) > 0$, then $C_n(\alpha)$ will eventually have an angle $(1-1/n)\mu(\alpha)$, and $C(\alpha)$ will be a crescent of angle $\mu(\alpha)$.

There is one point that is easy to miss. It is conceivable that, the crescent $C_m(\alpha)$ converges to a subset of $\mathbb{S}^1$. To address this concern, we refine the procedure for choosing Λ_n by insisting that, for each geodesic $\alpha \in \Lambda_n$, either $\mu(\alpha) > 0$ or α is a limit of geodesics of Λ on the far side of α from the basepoint x_0. It is clear that $C_m(\alpha)$ cannot tend to a subset of $\mathbb{S}^1$ if the limit $C(\alpha)$ has an angle $\mu(\alpha) > 0$. So we assume that $\mu(\alpha) = 0$. Let $\beta \in \Lambda_p$ lie on the far side of x_0 from λ, with $p > n$. Then $\mu(\beta, \alpha) > 0$.

Let $A_m \in \mathcal{A}_m$ be a leaf of $C_m(\alpha)$ and let $B_m \in \mathcal{A}_m$ be a leaf of $C_m(\beta)$, for m large. By taking m large enough, we may assume that the circular measure $\tau_m(A_m, B_m) > \mu(\beta, \alpha)/2$. By Proposition 14.3, either A_m stays on the same side of β as x_0 as m tends to infinity, or there is a crescent of angle $\mu(\beta, \alpha)/2$ with the same endpoints as α, lying between A_m and B_m. It follows that A_m must converge to a crescent leaf in $\mathbb{D}^2$.

So far we have defined a countable collection of disjoint crescents whose leaves lie in $\mathcal{A}$, and a countable collection of crescent leaves. To form $\mathcal{A}$ itself, we must also include any crescent leaf which is a limit of crescent leaves already defined.

It is straightforward to see that (Λ, μ) is the measured lamination derived from $\mathcal{A}$. $\blacksquare$

16 Generalizing the conditions for angle doubling

In Theorem 15.3 and Theorem 15.4 we have seen that the condition of being a crescent limit is equivalent to the ortho-condition. We recall the statements of these theorems. If we start with a lamination of $\mathbb{D}^2$ by a crescent limit (see Definition 14.5), we have an associated measured geodesic lamination (Λ, μ) which satisfies the ortho-condition. Conversely, given a measured geodesic lamination which satisfies the ortho-condition, we have shown how to construct a (non-unique) crescent limit whose associated measured geodesic lamination is the given one.

However, we have not yet established, except in the case of a finite lamination, that these conditions relate to the convex hull construction in $\mathbb{H}^3$. So far, our considerations in the case of general laminations, have been confined to $\mathbb{H}^2$. We now rectify this deficiency.

16.1 Theorem. *Let (Λ, μ) be a measured lamination satisfying the ortho-condition. Then this is the bending lamination associated to some simply connected open subset $\Omega \subset \mathbb{S}^2$. Moreover, Ω is uniquely determined by (Λ, μ), up to Möbius transformation.*

Proof: One problem that might occur is that the finite approximations may not satisfy the ortho-condition. To take care of this, the n-th approximant does not approximate (Λ, μ), but $(\Lambda, (1 - 1/n)\mu)$. By Lemma 13.8, this measured lamination satisfies the strong ortho-condition. We may therefore assume that all the finite approximants (Λ_n, μ_n) also satisfy the strong ortho-condition. By our previous considerations of finitely bent surfaces, we know that we obtain a finitely bent convex embedded surface S_n bounding a convex subset $Y_n \subset \mathbb{H}^3$, with corresponding bending lamination (Λ_n, μ_n). Let $X_n = \mathbb{S}^2 \cap \overline{Y_n}$ and let $\Omega_n = \mathbb{S}^2 \setminus X_n$. Ω_n is simply connected.

Now the Hausdoff limit of the closures in $\overline{\mathbb{D}^3}$ of hyperbolically convex subspaces Y_n is also the closure of a hyperbolically convex subspace $Y \subset \mathbb{D}^3$. It is easy to see that $\Omega = \mathbb{S}^2 \setminus \overline{Y}$ is simply connected, and that Y is the hyperbolic convex hull of $X = \mathbb{S}^2 \setminus \Omega$. It follows that (Λ, μ) is the bending lamination of Dome(Ω) ∎

Let $\Omega \subset \mathbb{S}^2$ be a non-empty simply connected subset whose complement $X \subset \mathbb{S}^2$ contains more than one point. Let $S = \text{Dome}(\Omega)$ be the hyperbolic convex hull of X. According to Theorem 2.1, S is isometric to $\mathbb{H}^2$. The bending lines of S give the measured lamination (Λ, μ) on $\mathbb{H}^2$.

The bending measure μ is a positive measure. However, we want to regard it as a complex measure with purely imaginary values, for the reasons explained at the beginning of Section 6.

16.2 Definition. In general, each complex-valued transverse Borel measure μ supported on Λ gives rise to a well-defined continuous map $\Phi_{\Lambda,\mu} : \mathbb{H}^2 \to \mathbb{H}^3$ and its image $S_\mu \subset \mathbb{H}^3$. These can be defined by approximating μ locally by a finitely supported measure. The details are given in [7]. The map $\Phi_{\Lambda,\mu}$ sends each rectifiable path to a rectifiable path of the same length. Note that the measure is not necessarily related to any convex hull construction. We also have the (usually) discontinuous complex earthquake map $E_{\Lambda,\mu} : \mathbb{H}^2 \to \mathbb{H}^3$, defined by taking the limit of the corresponding map for finite laminations defined at the end of Definition 6.1. □

Here is a technical point. A complex-valued measure can be written as the sum of a real-valued measure and a purely imaginary measure. The second of these is the result of multiplying a real-valued measure by i. A real-valued measure is the sum of a positive measure and a negative measure, supported on disjoint subsets. If both of these are allowed infinite values, we have trouble with defining the measure of the union of two disjoint sets, one with infinite positive measure and one with infinite negative measure. For this reason, infinite values are often excluded in the theory of complex valued measures. However, in our case, what we really need is a measure defined on each relatively compact open interval transverse to the lamination, which is invariant under leaf-preserving isotopy, and whose values agree whenever two transverse intervals have a common subinterval. Since Borel measures are finite on compact subsets, all these measures are finite, and we are never faced with the quandary of needing to define $\infty - \infty$.

16.3 Definition. Let μ be a complex-valued transverse Borel measure on a lamination Λ. We say that μ is *bounded* if there is a constant $m > 0$ such that, for each transverse arc A of length 1, the total transverse measure $\int_A d|\mu| < m$. □

16.4 Lemma. *For the boundary of the hyperbolic convex hull of the complement of any open topological disk, the bending measure is bounded. There is a universal bound for the constant as we vary over all simply connected Ω.*

Proof: Let S be the boundary of the convex hull.

In the appendix of [7] it is shown that, for a planar convex curve, the bending measure along a curve of length t is at most $\pi + t$. Therefore, if there is a geodesic which is orthogonal to all the bending lines, we can take $m = \pi + 1$.

In general, the estimate is a bit more complicated: it follows by approximating the 3-dimensional situation by the 2-dimensional situation just described. Readers who are prepared to accept this as obvious can skip the proof given below. A more precise result was proved in [7, Proposition 2.4.2], but the proof there is much too difficult to warrant the effort of understanding the less precise result we now require.

First we work in euclidean 3-space. Consider the ray from 0 in the (y,z)-plane with $z > 0$, at an angle $\alpha > 0$ to the y-axis. Now project the union of this ray and the y-axis orthogonally onto the plane orthogonal to the vector $\left(\sqrt{1-r^2}, r\cos(\phi), r\sin(\phi)\right)$. One can compute that the new bending angle β satisfies

$$\sin(\beta) = \sin(\alpha)\sqrt{1 + r^2\cos(\alpha)\cos(\alpha+2\phi) + \frac{r^4}{4}\cos^2(\alpha+2\phi)\sin^2(\alpha)}.$$

Taking account of the shape of the sin curve, this shows that, for r small, and independently of the size of $\alpha > 0$, $|\beta/\alpha - 1| < \epsilon$, where ϵ is small.

Next we approximate S by a finitely bent surface. Let $\gamma \subset S$ be a short geodesic arc with respect to the path metric on S which is transverse to the bending lamination. Let $x_0 \in \gamma$ be a point on a bending line ℓ. Now take a plane H orthogonal to ℓ through x_0. Since γ is short, all bending lines met by γ are almost orthogonal to H. (In this calculation, we can use universal constants, as is seen in [7].) It follows from the previous paragraph, if the length of γ, which we denote by t, is small, then the bending measure on γ is at most $(1+\epsilon)(\pi + t)$. Thus the measure is bounded as claimed. ∎

Since μ is fixed, we usually abbreviate $Phi_{\lambda+i\mu}$ to Φ_λ, as in § 10.4. As in Section 10, we define S_λ to be the image of Φ_λ.

Now we generalize the definition of quake-resistance given for finite laminations in Definition 10.2 to the current situation where the lamination is an arbitrary bending lamination.

16.5 Definition. We consider the convex hull boundary S as above. Let λ be an arbitrary earthquake of S supported on Λ. Suppose that for each such

λ, $\Phi_\lambda : \mathbb{H}^2 \to \mathbb{H}^3$ is a proper embedding or a fold (as in Definition 4.4). Then we say that S is *quake-resistant.* □

There are other definitions of "quake-resistant", which depend on the class of earthquakes permitted in the definition. For example, we could allow only finitely supported earthquakes; or only earthquakes which are absolutely continuous with respect to the bending measure; or only bounded earthquakes. Each of these gives a formally distinct class of quake-resistant surfaces. However, all of them turn out to be equivalent to each other. We will not systematically prove the equivalence, and will leave this to readers who may be interested. In each situation, the technique is to use some kind of approximation, including, for example, the slightly unorthodox approximation of a finitely supported earthquake by an earthquake which is absolutely continuous with respect to the bending measure.

16.6 Theorem*16.6.1. Suppose that (Λ, μ) is the bending lamination of a surface S which satisfies the ortho-condition. Then S is quake-resistant for arbitrary earthquakes supported on Λ.*

16.6.2. If S is a convex hull boundary which is quake-resistant for finitely supported earthquakes, then it satisfies the ortho-condition.

Proof: Suppose first that the ortho-condition is satisfied. It is easy to show that an earthquake does not change the metric ρ defined in § 12.1. It follows that, after an earthquake, the ortho-condition continues to be satisfied. By Theorem 16.1, the result of an earthquake λ is a properly embedded surface S_λ. This proves that S is quake-resistant.

Conversely, suppose that S is quake-resistant for finitely supported earthquakes. We want to prove the ortho-condition. Without loss of generality we suppose that Λ is a ladder. Given $\ell', \ell'' \in \Lambda$ and $\epsilon > 0$, let $\ell' = \ell_0, \ldots, \ell_N = \ell''$ be elements of Λ, such that $\rho(\ell', \ell'') > \sum_{i=1}^N d(\ell_{i-1}, \ell_i) - \epsilon$. Using the quake-resistance assumption, we can apply a finitely supported earthquake which enables us to assume that $\rho(\ell', \ell'') > d(\ell', \ell'') + 2\epsilon$.

Let $(\gamma_i)_{i\in\mathbb{N}}$ be a sequence of elements of Λ which is dense in Λ and such that each geodesic in the sequence is repeated an infinite number of times. Let λ_n be a finitely supported earthquake, such that, for $1 \le i < j \le n$, $d(\lambda_n(\gamma_i), \lambda_n(\gamma_j)) \ge \rho(\gamma_i, \gamma_j) > d(\lambda_n(\gamma_i), \lambda_n(\gamma_j)) - 1/n$. Let S_n be the image of S under λ_n. Then $S_n = \partial Y_n$, where Y_n is a closed, convex subset of $\mathbb{H}^3$.

Also, $Y_n = \mathcal{CH}(X_n)$, where X_n is a compact subset of $\mathbb{S}^2$ and $\Omega_n = \mathbb{S}^2 \setminus X_n$ is simply connected.

The limit of the Y_n in the Hausdorff metric on closed subsets of $\overline{\mathbb{D}^3}$ is a closed convex subset Y. Then ∂Y is a bent surface with boundary S, and its bending lamination is a ladder. S has the property that, for any two bending lines ℓ' and ℓ'', $d(\ell', \ell'') = \rho(\ell', \ell'')$, where d is the path metric on S. It follows readily that there is a single geodesic σ on S which is orthogonal to all the bending lines. Then $\sigma = H \cap S$, where $H \subset \mathbb{H}^3$ is a (flat) hyperbolic plane. The ortho-curve can be identified with σ. Our quake-resistant surface S therefore satisifes the ortho-condition. ■

16.7 Definition. Suppose we have Ω, $S = \mathrm{Dome}(\Omega)$, Λ as above and that S is quake-resistant. Let λ be a transverse measure Let ℓ be one of the bending lines, and let R be a hyperbolic plane containing ℓ. We say that R is a *separating plane* for ℓ if the conditions below are fulfilled. We denote by S_λ the image of the injective map $\Phi_\lambda : \mathbb{H}^2 \to \mathbb{H}^3$. The condition is that, for each λ, the closure of each component of $\mathbb{H}^3 \setminus R$ contains one of the two components of $S_\lambda \setminus \ell$. The folded case—see Definition 4.4—is an exception, in which the plane containing S is regarded as the separating plane.

We will say that S, or Ω, has the *separating planes property* if, for each bending line ℓ, there is a separating plane R containing ℓ.

It may be possible to vary R continuously by rotation about ℓ, while retaining its property of separation. The range of variation is a closed arc or a point. So there are either one or two extreme cases. We call the extreme cases *enveloping planes*.

We say that the *separating plane condition* is satisfied, if, for each bending line $\ell \subset S$, there is a separating plane R. □

This definition is the obvious generalization of Definition 10.6, to which it specializes when the lamination is finite. As in the definition of quake-resistance, we get various definitions of separating planes, according to the class of earthquakes used—finitely supported, bounded or absolutely continuous. As in the case of quake-resistance, we obtain equivalent concepts.

An enveloping plane, is divided into two halfplanes by ℓ, one half contained in $\mathcal{CH}(\mathbb{S}^2 \setminus \Omega)$ and the other half disjoint from this convex hull.

16.8 Definition. The halfplane contained in the convex hull is called an *enveloping halfplane*. If there are two enveloping halfplanes, we can distinguish them by specifying to which component of $S \setminus \ell$ they are nearest, using angle of rotation about ℓ to measure nearness. □

16.9 Theorem. *The ortho-condition, the condition that we have a crescent limit, quake-resistance and the separating plane condition are four equivalent conditions. Each of these conditions implies the existence of a 2-quasiconformal angle doubling map.*

Proof: We already know that the first three conditions are equivalent. By definition, the separating plane condition implies quake-resistance.

Let (Λ, μ) satisfy the ortho-condition. We already know that S is quake-resistant, and we need only show the existence of separating planes. We do this using the technique of finite approximation—see for example Theorem 16.1. By Theorem 10.7, we obtain, for each approximating finite lamination (Λ_n, μ_n), and, for each $\gamma \in \Lambda_n$, a separating plane $R_{\gamma,n}$ which contains the image of γ in $\mathbb{H}^3$.

Let γ be fixed. We now let R be a limit for the $R_{\gamma,n}$ as n tends to infinity. It is then easy to see that R is a separating plane for S, since we can approximate an arbitrary earthquake of S by a finite earthquake.

