Complex earthquakes and deformations of the unit disk

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

17 August 2004

Contents

1 Introduction

The central ingredients in our work are geodesic laminations (Λ, μ) in the hyperbolic plane $\mathbb{H}^2$, with a transverse measure μ, possibly invariant under a fuchsian group G representing a finite area quotient surface $\mathbb{H}^2/G$. Our investigations involve one-parameter deformations defined using the measured laminations, and the corresponding embedded pleated surfaces in hyperbolic space $\mathbb{H}^3$.

Our study is based on a generalization of the method of complex scaling (see [8] for an elementary exposition). The generalization is related to quakebends or complex earthquakes [7], [11]. In Chapters 2-4 we introduce a series of five functions which culminate in a holomorphic family $\Psi : \mathbb{C} \times \Omega_0 \to \mathbb{S}^2$. Here Ω_0 is (any) fixed quasidisk. For fixed $z \in \Omega_0$, $\Psi^z : \mathbb{C} \to \mathbb{S}^2$ is holomorphic. For fixed $t \in \mathbb{C}$, $\Psi_t : \Omega_0 \to \mathbb{S}^2$ is quasiregular. We show how to extract from $\Psi : \mathbb{C} \times \Omega_0 \to \mathbb{S}^2$ a holomorphic map from a certain simply connected domain containing the upper halfplane into the universal Teichmüller space $\mathcal{T}$. Its construction uses the boundary values of $\Psi_t : \Omega_0 \to \mathbb{S}^2$. If Ω_0 is invariant under a quasifuchsian group, Ψ_t is correspondingly equivariant.

The existence of Ψ is used to prove two theorems.

Theorem 5.1 is a new and more natural proof of Sullivan's theorem that there is a universal constant K with the following property. Given *any* simply connected region $\Omega \neq \mathbb{C}$, construct Dome (Ω), namely the relative boundary in $\mathbb{H}^3$ of the hyperbolic convex hull of $\mathbb{S}^2 \setminus \Omega$. Consider the family $\mathcal{F}_\Omega$ (homotopy class) of all quasiconformal mappings $f : \Omega \to \text{Dome}\,(\Omega)$ that pointwise fix the common boundary. This class is non-empty, and extremal maps in this class have maximal dilatation at most K. If in addition Ω is invariant under a group G of Möbius transformations, the functions f can be taken to be G-equivariant. If G represents a surface of finite area, the theorem shows there is likewise a universal bound K_{eq} for equivariant maps f. The new proof yields the estimate $K \leq 13.88$, and the same estimate $K_{eq} \leq 13.88$ is proved in the equivariant case. The equivariant estimate is much better than our original 82.8, but in the nonequivariant case it is not as good as Bishop's 7.8 found by an explicit construction. We will indicate below how we obtain a more precise characterization of $\mathcal{F}$ for the class of euclidean convex regions.

The second main theorem is the proof of Theorem 7.8 which we call the Disk Theorem, since it is related to a theorem of that name by McMullen [11]. Let $X \subset \mathbb{C}$ denote that (closed) set of parameters t for which the map $\Psi_t : \Omega_0 \to \mathbb{S}^2$ is injective. The Disk Theorem states that no component of $\mathbb{C} \setminus X$ is bounded. In particular, each component of the interior of X is simply connected. Numerous questions about this intriguing set remain open. The set X might be considered analogous to the Mandelbrot set.

We want to thank C. Earle for helpful discussions.

2 Defining the deformation: scaling for a discrete set of crescents

The basis of our work in this section is the *standard wedge of angle* α,

$$W_\alpha = \{z \in \mathbb{C} : 0 < \arg(z) < \alpha\},$$

where $0 < \alpha < 2\pi$. The wedge W_α is foliated by the set of its rays from 0 to ∞. In the usual orientation of W_α the ray $\arg(z) = 0$ will be called its *right edge*. See Figure 2.0.i.

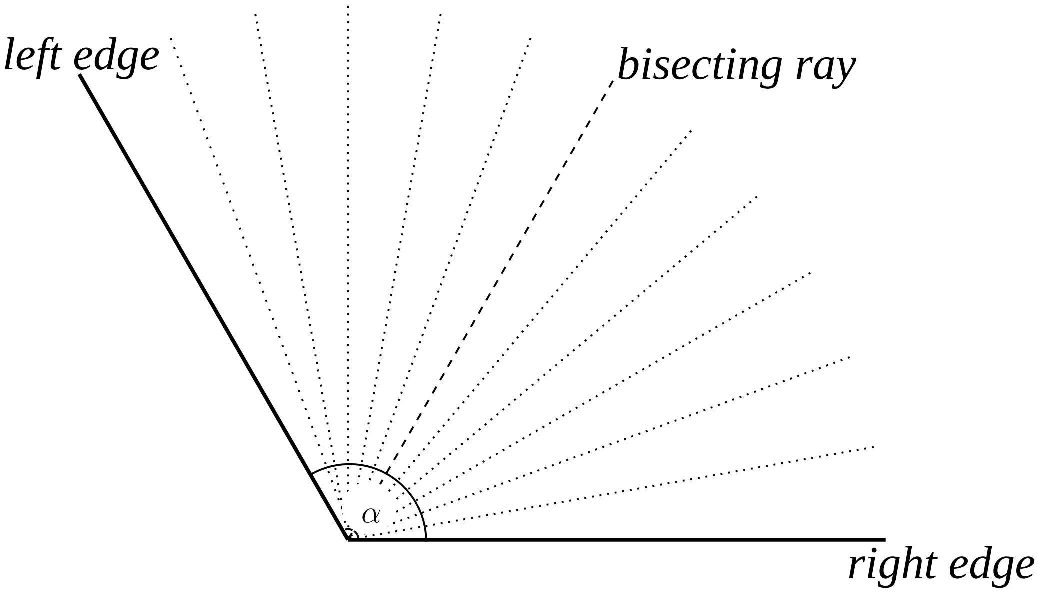

Figure 2.0.i. A standard wedge with angle α. Note the bisecting ray at angle $\alpha/2$. We indicate the foliation of the wedge by rays.

Let Ω be any simply connected region of $\mathbb{S}^2$ containing W_α, such that both $0, \infty \in \partial\Omega$. The bisecting ray $\{\arg z = \alpha/2\}$ of W_α separates Ω into two parts: let Ω_+ denote the component containing the positive real axis, and Ω_- the other component. See Figure 2.0.ii.

2.1 Definition. Given a standard wedge W_α, define the *scaling map*

$$E : (t, z) \in \mathbb{C} \times \Omega \mapsto E(t, z) \in \mathbb{C}$$

as follows:

- For $z \in \Omega_+ \setminus W_\alpha$ and $t \in \mathbb{C}$
$$E(t, z) = z.$$
- For $z = re^{i\theta\alpha} \in W_\alpha$ with $0 \leq \theta \leq 1$, and for $t = u + iv \in \mathbb{C}$,
$$E(t, z) = e^{t\theta\alpha} z = e^{u\theta\alpha} e^{iv\theta\alpha} z.$$
- For $z \in \Omega_- \setminus W_\alpha$ and $t \in \mathbb{C}$,
$$E(t, z) = e^{t\alpha} z = e^{u\alpha} e^{iv\alpha} z.$$

We define $E_t : \Omega \to \mathbb{C}$ by $E_t(z) = E(t, z)$. □

The effect of a scaling map is indicated in Figure 2.1.i.

The unit circle centred at 0 carries the obvious measure given by arclength along the circle, or equivalently by the angle subtended from 0. This gives a transverse measure on the foliated wedge, with total measure equal to α.