Doing this, for all $\gamma \in \bigcup_n \Lambda_n$, we find a separating plane for each geodesic in a dense subspace of Λ. Finally, if $\gamma \in \Lambda$ is some other geodesic, we take a limit of a sequence of geodesics $(\gamma_n \in \Lambda_n)_{n \in \mathbb{N}}$ which converges to γ. For a subsequence, the separating planes R_n for γ_n converge to a separating plane R for γ.

The fact that each of the conditions implies the existence of an angle doubling map is the content of Theorem 14.12 ∎

We still need to control boundary values for a general angle doubling. We know what the boundary map is in the case of a finite bending lamination, but the general case has not yet been established.

We next give the hypotheses for a lemma that will be useful in what follows. Suppose the separating planes condition is satisfied. Let ℓ_1 and ℓ_2 be two bending lines. These cut $S = \mathrm{Dome}(\Omega)$ into three pieces, P_0 with frontier ℓ_1 in S, P_1 with frontier $\ell_1 \cup \ell_2$ and P_2 with frontier ℓ_1. For $i = 1, 2$, let H_i be an enveloping halfplane (see Definition 16.8) with $\partial H_i = \ell_i \cup c_i$, where $c_i = \mathbb{S}^2 \cap \overline{H_i}$. Then c_i is an arc of circle in $\mathbb{S}^2$ with the same endpoints as ℓ_i. If there are two choices for the enveloping halfplane, we will choose that nearest P_0 for H_1 and that nearest $P_0 \cup P_1$ for H_2 (see Definition 16.8 for the meaning of "nearest" in this context).

16.10 Lemma. *$H_1 \cap \mathbb{H}_2 = \emptyset$. The only possible intersection of c_1 and c_2 is at a common endpoint.*

Proof: We will suppose that $H_1 \cap H_2 \neq \emptyset$, and prove a contradiction. Then $H_1 \cap H_2$ is either a geodesic or a geodesic ray ending at a point of $\ell_1 \cup \ell_2$. In either case, c_1 must cross c_2 and there is a crossing point x in the interior of each of these arcs.

By the definition of an enveloping halfplane, for any neighbourhood N_1 of c_1 in $\mathbb{S}^2$, there is some earthquake λ_1 supported on $\Lambda \cap \overline{P_0}$, such that $\lambda_1(\overline{P}_0) \cap N_1 \neq \emptyset$. An earthquake supported on ℓ_1 can be regarded as fixing H_1 and c_1 setwise. It moves $\overline{P_0} \cap \mathbb{S}^2$ away from one endpoint of c_1 towards the other. Taking N_1 suitably small, we see that we can use an earthquake like λ_1 to move part of $\overline{P_0} \cap \mathbb{S}^2$ from one side of c_2 to the other. But this contradicts the definition of H_2. ∎

17 Boundary values

The object of this section is to complete the proof that, if one of our equivalent conditions is satisfied, then the $K = 2$ conjecture is verified. The only point that remains to be proved is that the 2-quasiconformal map $\Omega \to \mathrm{Dome}(\Omega)$, obtained by angle doubling, extends continuously to the identity on $\partial\Omega$. This will follow from the general result Theorem 17.7 about quasi-isometries. Before we can apply results about quasi-isometries to our particular situation, we need to show that quasiconformal homeomorphisms are quasi-isometries. This very useful general fact is established in Theorem 17.1.

An indication of this phenomenon is seen in the the Ahlfors-Beurling and Douady-Earle extensions to $\mathbb{D}^2$ of quasisymmetric homeomorphisms of $\mathbb{S}^1$. These are known to be quasi-isometries, with constants determined by the quasisymmetric constants. Our work also controls constants, and we obtain a sharp result for the multiplicative constant.

In the converse direction, it is clear that we cannot hope for a quasi-isometry to be quasiconformal: a quasi-isometry doesn't even need to be continuous, and quasiconformality is a local property.

17.1 Theorem. *Suppose $f : \mathbb{D}^2 \to \mathbb{D}^2$ is a K-quasiconformal map. Let d denote the hyperbolic metric. Then*

$$(d(z,w) - a)/K \leq d(f(z), f(w)) \leq K d(z,w) + a, \quad \forall\, z, w \in \mathbb{D}^2, \quad (17.1.5)$$

where $a = a(K) > 0$. In particular, we can take $a(K) = K \log 2$ if $1 < K \leq 2$. An estimate for $a(K)$ is given in the proof when $K > 2$.

Proof: We may assume that $f(0) = 0$. It suffices to estimate $d(0, f(z))$. By the Hersch-Pfluger distortion theorem (see [11, p. 64]),

$$|f(w)| \leq \phi_K(|z|), \quad \forall\, z \in \mathbb{D}^2. \qquad (17.1.6)$$

Here ϕ_k is the distortion function which only depends on the maximal dilatation K. In particular for $K = 2$,

$$\phi_2(t) = \frac{2\sqrt{t}}{1+t}, \quad 0 \leq t \leq 1.$$

For a 2-quasiconformal map f, from Inequality 17.1.6 follows

$$|f(z)| \leq \frac{2\sqrt{|z|}}{1+|z|}. \qquad (17.1.7)$$

We will find it useful to define the function $g : [0,1] \to [1,2]$, given by $g(t) = \frac{(1+x)^2}{/}1 + x^2$. We recall that the hyperbolic distance $d(0,z) = \log\left(\frac{1+|z|}{1-|z|}\right)$. Substituting Inequality 17.1.7, we find

$$d(0, f(z)) - 2d(0,z) \leq 2\log\left(g\left(\sqrt{|z|}\right)\right) < 2\log 2.$$

From this we conclude that

$$d(0, f(z)) \leq 2d(0,z) + 2\log 4. \tag{17.1.8}$$

This proves the right side of Inequality 17.1.5 when $K = 2$. The left side is proved by applying Inequality 17.1.8 to f^{-1}.

To continue, we need the following lemma.

17.2 Lemma. *Given $L \geq 1$, and a point p on a hyperbolic geodesic γ, there is an L-quasiconformal diffeomorphism $f_L : \mathbb{H}^2 \to \mathbb{H}^2$ which fixes p and sends γ to itself, stretching hyperbolic distances along γ by a factor L, and stretches all other distances by a factor between 1 and L.*

Proof: We use the upper halfplane model. Take $i = p$ and γ be vertical geodesic through i. We define $f_L(r,\theta) = (r^L, \theta)$. ∎

We continue with the proof of Theorem 17.1, in the case $1 < K < 2$. We apply Lemma 17.2 to $\mathbb{D}^2$, with $L = 2/K$, $p = 0$ and γ equal to the geodesic passing through 0 and z. Note that $f_L f$ is 2-quasiconformal and fixes 0. We obtain

$$Ld(0, f(z)) = d(0, f_L f(z)) \leq 2d(0,z) + 2\log 2.$$

Therefore $d(0, f(z)) \leq Kd(0,z) + K\log 2$.

Now let $K > 2$, and choose $n > 1$ as small as possible so that $P = K^{1/n}$ satisfies $1 < P \leq 2$. We note that any K-quasiconformal map $g : \mathbb{D}^2 \to \mathbb{D}^2$ can be written as a composition of P-quasiconformal maps $f_i : \mathbb{D}^2 \to \mathbb{D}^2$, $1 \leq i \leq n$ (see [1, p.100]):

$$g = f_1 \circ f_2 \circ \cdots \circ f_n.$$

Applying Inequality 17.1.5, with K replaced by P and $1 < P \leq 2$, to the f_i we find a constant a_n such that, for all $z, w \in \mathbb{D}^2$, we have $d(g(z), g(w)) \leq Kd(z,w) + a_n$. The inductive procedure above allows us to find a value for a_n in terms of K:

$$a_n = \frac{P\log(2)\,(P^n - 1)}{P-1} \leq \frac{(K-1)\sqrt{2}\log 2}{\sqrt{2}-1}.$$

So, if $K > 2$, we can take $a(K) = 2.37(K-1)$. ∎

Suppose $g : R \to S$ is a homeomorphism between complete hyperbolic surfaces of finite type. Let L_g be the Lipschitz constant of g if g is Lipschitz and ∞ otherwise. Let K_g be the dilatation of g if g is quasiconformal and ∞ otherwise.

17.3 Corollary. *Let $f : R \to S$ be a homeomorphism between complete hyperbolic surfaces of finite type. Set $K(f) = \inf\{K_g : g \simeq f\}$, and $L(f) = \inf\{L_g : g \simeq f\}$, where g varies over all homeomorphisms which are homotopic to f. Then $L(f) \leq K(f)$.*

Proof: . We may assume without loss of generality that f is the Teichmüller map with $K_f = K(f)$. According to Thurston's unpublished manuscript [18],

$$L(f) = \inf_{\alpha} \frac{\text{length}[f(\alpha)]}{\text{length}[\alpha]},$$

where α is a simple closed geodesic on R of hyperbolic length length$[\alpha]$ and length$[f(\alpha)]$ is the length on S of the geodesic freely homotopic to $f(\alpha)$.

Choose fuchsian representations G_R, G_S for R and S so that $\pi_R : \mathbb{D}^2 \to \mathbb{D}^2/G_R$ and $\pi_S : \mathbb{D}^2 \to \mathbb{D}^2/G_S$. Let $f^* : \mathbb{D}^2 \to \mathbb{D}^2$ be a lift of f. Let α be an arbitrary simple closed geodesic on R. We parametrize α by pathlength and lift to a geodesic $\beta : \mathbb{R} \to \mathbb{D}^2$, with $\pi_R(\beta) = \alpha$.

Let $A = \text{length}(\alpha)$. Since f is a quasi-isometry by Theorem 17.1, we know that $f^* \circ \beta : \mathbb{R} \to \mathbb{D}^2$ has the same endpoints as a geodesic γ. Projecting onto γ reduces lengths, and so, from Theorem 17.1 we have $n\text{length}(f(\alpha)) \leq d(f(\beta(0)), f(\beta(nA))) \leq K_f nA + a(K_f)$. Dividing by nA, and applying Thurston's result, referred to above, we find that $L(f) \leq K_f + \frac{a(K_f)}{nA}$. Corollary 17.3 follows. ∎

The following is a crucial fact.

17.4 Lemma. *There are universal constants C_1 and C_2, independent of Ω, such that, for any geodesic γ of length t on S, the bending measure along γ is bounded by $C_1 t + C_2$.*

17.5 Question. Perhaps one can take $C_1 = 1$. If that were the case, then the nearest point retraction $r : \Omega \to S$ would be a $(2, \epsilon)$-quasi-isometry, for some ϵ. □

Proof (of Lemma 17.4): In [7, A5.2, p252], it is shown that if γ is orthogonal to all the bending lines it meets, then we can take $C_2 = \pi$ and $C_1 = 1$, and this is best possible. In general, γ does not meet the bending lines orthogonally, but there is a universal δ, such that, if $t < \delta$, then the bending lines are almost orthogonal to γ. We need to show that there is a universal bound to how far from orthogonal to γ the bending lines can get. An explicit computation is given in [7, Proposition 2.14.2], but this is unnecessarily complicated for our present purposes. It then follows that we may choose δ so that the total bending measure along γ is less than 2π if $t < \delta$. We can then take $C_1 = 2\pi/\delta$ and $C_2 = 2\pi$. ∎

Using Lemma 17.4, we can prove the following result, pointed out explicitly in [4], and implicitly contained in [7].

17.6 Lemma. *For any simply connected open subset $\Omega \subset \mathbb{S}^2$, the nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$ is a quasi-isometry.*

Proof: In Theorem 23.9, we will show that r is 2-lipschitz with respect to the hyperbolic metric on Ω and the pathmetric on S. Here we present a proof of a much worse estimate. Let $S = \mathrm{Dome}(\Omega)$ and let $\log 2 > \epsilon > 0$. We factor r into $r_\epsilon : \Omega \to S_\epsilon$ and $r_2 : S_\epsilon \to S$. Here $S_\epsilon \subset \mathbb{H}^3$ is the surface ϵ-distant from S, and r_1 and r_2 are nearest point retractions.

By [7, Theorem 2.3.1], r_1 is $4\cosh(\epsilon)$-lipschitz and r_-^{-1} is $\mathrm{cosech}(\epsilon)$-lipschitz. So we need only check that r_2 is a quasi-isometry. Since r_2 decreases the path-length, we need only show that its inverse satisfies appropriate inequalities. We prove this for a finitely bent surface S and then take limits.

Applying Lemma 17.4, we see that a geodesic γ on S of length t has an inverse image $r_2^{-1}(\gamma)$, which consists of a finite number of pieces of total length $\cosh(\epsilon)t$ between the crescents, and crossing crescents through a total angle of at most $\sinh(\epsilon)(C_1 t + C_2)$.

Let $x, y \in S_\epsilon$, and let $t = d_S(r_2 x, r_2 y)$. Then

$$d_{S_\epsilon}(x,y) \le (\cosh(\epsilon) + C_1 \sinh(\epsilon)) d_S(r_2 x, r_2 y) \text{ and } d_S(r_2 x, r_2 y) \le d_{S_\epsilon}(x,y).$$

So r_2 is a quasi-isometry, and so is the composition $r = r_2 r_1$. ∎

Let $\Omega \subset \mathbb{S}^2$ be an arbitrary simply connected open region, $S = \mathrm{Dome}(\Omega)$, and Λ the bending lamination on S. Let $A \subset \partial\Omega$ be the set of points which are endpoints of bending lines in Λ, or limit points in $\mathbb{S}^2$ of flat pieces for Λ. We have seen before that, for each point $p \in A$, there is a maximal disk $D \subset \Omega$, such that $p \in \partial D \cap \partial\Omega$ and $\partial D \cap \partial\Omega$ contains at least two points.

Let $g : \Omega \to \Omega$ be a (K, ϵ)-quasi-isometry with respect to the hyperbolic metric on Ω. Although g need not be continuous on Ω, we assume it has the following continuity property on $\partial\Omega$. Suppose $p \in A$. We assume that there is a sequence $(x_n)_{n\in\mathbb{N}}$ converging in the spherical metric to p, which is quasigeodesic for the metric d_Ω, such that $(g(x_n))_{n_\mathbb{N}}$ also tends to p. We recall that $(x_n)_{n\in\mathbb{N}}$ is said to be *quasigeodesic* if there are constants $k \geq 1$ and $\epsilon \geq 0$, such that, for all $m, n \in \mathbb{N}$,

$$(d_\Omega(x_n, x_m) + \epsilon)/k \leq |n - m| \leq k d_\Omega(x_n, x_m) + \epsilon.$$

17.7 Theorem. *Under the above assumptions, if we extend g to be the identity on $\partial\Omega$, then it is continuous on $\partial\Omega$ (even though it may be discontinuous at all points of Ω).*

By Theorem 17.1, this theorem applies in particular when g is quasiconformal.

Proof: It is clear that, any d_Ω-quasigeodesic α converging to p in the spherical metric is at a finite Ω-distance from the sequence $(x_n)_{n\in\mathbb{N}}$. The image sequence under g is also a quasigeodesic sequence converging to p. By the Koebe one quarter theorem, infinitesimal Ω-distances are of the order of the corresponding infinitesimal spherical distance divided by the spherical distance from $\partial\Omega$. It follows that, if α is any Ω-quasigeodesic converging to p, then $g\alpha$ is also an Ω-quasigeodesic converging to p.