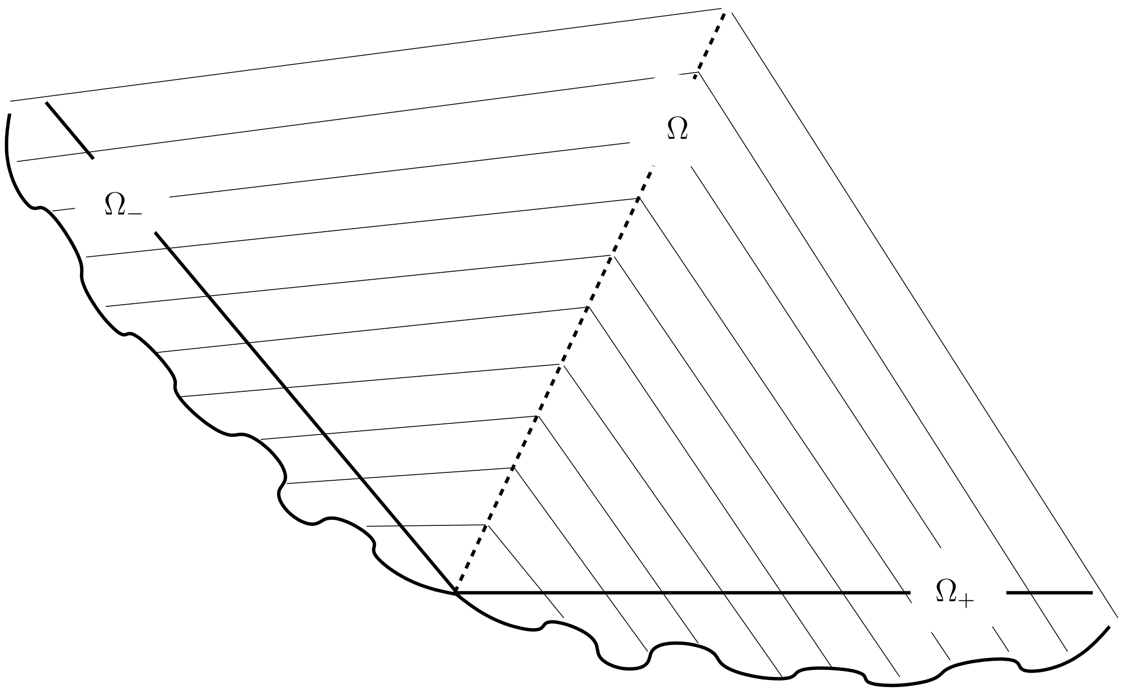

Figure 2.0.ii. This is the same wedge as that shown in Figure 2.0.i. We indicate the subsets Ω, Ω_+ and Ω_- referred to in Definition 2.1, by means of hatching.

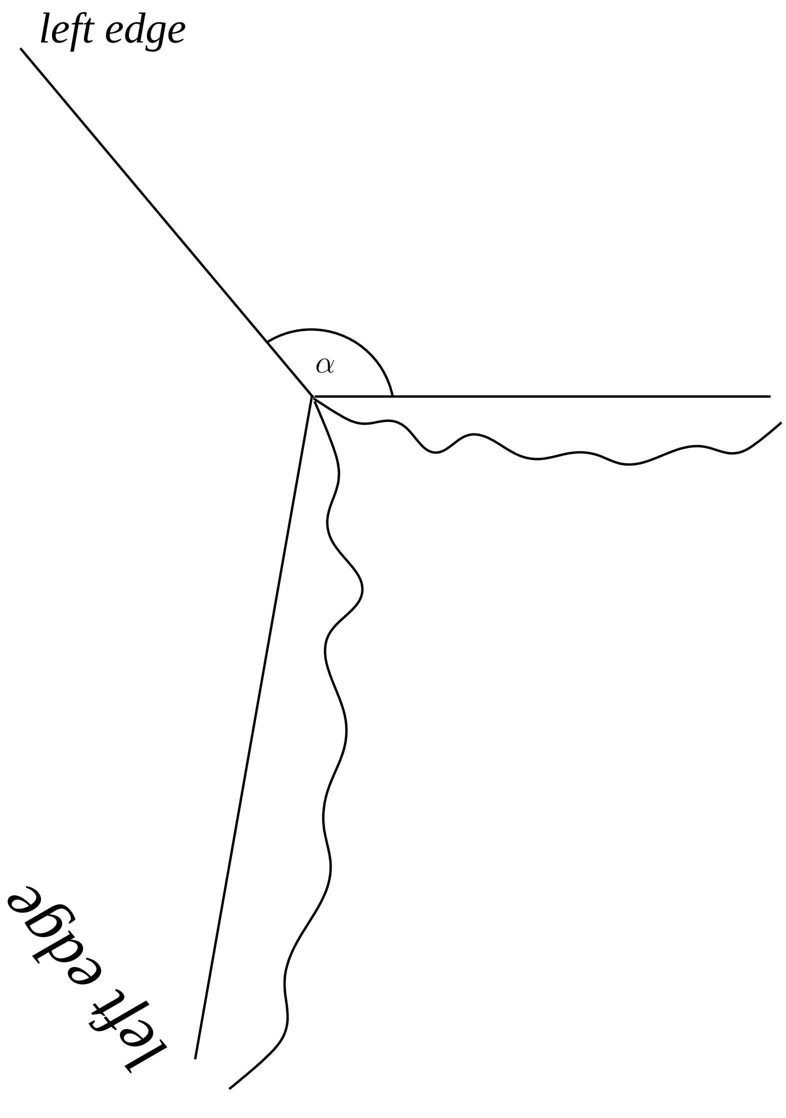

Figure 2.1.i. This picture shows the effect of scaling. The particular scaling parameter t has been chosen so that $\alpha.\mathrm{Re}(t) = \log(2)$ and $\mathrm{Im}(t) = 1$. The effect is to magnify at the lefthand edge by a factor 2 and to rotate by doubling the angle α. To help the reader understand the effect of the scaling, it is applied also to the label "left edge". At the righthand edge, the scaling map is the identity. In between, the magnification and the amount of rotation vary as a linear function of the original angle.

2.2 Lemma.

- *For fixed $z \in \Omega$, $E(t,z)$ is a holomorphic function of t.*
- *For fixed $t \in \mathbb{C}$, the restriction of E_t to a c-leaf $\{\arg z = \theta\alpha\}$ of the foliation of W_α, or to a component of $\Omega \setminus W_\alpha$, is a Möbius transformation.*

For each $t = u + iv \in \mathbb{C}$, the image of W_α under E_t is the wedge $W_{(v+1)\alpha}$. The image of Ω is a new region $E_t(\Omega)$. In general $E_t : \Omega \to \mathbb{C}$ is not an embedding. If $\mathrm{Im}(t) > -1$, then the scaling map E_t is locally injective and is injective on W_α.

There is a one parameter family of hyperbolic Möbius transformations, namely $\left\{z \mapsto e^\lambda \cdot z\right\}_{\lambda\in\mathbb{R}}$, that fixes each vertex $0, \infty$ and maps W_α onto itself. For $t \in \mathbb{C}$, we write $A_t : z \in \mathbb{C} \mapsto e^\lambda \cdot z \in \mathbb{C}$. The Möbius transformation B defined by $B(z) = e^{i\alpha}/z$ also sends W_α to itself, interchanging the two vertices and the two edges. Note that $B^{-1} = B$. The group of Möbius transformations preserving W_α consists of the A_λ, together with all products of the form $A_\lambda.B$. We have

$$E_t \circ A_\lambda = A_\lambda \circ E_t \text{ and } E_t \circ B = A_{t\alpha} \circ B \circ E_t.$$

The A_λ and B also preserve the foliation of W_α and its transverse measure.

It follows that the conjugate of E_t by A_λ is equal to E_t. The conjugate of E_t by B is given by

$$B \circ E_t \circ B^{-1} = A_{-t\alpha} \circ E_t.$$

This keeps the left edge of W_α fixed, rather than the right edge. To keep the right edge fixed, we normalize by postmultiplying by $A_{t\alpha}$, obtaining the same scaling map E_t as the one we started with.

We now generalize the use of the term *scaling map* in a way that will be more convenient to use. In particular we extend the concept to any crescent C_α of vertex angle $\alpha > 0$. Here we are using the term *crescent* to denote a region in $\mathbb{S}^2$ bounded by two circular arcs.

There is a Möbius transformation T such that $T(W_\alpha) = C_\alpha$. The scaling map $\tilde{E}_t = T \circ E_t \circ T^{-1}$ fixes the image of the right edge of W_α under T. If instead one wants to fix the other edge of C_α, one needs to replace T by TB. These two choices are related by postmultiplication by $T \circ A_{t\alpha} \circ T^{-1}$. Replacing T by $T \circ A_\lambda$ does not change $\tilde{E}$. (The domain $T(\Omega)$ may change, but this makes no difference to any of our arguments.)

More generally, we may not want to fix either edge of C_α. Instead, we may want the image under T of the right edge of W_α to be mapped by the Möbius transformation

S_t. Here S_t is a Möbius transformation depending holomorphically on t, with $S_0 = \mathrm{Id}$. In this case, the scaling map is given by

$$\tilde{E}_t = S_t \circ T \circ E_t \circ T^{-1}.$$

The left edge of C_α is then mapped by $S_t \circ T \circ A_{t\alpha} \circ T^{-1}$. We summarize the situation in the following lemma.