Now transfer the lamination Λ to $\mathbb{D}^2$. The set A' of points in $\mathbb{S}^1 = \partial\mathbb{D}^2$, corresponding to A defined above, is dense in $\mathbb{S}^1$. Given any point $z \in \mathbb{D}^2$, it is arbitrarily close to the intersection of two geodesics each joining a pair of points in A', So the same applies to $z \in S$, with geodesics joining points of A. Since r is a quasi-isometry, it follows that there is an upper bound for the distance of any $z \in \Omega$ from the intersection of two Ω-geodesics joining points of A, this upper bound depending only on the quasi-isometry constants of r. Under g these two Ω-geodesics are sent to Ω-quasigeodesics with the same endpoints as the geodesics, and with constants depending only on the constants of g. It follows that there is a u such that, for all $z \in \Omega$, $d_\Omega(z, g(z)) < u$. The size of u can be computed from the quasi-isometry constants of r and g.

Now let $x \in \partial\Omega$, let $y \in \Omega$. Let $\beta : [0, v] \to \Omega$ be the Ω-geodesic from y to $g(y)$ parametrized by pathlength. We have

$$u \geq v = \int_0^v \|\beta'(t)\|\, dt \geq \int_0^v 1/2d(\beta(t), \partial\Omega)\, dt \geq \int_0^v 1/2d(\beta(t), x)\, dt \geq \left|\log\left(\frac{d(x, gy)}{d(x, y)}\right)\right|.$$

This means that $d(x, gy) \leq e^u d(x, y)$, which proves our claim. ∎

We can now prove the following theorem.

17.8 Theorem. *Suppose that $\mathcal{A}$ is a crescent limit in $\mathbb{D}^2$, and that $f : \mathbb{D}^2 \to \Omega$ is the associated angle doubling map. Or, equivalently by Theorem 16.9, suppose that one of the other three conditions found there are satisfied. Then $h = f \circ \iota : \mathrm{Dome}(\Omega) \to \Omega$ is a 2-quasiconformal mapping that extends continuously to fix $\partial\Omega$ pointwise.*

Proof: We plan to apply Theorem 17.7 to $g = f \circ \iota \circ r$. We need to find sufficiently many quasigeodesic sequences with which g commutes. There are several cases to consider.

Suppose P is a flat piece of the surface $S = \mathrm{Dome}(\Omega)$. Let $P' = \mathbb{S}^1 \cap \overline{P}$. Then P' is a finite or countable disjoint union of points and closed arcs.

Suppose that p is in the interior of one of the closed arcs that is a component of P'. One possibility is that there is a neighbourhood of p on which g is the identity. In that case, the hypotheses of Theorem 17.7 are satisfied for g.

If there is no such neighbourhood, then P must be a hyperbolic halfplane. Let γ be the bounding geodesic of the halfplane, and let $\alpha \subset \mathbb{S}^1$ be the remainder of the boundary of P. With appropriate normalizations, r extends continuously to the identity on α, ι extends continuously to α, and the angle doubling map f extends continuously to $\iota(\alpha)$ so as to be the inverse of ι. The last statement can be proved by using finite approximations. Note that ι is defined on compact subsets of S by a limiting process, and λ will meet a large enough compact set.

Now suppose that $p \in \partial\Omega$ is an endpoint of a geodesic $\gamma \in \Lambda_S$. It is not the case that every γ in Λ has a corresponding $\alpha \in \mathcal{A}$ with the same endpoints. The previous paragraph discusses the situation where there is no such α. (The α of the previous paragraph was in the boundary of $\mathbb{D}^2$, not in its interior.) We confine our attention to the situation where there is such an $\alpha \in \mathcal{A}$.

For any point $x \in S$, $r^{-1}(x)$ is either a point or an arc. The latter occurs when x lies on a bending line of S with strictly positive atomic measure. Therefore, if we take any element α of $\mathcal{A}$, its inverse image under ι is a curve on S. The inverse image under r of this curve on S is a curve in Ω. (If α crosses a geodesic in Λ at a point, that point is replaced by an interval. If α

is tangent at a point x in $\mathbb{D}^2$ to a geodesic $\gamma \in \Lambda$, then we use only a single point of the (possible) interval $r^{-1}(x)$ when constructing the curve $r^{-1}(\alpha)$.

We return to the situation above, where some $\alpha \in \mathcal{A}$ converges to p. By Lemma 17.6, $\alpha' = r^{-1}(\alpha)$ is a quasigeodesic in Ω which also converges to p. Then $g(\alpha')$ is a quasigeodesic in Ω which converges to p.

Now take any sequence $(x_i)_{i\in\mathbb{N}}$ of points of $S = \mathrm{Dome}(\Omega)$ converging to a point $p \in \partial\Omega$. We choose a sequence $(y_i)_{i\in\mathbb{N}}$ in Ω, so that $r(y_i) = x_i$. The sequence $(y_i)_{i\in\mathbb{N}}$ in Ω must converge to p. For otherwise there is a subsequence which converges to some $q \in \overline{\Omega}$, where $q \neq p$. Applying r, we find that $p = r(q)$. Therefore $q \in \partial\Omega$ and $p = q$, a contradiction.

It follows from Theorem 17.7 that g extended by the identity map on $\partial\Omega$ is continuous. So $(g(y_i))_{i\in\mathbb{N}} = (h(x_i)_{i\in\mathbb{N}}$ converges to p. Theorem 17.8 follows. ∎

We also have the following more standard result.

17.9 Theorem. *Let $f : \mathbb{D}^2 \to \mathbb{D}^2$ be a quasi-isometry, and let Λ be a geodesic lamination on $\mathbb{D}^2$. Suppose that, as x tends to and point $y \in \mathbb{S}^1$ along either a closed flat piece or along a geodesic of Λ, $f(x)$ also tends to y. Then the natural extension of f to a homeomorphism of $\mathbb{S}^1$ is the identity map.*

Proof: It is well-known that f extends to a homeomorphism of $\mathbb{S}^1$. This homeomorphism is the identity on endpoints of geodesics and on the limit point of any flat piece. This means that $f|\mathbb{S}^1$ is the identity on a dense subset. Therefore it is the identity throughout $\mathbb{S}^1$. ∎

18 Continuous extension to the boundary

The main result of this section is the following.

18.1 Theorem. *Let $\Omega \subset \mathbb{S}^2$ be a simply connected region satisfying the equivalent conditions of Theorem 16.9. Then the isometry $\iota^{-1} : \mathbb{D}^2 \to \mathrm{Dome}(\Omega)$ extends to a continuous map $\overline{\mathbb{D}^2} \to \overline{\mathbb{D}^3}$. Let $f : \mathbb{D}^2 \to \Omega$ be an angle doubling map. Then f extends continuously to the boundary $\mathbb{S}^1$. In particular, $\partial\Omega$ is locally connected.*

We will carry out our proof in the ball model $\mathbb{D}^3$ of $\mathbb{H}^3$. In the proof, d will denote euclidean distance and $\mathrm{diam}(X)$ the euclidean diameter of a set X. We will denote by $I : \mathbb{D}^2 \to \mathrm{Dome}(\Omega)$ the (hyperbolic) isometry ι^{-1} normalized as follows. Regarding $\mathbb{D}^2$ as the equatorial plane in $\mathbb{D}^3$, I fixes the vertical diameter $[-i, i] \in \mathbb{D}^2$. By assumption, this geodesic is assumed to be an element of the bending lamination Λ, or to be disjoint from the bending lamination in $\mathbb{D}^2$. I is further normalized so that the limit support plane of the convex hull $\mathcal{CH}(\mathbb{S}^2 \setminus \Omega)$ on the left of the geodesic $[-i, i]$ is equal to the equatorial plane $\mathbb{D}^2$. We adopt a convention for bending that ensures that the convex hull lies above $\mathbb{D}^2$.

We first note the following result.

18.2 Lemma. *Let $(\ell_n)_{n\in\mathbb{N}}$ be a sequence of geodesics in Λ. Then $\mathrm{diam}(\ell_n)$ converges to 0 if and only if $\mathrm{diam}(I(\ell_n))$ converges to 0.*

Proof: In both $\mathbb{D}^2$ and $\mathrm{Dome}(\Omega)$, a geodesic is very small if and only if it is disjoint from a large compact set. But compact sets in $\mathbb{D}^2$ and $\mathrm{Dome}(\Omega)$ are in one-to-one correspondence since $\iota : \mathrm{Dome}(\Omega) \to \mathbb{D}^2$ is a homeomorphism.

∎

The main lemma required is as follows.

18.3 Lemma. *We assume the hypotheses of Theorem 18.1. Suppose that $\{\ell_n\} \subset \Lambda$ is a sequence of leaves for which $\lim \mathrm{diam}(\ell_n) = 0$. Let $L_n \subset \mathbb{D}^2$ be the halfplane not containing $0 \in \mathbb{D}^2$ and bounded by ℓ_n. Then $\lim \mathrm{diam}(I(L_n)) = 0$.*

Proof: Suppose to the contrary there is a sequence $(\ell_n)_{n\in\mathbb{N}}$ for which the lemma is false. We will prove a contradiction. We denote by X_n the space $\mathbb{S}^2 \cap \overline{I(L_n)}$.

We may assume that all the endpoints of the ℓ_n are distinct, for otherwise we may take a subsequence where there is a common endpoint. From our assumptions, the total bending at this common endpoint is at most π. This is sufficiently explicit so that we can check the conclusion of our lemma directly. (It might help to temporarily conjugate to the situation where the common endpoint is the point at infinity in the upper halfspace model.) We will assume not only that all the endpoints are distinct, but also that none of them is equal to $\pm i$, the points which were used to normalize I.

Our normalization ensures that $\mathcal{CH}(\mathbb{S}^2 \setminus \Omega)$ lies in the upper half of $\mathbb{D}^3$. Therefore the smallest disk on $\mathbb{S}^2$ that contains X_n is contained in the northern hemisphere. The hyperbolic convex hull of this disk contains L_n, and so $\operatorname{diam}(X_n) > \operatorname{diam}(I(L_n))/2$. Going to a subsequence, we may assume that there is an $m > 0$ such that, for all n, $\operatorname{diam}(X_n) \geq 2m > 0$.

Let H_n denote an enveloping halfplane bounded by $I(\ell_n)$ (see Definition 16.8. Recall that $H_n \subset \mathcal{CH}(\mathbb{S}^2 \setminus \Omega)$, So that H_n lies in the upper half of $\mathbb{D}^3$. If there are two such halfplanes we choose the one nearest to $I(L_n) \subset \operatorname{Dome}(\Omega)$ (see Definition 16.8 for the meaning of "nearest" in this context).

The circular arc $c_n = \partial H_n \cap \mathbb{S}^2$ has as the same endpoints as $I(\ell_n)$ and lies in the upper half of $\mathbb{D}^3$. The arc c_n separates $\mathbb{S}^2 \setminus \Omega$ into two components, whose closures contain $\pm i$ and $\mathbb{S}^2 \cap \overline{I(L_n)}$ respectively. The two endpoints of c_n are very close together in the northern hemisphere. The arc that completes c_n to a circle C_n lies in Ω, and may not lie completely in the northern hemisphere. This arc is very short. One way to get a picture of the situation is to use stereographic projection of $\mathbb{S}^2$ from the south pole to the plane. C_n encloses a disk (mostly) in the northern hemisphere which contains X_n.

Therefore $\operatorname{diam}(c_n) \geq m$. Passing to a subsequence we may assume that $C = \lim c_n = \lim C_n$ exists. We may also assume that ζ, the limit of the endpoints of $I(\ell_n)$, exists. C is a circle in the northern hemisphere which is disjoint from Ω. Since Ω contains the southern hemisphere, the disk in the northern hemisphere bounded by C is also disjoint from Ω and therefore from $\partial\Omega$. We have $\operatorname{diam}(C) \geq m$.

Since X_n is connected, it must lie in a small neighbourhood of the circle C. Any bending line in L_n has smaller diameter than ℓ_n. By Lemma 18.2, for each bending line $\ell \subset L_n$, $I(\ell)$ has small diameter. It follows that, for any flat piece $P \subset L_n$, $I(P)$ is contained in a small neighbourhood of C in $\mathbb{R}^3$. Therefore $I(L_n)$ is contained in a small neighbourhood of C.

Each bending line $\ell \subset L_n$ cuts L_n into two pieces. Let F_ℓ be the piece which is further from ℓ_n in the pathmetric on Dome(Ω). Let K_ℓ be the enveloping halfspace bounded by ℓ which is nearest to F_ℓ. By Lemma 16.10, $H_n \cap K_\ell = \emptyset$. Also, if ℓ' is another bending line in L_n, then $K_\ell \cap K_{\ell'} = \emptyset$.

Take a point $p \in L_n$ with $d(p, \zeta) = m/3$. A hyperbolic geodesic γ for the pathmetric on Dome(Ω) from p to a point of ℓ_n passes through many, possibly infinitely many, bending lines of L_n. Now consider the set of bending lines $\ell \subset L_n$ which meet γ, and with $m/8 \geq d(\ell, \zeta) \geq m/16$. (Remember that ℓ is very small with respect to euclidean coordinates.) Let $c_\ell = \mathbb{S}^2 \cap \overline{K_\ell}$. We can extend c_ℓ to a complete circle C_ℓ by adding a small arc of circle in Ω.

Now K_ℓ separates ℓ_n from F_ℓ in the convex hull of $\mathbb{S}^2 \setminus \Omega$, and so the C_ℓ separates the endpoints of ℓ_n from $\mathbb{S}^2 \cap \overline{F_\ell}$ in $\mathbb{S}^2$. By Lemma 16.10, c_ℓ can only meet c_n tangentially, if at all. So C_ℓ must be a circle which is very close to C_n and to C.

We are near to obtaining a contradiction. The details are explained most easily if we move to the upper halfspace model, with ζ at infinity. C becomes the x-axis. In $\overline{\mathbb{D}^3}$, we had $m/8 \geq d(\ell, \zeta) \geq m/16$. In the current situation, we can assume that such geodesics ℓ lie near the interval $[0, 1] \subset C$ and that, for a fixed n, every point in $[0, 1]$ is very near some geodesic $\ell \subset I(L_n)$. C_ℓ is a straight line or circle of very large radius which is almost equal to $[0, 1]$ near $[0, 1]$.

We take a finite chain of maximal disks, each containing at least one of the geodesics ℓ which we are working with. (These are the bending lines which are crossed by the geodesic γ and are very near $[0, 1]$.) The maximal disks are the boundaries of support planes for Ω. We choose them so that the support planes meet γ along a sequence of points whose successive distances from each other are extremely small. Also, whenever there is atomic bending along γ which is greater than some very small angle, we use two support planes, one on each side of the atom. Given two successive maximal disks, they meet at two points. We assume that our maximal disks are chosen so that the interval joining these two points is almost horizontal.

Now consider three successive maximal disks D_1, D_2 and D_3 in our chain. The two intervals resulting from the intersections as at the end of the preceding paragraph are both nearly horizontal. A point that is rather easy to overlook is that D_1 and D_3 cannot lie "on the same side" of D_2. That is the geodesic γ does not make a U-turn in D_2. The reason is that the two corresponding enveloping halfplanes are disjoint by Lemma 16.10. The path connecting successive centres of maximal disks is almost vertical. However, it needs to travel horizontally from near 0 to near 1. This contradiction proves Lemma 18.3. ∎

Proof (of Theorem 18.1): It suffices to prove that the isometry $I : \mathbb{D}^2 \to P$ has a continuous extension to $\partial\mathbb{D}^2$. For this implies that $\partial P = \partial\Omega$ is a parameterized curve and is hence locally connected. The angle doubling map $f : \mathbb{D}^2 \to \Omega$ then has a continuous extension as well.