Suppose the crescent C_α with a designated right edge lies in a simply connected region $\Omega \subset \mathbb{S}^2$, with the vertices of C_α in $\partial\Omega$. Let Ω_+, Ω_- be the two components of Ω determined by the c-leaf which bisects C_α, so labelled that the right edge of C_α lies in Ω_+. Let S_t be a holomorphic family of Möbius transformations, with $S_0 = \mathrm{Id}$.

2.3 Lemma. *Let $E : \mathbb{C} \times \Omega \to \mathbb{S}^2$ be the scaling map associated to the crescent $C_\alpha \subset \Omega$, such that $E_t = S_t$ on the right edge of C_α. The following properties hold:*

2.3.1. E is continuous.

2.3.2. For fixed $z \in \Omega$, $E(\cdot, z) : \mathbb{C} \to \mathbb{S}^2$ is holomorphic.

2.3.3. $E_0 : \Omega \to \mathbb{S}^2$ is the identity map.

2.3.4. For fixed $t \in \mathbb{C}$, the restriction of E_t to any c-leaf $\ell \subset C_\alpha$, or to Ω_+, or to Ω_-, is a Möbius transformation.

2.3.5. For each $t = u + iv \in \mathbb{C}$ with $v > -1$, the scaling map is a locally injective K_t-quasiregular mapping with maximal dilatation

$$K_t = \frac{1 + |\kappa(t)|}{1 - |\kappa(t)|}, \text{ where } \kappa(t) = \frac{-t}{2i + t} \cdot \frac{z}{\bar{z}}.$$

2.3.6. The restriction of E_t to the crescent C_α itself is injective if and only if $0 < (v + 1)\alpha < 2\pi$.

We next extend the definition of scaling maps to the case of a finite number of crescents $\{C_1, \ldots, C_k\}$. It is convenient to take the crescents as open and we assume they are mutually disjoint in a given simply connected region $\Omega \subset \mathbb{S}^2$ (however the closures of two crescents may touch tangentially). As before we assume that the vertices lie on $\partial\Omega$. We allow distinct crescents to share the same pair of vertices. Each crescent is foliated by arcs of circle going through the two vertices.

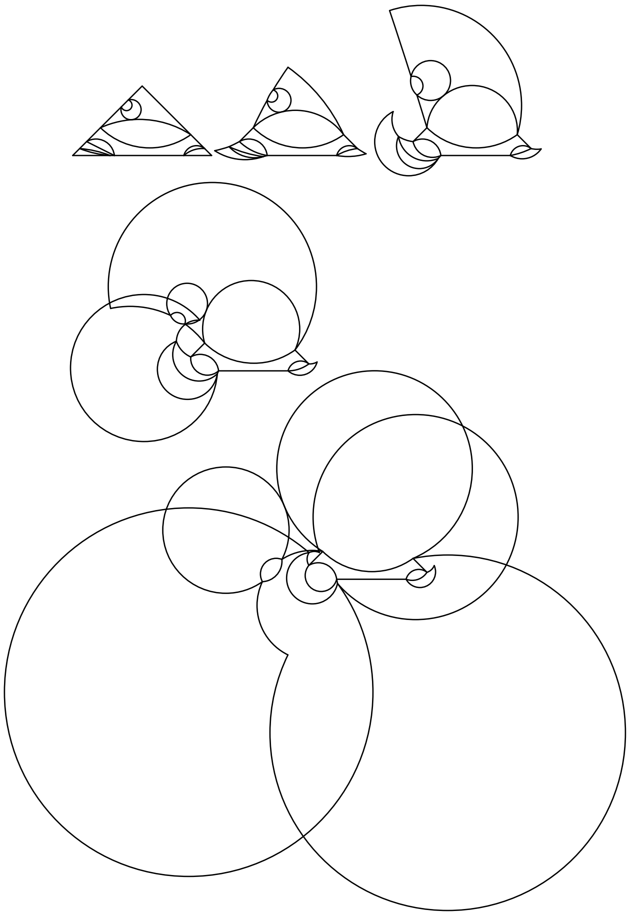

Figure 2.4.i. We start with a number of crescents in a topological disk. In this case, the topological disk is a triangle Δ. We use the notation of Lemma 2.5. The first diagram shows Δ and its crescents. The other four diagrams show the images of the various intervals and arcs in Δ under the scaling map E_t for the four values of t in $\{-0.06+0.3i, -0.18+0.9i, -0.4+1.2i, -0.18+1.8i\}$. The last two diagrams show that $E_t : \Delta \to \mathbb{S}^2$ is not necessarily injective. In the last diagram, several of the finite regions in the domain triangle have an image which includes ∞. The scaling map E_t is surjective for many values of t. E_t is normalized by taking it to be the identity map on a six-sided region, one edge of which lies on the bottom edge of the triangle. This shape can be observed in each of the five diagrams. All scaling maps with $\mathrm{Im}(t) > 0$ are locally injective.

2.4 Definition. By a *crescent leaf*, we mean one of the arcs of circle just described. We abbreviate this to c-leaf. □

The components of $\Omega \setminus \cup C_i$ are called *gaps*. Thus gaps are closed sets in Ω. The interior of a gap may not be connected; this will happen when there are tangencies between various crescents. A gap has at least three boundary points on $\partial\Omega$, unless the gap is itself a crescent.

The following lemma is proved by an induction argument, interpolating in the gaps between crescents by the Möbius transformations which act on the successive edges of the crescents.

2.5 Lemma. *There is a continuous map $E : \mathbb{C} \times \Omega \to \mathbb{S}^2$ (which we call the scaling map) with the following properties:*

2.5.1. For fixed $z \in \Omega$, $E(\cdot, z) : \mathbb{C} \to \mathbb{S}^2$ is holomorphic.

2.5.2. $E_0 : \Omega \to \mathbb{S}^2$ is the identity.

2.5.3. For each i, the restriction of E to the crescent C_i is a scaling map.

2.5.4. The restriction of E_t to a gap, or to a c-leaf of the foliation of a crescent C_i, is a Möbius transformation.

E_t is uniquely determined by the above properties up to postcomposition by Möbius transformations S_t, provided S_t depends holomorphically on t.

Figure 2.4.i illustrates the effect of a scaling map.

2.6 Normalization. We will invariably normalize the scaling map. For example, we choose three distinct points $\{x_1, x_2, x_3\}$ in $\overline{\Omega}$, where, for $i = 1, 2, 3$, if $x_i \notin \Omega$, then $x_i \in \partial V_i$, for some gap or c-leaf V_i. We then define $E_t(x_i) = x_i$ for all t. By choosing $\{x_1, x_2, x_3\} \subset V$, where V is a gap or a c-leaf, we can ensure that, for all t, $E_t|V = \mathrm{Id}\,|V$. We will normally use this slightly more special style of normalization.

There is no loss in generality in assuming that the three points fixed are $-1, 1, i$. Alternatively, if we want V to remain fixed, we may assume that V is placed in some special position. □

The following corollary can be proved by a countable induction and limiting process.

2.7 Corollary. *Lemma 2.5 continues to hold in the case of a set of disjoint crescents which is locally finite in Ω.*

For future reference, we record the following result in the case of a family of disjoint crescents which is locally finite in Ω. We use the same assumptions and notation as in Lemma 2.5 and Corollary 2.7.

2.8 Lemma. *Let Ω be a Jordan region and let W be the union of all the crescents in the locally finite family. The scaling map E_t on Ω extends continuously to those points $\zeta \in \partial\Omega$ which are in the boundary of a crescent, or in the boundary of a component of $\Omega \setminus W$. For each such fixed ζ, the map $E(\cdot, \zeta) : \mathbb{C} \to \mathbb{S}^2$ is holomorphic.*

If there are only a finite number of crescents, Lemma 2.8 shows that the scaling map can be extended to the whole of $\mathbb{C} \times \overline{\Omega}$. However this is not possible in general for an infinite family, even if the family is locally finite in the interior of Ω, as the following example shows. (There is no problem if the family is locally finite in $\overline{\Omega}$.)

2.9 Example. Consider an infinite "chain of beads" which covers the open interval $(-1, 1)$ in the x-axis in the plane. Each bead (open disk) meets exactly two other beads. Successive beads overlap, but the overlap is a tiny amount so that successive beads are almost tangent to each other. As we go to one of the endpoints of the interval, the successive exterior angles of intersection increase to π, that is, the beads become more and more nearly tangent. The situation is illustrated in Figure 2.9.i.

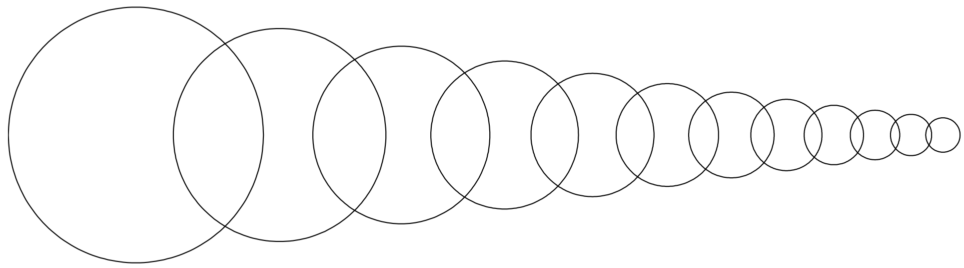

Figure 2.9.i. The picture illustrates Example 2.9. It represents a countable union of round disks, getting smaller and smaller. The union of the interiors is simply connected. The angle between successive circles gets smaller and smaller. In other words, successive disks become more and more nearly tangent to each other.

The union Ω of the beads is an open simply connected region. Its dome corresponds to a discrete lamination, each g-leaf carrying a bending measure which is almost equal to π. We can find a geodesic τ, lying on the dome, that is orthogonal to all the leaves. We can assume that the hyperbolic distance between successive leaves increases to infinity. Let (Λ, μ) be the corresponding measured geodesic lamination of $\mathbb{H}^2 \cong \mathbb{D}^2$. Let $r : \Omega \to \mathrm{Dome}\,(\Omega)$ be the nearest point retraction. For each g-leaf ℓ, we have a crescent $r^{-1}(\ell)$ with angle nearly π. From Corollary 2.7, we see that we can define a scaling map $E_t : \mathbb{D}^2 \to \mathbb{S}^2$ for all $t \in \mathbb{C}$. We have $\Omega = E_i(\mathbb{D}^2)$, and Ω is a Jordan region.

For values of $t \in \mathbb{C}$ with small imaginary part (the exact condition is explained in [9, Section 4]), we can extend E_t continuously to $\mathbb{S}^1 = \partial\mathbb{D}^2$, so that, for $z \in \mathbb{S}^1$, $E(t, z)$ is a holomorphic function of such t.

However, it is not possible to extend E_t continuously to either endpoint of τ, if t lies on the imaginary axis, with $\mathrm{Im}(t)$ just bigger than 1. On the other hand, E_t *can* be extended continuously to the closed disk when $t = i = \sqrt{-1}$.

In a similar way, consider a region Ω with a prime end which has an impression consisting of more than one point (see Definition 5.2). Suppose Ω is the image of $\mathbb{D}^2$ under the scaling E_i, where $i = \sqrt{-1}$. The ability to extend E_i to a point of $\partial\mathbb{D}^2$ implies that the impression of its image is a single point. Therefore we cannot extend E_i to $\partial\mathbb{D}^2$ in such a situation. $\square$

3 Defining the deformation: complex earthquakes

In [7, Chapter 3], there is a thorough discussion of measured laminations, with complex valued transverse measures. We discussed associated surfaces in $\mathbb{H}^3$, using what we called "quakebends". These surfaces are special cases of pleated surfaces, where the pleating happens to be associated to a transverse measure. The concept is due to Bill Thurston, who used the idea in computer programs. Subsequent authors rechristened the notion of quakebend with the name "complex earthquake", and we will use the newer terminology. A complex earthquake is a map from $\mathbb{H}^2$ (which we shall usually represent as the Poincaré disk $\mathbb{D}^2$) into $\mathbb{H}^3$ (which we shall usually represent as the Poincaré ball $\mathbb{D}^3$). The image may be embedded, but typically it is not; the image may even be dense in $\mathbb{H}^3$.

In essence, a complex earthquake is the composition of a pure earthquake in $\mathbb{H}^2$ (the real part) followed by a pure bending of $\mathbb{H}^2 \subset \mathbb{H}^3$ within $\mathbb{H}^3$ (the imaginary part). It is the first factor that is discontinuous. We will also bring in associated maps, called *scaling maps*, $\mathbb{D}^2 \to \mathbb{S}^2 = \partial\mathbb{D}^3$.

The purpose of this section is to recall the notion of complex earthquakes and its properties.

In [7, Section 3.11], measured laminations were discussed in terms which made easier the link to standard measure theory ideas. The space of geodesics in $\mathbb{H}^2$ can be identified with the open Möbius band X of pairs of distinct points of $\partial\mathbb{H}^2 = \mathbb{S}^1$. A measured lamination was defined as a Borel measure on X with the property that, given any two pairs of points in the support of the measure, the corresponding geodesics do not cross. The space of measured laminations was topologized as a subspace of all Borel measures on X, using the weak topology from continuous functions $f : X \to \mathbb{R}$ with compact support.

Suppose we are given a given a measured geodesic lamination (Λ, μ). A component of $\mathbb{D}^2 \setminus \bigcup \Lambda$ is called a *flat*, because its image under a complex earthquake is always flat, in the sense that it is isometrically embedded in a hyperbolic plane which is a subspace of hyperbolic 3-space.

3.1 Definition. We will usually refer to an element of Λ as a *geodesic leaf* or *g-leaf*. We will refer to a crescent leaf, as in Definition 2.4, as a *c-leaf*. It will turn out that there is an intimate relationship between these two types of leaf. □

3.2 Orientation. The flats determined by Λ are oriented from the orientation of $\mathbb{D}^2$. Choose a flat or a line $\ell \in \Lambda$ with $\mu(\ell) = 0$ to serve as the basepoint V_0 of our construction. Except in a very special case, we can always choose V_0 to be a flat, and for simplicity we will assume this can be done. Orient the leaves bounding V_0 so that V_0 lies to their left.

There is a geodesic segment σ from V_0 to any g-leaf ℓ. Orient ℓ so that it is consistent with the direction of the boundary g-leaf ℓ_0 of V_0 that σ first crosses. The orientation of ℓ is independent of the particular σ chosen. This process gives a consistent orientation to all leaves of Λ. We can now do earthquakes consistently along Λ. An earthquake by a positive amount along an oriented g-leaf ℓ moves its right side in the positive direction of ℓ, with respect to its left side.

We will regard $\mathbb{D}^2$ as the equatorial plane in the Poincaré ball $\mathbb{D}^3$. Then we designate the *positive direction* of rotation about an oriented g-leaf ℓ as the direction of rotation of a corkscrew motion. This is the usual direction of increasing argument of a complex number in the (x, y)-plane, when rotating about the positively oriented z-axis. □

Although the general theory of earthquakes applies to all complex-valued transverse Borel measures, we simplify the discussion by restricting our attention to measures of the form $t\mu$, where $t \in \mathbb{C}$ and μ is a non-negative measure with support on Λ.

Let V denote a flat or a g-leaf, and let $0 \leq \theta \leq 1$. The complex earthquake $\mathbb{C}E_t$ associated to $(\Lambda, t\mu)$ assigns to each such pair (V, θ) a unique Möbius transformation, which we denote by $M(t, V, \theta)$. The notation is deceptive: $M(t, V, \theta)$ and $\mathbb{C}E_t$ depend on t only via their dependence on Λ and on $t\mu$. Thus if t is multiplied by a positive real number and μ is divided by the same number, then $M(t, V, \theta)$ and $\mathbb{C}E_t$ are unchanged. For the base element V_0, we assign, for each $t \in \mathbb{C}$, $M(t, V_0, \theta) = \mathrm{Id}$. $M(t, V, \theta)$ depends on θ if and only if $V \in \Lambda$ is a g-leaf with positive measure. We will often omit the parameter θ where it is unnecessary.

3.3 The finite case. Suppose that Λ is a finite measured lamination, in which case every g-leaf has strictly positive μ-measure. Let V' and V'' be two flats, and choose a geodesic arc γ from a point of V' to a point of V''. Denote the successive leaves of Λ crossed by σ by $\ell_1, \ldots, \ell_n$, oriented according to the convention in (3.2). Let $A_{t,i}$ be the loxodromic transformation with axis ℓ_i, translating the signed distance $\mathrm{Re}(t)\mu(\ell_i)$ along ℓ_i and bending through the signed angle $\mathrm{Im}(t)\mu(\ell_i)$ around ℓ_i.