We claim that given $\varepsilon > 0$ there exists $\delta > 0$ such that for every $z, w \in \mathbb{D}^2$ with $d(z, w) < \delta$, $d(I(z), I(w)) < \varepsilon$. That is, I is uniformly continuous on $\mathbb{D}^2$ with respect to the euclidean metric on $\mathbb{D}^3$. Once we prove this, our theorem is an immediate consequence.

According to Lemma 18.3, given $\varepsilon > 0$ there exists $\delta > 0$ such that for any leaf $\ell \subset \Lambda$ bounding a halfplane $L \subset \mathbb{D}^2$ with $\operatorname{diam}(L) < \delta$, the diameter of the image $\operatorname{diam}(I(L)) < \varepsilon$. We will use the set $\mathfrak{L} = \{L : \operatorname{diam}(L) < \delta\}$. The set $\mathfrak{L}$ is a countable or finite set of mutually disjoint closed halfplanes. When two points z, w lie in an element of $\mathfrak{L}$, the pair satisfies $d(z, w) < \delta$ and $d(I(z), I(w)) < \varepsilon$.

Form a sublamination $\Lambda_\delta \subset \Lambda$ by taking only those leaves with diameter $\geq \delta$. Set $V_\delta = \mathbb{D}^2 \setminus \bigcup \mathfrak{L}$; V_δ is a connected, hyperbolically convex set. Let I_δ be the isometry $\mathbb{D}^2 \to \mathbb{D}^3$ determined by the bending lamination (Λ_δ, μ). Then $I_\delta(z) = I(z)$ for $z \in V_\delta$.

We first show that I_δ extends continuously to a map $\overline{\mathbb{D}^2} \to \overline{\mathbb{D}^3}$. We need only show that, for any sequence $(x_i)_{i\in\mathbb{N}}$ in $\mathbb{D}^2$, converging to $x \in \mathbb{S}^1$, the sequence $(I_\delta(x_i)_i)_{i\in\mathbb{N}}$ converges to a point $y \in \partial\Omega$. If x lies in the interior of a boundary interval of a flat piece, then, since the flat piece is mapped by one fixed Möbius transformation, the result is clear. So we assume x is not such a point. Since we are using Λ_δ instead of Λ, x must be the endpoint of a bending line ℓ.

There is a Möbius transformation E_ℓ such that, for $z \in \ell$, $I(z) = E_\ell(z)$. We claim that $(I_\delta(x_i)_i)_{i\in\mathbb{N}}$ converges to $y = E_\ell(x)$. Standard reasoning shows that it is sufficient to prove that every subsequence has a further subsequence converging to y.

If the sequence eventually lies in a fixed flat piece or geodesic P, then P must be at zero hyperbolic distance from ℓ. Since I and I_ℓ preserve the lengths of rectifiable paths, our claim is then obvious.

The only other possibility is that there is a sequence P_i of distinct flat pieces or geodesics, such that $x_i \in P_i$. Since we are working in Λ_δ, $\overline{P_i}$

converges to $\overline{\ell}$ in the Hausdorff topology. Let E_i be a Möbius transformation such that $E_i|P_i = I|P_i$. Let B be a compact ball in $\mathbb{D}^2$, centred at 0, and meeting each geodesic in Λ_δ in an interval of hyperbolic length at least one. Working in B, we see that $E_i|\ell$ converges to $E_\ell|\ell$. Taking a subsequence, we may assume that E_i converges to a Möbius transformation E', where $E'|\ell = E_\ell|\ell$.

By continuity for the euclidean distance, $I(x_i) = E_i(x_i)$ converges to $E'(x) = E_\ell(x) = y$. This completes the proof of our claim. In particular, we now know that $I|\overline{V_\delta}$ is uniformly continuous.

Now we can finish the proof that I is uniformly continuous. We know that for any z, w in an element of $\mathfrak{L}$, $d(I(z), I(w)) < \varepsilon$. There is a $\delta_1 \leq \delta$ such that $d(I(z), I(w)) < \varepsilon$ when $z, w \in V_\delta$ satisfy $d(z, w) < \delta_1$.

Finally suppose $d(z, w) < \delta_1$ where z lies in a component of $\mathfrak{L}$ but w lies in a different component, or lies in V_δ. The geodesic $[z, w]$ is a union of two or three segments one of which lies in V_δ. Therefore $d(I(z), I(w)) < 3\varepsilon$, at most. This completes the proof of uniform continuity. It also shows that $\partial\Omega$ is the continuous image of a circle and is therefore locally connected.

We have now proved the first part of the statement of Theorem 18.1. To prove the same result about the angle doubling map $f : \mathbb{D}^2 \to \Omega$, we use the fact that f is 2-quasiconformal. Since $\partial\Omega$ is locally connected, f can be continuously extended to the closed disk. A proof from first principles of this can be given using prime end theory—prime ends are particularly well-behaved when the boundary is locally connected. Alternatively, if one is prepared to assume this result for the Riemann mapping $h : \mathbb{D}^2 \to \Omega$, we can write $f = h \circ (h^{-1} \circ f)$. The extension result is standard for the quasiconformal homeomorphism $h^{-1} \circ f : \mathbb{D}^2 \to \mathbb{D}^2$. ∎

19 The strong ortho-condition and Jordan curves

In this section, we will prove the following result.

19.1 Theorem. *Let Ω be a simply connected region in the plane satisfying one of the four equivalent conditions of Theorem 16.9, and suppose $\partial\Omega$ is not a Jordan curve. Then there is a ladder in the corresponding bending lamination Λ, and a geodesic γ in the convex bent surface S which is orthogonal to all the bending lines in the ladder. Furthermore the ladder satisfies the ortho-condition but not the strong ortho-condition. It follows that if the strong ortho-condition is satisfied, then $\partial\Omega$ is a Jordan curve.*

Proof: We know from Theorem 18.1 that the angle doubling map $f : \mathbb{D}^2 \to \Omega$ extends to a continuous map on the closed disk, also denoted by $f : \overline{\mathbb{D}^2} \to \overline{\Omega}$. So, if $\partial\Omega$ is not a Jordan curve, there are distinct points, $x, y \in \mathbb{S}^1$, such that $f(x) = f(y)$. Let γ be a geodesic from x to y. Without loss of generality, we may assume that all geodesics in the the bending lamination Λ are crossed by γ. (This might give us a lamination which is not closed in $\mathbb{D}^2$, in which case we include the one or two geodesics in the closure that had previously been omitted, assigning them zero atomic measure.)

For each bending line ℓ, the boundary c_ℓ of the corresponding enveloping halfplane H_ℓ has to pass through $f(x) = f(y)$. It follows that there is a unique enveloping halfplane for each bending line. (In general, there may be two.) By Lemma 16.10, the arcs of circle c_ℓ all pass through $f(x)$ with a common tangent.

We look at the situation in the upper halfspace model, with $f(x)$ being the point at infinity. The c_ℓ become disjoint straight segments of lines. Since they have a common tangent at infinity, the straight lines are parallel, and we take them to be vertical.

Since c_ℓ meets $\partial\Omega$ only at the endpoints of ℓ, we see that $\partial\Omega$ consists of two curves Ω_- and Ω_+, with Ω_- below Ω_+. Then c_ℓ starts at Ω_+, travels upwards, goes through ∞ at the top of the picture and re-emerges at the bottom of the picture, when it travels upwards to Ω_-. The two angles with which c_ℓ meets Ω_+ are equal to the two angles with which c_ℓ meets Ω_-.

We fix some c_ℓ and call it the y-axis. We take the origin to be the point halfway between the two intersections with $\partial\Omega$. Let Ω_+ be the graph of a function $g_+ : \mathbb{R} \to \mathbb{R}$ and similarly for g_-. These two functions have well-defined left and right derivatives at all points of the domain. The statement that the angles are the same at Ω_+ and Ω_- means that the left and right derivatives for g_+ and g_- are equal, but with opposite sign.

It follows that $g_+ + g_- : \mathbb{R} \to \mathbb{R}$ is a continuous function whose left and right derivatives exist everywhere and are zero. It follows that the function is identically zero. This proves that the vertical half plane above the x-axis is orthogonal to each bending line.

The ortho-curve is the intersection of $\mathrm{Dome}(\Omega)$ with the vertical halfplane above the x-axis. The two ends of the ortho-curve are equal at infinity. It follows that the strong ortho-condition is not satisfied. ∎

20 Angle-scaling mappings

In Definition 16.2, we discussed a map $\Phi_{\Lambda,\nu} : \mathbb{H}^2 \to \mathbb{H}^3$ which sends rectifiable paths to rectifiable paths of the same length, in the case where ν is a complex valued measure. In this section, we will denote the same map by $\Phi(\Lambda,\nu)$, and, by abuse of notation, we will denote the image of this map also by $\Phi(\Lambda,\nu)$, when the map is an embedding.

Let μ be a real-valued, non-negative measure and take $\nu = i\mu$. To normalize the construction, in say the ball model of $\mathbb{H}^3$, we may fix the equatorial plane P to represent $\mathbb{D}^2$. Then positive bending will be bending in the upper halfball. The pleated surface will then be locally convex (toward the north for small amounts of bending). We will normalize S by choosing a leaf $\ell \subset \Lambda$ to lie in P such that P itself becomes the medial support plane at ℓ. See [7] for details. If $\Phi(\Lambda,\nu)$ is an embedding, its image S is the dome over a region Ω containing the lower hemisphere.

Let $S' = \Phi(\Lambda,-\nu)$. For all real Borel measures μ, S' is the image of S under reflection in P. S is embedded if and only if S' is embedded, and in this case, S lies in the upper half ball, and S' in the lower half ball. If Ω' is the reflection of Ω under P, then S' is the dome over Ω'.

Given $t \in \mathbb{C}$, we define $\Phi_t(\Lambda,\mu) = \Phi(\Lambda,t\mu)$. The surfaces satisfy the symmetry that the surface Φ_{u-iv} is the reflection of Φ_{u+iv} in the equatorial plane P, since we first apply the earthquake corresponding to u, and then bend either upwards or downwards.

20.1 The bending slice of a measured lamination All the measured laminations we use in this section are bounded (see Definition 16.3). In fact there are only rare occasions anywhere in this paper when unbounded earthquakes are used, and we always point these out.

20.2 Proposition. . *There is a universal $\epsilon > 0$, with the property that, if (Λ,μ) is any measured lamination of the hyperbolic plane, such that, for every transverse interval J of length one, $\mu(J) \le \epsilon$—then $\Phi(\Lambda,\mu)$ is an embedding to a dome of a simply connected region.*

We first prove the following lemma.

20.3 Lemma. *Let γ be a piecewise geodesic in $\mathbb{H}^3$ of length L, such that the sum of the bending angles along each subinterval of length 1 is at most $\epsilon > 0$. Suppose $\theta_0 > 0$ and $\sin(\theta_0) > \epsilon$ and $\theta_0 + \epsilon < \pi/2$. Then the endpoints of γ*

are at least $L\cos(\epsilon+\theta_0)$ *apart. (Explicit possible values are* $\theta_0 = \pi/4$ *and* $\epsilon = 1/2$*, in which case* $\cos(\epsilon+\theta_0) \sim 0.28154$*. In fact, this method will prove Proposition 20.2 with* $\epsilon = .739$*.)*

Proof: Let $\alpha : \mathbb{R} \to \mathbb{D}^2$ be a geodesic parametrized by hyperbolic arclength. Let p be a point in $\mathbb{S}^1$. Suppose that the angle that oriented geodesic from p through $\alpha(s)$ makes an angle $\theta(s)$ with $\alpha'(s)$, with $0 < \theta(s) < \pi/2$. Then $\theta'(s) = \sin(\theta(s))$. (Compare the proof of Lemma 14.2.)

Consider a closed subinterval $\alpha[s_1, s_2]$, such that, for $s_1 \leq s \leq s_2$, $0 < \theta_0 \leq \theta(s) \leq \theta_0 + \epsilon < \pi/2$. Then $\theta(s_1) - \theta(s_2) > \sin(\theta_0)(s_2 - s_1)$. We have strict inequality because $\theta(s)$ is not constant on $[s_1, s_2]$.

If p is a finite point, rather than a point at infinity, then the inequality continues to hold. In fact, if $\theta(s_1)$ remains the same, then $\theta(s_2)$ becomes smaller as a result of moving p from a point at infinity to a finite point.

Now suppose we have a piecewise geodesic $\beta : [0,\infty] \to \mathbb{H}^3$, parametrized by pathlength, such that, on any subinterval of length 1, the sum of bending angles is at most ϵ. Let $p = \beta(0)$. We define $\theta(s)$ as the angle which the geodesic ray from p to $\beta(s)$ makes with the tangent vector $\beta'(s)$. This is not well-defined at a corner point of β, where we use a limiting value either from the left or the right, as seems appropriate. The initial segment of β has $\theta(s) = 0$.

Take a maximal subinterval $[s_1, s_2]$, such that if $s \in [s_1, s_2]$ is a smooth point of β, then $\theta(s) \geq \theta_0$. We claim that $s_2 < s_1 + 1$. To see this, we suppose $s_2 \geq s_1 + 1$ and prove a contradiction.

The initial segment of β has $\theta(s) = 0$, so our maximal interval is disjoint from the initial segment. Using the fact that $\theta(s)$ is strictly decreasing if it is positive, we see that there is an interval (s_0, s_1) on which $\theta(s) < \theta_0$.

Now $\theta'(s) \leq -\sin(\theta_0)$ on $[s_1, s_2]$. Also, the total bending angle on $[s_1, s_1 + 1]$ is at most ϵ. As a third ingredient, note that $\mathbb{S}^2$ has a metric in which the distance between any two points is equal to the angle subtended at the centre. It follows that such angles satisfy the triangle inequality, and it is therefore justified to add angles while deducing the following inequality. We have, for $0 < s < 1$,

$$\theta(s_1 + s) \leq \theta_0 + \epsilon - s\sin(\theta_0). \tag{20.3.9}$$

For s near 1, this gives a contradiction, so we now know that $s_2 < s_1 + 1$. From Inequality 20.3.9, we see that, for $s \in [s_1, s_2]$, $\theta_0 \leq \theta(s) \leq \theta_0 + \epsilon$.

It now follows that, for all $s \in [0,\infty)$, $0 \leq \theta(s) \leq \theta_0 + \epsilon$. ∎

Proof (of Proposition 20.2): Suppose that $\Phi = \Phi(\Lambda, \mu)$ is not an embedding. Then there must be two points in $\mathbb{H}^2$ which are sent to the same point of $\mathbb{H}^3$. Let γ be a geodesic arc in $\mathbb{H}^2$ joining the two points. The endpoints of γ cannot lie on the same bending line, since the bending line is embedded in $\mathbb{H}^3$. In fact, no segment of γ can lie on a bending line. When we approximate the measured lamination by a finite measured lamination, we leave $\gamma \subset \mathbb{H}^2$ fixed, erase Λ and replace it by a new finite measured lamination, which, by abuse of notation, we will continue to call Λ. The image $\Phi\gamma$ changes from a closed loop in $\mathbb{H}^3$ to a path that goes some distance from its starting point, and then returns to a point very close to its starting point. We will work with this approximation.