A defining property of a complex earthquake is that, on the flats,

$$M(t, V')^{-1} \circ M(t, V'') = A_{t,1} \circ A_{t,2} \circ \ldots \circ A_{t,n}. \tag{3.3.a}$$

In fact (3.3.a) itself determines the complex earthquake on the flats up to postcomposition with a Möbius transformation. In the present context, (see § 2.6) we have already normalized the map on V_0 as $M(t, V_0, \theta) = \mathrm{Id}$, so the earthquake is determined without ambiguity on the flats. If V is a flat, then the image $M(t, V)(V) \subset \mathbb{D}^3$ is also called a flat. See Figure 3.3.i for an illustration of the situation.

We obtain a map $\mathbb{C}E_t : \mathbb{D}^2 \to \mathbb{D}^3$ by defining, for each flat V, $\mathbb{C}E_t|V = M(t, V)|V$. This defines $\mathbb{C}E_t$ except on the g-leaves, where it is so far undefined. □

3.4 G-leaves with positive measure. We now examine the situation at a g-leaf $\ell_j \subset \Lambda$. There is a flat V_- bordering ℓ_j on the left, and a flat V_+ bordering on the right. We know that $M(t, V_+) = M(t, V_-)A_{t,j}$ where the rotational part of $A_{t,j}$ rotates about the g-leaf ℓ_j with signed angle $\mathrm{Im}(t)\mu(\ell_j)$, and the translational part translates along ℓ_j a signed distance $\mathrm{Re}(t)\mu(\ell_j)$. In particular, $A_{t,j}$ preserves ℓ_j and therefore $M(t, V_+)(\ell_j) = M(t, V_-)(\ell_j)$. We denote this geodesic in $\mathbb{D}^3$ by $\mathbb{C}E_t(\ell_j)$, even though $\mathbb{C}E_t$ is not well-defined on ℓ_j. The flats $\mathbb{C}E_t(V_-)$ and $\mathbb{C}E_t(V_+)$ share the edge $\mathbb{C}E_t(\ell_j)$.

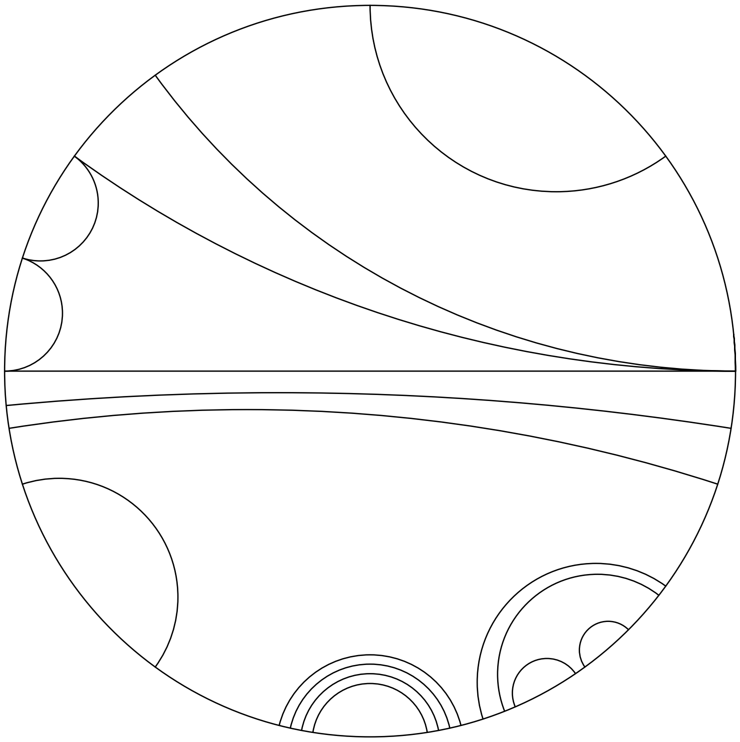

Figure 3.3.i. This is a picture of a finite geodesic lamination. We are using the Poincaré metric on the unit disk. To each of the regions shown is assigned a Möbius transformation. The transformations on different sides of a geodesic differ by a Möbius transformation having that geodesic as an axis.

We define the *pleated surface* $P_{\Lambda,t\mu}$ to be the union of the sets $\mathbb{C}E_t(V)$ as V varies over flats and g-leaves.

We have already associated a Möbius transformation $M(t,V)$ to each flat V of Λ. Now let $V=\ell_j\in\Lambda$, and suppose $0\le\theta\le 1$. We define $M(t,V,\theta)=M(t,V_-,\theta)B_{\theta t\mu(V),j}$ where $B_{s,j}$ is the Möbius transformation with axis ℓ_j which translates signed hyperbolic distance $\mathrm{Re}(s)$ and rotates signed angle $\mathrm{Im}(s)$ about ℓ_j. Unless otherwise stated, if we have to make a choice, we set $\theta=1/2$. Using θ allows us to interpolate continuously between $M(t,V_-)$ and $M(t,V_+)$.

There is another way of looking at the Möbius transformations $M(t,V,\theta)$, which demonstrates that their intrinsic significance goes further than simply providing a continuous interpolation. Let $\Omega\subset\mathbb{C}$ be a simply connected proper open subset. Let Λ be the bending lamination on Dome(Ω) and let μ be the bending measure. Suppose that Λ is finite. We think of (Λ,μ) as a geodesic lamination of $\mathbb{D}^2$ with a transverse measure, and compute the Möbius transformations $M(t,V,\theta)$ for this situation. Let X be the hyperbolic convex hull of $\mathbb{S}^2\setminus\Omega$. The set of hyperbolic support planes P for X is then equal to the set of planes $\{M(i,V,\theta)(\mathbb{D}^2)\}$, where V varies over all g-leaves and flats, and $0\le\theta\le1$. The correspondence between the set of support planes P and the set of such Möbius transformations is bijective. However, the correspondence between the set of such Möbius transformations and the set of such pairs (V,θ) is not bijective—one Möbius transformation may arise from many distinct pairs (V,θ). □

3.5 The general case. The complex earthquake associated to $(\Lambda,t\mu)$ is defined by approximating (Λ,μ) by finite measured laminations $\{(\Lambda_n,\mu_n)\}$. We will think of the approximation in terms of the Möbius transformations defined on the successive flats and leaves of each Λ_n.

Let $0\in\mathbb{D}^2$ lie in V_0, the flat or g-leaf of measure zero that we have chosen as our base. We can also assume that, for each n, 0 lies in a flat of Λ_n. Let $z\in V$, where V is a flat or a g-leaf of Λ. It is shown in [7] that, as (Λ_n,μ_n) gets nearer to (Λ,μ) in the topology described at the beginning of this section, the Möbius transformation at z for $t\mu$ converges to a Möbius transformation $M(t,V)$, provided that, when V is a g-leaf, $\mu(V)=0$.

For a g-leaf ℓ, with $\mu(\ell)>0$, the approximating laminations Λ_n can have the property that ℓ splits up into a finite number of disjoint leaves, each very near to ℓ, with the measure $\mu(\ell)$ distributed between the finite number of leaves. By changing the approximation a little, we can ensure that, for each n, z lies in a flat V_n of Λ_n and not on a g-leaf. The corresponding sequence $(M(t,V_n))_{n\in\mathbb{N}}$ lies in a compact set of Möbius transformations. Any convergent subsequence converges to a Möbius transformation of the form $M(t,V,\theta) = M(t,V_-) \circ B_{\theta t\mu(\ell)}$, where B_s has axis ℓ and $s \in [0, t\mu(\ell)]$. Moreover, given θ with $0 \le \theta \le 1$, we can choose our approximating laminations so that the limit of the sequence is exactly $M(t,V_-) \circ B_{\theta t\mu(\ell)}$. For $V = \ell$ the image $M(t,V,\theta)(\ell)$ is a g-leaf of the pleated surface $P_{\Lambda,t\mu}$ which is independent of θ.

The bottom line is that given a measured lamination $(\Lambda, t\mu)$, and given V, a g-leaf or a flat, there is a uniquely defined Möbius transformation $M(t,V,\theta)$. These satisfy the composition formula (3.3.a). For leaves ℓ with non-negative measure, there is a range of possible choices for $M(t,\ell,\theta)$, and we usually choose to use the middle point $\theta = 1/2$ of this range of choices.