If a subinterval of γ coincides with a bending line, then the whole of γ must lie on a bending line. But then $\Phi\gamma$ is a geodesic in $\mathbb{H}^3$ which is impossible. Let p be a point where γ crosses a bending line ℓ. Let P_1 and P_2 be the two flat pieces of the approximation on either side of ℓ and let $\gamma_1 \subset P_1$ and $\gamma_2 \subset P_2$ be small subintervals of γ with the same length, each having p as an endpoint. Then γ_1 and γ_2 meet ℓ at p at the same angle.

Let $\ell^\perp$ be the line through $\Phi(p)$ perpendicular to $\Phi(\ell)$, exactly half way between $\Phi(P_1)$ and $\Phi(P_2)$. Rotation about $\ell^\perp$ through an angle π interchanges $\Phi(P_1)$ and $\Phi(P_2)$. Since the angles are the same, it also interchanges $\Phi(\gamma_1)$ and $\Phi(\gamma_2)$. It follows that these subintervals lie in a plane containing $\ell^\perp$. Therefore the bending angle for the curve $\Phi\gamma$ at this point is less than or equal to $\mu(\ell)$.

Lemma 20.3 now shows both that Φ is an embedding, and that it is proper. It therefore separates $\mathbb{H}^3$ into two components, one of which is a convex subspace Y (see Lemma 4.6). The boundary of Y in $\mathbb{H}^3$ is the image of Φ. It follows that $\mathbb{S}^2 \setminus \overline{Y}$ is a topological disk. (For example, we can project radially from a point in the interior of Y to obtain a homeomorphism with the boundary of Y.) ∎

Given a measured lamination (Λ, μ) in $\mathbb{D}^2$, let $\mathcal{D} = \mathcal{D}(\Lambda, \mu)$ be the set of $t \in \mathbb{C}$ such that $\Phi_t(\Lambda, \mu) : \mathbb{H}^2 \to \mathbb{H}^3$ gives the dome over some simply connected region $\Omega_t \subset \mathbb{S}^2$. $\mathcal{D}(\Lambda, \mu)$ is called the (Λ, μ)-*bending slice*. It is a closed subset of $\mathbb{S}^2$ and is symmetric under the reflection $t \mapsto \bar{t}$ if the pleated surfaces are suitably normalized. This is because of the symmetry between positive and negative bending.

It follows from Proposition 20.2 that $\mathcal{D}$ contains a neighbourhood of the real line. It is possible that there is no $\epsilon > 0$ such that it contains an

ϵ-neighbourhood of the real line.

20.4 Questions. Is $\mathcal{D}(\Lambda, \mu)$ necessarily connected? What is the relationship between the subset of t such that Ω_t is a quasidisk and the bending slice itself? Is it an open and dense subspace of the bending slice? If $\partial\Omega = \partial\Omega_i$ is not locally connected, can $\partial\Omega_t$ be locally connected for some value of $t \in \mathcal{D}(\Lambda, \mu)$ with $\Im t \neq 0$. □

20.5 Angle scaling and grafting. We recall from Definition 5.1 that angle scaling with factor $(1+t)$ can be defined for any finite union of disjoint crescents, where $t = u + iv \in \mathbb{C}$.

In order to fix the notation, we repeat the definitions. We return again to our model wedge of angle α, $W = \{z : 0 < \arg z < \alpha\}$. For any $t = u+iv \in \mathbb{C}$

$$z \mapsto ze^{it\theta} \tag{20.5.10}$$

is a well-defined map $W \to \mathbb{C}$. If $-1 < v \leq 1$, the map is a K_v-quasiconformal map from W to a new wedge of vertex angle $(1+v)\alpha$. Here

$$K_v = \max\left\{(1+v), \frac{1}{1+v}\right\}.$$

The limiting case $v = -1$ collapses W while $v = 1$ corresponds to angle doubling.

Consider also the reflection W' of W in the imaginary axis $W' = \{z : \pi - \alpha < \arg z < \pi\}$. Using the angle φ in W' the formula is $z = r \mapsto ze^{-it\varphi}$.

Note that that combined mapping on $W \cup W'$ is the identity on one pair of symmetric edges. On the other pair of symmetric edges (obtained by reflection in the imaginary axis) we have the map $z \mapsto ze^{t\alpha}$ Denote the combined mapping of $W \cup W'$ by B_t.

We denote by $\mathcal{E}\mathbb{D}^2$ the exterior of $\mathbb{D}^2$, namely the open disk $\mathbb{S}^2 \setminus \overline{\mathbb{D}^2}$. Let $J : \mathbb{S}^2 \to \mathbb{S}^2$ be given by inversion in $\mathbb{S}^1$.

We will now extend the above angle scaling maps to the whole of $\mathbb{S}^2$ when working with a finite lamination (Λ, μ) satisfying the ortho-condition. We can insert associated crescents in $\mathbb{D}^2$ and, by the inversion J, in $\mathcal{E}\mathbb{D}^2$. To each pair of these apply a suitable conjugate of B_t as above. In each gap we use a Möbius transformation, putting the maps together so that we obtain a continuous map $A_t : \mathbb{S}^2 \to \mathbb{S}^2$ as was done for angle doubling maps in the lead up to Theorem 5.3. We can check that $A_t = J \circ A_{-t} \circ J$.

When $-1 < v < 1$, A_t is a homeomorphism of $\mathbb{S}^2$. When $1 > v \geq 0$, the dome over $\Omega_t = A_t(\mathbb{D}^2)$ is found by first earthquaking the flat hyperbolic plane along Λ using the measure $u\mu$, and then bending, using the measure $v\mu$. The dome of Ω_t can be thought of as a negatively bent (and architecturally unsound) dome over the simply connected region $A_t(\mathcal{E}\mathbb{D}^2)$. It is associated with the earthquake $-uJ\mu$ on $J\Lambda$ in $\mathcal{E}\mathbb{D}^2$, followed by bending $-vJ\mu$ on $J\Lambda$.

When $0 \geq v > -1$, the bending can be thought of as negative over $\mathbb{D}^2$, and positive over $\mathcal{E}\mathbb{D}^2$.

Most of the preceding discussion can be generalized to the case of a measured lamination satisfying the ortho-condition, even when the lamination is not finite. On both $\mathbb{D}^2$, and on $\mathcal{E}\mathbb{D}^2$, a sequence of finite approximations (Λ_n, μ_n) converging to (Λ, μ) will lead to the sequence of bending maps $A_{n,t}$ converging to a limit A_t. The convergence is uniform on compact subsets of $(\mathbb{S}^2 \setminus \mathbb{S}^1) \times \mathbb{C}$, as we see from **??**. Since the function is analytic as a function of t for each n, and the convergence is uniform, $A_t(z)$ is an analytic function of t for $z \in (\mathbb{S}^2 \setminus \mathbb{S}^1)$.

We summarize our discussion in the following lemma.

20.6 Lemma. *Suppose (Λ, μ) is a measured lamination satisfying the ortho-condition. For all $t = u + iv \in \mathbb{C}$, we have a map $A_t : (\mathbb{S}^2 \setminus \mathbb{S}^1) \to \mathbb{S}$. As a function of two variables (t, z), with $z \in (\mathbb{S}^2 \setminus \mathbb{S}^1)$, b defines a continuous function. If z is fixed, then $t \mapsto A_t(z)$ is an analytic function of t.*

When $-1 < v < 1$, there is a K_v-quasiconformal homeomorphism $A_t : \mathbb{S}^2 \to \mathbb{S}^2$ determined by the crescent limit $\mathcal{A}$, and its reflection. The values of A_t on $\mathbb{S}^1$ do not depend on the particular crescent limit $\mathcal{A}$, but only on the measured lamination (Λ, μ).

In the limiting case $v = 1$, A_t is the 2-quasiconformal angle doubling map $\mathbb{D}^2 \to \mathbb{S}^2$ and the angle collapsing map in $\mathcal{E}\mathbb{D}^2$. When $v = -1$, A_t is equal to angle collapsing in $\mathbb{D}^2$ and to angle doubling in $\mathcal{E}\mathbb{D}^2$.

Proof: Most of the above statements have already been proved. To show that A_t is defined on all of $\mathbb{S}^2$ when $-1 < v < 1$, we use the fact that every sequence of K_v-quasiconformal homeomorphisms of $\mathbb{S}^2$, normalized so that the i, $-i$ and 0 are fixed, ∎

We now generalize our work still further, so as to incorporate the process for changing conformal structures on surfaces known as *grafting*. In $\mathbb{D}^2$ (or on a Riemann suface covered by $\mathbb{D}^2$) grafting can be described as follows.

Cut $\mathbb{D}^2$ along a hyperbolic geodesic ℓ (which descends to a simple closed geodesic in the Riemann surface case), insert a crescent of angle α, and sew back together. When this is applied to the finite lamination (Λ, μ), with bending angles as assigned by μ, the effect on $\partial\mathbb{D}^2$ is the same as angle doubling on any inserted crescents corresponding to (Λ, μ). Inserting angles at a constant ratio to the angles given by the measure, we obtain angle scaling. The natural structure obtained is a geometric structure modelled on $\mathbb{S}^2$, with group $\mathrm{PSL}(2, \mathbb{C})$, and we obtain such a structure by grafting in a crescent of any angle $\alpha > 0$, for example $\alpha = 4\pi$.

Grafting is also incorporated into the definition of complex earthquakes given in [12]. Our maps A_t defined above are not the same as these complex earthquakes, but have the same values on $\mathbb{S}^1$.

The following lemma can be proved first in the case of a finite disjoint union of crescents, and then, by a limiting process for any crescent limit. Before stating it, we recall a standard definition.

20.7 Definition. A *quasiregular function* has many properties of a quasiconformal mapping except it is not necessarily a homeomorphism. A K-quasiregular map F is characterized by the fact that it can be factored as $F = G \circ H$ where H is K-quasiconformal and G is holomorphic on the range of H. □

20.8 Lemma. *For any crescent limit $\mathcal{A}$ the angle scaling map $A_t(z) = A_t(z; \mathcal{A})$ is defined for $z \in \overline{\mathbb{D}^2}$ and $t \in \widetilde{R} = \{t \in \mathbb{C},\ \Im t > -1\}$. It is defined in $\overline{\mathcal{E}\mathbb{D}^2}$ for $t \in -\widetilde{R} = \{t : \Im t < 1\}$. It has the following properties.*

20.8.1. For any $t \in \mathbb{C}$, $A_t : \mathbb{S}^2 \to \mathbb{S}^2$ is an angle scaling map with scaling factor $t + 1$.

20.8.2. For each $t \in R = \{t : -1 < \Im t < 1\}$, $A_t : \mathbb{S}^2 \to \mathbb{S}^2$ is a K_t-quasiconfomal homeomorphism where

$$K_t = \frac{1 + |\kappa(t)|}{1 - |\kappa(t)|}, \quad \kappa(t) = \frac{t}{t + 2i}.$$

20.8.3. For each $t \in \widetilde{R}$, A_t is a K_t-quasiregular function and a locally injective developing mapping $\mathbb{D}^2 \to \mathbb{S}^2$.

20.8.4. For $v = 1$ and any $u \in \mathbb{R}$, the restriction of A_t to $\mathbb{D}^2$ is the angle doubling map for the image in the flat $\mathbb{D}^2$ of (Λ, μ) under the earthquake A_u.

20.8.5. For $v = \Im t = 0$, $A_t : \mathbb{D}^2 \to \mathbb{D}^2$ has the boundary values of the $e_{u\mu}$-earthquake of Λ.

20.8.6. For each $z \in \overline{\mathbb{D}^2}$, A_t is holomorphic for $t \in \widetilde{R}$. For each $z \in \mathbb{S}^2$, A_t is holomorphic for $t \in R$.

20.8.7. $A_0(z) = z$ for $z \in \mathbb{S}^2$.

20.8.8. The corresponding statements hold for $z \in \mathcal{E}\mathbb{D}^2$ when $t \in -\widetilde{R}$.

Although we have now captured the measure μ as the limit of crescent angles, unlike the situation when $\Omega = \mathbb{D}^2$, there is no guarantee that under a pure earthquake ${}_t(z)$, $t \in \mathbb{R}$, the image of Ω remains embedded in $\mathbb{C}$. Nor is there a natural way to extend the earthquake to the complement of Ω.

21 The λ Lemma

We recall the λ-lemma, as treated in Mané–Sad–Sullivan [13], Sullivan–Thurston [17], Bers–Royden [3] and Slodkowski [15]. The strongest of these is Slodkowski's, and this is the version we now state.

21.1 Theorem ($\boldsymbol{\lambda}$-Lemma). *Let $B \subset \mathbb{S}^2$ contain at least three points and let $\{f_\lambda(z) \in \mathbb{S}^2\}$ denote a family of functions with arguments $\lambda \in \mathbb{D}^2$ and $z \in B$. Suppose these functions have the following properties.*

- *For each fixed $\lambda \in \mathbb{D}^2$, $f_\lambda : z \in B \mapsto f_\lambda(z) \in \mathbb{S}^2$ is injective.*
- *For each fixed $z \in B$, the map $\lambda \mapsto f_\lambda(z)$ is a complex analytic map $\mathbb{D}^2 \to \mathbb{S}^2$.*
- *For each $z \in B$, $f_0(z) = z$.*

Then

21.1.1. $f_\lambda(z)$ is jointly continuous in λ and z.

21.1.2. For fixed $\lambda \in \mathbb{D}^2$, $f_\lambda(z)$ $(z \in B)$ is the restriction to B of a K^*_λ-quasiconformal mapping $f^*_\lambda : \mathbb{S}^2 \to \mathbb{S}^2$ where

$$K^*_\lambda = \frac{1+|\lambda|}{1-|\lambda|}.$$

21.1.3. With $\lambda \in \mathbb{D}^2$ and $z \in \mathbb{S}^2$, $\{f^*_\lambda(z)\}$ gives a holomorphic motion of $\mathbb{S}^2$.

A family with the three properties of $\{f_\lambda\}$ is called a *holomorphic motion* of B. Note that continuity in z is *not* assumed; instead, continuity is a conclusion. In particular, if $B = \mathbb{D}^2$, then $f_\lambda(B)$ is a quasidisk.

We will show how the λ-Lemma leads to different proofs of some results we have already proved, and leads to a better understanding of the situation.

21.2 Angle-scaling mappings as holomorphic motions. Let $R \subset \mathbb{C}$ be the set of points $t \in \mathbb{C}$ with $-1 < \Im t < 1$.

21.3 Theorem. *Let $\mathcal{A}$ be a crescent limit in $\mathbb{D}^2$.*

21.3.1. For each $t = u + iv \in R$, the angle scaling map $A_t(z) = A_t(z; \Lambda, \mu)$ for $z \in \mathbb{D}^2 \cup \mathcal{E}\mathbb{D}^2$ has a continuous extension to $\partial\mathbb{D}^2$.

21.3.2. *With* $(t, z) \in R \times \mathbb{S}^2$, $\{A_t(z)\}$ *gives a holomorphic motion of* $\mathbb{S}^2$*. In particular* $A_t(z)$ *is continuous for* $(t, z) \in R \times \mathbb{S}^2$*.*

21.3.3. $A_t(z)$ *is (at worst) a* K_t^**-quasiconformal map of* $\mathbb{S}^2$ *where*

$$K_t^* = \frac{1 + |\lambda(t)|}{1 - |\lambda(t)|}, \text{ with } \lambda(t) = \tanh \frac{\pi t}{4} \in \mathbb{D}^2, \;\; t \in R.$$

21.3.4. *For fixed* $z \in \mathbb{D}^2$, $A_t(z)$ *has an analytic continuation to* $t \in \widetilde{R} = \{t \in \mathbb{C} : -1 < \Im t\}$*; for fixed* $z \in \mathcal{E}\mathbb{D}^2$, $A_t(z)$ *has an analytic continuation to* $t \in -\widetilde{R} = \{t : \Im t < 1\}$*.*

Proof: Lemma 20.8 implies that $A_t(z)$ is a holomorphic motion of $B = \mathbb{D}^2 \cup \mathcal{E}\mathbb{D}^2$. According to the λ-Lemma, it is the restriction of a holomorphic motion on $\mathbb{S}^2$. Hence $A_t(z)$ has a continuous extension to $\partial\mathbb{D}^2$ and becomes a quasiconformal map of the whole sphere $\mathbb{S}^2$. The map $z \mapsto \tanh(\pi z/4)$ is the Riemann mapping $R \to \mathbb{D}^2$. Thus the first three items are consequences of the λ-lemma. The fourth item follows from Lemma 20.8 ∎

21.4 Remark. In 21.3.3, we cannot dispense with the assumption that $-1 < \Im t < 1$, even if (Λ, μ) itself satisfies the strong ortho-condition. Consider the following example. In $\mathbb{D}^2$ insert a sequence of mutually disjoint crescents $\{C_n\}$ such that C_n has vertex angle $\pi - \frac{1}{n}$. We may take the crescents to be centered on geodesics with a common orthogonal. Taking $t = (1+\varepsilon)i$ corresponds to taking a bending angle $(1+\varepsilon)(\pi - \frac{1}{n})$ on the n-th geodesic. If n is large, this bending angle is larger than π, and so the angle scaling map $A_t : \mathbb{S}^1 \to \mathbb{S}^2$ will not be an embedding and $A_t : \mathbb{D}^2 \to \mathbb{D}^2$ is also not an embedding, although it is locally injective. □

22 Equivariance

22.1 Suppose $R = \mathbb{D}^2/G$ is a Riemann surface of finite hyperbolic area represented in its universal covering surface by the fuchsian group G. Suppose Λ is a finite geodesic lamination on R with leaves ℓ_i each of which is a simple loop. Choose a set of nonoverlapping tubular neighbourhoods about the geodesic loops which is maximal in the sense that no neighbourhood can be enlarged. A lift of a tubular neighbourhood about ℓ_i is a crescent in $\mathbb{D}^2$ and thereby has a vertex angle β_i. Let Λ^* denote the measured lamination that comes by lifting Λ to $\mathbb{D}^2$ and assigning the angle β_i to all lifts of ℓ_i, for each index i. The leaves of the measured lamination Λ^* are discrete in $\mathbb{D}^2$ and the set of crescents about them, which is also invariant under G, are nonoverlapping and form a maximal set in the sense that none can be enlarged.

Therefore Λ^* satisfies the ortho-condition. In consequence we can apply our theory to Λ^*. In detail,

22.2 Theorem. *Let Λ be a finite lamination of closed leaves $\{\ell_i\}$ on the hyperbolic surface $R = \mathbb{D}^2/G$. Find a maximal set of tubular neighbourhoods about the geodesics and denote the corresponding angles by β_i. Let $\Lambda^* = \{(\ell_i, \beta_i)\}$ be the corresponding G-invariant lifted measured lamination to $\mathbb{D}^2$. There is a 2-quasiconformal angle doubling map $f : \mathbb{D}^2 \to \Omega$ where Ω is simply connected and invariant under a group of Möbius transformations G_1 isomorphic to G. Moreover the 2-quasiconformal map $f \circ \iota : \mathrm{Dome}(\Omega) \to \Omega$ is equivariant under G_1 and extends to be the identity on $\partial\Omega$.*

The following result is a direct consequence of Corollary 17.3.

22.3 Corollary. *Suppose $f : \mathrm{Dome}(\Omega) \to \Omega$ is a K-quasiconformal map that pointwise fixes $\partial\Omega$. Assume f is equivariant under a kleinian group preserving Ω that represents a Riemann surface with finite hyperbolic area. Then in the respective hyperbolic metrics, there are K-Lipschitz maps in the homotopy classes of f and of f^{-1}. In particular if f is associated with an angle doubling mapping, then in the homotopy classes of f and f^{-1} there 2-Lipschitz maps.*

22.4 The equivariant constant Given a simply connected region $\Omega \subset \mathbb{S}^2$, let G_Ω be the group of conformal equivalences of Ω with itself. Let K_{eq} denote the smallest constant, such that, for each simply connected region Ω, there is a K_{eq}-quasiconformal, equivariant map dome$(\Omega) \to \Omega$ which is the identity on $\partial\Omega$.

22.5 Theorem.

$$K_{eq} \leq K^3.$$

Proof: This is an immediate consequence of a result of Lehto [10, Theorem 4.7, p. 201] in Teichmüller theory. He presents it in the sharp infinitesimal form. ∎

In particular if $K \leq 2$ then $K_{eq} \leq 8$. We will see in 22.8 that $K_{eq} > 2$.

22.6 One dimensional Teichmüller spaces. Consider the square once-punctured torus T; this means that there is a pair of simple geodesics crossing each other once, and mutually orthogonal at their point of intersection. Denote these by β and γ.

Let C_γ be the maximal collar about γ, centered on γ. It is bounded by a pair of equidistant lines of hyperbolic distance d from γ.

Consider the totality of lifts of C_γ to the universal cover $\mathbb{D}^2$. This is a maximal set of crescents, each of vertex angle α, where $\cosh d = \sec(\alpha/2)$. The geodesics over γ form into infinite discrete ladders. Each lift of β determines such a ladder for which it is the common orthogonal.

Consider the angle doubling map $A_i : \mathbb{D}^2 \to \Omega_i$. According to Theorem 19.1, $\partial\Omega_i$ is not a Jordan curve. Rather there are a countable number of double points. The two endpoints of each lift of β are mapped to the same point in $\mathbb{S}^2$.

Let G denote the fuchsian group of covering transformations for the square torus, and let Γ be the image group determined by A_i. We can think of the parametrized family A_{iv}, for $0 \leq v \leq 1$. For $v < 1$, we have a quasifuchsian group whose domain of discontinuity consists of the two components $A_{iv}(\mathbb{D}^2)$ and $A_{iv}(\mathcal{E}\mathbb{D}^2)$. The behaviour of A_{iv} on $\mathcal{E}\mathbb{D}^2$ corresponds under inversion in $\mathbb{S}^1$ to A_{-iv} on $\mathbb{D}^2$. As v tends to 1, the component which is the image of $\mathbb{D}^2$ survives as a single component. This shows that the conformal structure at the boundary continues to be a once punctured torus. The image of $\mathcal{E}\mathbb{D}^2$, however, is the result of an angle collapse, and we get countably many complementary components. The collapse corresponds to pinching the geodesic β to a point. Therefore the conformal structure at infinity of the other side is that of a thrice punctured sphere. The kleinian group Γ is a boundary point of the deformation space of G and represents a cusp. The conformal boundary $\partial(\mathbb{H}^3/\Gamma)$ is the union of a once-punctured torus and the 3-punctured sphere.

Return to the discrete geodesic lamination Λ of $\mathbb{D}^2$ consisting of the lifts of γ. We will assign the bending measure μ which is just the atomic measure assigning the bending angle α to each leaf. Correspondingly we have the maximal set of crescents of angle α. With this data, consider the angle scaling maps $A_t : \mathbb{D}^2 \to \mathbb{S}^2$ for $t \in \widetilde{S} = \{t \in \mathbb{C} : -1 < \Im t < +\infty\}$.

When $\Im t$ is sufficiently large, $A_t : \mathbb{D}^2 \to \mathbb{S}^2$ fails to be injective, though it is a locally injective. Furthermore, A_t conjugates the fuchsian covering group G of T to a group Γ_t acting on Ω_t: Ω_t/Γ_t is also a once-punctured torus. Here Ω_t is a simply connected surface with a projective structure. The developing map $\Omega_t \to \mathbb{S}^2$ is injective if $-1 < \Im t \leq 1$ and then Γ_1 is a discrete group of Möbius transformations. But if $\Im t > 1$, then the developing map will not always be injective, and the image of the holonomy homomorphism $\Gamma_t \to \mathrm{PSL}(2, \mathbb{C})$ will not always be discrete.

Conjugation by a Riemann map $\Psi_t : \Omega_t \to \mathbb{D}^2$ will send Γ_t to a fuchsian group Γ'_t. Then $\mathbb{D}^2/\Gamma'_t$ is conformally equivalent to Ω_t/Γ_t. The composition $\Psi_t \circ A_t$ is a K_t-quasiconformal homeomorphism of $\mathbb{D}^2$, which conjugates G to Γ'_t.

Set $L = \mathrm{length}(\gamma)$. Recall that A_t is the scaling map associated to the complex scaling for the measure $t\mu$. Therefore $A_{nL/\alpha}$ is the angle scaling which has the same effect on $\mathbb{S}^1$ as the earthquake along each lift of γ through a distance nL. This corresponds to an n-fold Dehn twist.

Each $A_{nL/\alpha}$ sends T (which corresponds to $t = u + iv = 0$) to another square torus $A_{nL/\alpha}(T)$ which is the same underlying torus but with a new marking. Acting in $\mathbb{D}^2$, A_{nL} sends Λ to $A_{nL}(\Lambda)$. Again, this is the same lamination as Λ, but with a different marking.

For $n \in \mathbb{Z}$, we set $t = (nL/\alpha) - i$. As before, A_t collapses the boundary torus to a thrice punctured sphere (but in a different way).

On the other hand consider the effect of an earthquake A_u with $u \neq nL/\alpha$. In $\mathbb{D}^2$, the leaves $A_u(\Lambda)$ will no longer be organized in ladders with common orthogonals, although there will still be an associated maximal set of crescents of angle α.

The angle scaling map $A_t : \mathbb{D}^2 \to \mathbb{S}^2$, with $t = u + i$, maps $\mathbb{D}^2$ to a Jordan region Ω_t, by Theorem 19.1. Therefore the collapsing map A_{u-i} will not collapse the punctured torus, but map it to another proper once-punctured torus. For these tori, there is a holomorpic extension of $A_t(z)$ to values of $t = u + iv$ with $v < -1$. That is for a while longer, collapsing can be continued to result in further proper tori. The possibility of defining structures for certain values of t below the line $\Im t = -1$ is also proved in [12].

Let $\mathcal{T} \supset \widetilde{S}$ be the maximal connected open set of values of t to which the functions $t \mapsto A_t(z)$ can be simultaneously analytically continued for all $z \in \mathbb{D}^2$. The set $\mathcal{T}$ is invariant under the horizontal translation $t \mapsto t+(L/\alpha)$.

We next cite an important result due to McMullen ([12, Theorem 1.3]).

22.7 Theorem (McMullen's Disk Theorem). *$\mathcal{T}$ is biholomorpically equivalent to the Teichmüller space of once-punctured tori. It is not a euclidean halfplane.*

In particular, the Teichmüller metric on $\mathcal{T}$ coincides with its hyperbolic metric. See [12, Figure 8] for a picture of $\mathcal{T}$. $\mathcal{T}$ is originally defined as a certain subset of $\mathbb{C}$. In the above result, we think of $t \in \mathcal{T}$ also as the structure on the punctured torus given by the homeomorphism $\Psi_t \circ A_t : \mathbb{D}^2 \to \mathbb{D}^2$. We also talk of this point as being given by the angle scaling map $A_t : \mathbb{D}^2 \to \Omega_t$, which is a homeomorphism to a surface with a projective structure. In [12], the corresponding complex earthquake (which is not continuous) is used. The initial square punctured torus T serves as the origin $t = 0$. The fuchsian group G for T is conjugated by A_t to a group Γ_t acting on the manifold Ω_t with a projective structure.

Points $t \in \mathcal{T}$ with $\Im t < 0$ correspond to negative bending of the equatorial plane $\mathbb{D}^2 \subset \mathbb{D}^3$. Bending beyond an angle π along the simple closed geodesic γ is not possible. It follows that we cannot have t with $\Im t < -\pi/\alpha$. Let χ be the supremum of those real numbers s, such that the horizontal line $y = s$ lies below $\mathcal{T}$. Then $-\pi/\alpha \le \chi < -1$.

Let J denote reflection in the real axis in the t-plane. Since positive and negative bending are related by J, we know that, for $n \in \mathbb{Z}$, $-i + n(L/\alpha) \in \partial\mathcal{T}$.

Let $\widehat{\mathcal{T}}$ denote the open horizontal strip in $\mathbb{C}$ bounded below by the real axis and above by $J(\partial\mathcal{T})$. For $n \in \mathbb{Z}$, the strip $\widehat{\mathcal{T}}$ is invariant under the translation $t \mapsto t + n(L/\alpha)$. For any $t \in \overline{\widehat{\mathcal{T}}}$, $\Omega_t \subset \mathbb{S}^2$ and the earthquake $A_t : \mathbb{D}^2 \to \Omega_t$ is a K_t-quasiconformal homeomorphism.

Let $\xi = \phi - i\chi$ be a point on the upper boundary of $J(\widehat{\mathcal{T}})$. Given a neighborhood U' of the point ξ in which $\Im t > 1$, set $U = U' \cap \widehat{\mathcal{T}}$.

Let T_t be the torus obtained from the from the square torus by applying the angle scaling map $A_t : \mathbb{D}^2 \to \Omega_t$ and the Riemann mapping $\Psi_t : \Omega_t \to \mathbb{D}^2$.

22.8 Theorem. *If U' is small enough, then, for any $t = x + iy \in U'$, the Teichmüller distance from T_x to T_t is strictly greater than* $\log 2$. *For any equivariant K-quasiconformal homeomorphism $\Omega_t \to \mathrm{Dome}(\Omega_t)$ which extends to the identity on $\partial\Omega_t$, we have $K > 2$. Therefore $K_{eq} > 2$.*

Proof: Let $d_{\mathcal{T}}$ denote the hyperbolic metric in $\mathcal{T}$ and let d_- denote the hyperbolic metric in the halfplane $\{t \in \mathbb{C} : \Im t > \chi\} \supset \mathcal{T}$. In particular for $\xi = \phi + \chi i$,

$$d_-(\phi, \xi) = \log 2 < d_{\mathcal{T}}(\phi, \xi).$$

We choose $\epsilon > 0$ and U' small enough, so that, for $\tau \in U$, $\log 2 + \epsilon < d_{\mathcal{T}}(\Re\tau, \tau)$.

By the definition of the Teichmüller distance, the dilatation of any quasiconformal homeomorphism between T_u and T_τ is greater than $2e^\epsilon$. Let $f_t : \mathrm{Dome}(\Omega_t) \to \Omega_t$ be any equivariant quasiconformal homeomorphism which extends to the identity on $\partial\Omega_t$. Let $I_t : \mathbb{D}^2 \to \mathrm{Dome}(\Omega_t)$ be the normalized isometry. Let $\Lambda' = A_u(\Lambda)$. Then $f_t \circ I_t : \mathbb{D}^2 \to \Omega_t$ has the same boundary values as $A_{iv,\Lambda'} : \mathbb{D}^2 \to \Omega_t$, the map obtained by bending through an angle $v\alpha$ along Λ', since the values are the same at endpoints of geodesics in Λ and limits in $\mathbb{S}^1$ of points in a fixed flat piece, and these are dense in $\mathbb{S}^1$. The Riemann mapping $\Psi_t : \Omega_t \to \mathbb{D}^2$ can be extended to a homeomorphism of the boundary, because $\partial\Omega_t$ is a Jordan curve. So $\Psi_t \circ f_t \circ I_t$ and $\Psi_t \circ A_{iv,\Lambda'}$ have the same boundary values. These maps are equivariantly homotopic with the boundary fixed, and can therefore be thought of as defining the same point of Teichmüller space. Furthermore $A_{iv,\Lambda'}$ sends T_u to T_t.