There results a complex earthquake $\mathbb{C}E_t$ which maps $\mathbb{D}^2$ onto a pleated surface in $\mathbb{D}^3$. This map is continuous except possibly along a g-leaf ℓ, where it is discontinuous if and only if $\mathrm{Re}(t)\mu(\ell) \neq 0$. □

3.6 Definition. Let (Λ, μ) be a geodesic lamination of $\mathbb{D}^2$, with non-negative transverse measure μ. Let $t \in \mathbb{C}$. We define $\mathbb{C}E_t : \mathbb{D}^2 \to \mathbb{D}^3$ to be the corresponding *complex earthquake*. More precisely $\mathbb{C}E_t$ is defined except on g-leaves ℓ with $\mathrm{Re}(t)\mu(\ell) \neq 0$, and it is discontinuous at such g-leaves. □

We will be particularly interested in how, for a fixed geodesic lamination Λ, the complex earthquake changes with the parameter t. The complex earthquake $\mathbb{C}E_t$ has its image in hyperbolic 3-space, so it does not have direct sense to say that it depends holomorphically on t. However the following lemma shows that the component parts of the complex earthquake do depend holomorphically on t.

3.7 Lemma. *Let (Λ, μ) be a measured lamination with non-negative measure. Then, for each flat or g-leaf V and each θ with $0 \le \theta \le 1$, the Möbius transformation $M(t,V,\theta)$ depends holomorphically on t. (Our earthquakes are assumed to be normalized, for example so that one particular gap or c-leaf is pointwise fixed—see § 2.6.) $M(t,V,\theta)$ depends on θ if and only if V is a g-leaf with $\mu(V) > 0$.*

To prove the lemma, we use the fact that these transformations are holomorphic for a finite lamination, and then use the fact that a uniform limit of holomorphic functions is holomorphic.

3.8. Another way to construct $P_{\Lambda,t\mu}$ is to first construct the signed earthquake using $\mathrm{Re}(t)\mu$. This sends $\mathbb{D}^2$ to $\mathbb{D}^2$, moves Λ to a new geodesic lamination Λ^*, and correspondingly transfers μ to μ^*. Then bend along Λ^*, using $\mathrm{Im}(t)\mu^*$. Since bending is a continuous operation, we can consider the pleated surface as being the image of a continuous map $\mathbb{D}^2 \to \mathbb{D}^3$, using only the bending in the construction just described.$\square$

4 Defining the deformation: scaling maps and non-atomic measures

In this section we will consider the construction of particular maps $\mathbb{C} \times \mathbb{D}^2 \to \mathbb{S}^2$, although sometimes $\mathbb{C}$ is replaced by some open connected subset of $\mathbb{C}$. We do not have axioms which characterize the kind of maps that interest us, but they always have the property that if $z \in \mathbb{D}^2$ is fixed, the corresponding map $\mathbb{C} \to \mathbb{S}^2$ is holomorphic. They will also have the property that, for every parameter t in a certain open connected subset of $\mathbb{C}$, the corresponding map $\mathbb{D}^2 \to \mathbb{S}^2$ is locally injective; it will be either quasiconformal or quasiregular.

We recall that a *quasiregular map* is a map of the form $f = h \circ g$ where g is quasiconformal and h is meromorphic on the range of g. The map f is locally injective if and only if h is locally injective. The reason such maps are important is that a locally injective quasiregular map $\mathbb{D}^2 \to \mathbb{S}^2$, like a quasiconformal map, determines a point of universal Teichmüller space. In § 6.1 we will briefly remind the reader of some of the basic facts in this theory.

The data used to define our mappings will be a pair $(\Lambda, t\mu)$, where Λ is geodesic lamination on $\mathbb{D}^2$, $t \in \mathbb{C}$ and μ is a non-negative transverse measure on Λ.

4.1 Definition. We introduce the norm

$$\|\mu\| = \sup_{\tau} \mu(\tau),$$

where τ ranges over all transverse open geodesic intervals of length one. $\square$

4.2 The mapping $\mathfrak{T}_{\mathfrak{o}} \times \mathbb{D}^2 \to \mathbb{S}^2$. In Section 4 of [9], we introduced the function

$$f(u,x) = \min\left(\operatorname{arcsinh}(e^{|x|}\sinh(u)), e^{|x|/2}u\right). \tag{4.2.a}$$

In the cited paper, we defined a "strip" U about the real axis. Here we will use a somewhat larger neighbourhood $\mathfrak{T}_{\mathfrak{o}}$ which will give better numerical estimates.

We also defined the constant c_2, which was the largest positive number with the following property.

Let (Λ,μ) be any measured geodesic lamination with $\|\mu\| < c_2$. We proved that the corresponding pleated surface map $\mathbb{C}E_t : \mathbb{D}^2 \to \mathbb{D}^3$ is an embedding when $t = i = \sqrt{-1}$. We showed that $.73 \le c_2 \le 2\arcsin(\tanh(1/2)) \approx .9607$. Unpublished work by David Epstein and Dick Jerrard should prove that $c_2 > .948$, though detailed proofs have not yet been written. We conjecture that the correct value is $c_2 = 2\arcsin(\tanh(1/2))$.

If $p \in \mathbb{R}$, let $\lceil p \rceil$ denote the smallest integer greater than or equal to p. We will consider the following two simply connected subregions of the parameter t-plane, a "strip" and a "halfplane":

4.3 Definition.

$$\mathfrak{T}_o = \text{Interior}\left\{t = u + iv : |v| < \frac{c_2}{\lceil f(1,|u|)\rceil}\right\},$$
$$\mathfrak{T} = \text{Interior}\left\{t = u + iv : v > -\frac{c_2}{\lceil f(1,|u|)\rceil}\right\}.$$

These subspaces are illustrated in Figure 4.3.i and Figure 4.3.ii. □

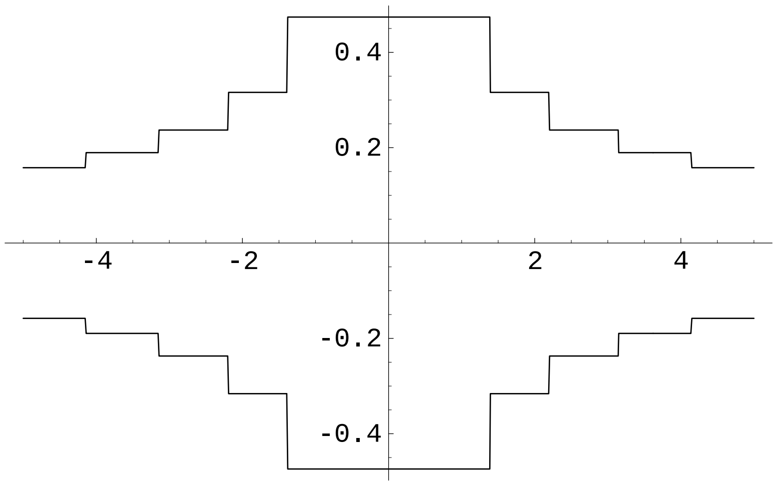

Figure 4.3.i. This is an illustration of $\mathfrak{T}_o$. This is the set of values of $t \in \mathbb{C}$ for which we can show that the map Φ_t of Theorem 4.4 is injective. It is the set of points lying between the lower and upper curves shown in the illustration.

The following result is proved as Theorem 4.14 of [9] as an application of the extended λ-lemma. We will work with the Poincaré models of 2- and 3-dimensional hyperbolic space $\mathbb{D}^2, \mathbb{D}^3$.

4.4 Theorem. *Suppose (Λ, μ) is a measured geodesic lamination with $\|\mu\| = 1$. The following properties hold for points $t \in \mathfrak{T}_o$.*

4.4.1. The complex earthquake map $\mathbb{C}E_t : \mathbb{D}^2 \to \mathbb{D}^3$ is a quasi-isometry.

4.4.2. $\mathbb{C}E_t$ extends to a holomorphic motion of $\mathbb{S}^1 = \partial\mathbb{D}^2$ in $\mathbb{S}^2 = \partial\mathbb{D}^3$.

4.4.3. The extension of $\mathbb{C}E_t$ to $\mathbb{S}^1$ is itself the restriction to $\mathbb{S}^1$ of a holomorphic motion of $\mathbb{S}^2$ in $\mathbb{S}^2$, which we denote by $\Phi_t = \Phi_{(\Lambda,t\mu)} : \mathbb{S}^2 \to \mathbb{S}^2$.