It follows that the quasiconformal constant of f_t must be greater than $2e^\epsilon$. ∎

23 Further properties of mappings dome to floor

23.1 The introduction of slits. Suppose $\Omega' \subset \mathbb{S}^2$ is simply connected and that $0 \in \Omega'$. Denote by Ω the result of removing the line segment along the positive real axis running from $z = 0$ to the first boundary point.

23.2 Lemma. *Any K-quasiconformal mapping $\Phi : \mathrm{Dome}(\Omega) \to \Omega$ that extends to the identity on $\partial\Omega$ satisfies $K \geq 2$. Therefore if the $K = 2$ conjecture is true, an extremal quasiconformal map $\Phi : \mathrm{Dome}(\Omega) \to \Omega$ has maximal dilation exactly two.*

Proof: Let $\{c_n\}$ be an increasing sequence of positive numbers converging to infinity and set $C_n : \mathbb{C} \to \mathbb{C}$ be defined by $C_n(z) = c_n z$. The sequence of regions $\Omega_n = C_n(\Omega)$ converges in the sense of Carathéodory to $\Omega_\infty = \{z \in \mathbb{C} : z \notin [0, +\infty)\}$. The slits become progressively longer and in the limit become $\{[0, +\infty)\}$. Also $\lim \mathrm{Dome}(\Omega_n) = \mathrm{Dome}(\Omega_\infty)$, where the latter is taken to be the union of the two sides of the vertical wall erected on $\{[0, +\infty)\}$. We are working in $\mathbb{U}^3 \subset \mathbb{R}^3 \subset \mathbb{S}^3$ and we use the spherical metric in $\mathbb{S}^3$.

Consider a sequence of K-quasiconformal mappings $g_n : \Omega_n \to \mathrm{Dome}(\Omega_n)$ where g_n fixes each point on $\partial\Omega_n$. We claim that for a subsequence $g = \lim g_m$ is a K-quasiconformal map $g : \Omega_\infty \to \mathrm{Dome}(\Omega_\infty)$.

We will apply the main convergence theorem for quasiconformal mappings. The situation is nonstandard since the image regions are not in $\mathbb{C}$ but rather are embedded in $\mathbb{R}^3 \subset \mathbb{S}^3$. Nevertheless the same results are valid. That the limit function may be a constant is ruled out by the fact that g_m pointwise fixes a segment of a horizontal slit. Thus g is a K-quasiconformal map which necessarily fixes every point on $\partial\Omega_\infty$. ■

23.3 Remark. Lemma 23.2 also holds if the euclidean slit is replaced by a circular slit, or by an inner pointing "cusp" consisting of two smooth arcs meeting at a zero angle. □

23.4 Example. Consider the disks $D_n = \{z : |z - 2n| < 1\}$ and $D'_n = \{z : |z - 2n - 1 + i| < 1, \ -\infty < n < \infty\}$. The region $\Omega = \cup_{n=-\infty}^{\infty}(D_n \cup D'_n)$ is simply connected and its boundary has inward pointing cusps. Therefore $K \geq 2$ and the boundary dilatation is 2 for Ω. On $\mathrm{Dome}(\Omega)$ the bending angles are all $\frac{\pi}{2}$. The leaves of the corresponding lamination in $\mathbb{D}^2$ pairwise share end points. There is no ortho-curve and the lamination allows no crescents at all to be inserted. In particular Ω cannot be obtained by angle doubling. If we apply a right earthquake and try to correspondingly rebuild the dome, the underlying region will no longer be embedded in $\mathbb{S}^2$.

On the other hand consider the region formed by three disks $\Omega_1 = D_0 \cup D_1' \cup D_1$. Its dome has two bending lines with bending angles $\frac{\pi}{2}$ and back in $\mathbb{D}^2$ the two geodesics have a common endpoint. The ortho-curve is the limiting case where the bending angle is π. Since Ω_1 can be obtained by angle doubling, $K = 2$. □

23.5 Extremality of the angle doubling map. The angle doubling map is extremal only in the degenerate cases described in the next theorem.

23.6 Theorem. *Let Λ be a finite measured lamination which satisfies the ortho-condition. The corresponding angle doubling map $f : \mathbb{D}^2 \to \Omega$ is an extremal quasiconformal mapping for its boundary values in and only in the following degenerate case. There exists $\zeta \in \partial\mathbb{D}^2$ which is the common endpoint of leaves of Λ such that the corresponding bending angles sum to exactly π.*

Proof: If there is a vertex ζ such that the sum of bending angles at ζ is equal to π then the boundary dilatation H necessarily satisfies $H \geq 2$. On the other hand since the angle doubling map f is 2-quasiconformal we have that $H \leq K = 2$. Therefore $H = K = 2$ and consequently f is extremal (there cannot be another quasiconformal map with the same boundary values but smaller maximal dilatation).

If the sum of angles at each vertex is less than π then $H < 2$. For according to [9], if the local boundary dilatation $H_q < 2$ at each boundary point q then the boundary dilatation also satisfies $H < 2$. We claim that f then cannot be extremal. For otherwise $K = 2$ for the extremal map and we would have $H < K = 2$. The Strebel frame condition applies ([16]; see also [14]) to guarantee a uniquely extremal Teichmüller map with the same boundary values as f. But the angle doubling map f is not of Teichmüller type except in the case that Λ has one line with bending angle π. Except in this case, f is conformal on some open subset of $\mathbb{D}^2$. This is a contradiction. ■

Figure 27: Ω as a union of gaps and crescents

23.7 The nearest point retraction is 2-Lipschitz. Assume that $\Omega \subset \mathbb{C}$ is simply connected and is not equal to $\mathbb{C}$. Assume also that that $S = \mathrm{Dome}(\Omega) = \mathcal{CH}(\Omega^c)$ is finitely bent.

Consider a flat piece F of S; it is a polygon in a hyperbolic plane $H \subset \mathbb{H}^3$. The boundary $\partial H \subset \mathbb{S}^2$ is a circle. D, one of the two disk components of $\mathbb{S}^2 \setminus \partial H$, lies in Ω. The disk D is a *maximal disk* in the sense that it is not properly contained in any other open disk lying in Ω. Since D is associated with a flat piece, $\partial D \cap \partial\Omega$ consists of at least three distinct points.

The nearest point map uniquely determines a region $G \in D$ by the property that $r : G \to F$ is an isometry if we use the hyperbolic metrics on S and D. The inverse images in Ω of the bending lines are disjoint crescents.

The relative boundary $\partial G \cap D$ is a union of geodesics in the hyperbolic metric of D; there is at least one of them. Each component c of $\partial G \cap D$ is an edge of a crescent $C \subset \Omega$. The other edge c' is a geodesic in another maximal disk D', corresponding to a flat piece F' that is adjacent to F along a bending line ℓ; the exterior bending angle satisfies $0 < \alpha < \pi$ (if S is not folded). The vertices of C are $\partial D \cap \partial D'$ which are also the endpoints of ℓ. The vertex angle of C is α. The nearest point retraction r sends C onto ℓ.

Overall, Ω is the union of gaps and crescents.

23.8 Lemma (The gap lemma). *Suppose $\Omega \subset \mathbb{C}$ is simply connected $\neq \mathbb{C}$ and Dome(Ω) is finitely bent. Let $D \subset \Omega$ be a maximal disk and $G \subset D$ be a gap. Then the hyperbolic metrics $\rho_D|dz|$ of D and $\rho_\Omega|dz|$ of Ω satisfy*

$$\rho_\Omega(z) < \rho_D(z) < 2\rho_\Omega(z), \quad \forall z \in G. \tag{23.8.11}$$

Proof: The lefthand inequality of Inequality 23.8.11 is obvious; we need to prove the righthand inequality. It suffices to consider the case that the bending lines of S have no common endpoints, for we can approximate S (and Ω) by finitely bent surfaces with such a property.

Consider then the gap $G \subset D$.

Let $\zeta \in \partial G$ be a point that lies on an edge c of a crescent C associated with the intersecting maximal disks D, D'. Choose a Möbius transformation T that sends one endpoint of c to $z = 0$, the other to ∞. T maps D and D' to euclidean halfplanes $T(D)$ and $T(D')$. Thus $T(\Omega)$ contains $T(D) \cup$

$T(C) \cup T(D')$ which subtends angle $\pi + \alpha$ in $T(\Omega)$. Therefore for the shortest euclidean distances,

$$d(T(\zeta), \partial T(\Omega)) = d(T(\zeta), \partial T(D)) = |T(\zeta)| = d(T(\zeta), 0).$$

Since $\infty \notin T(\Omega)$, by the Koebe inequality (see [2]),

$$\rho_{T(\Omega)}(T(\zeta)) \geq \frac{1}{2d(T(\zeta), \partial T(\Omega))}.$$

Since $T(D)$ is a halfplane,

$$\rho_{T(D)}(T(\zeta)) = \frac{1}{d(T(\zeta), \partial T(D))} = \frac{1}{d(T(\zeta), \partial T(\Omega))}.$$

We conclude that

$$\frac{\rho_D}{\rho_\Omega}(\zeta) = \frac{\rho_{T(D)}}{\rho_{T(\Omega)}}(T(\zeta)) \leq 2.$$

This holds for all points $\zeta \in \partial G \cap D$.

Next consider a component c of $\partial G \cap \partial D \subset \partial\Omega$. Choose a Möbius transformation T sending D to the upper halfplane, and the endpoints of c, which are endpoints of crescent edges, to $0, \infty$. In light of our considerations of points on crescent edges, there exists a uniform strip along $\mathbb{R}$ in the euclidean metric with the following property. For all points $\zeta \in G$ sufficiently close to c (such that $T(\zeta)$ lies in the strip), the orthogonal projection of $T(\zeta)$ to $\mathbb{R}$ is the closest point of $T(\zeta)$ to $\partial T(\Omega)$. We conclude as before that for all ζ sufficiently close (in the euclidean sense) to c,

$$\rho_{T(\Omega)}(T(\zeta)) \geq \frac{1}{2d(T(\zeta), \partial T(\Omega))} = \frac{1}{2d(T(\zeta), \partial T(D))} = \frac{\rho_{T(D)}(T(\zeta))}{2}.$$

We have shown, for all points ζ on or near ∂G, that

$$\rho_D(\zeta)|d\zeta| \leq 2\rho_\Omega(\zeta)|d\zeta|,$$

since this expression is invariant under renormalization.

Now for all points $z \in D$,

$$\Delta \log \rho_\Omega(z) = {\rho_\Omega}^2, \quad \Delta \log \rho_D = {\rho_D}^2.$$

Therefore if we set $\tau = \log \frac{\rho_D}{\rho_\Omega}$, we find $\Delta\tau(z) > 0$. That is, τ is subharmonic in D. Consequently τ cannot have a maximum in D. So it can't have a maximum in the interior of G. Since $\tau(\zeta) \leq \log 2$ on $\partial G \cap D$ and near $\partial G \cap \partial\Omega$, $\tau(z) \leq \log 2$ for all $z \in G$. This establishes Inequality 23.8.11 for G. ∎

23.9 Theorem. *The nearest point retraction* $r : \Omega \to \mathrm{Dome}(\Omega)$ *is* 2*-Lipschitz in the respective hyperbolic metrics. The result is best possible.*

Proof: First consider the finitely bent case.

Suppose G is a gap. We can renormalize so that $r(G) \subset S = \mathrm{Dome}(\Omega)$ is contained in a vertical halfplane in $\mathbb{U}^3$. Then r is a euclidean rotation and, for $z \in G$, $|r'(z)| = 1$. Consequently, from Inequality 23.8.11, for $z \in G$ we have

$$\rho_S(r(z))|r'(z)||dz| = \rho_S(r(z))|dz| = \rho_D(z)|dz| \leq 2\rho_\Omega(z)|dz|,$$

where the extreme terms give a form that is invariant under renormalization.

Next consider a crescent C. Renormalize so that its endpoints are $0, \infty$. Then C is a thin wedge of vertex angle $\alpha < \pi$. The euclidean halfplanes D, D' are adjacent along its two edges, using the earlier notation. Therefore given $z \in C$ the closest euclidean distance $d(z, \partial\Omega)$ is $|z|$, the distance to $z = 0$. Consequently

$$\rho_\Omega(z) \geq \frac{1}{2d(z, \partial\Omega)} = \frac{1}{2|z|}.$$

The bending line ℓ corresponding to C becomes the vertical halfline ending at $0 \in \mathbb{U}^3$. The nearest point retraction $r : C \to \ell$ is a euclidean isometry on each line in C through 0. In particular $r : C \to \ell$ preserves euclidean distances to $z = 0$. Also, $|r'(z)| \leq 1$. Consequently

$$\rho_S(r(z))|r'(z)| < \rho_S(r(z)) = \frac{1}{|z|} \leq 2\rho_\Omega(z).$$

We conclude that

$$\rho_S(r(z))|r'(z)| \leq 2\rho_\Omega(z), \quad \forall z \in \Omega. \tag{23.9.12}$$

Upon integration,

$$d_S(r(x), r(y)) \leq 2d_\Omega(x, y). \tag{23.9.13}$$

Inequality 23.9.13 applies equally to general simply connected regions $\Omega \neq \mathbb{C}$. For we can approximate a general region Ω by regions Ω_n with finitely bent domes so that Ω_n converges to Ω in the sense of Carathéodory, and $\mathrm{Dome}(\Omega_n)$ converges to $\mathrm{Dome}(\Omega)$ in the sense of Hausdorff. The corresponding nearest point retractions, considered as maps into $\mathbb{H}^3$, converge uniformly on compact subsets of Ω. ∎

We recall that r is automatically equivariant with respect to a kleinian group acting on Ω. To see that Theorem 23.9 is best possible, we take Ω to be the complement of a half-infinite line.

Our result improves the original result in [7, Theorem 2.3.1], repeated in [5, Theorem 2], that r is 4-Lipschitz. In [6, Corollary 4.4] it is shown that the nearest point retraction is homotopic to a $2\sqrt{2}$-Lipschitz, equivariant map.

23.10 Angle doubling is 4-bilipschitz. It has been suggested that in parallel to the $K = 2$ conjecture there should be a $L = 2$ conjecture for the minimal bilipschitz constant L of maps $\mathrm{Dome}(\Omega) \to \Omega$ that pointwise fix $\partial\Omega$. Bishop showed [4, Theorem 1.4] that the nonequivariant bilipshitz constant may be taken as the number 13.3. In contrast we offer the following result for the angle doubling mappings.

23.11 Proposition. *The mapping* $\mathrm{Dome}(\Omega) \to \Omega$ *associated with an angle doubling mapping* $f : \mathbb{D}^2 \to \Omega$ *is* 4*-bilipschitz in the respective hyperbolic metrics.*

Proof: It is enough to show that $f : \mathbb{D}^2 \to \Omega$ is 4-bilipschitz. By approximation, we may assume that the measured lamination $\Lambda \subset \mathbb{D}^2$ is finite.

Insert a set of cresents in $\mathbb{D}^2$ that are associated with Λ. The mapping f is differentiable except on the boundaries of the crescents so it suffices to prove f is locally 4-bilipschitz. In addition the property is invariant under renormalization of f by conjugation by Möbius transformations so we have the freedom to choose the most propitious one.

First consider a component G of the complement of the crescents in $\mathbb{D}^2$. Choose a point $p \in G$. By renormalization of f, we map assume that p is the point $z = 0$ and that the restriction of f to G is the identity.

Since G contains an endpoint of a leaf of Λ in its boundary,

$$d(f(0), \partial\Omega) = d(0, \partial\Omega) = d(0, \partial\mathbb{D}^2 \cap \partial G) = 1.$$

Here d refers to the euclidean distance. Therefore the standard estimate for the hyperbolic density ρ_Ω of Ω gives

$$\rho_\Omega(0) \geq \frac{1}{2d(0, \partial\Omega)} = \frac{1}{2}.$$

On the other hand, $\mathbb{D}^2 \subset \Omega$ so that

$$\rho_\Omega(0) > \rho_{\mathbb{D}^2}(0) = 2.$$

Combining the two estimates, we get

$$\frac{\rho_{\mathbb{D}^2}(0)}{4} \le \rho_\Omega(f(0))|f'(z)| \le \rho_{\mathbb{D}^2}(0).$$

This shows that f is 4-bilipschitz in the gaps.

Now suppose p is a point in a crescent C. We normalize with $D = \mathbb{U}^2$, taking the crescent to have endpoints at 0 and ∞, with p at height 1. Denote by ℓ the ray from 0 through p. By further renormalization, we may assume that the restriction of f to ℓ is the identity. Then the directional derivative $D_\beta f$ satisfies

$$1 \le |(D_\beta f)(0)| \le 2.$$

The value 1 occurs in the direction of ℓ and 2 occurs in the orthogonal direction.

Consequently,

$$\frac{\rho_D(p)|dz|}{2} = |dz|/2 \le \rho_\Omega(p)|dz| \le \rho_\Omega(p)|df| \le 2\rho_D(p)|dz|.$$

We conclude that f is 4-bilipschitz as claimed. ∎

23.12 Euclidean convex regions. We will prove the following fact announced in Section 1.

23.13 Proposition. *Suppose $\Omega \subset \mathbb{C}$ is a euclidean convex region. The nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$ is a 2-quasiconformal mapping.*

Proof: . We may assume by approximation that Ω is a convex polygon. As pointed out in Section 1, $\mathrm{Dome}(\Omega)$ is smooth. It is the union of a finite number of flat pieces together with a portion of a half-cone near each vertex and possibly half-cylinders lying over rectangular subregions. The nearest point retraction r is conformal to the flat pieces, and from § 3.6 and § 3.5 we know r does not exceed 2-quasiconformal near the corners. ∎

23.14 Boundary dilatations and the Brennan conjecture In this section we will present a weaker condition that also implies Brennan's conjecture. The example in § 8.11 suggests how this new condition may be easier to apply.

Let $\Omega \ne \mathbb{C}$ be a simply connected region. Two quasiconformal mappings $f_1, f_2 : \Omega \to \Omega^*$, are said to have *the same boundary values* if ${h^*}^{-1} f_1 h, {h^*}^{-1} f_2 h$:

$\mathbb{D}^2 \to \mathbb{D}^2$ have the same boundary values. Here $h : \mathbb{D}^2 \to \Omega$, $h^* : \mathbb{D}^2 \to \Omega^*$ are Riemann maps. (Recall that a quasiconformal homeomorphism between two open unit disks can always be continuously extended to the closed disks, so the "boundary values" of these quasiconformal homeomorphisms makes sense.)

Given a quasiconformal map f of Ω and a point $p \in \partial\Omega$, the notion of *substantial boundary dilatation* $H_p(f)$ of f at p is defined as follows. Denote the disk of radius r about p by $B_r(p)$. Suppose we have a quasiconformal map g of Ω with the same boundary values as f (write $g \sim f$). We recall that we are not talking here of the boundary of Ω, but of the boundary of $\mathbb{D}^2$. We denote the maximal dilatation of g restricted to $B_r(p) \cap \Omega$ by $K_r(g)$. We make the following sequence of definitions:

$$\tilde{H}_p(g) = \lim_{r \searrow 0} K_r(g), \quad H_p(f) = \inf_{g \sim f} \tilde{H}_p(g),$$

$$H(f) = \max_{p \in \partial\Omega} H_p(f).$$

The number $H(f)$ is called the *boundary dilatation* of f. It is known [8, 9] that the maximum is attained.

Automatically $H(f) \leq K_0$ where K_0 is the maximal dilatation of an extremal quasiconformal map with the boundary values of f. Typically, there is strict inequality. See [14] for an exposition of this topic.

If $\partial\Omega$ is smooth and f has C^1-boundary values, then $H = 1$. Suppose instead $\partial\Omega$ and $\partial f(\Omega)$ are smooth except for a finite number of corners of interior angle α_i, $0 < \alpha_i \leq 2\pi$. Suppose f is correspondingly piecewise C^1 with non-zero derivative on $\partial\Omega$. Then $H_p = 1$ except at the points corresponding to the corners where $H_{p_i} = \max(\frac{\alpha_i}{\pi}, \frac{\pi \sin(\alpha_i/2)}{\alpha_i})$.

When $\partial\Omega$ is smooth except for a finite number of corners with interior angles $0 \leq \alpha_i \leq 2\pi$, any map $\Psi : \mathrm{Dome}(\Omega) \to \Omega$ which is the identity on $\partial\Omega$ has the following property. The substantial boundary dilatations satisfy $H_p(\Psi) = 1$ for p on the smooth pieces of $\partial\Omega$, $H_{p_i}(\Psi) = \frac{\alpha_i}{\pi}$ at vertices with $\pi \leq \alpha_i \leq 2\pi$, and $H_{p_i}(\Psi) = \frac{\pi \sin(\alpha_i/2)}{\alpha_i}$ at vertices with $0 \leq \alpha_i \leq \pi$ ($\alpha_i = 0$ indicates a cusp).

23.15 Theorem. *If for all simply connected regions $\Omega \neq \mathbb{C}$ and for some map $\Psi : \mathrm{Dome}(\Omega) \to (\Omega)$ which is the identity on $\partial\Omega$, the boundary dilatation satisfies $H(\Psi) \leq 2$, then Brennan's conjecture is true. If $H(\Psi) \leq 2$ for a particular region Ω, then Brennan's conjecture holds for that particular region.*

The theorem immediately implies the known result that Brennan's conjecture holds for regions with smooth boundaries except for possibly a finite number of corners. Like Bishop's theorem this is a consequence of a result of Astala. We will give a proof in § 23.14. We will also apply this result in an example.

The following conjecture is weaker than the $K = 2$ conjecture (Conjecture 2.4), and perhaps has more chance of being true (see Example 23.4 and § 8.11).

23.16 Conjecture. *Given any simply connected region $\Omega \neq \mathbb{C}$, there exists a quasiconformal map $\Psi : \mathrm{Dome}(\Omega) \to \Omega$ which is the identity on $\partial\Omega$ and for which $H(\Psi) \leq 2$.*

23.17 Remark. Bishop showed that Brennan's conjecture would be true if the following held: For any simply connected $\Omega \neq \mathbb{C}$ there exists a 2-quasiconformal map $H : \Omega \to \mathbb{D}^2$ with bounded euclidean gradient. This requires no boundary assumptions. Brennan's conjecture would also follow if a quasiconformal H only had its boundary dilatation less than or equal to two, as well as bounded euclidean gradient. It follows that Conjecture 23.16 is sufficiently strong so that one could infer from its truth the truth of Brennan's Conjecture. The proof follows the pattern of Bishop's proof, making use of the proof of Theorem 23.15 given in § 23.14. □

Proof: Let $\Omega \subset \mathbb{S}^2$ be a simply connected region. Assume that $H(\Phi) \leq H \leq 2$ holds for all quasiconformal maps $\Phi : \Omega \to \mathrm{Dome}(\Omega)$ that fix $\partial\Omega$. Denote by g the associated function $g = \iota \circ \Phi : \Omega \to \mathbb{D}^2$, where as usual $\iota : \mathrm{Dome}(\Omega) \to \mathbb{D}^2$ is an isometry.

By theorems of [8, 9], for each $\varepsilon > 0$ there is a compact set $A_\varepsilon \subset \Omega$ with the following property. The map $\Phi : \Omega \to \mathrm{Dome}(\Omega)$ may be taken so that the restriction of the corresponding g to $\Omega \setminus A_\varepsilon$ is $(H + \varepsilon)$-quasiconformal, or in our case, $(2 + \varepsilon)$-quasiconformal.

Bishop [4] showed in particular that $g : \Omega \to \mathbb{D}^2$ is Lipschitz in the euclidean metric, *i.e.* $|g'(z)| < C_1$ for some constant $C_1 < \infty$. Choose a Riemann map $f : \Omega \to \mathbb{D}^2$ and factor f as $f = g_1 \circ g$. The mapping $g_1 : \mathbb{D}^2 \to \mathbb{D}^2$ is $(2 + \varepsilon)$-quasiconformal on $\mathbb{D}^2 \setminus B_\varepsilon$. Here $B_\varepsilon = g(A_\varepsilon)$ is compact in $\mathbb{D}^2$.

We can in turn factor g_1 as $g_1 = g_3 \circ g_2$ where $g_3 : \mathbb{D}^2 \to \mathbb{D}^2$ is $(2 + \varepsilon)$-quasiconformal on the whole disk $\mathbb{D}^2$ and $g_2 : \mathbb{D}^2 \to \mathbb{D}^2$ is quasiconformal in

$\mathbb{D}^2$ and conformal in $\mathbb{D}^2 \setminus B_\varepsilon$. This is done by solving the Beltrami equations for μ_3 which is the complex dilatation of g_1 on $\mathbb{D}^2 \setminus B_\varepsilon$ and zero in the rest of $\mathbb{D}^2$, and for μ_2 which is the complex dilatation of g_1 on B_ε and zero in the rest of $\mathbb{D}^2$.

The map g_2 extends to a quasiconformal map of $\mathbb{S}^2$ which is conformal in the complement of $B_\varepsilon \cup B'_\varepsilon$, where B'_ε is the reflection of B_ε in the unit circle $\partial\mathbb{D}^2$. In particular there is a compact set $\widetilde{B_\varepsilon}$, such that $\mathbb{D}^2 \supset \widetilde{B_\varepsilon} \supset B_\varepsilon$ and such that $|g_2'(z)| < C_2$ for $z \in \mathbb{D}^2 \setminus \widetilde{B_\varepsilon}$.

We finish by following Bishop's argument [4, p.8]. For $2 < t < 4$, we want to show that

$$\iint_\Omega |f'(z)|^t \, dxdy < \infty.$$

Since $\widetilde{A_\varepsilon} = g^{-1}(\widetilde{B_\varepsilon})$ is a compact subset of Ω, we have, for all $0 < t < \infty$,

$$\iint_{\widetilde{A_\varepsilon}} |f'(z)|^t \, dxdy < \infty.$$

Therefore it suffices to estimate

$$\iint_{\Omega \setminus \widetilde{A_\varepsilon}} |f'(z)|^t \, dxdy = \iint_{\Omega \setminus \widetilde{A_\varepsilon}} |g_3' \circ (g_2 \circ g)|^t |(g_2 \circ g)'|^2 |(g_2 \circ g)'|^{t-2} \, dxdy.$$

Since $|g'| < C_1$ and $|g_2'| < C_2$ on $\widetilde{B_\varepsilon} = g(\widetilde{A_\varepsilon})$, we obtain

$$\iint_{\Omega \setminus \widetilde{A_\varepsilon}} |f'|^t \, dxdy < (C_1 C_2)^{t-2} \iint_{(g_2 \circ g)(\Omega \setminus \widetilde{A_\varepsilon})} |g_3'|^t \, dudv.$$

According to Astala's theorem applied to the $(2+\varepsilon)$-quasiconformal map $g_3 : \mathbb{D}^2 \to \mathbb{D}^2$, for all t satisfying

$$2 \le t \le \frac{2(2+\varepsilon)}{2+\varepsilon-1} = \frac{4+2\varepsilon}{1+\varepsilon}, \text{ we have } \iint_{\mathbb{D}^2} |g_3'(w)|^t \, dudv < \infty.$$

Since $(g_2 \circ g)(\Omega \setminus \widetilde{A_\varepsilon}) \subset \mathbb{D}^2$, we conclude that

$$\iint_{\Omega \setminus \widetilde{A_\varepsilon}} |f'|^{\frac{4+2\varepsilon}{1+\varepsilon}} \, dxdy < \infty.$$

Therefore

$$\iint_\Omega |f'|^{\frac{4+2\varepsilon}{1+\varepsilon}} \, dxdy < \infty.$$

The exponent increases strictly to 4 as ε decreases to 0. Therefore $f' \in L^p$, $\forall p,\ 2 < p < 4$. Therefore the upper bound in Brennan's conjecture is 4 as was to be proved. ∎

References

[1] L.V. Ahlfors. *Lectures on Quasiconformal Mappings.* D.Van Nostrand, 1966.

[2] L.V. Ahlfors. *Conformal Invariants: Topics in Geometric Function Theory.* McGraw-Hill Book Co., 1973.

[3] L. Bers and H.L. Royden. Holomorphic families of injections. *Acta Math.*, 157:259–286, 1986.

[4] C. Bishop. Quasiconformal Lipschutz maps, Sullivan's convex hull theorem, and brennan's conjecture. *Stony Brook preprint*, 2000.

[5] M. Bridgeman and R. Canary. From the boundary of the convex core to the conformal boundary. *preprint, www.math.lsa.umich.edu/ canary*, 2000.

[6] R. Canary. The conformal boundary and the boundary of the convex core. to appear in Duke Math. J.

[7] D.B.A. Epstein and A. Marden. *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, pages 113–253. Cambridge University Press, 1987.

[8] R. Fehlman. Ueber extremale quasikonforme abbildungen. *Comment. Math. Helv.*, 56:558–580, 1981.

[9] N. Lakic. Substantial boundary points for plane domains and Gardiner's conjecture. *Ann. Acad. Sci. Fenn.*, 25:285–306, 2000.

[10] O. Lehto. *Univalent Functions and Teichmüller Spaces.* Springer-Verlag, 1987.

[11] O. Lehto and K.I. Virtanen. *Quasiconformal Mappings in the Plane, 2nd ed.* Springer-Verlag, 1973.

[12] C. McMullen. Complex earthquakes and Teichmüller theory. *Jour. A.M.S.*, 11, 1998.

[13] D. Sullivan R. Mañé P. Sad˙ On the dynamics of rational maps. *Ann. Sci. École Norm. Sup*, 16:193–217, 1983.

[14] E. Reich. Extremal quasiconformal mappings of the disk. In R. Kühnau, editor, *Geometric Function Theory, vol.1, Handbook of Complex Analysis*. Elsevier, 2001.

[15] Z. Slodkowski. Holomorphic motions and polynomial hulls. *Proc. A.M.S.*, 111:347–355, 1991.

[16] K. Strebel. On the existence of extremal Teichmüller mappings. *Jour. d'Anal. Math.*, 30:441–447, 1976.

[17] D. Sullivan and W.P. Thurston. Extending holomorphic motions. *Acta Math.*, 157:243–257, 1986.

[18] W.P. Thurston. Minimal stretch maps between hyperbolic surfaces. unpublished manuscript.