4.4.4. The restriction $\Phi_t|\mathbb{S}^1$ is injective. Its image is a Jordan curve bounding the region $\Omega_t = \Phi_t(\mathbb{D}^2)$. The bending measure of $\mathrm{Dome}(\Omega_t)$ is $\mathrm{Im}(t) \cdot \mu^*$, where μ^* is defined in § 3.8.

4.4.5. In particular, $\Phi_t : \mathbb{D}^2 \to \mathbb{S}^2$ is a quasiconformal homeomorphism with maximal

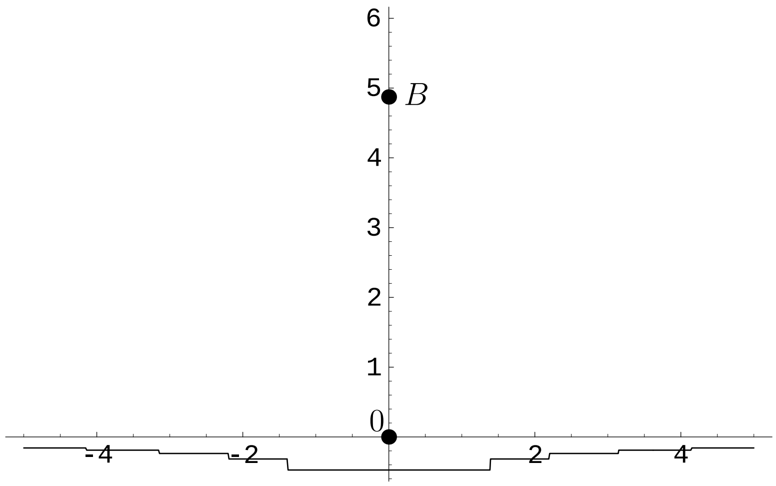

Figure 4.3.ii. Here we illustrate $\mathfrak{T} = \mathfrak{T}_o \cup \mathbb{U}^2$. This is the set of points in $\mathbb{C}$, lying above the curve shown. This is the same as the curve shown in Figure 4.3.i, but drawn to a different vertical scale. As explained in § 4.5, there is a holomorphic map from $\mathfrak{T}$ to universal Teichmüller space. The point $B = (0, c_1)$ (B stands for Bridgeman) is marked in the diagram. Here c_1 comes from Definition 6.10. The Poincaré metric on $\mathbb{D}^2$ is pulled back to $\mathfrak{T}$ using the Riemann map $\mathfrak{T} \to \mathbb{D}^2$. The proof of Theorem 5.1 needs a computation of the hyperbolic distance from 0 to B with respect to this metric on $\mathfrak{T}$.

dilatation K_t given by

$$K_t \leq \frac{1+|h(t)|}{1-|h(t)|}.$$

Here $h : \mathfrak{T}_\circ \to \mathbb{D}^2$ is a Riemann map taking 0 *to* 0.

4.4.6. If there is a group Γ of Möbius transformations which preserves (Λ, μ), then the holomorphic motion Φ_t can be chosen so that there is a homomorphism $\varphi : \Gamma \to \Gamma'$, with

$$\Phi_t(\gamma(z)) = \varphi(\gamma) \circ \Phi_t(z), \quad \textit{for all } \gamma \in \Gamma \textit{ and } z \in \mathbb{D}^2.$$

Here Γ' is another group of Möbius transformations.

Our normalization convention (§ 2.6) ensures that $\mathbb{C}E_t$ fixes a flat or g-leaf and Φ_t fixes the corresponding gap or c-leaf. □

4.5 The map Ψ. We continue to assume that $\|\mu\| = 1$. So far we have defined the holomorphic motion of $\mathbb{S}^2$ in $\mathbb{S}^2$ $\Phi_t : \mathbb{S}^2 \to \mathbb{S}^2$ for $t \in \mathfrak{T}_\circ$. We will be able to deduce from this a holomorphic map from $\mathfrak{T}_\circ$ into universal Teichmüller space $\mathcal{T}$, as we will see later in Section 6. It turns out that Sullivan's Theorem (Theorem 5.1) can be deduced from the existence of such a map. However, the larger the domain of such a holomorphic map into universal Teichmüller space, the better the estimate we will get for the constant K in Sullivan's Theorem. In particular, it seems a reasonable objective to extend the holomorphic motion Φ_t so that it applies for all $t \in \mathfrak{T}_\circ \cup \mathbb{U}^2$. If we take this objective literally, then it is unattainable. Instead we will proceed by constructing another map, denoted by Ψ_t and defined for $t \in \mathbb{U}^2$. Ψ will give rise to a holomorphic map $\mathbb{U}^2 \to \mathcal{T}$ which agrees on $\mathbb{U}^2 \cap \mathfrak{T}_\circ$ with the map induced by Φ.

Our first task is to give the domain of Ψ. We fix a point $t_0 = iv_0$ on the positive imaginary axis, so that $t_0 \in \mathfrak{T}_\circ$. The map $\Phi_{t_0} : \mathbb{S}^2 \to \mathbb{S}^2$ restricts to map the unit disk onto a quasidisk that we will denote by Ω_0 (or by Ω_{t_0}). In particular, Ω_0 is a Jordan domain. The complex earthquake $\mathbb{C}E_{t_0} : \mathbb{D}^2 \to \mathbb{D}^3$ has image $\mathrm{Dome}\,(\Omega_0)$, which is the boundary in $\mathbb{D}^3$ of the hyperbolic convex hull of $\mathbb{S}^2 \setminus \Omega_0$. In Figure 4.5.i, we have a picture of Ω_0.

We wish to define $\Psi : \mathbb{C} \times \Omega_0 \to \mathbb{S}^2$ in such a manner that the boundary values of Ψ_t are related to $\Phi_t|\mathbb{S}^1$. We first explain how this is done when (Λ, μ) is a finite lamination. In that case Ω_0 contains some intrinsically defined crescents. To construct these, let $r : \Omega_0 \to \mathrm{Dome}\,(\Omega_0)$ be the nearest point retraction. Then for each bending line ℓ_i, we have the crescent $C_i = r^{-1}(\ell_i)$, of angle $v_0\mu(\ell_i)$.

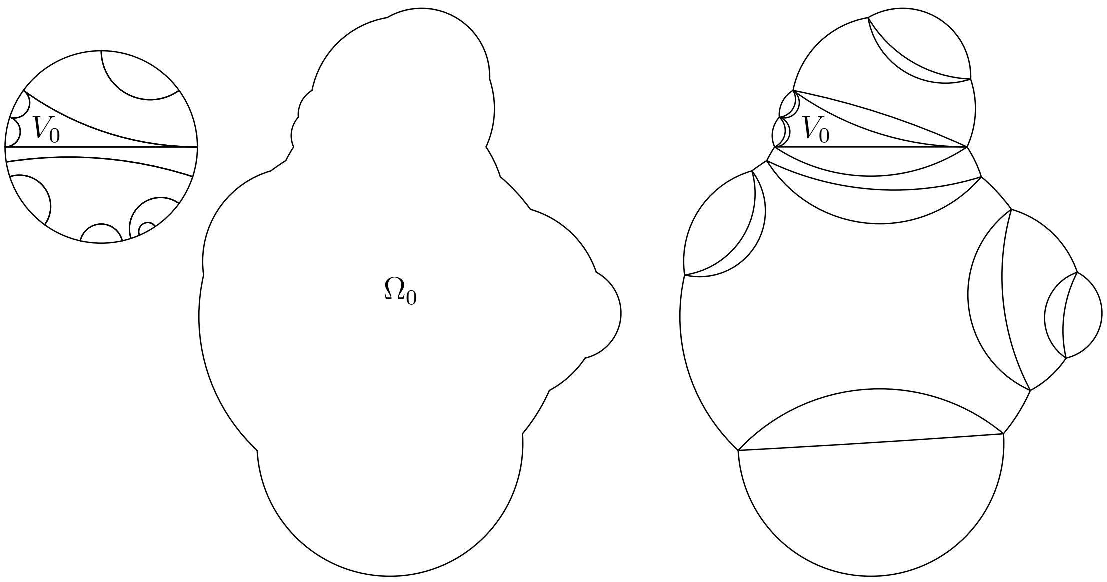

Figure 4.5.i. The lefthand picture shows the unit disk, thought of as a hyperbolic plane with the Poincaré metric, together with a measured geodesic lamination. The measure of each geodesic is proportional to the angle of the corresponding crescent in the righthand picture. The middle picture shows Ω_0, the result of applying Φ_{t_0} to the unit disk. The righthand picture shows Ω_0 with its canonical crescents. Each crescent is the inverse image of a bending line under the nearest point retraction onto $\mathrm{Dome}\,(\Omega_0)$. These crescents are, by their definition, disjoint, and they can therefore be used to "extend" the holomorphic motion Φ_t to all $t \in \mathbb{C}$. The reason for the inverted commas round "extend", is that the extension is not really an extension, but it is close enough to give a well-defined map into universal Teichmüller space, as we will see in due course. We have normalized earthquakes and scaling maps so that the region V_0 shown in the lefthand and righthand pictures is fixed.

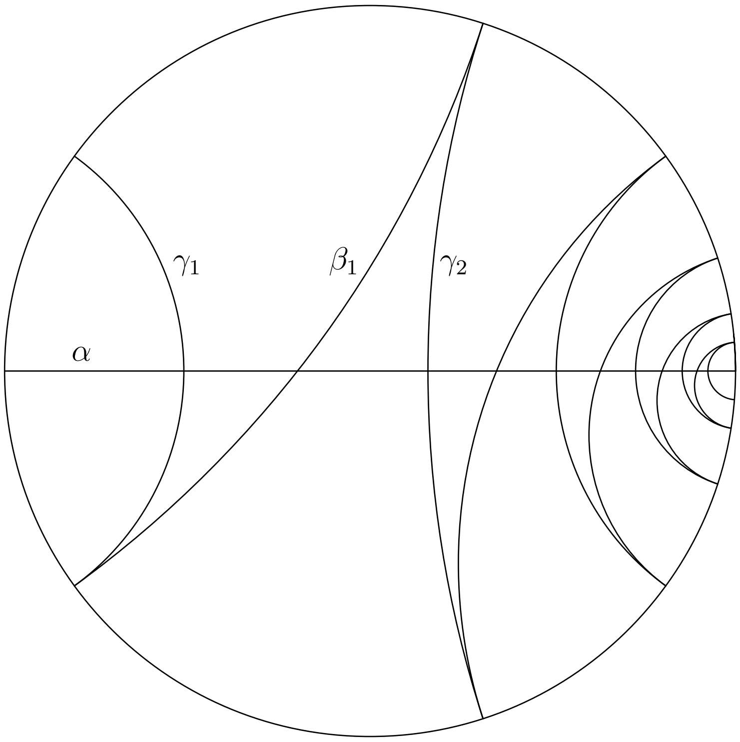

Figure 4.5.ii. For a finite measured lamination of $\mathbb{D}^2$, it is always possible to fit disjoint crescents, with vertices at the endpoints of geodesics in the lamination, and crescent angles proportional to the measure of the corresponding leaf. The definition of Φ_t—in our case carried out using the Λ-lemma—can then be replaced by that of Ψ_t. The question arises of why this is not always possible. Here we provide an example of a countable lamination, where it is not possible to insert such disjoint crescents. We take a countable set of disjoint geodesics $(\gamma_i)_{i\in\mathbb{N}}$, all orthogonal to a fixed geodesic α. We then join the geodesics $(\gamma_i)_{i\in\mathbb{N}}$ with geodesics $(\beta_i)_{i\in\mathbb{N}}$, where β_i joins the end of γ_i with the beginning of γ_{i+1}. If the measure on each of the geodesics γ_i and β_i is the same, then we would have to insert disjoint crescents, each with crescent angle $\epsilon > 0$. It is easy to show that the angle of the crescent associated to γ_i to the boundary circle is at least ϵ greater than the angle of the crescent associated to γ_{i+1} to the boundary circle. So we would be able to fit at most $[2\pi/\epsilon]$ crescents associated to the γ_i. This is the explanation of why we construct our crescents in Ω_0 and not in $\mathbb{D}^2$.

Let E_t be the scaling map defined in Lemma 2.5 using Ω_0 with the crescents $\{C_i\}$. By Lemma 2.8, for each $t \in \mathbb{C}$, E_t extends to a continuous map $E_t : \overline{\Omega_0} \to \mathbb{S}^2$. $E_t|\partial\Omega_0$ is closely related to $\Phi_t|\mathbb{S}^1$. It is even more closely related to $(\Phi_t|\mathbb{S}^1) \circ (\Phi_{t_0})^{-1}|\partial\Omega_0) : \Omega_0 \to \mathbb{S}^2$. To make them match precisely, we perform an affine transformation on the parameter t. We set $\Psi_t = E_{i(t-t_0)/t_0}$. It is easy to see, for example by induction on the number of crescents and using Definition 2.1, that

$$\Psi_t \circ \Phi_{t_0}|\mathbb{S}^1 = \Phi_t|\mathbb{S}^1, \tag{4.5.a}$$

provided $t \in \mathfrak{T}_\circ$ (to ensure that the righthand side of 4.5.a makes sense). A quick check, which is not a formal proof, is obtained by looking at the special cases $t = 0$ and $t = t_0$.

We want to carry out a similar construction for Ψ_t when (Λ, μ) is a general measured lamination with non-negative transverse measure. We require that 4.5.a should continue to be satisfied, provided that $\|\mu\| = 1$ and $t \in \mathfrak{T}_\circ$. The natural way to approach this is to take a sequence of finite measured laminations (Λ_n, μ_n) which converges to (Λ, μ) in an appropriate sense, and then show that the corresponding sequence of maps $(t, z) \mapsto \Psi_{n,t}(z)$ converges uniformly on compact subsets of $\mathbb{C} \times \Omega_0$.

Such a proof by approximation, though presumably possible, is not so easy to carry out entirely within $\mathbb{S}^2$, because Möbius transformations do not preserve a metric on $\mathbb{S}^2$. It is much easier to carry out the proof in hyperbolic 3-space, since the hyperbolic metric is preserved by Möbius transformations. Even better, the necessary analysis has already been carried out in this case, in [7, Chapter 3], and so, if we can make the appropriate connection with this previous work, all we need to do is to cite it. We now proceed to set up the apparatus which enables us to carry out the construction of Ψ_t in hyperbolic 3-space instead of in $\mathbb{S}^2$. But first we need to make a point about approximations by finite measured laminations. □

As explained near the beginning of Section 3, there is a topology on the space of measured laminations, so that we know what is meant by an approximating finite measured lamination.

4.6 Lemma. *Let (Λ, μ) be a geodesic lamination with non-negative transverse measure. Suppose that $\|\mu\| = 1$. Then μ can be approximated by finite measured laminations μ_n with $\|\mu_n\| = 1$.*

Proof: The danger we have to avoid is that, in some situations, an obvious approximation can have norm around two, instead of around one. Suppose, for example, that γ is a geodesic. We mark on γ a basepoint 0 and all points on γ at an integer distance from 0. We set Λ equal to the set of geodesics orthogonal to γ and passing through some marked integral point. We take $\mu(\ell) = 1$, for each $\ell \in \Lambda$. We define Λ_n to consist of $2n+1$ geodesics, all orthogonal to γ, through points on γ at distance $j(1 - \frac{1}{n^2})$ from 0, for $0 \leq j \leq n$. We set $\mu_n(\ell) = 1$ if $\ell \in \Lambda_n$. Then (Λ_n, μ_n) converges to (Λ, μ). But, for each n, $\|\mu_n\| = 2$. The situation is shown in Figure 4.6.i.

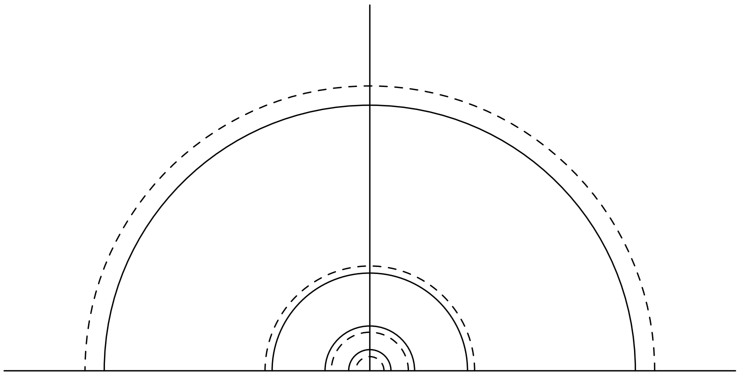

Figure 4.6.i. We show a picture in the upper halfplane with the hyperbolic metric. The dotted semicircles are geodesics a hyperbolic distance one apart. The solid semicircles are approximating geodesics. The measure of each leaf in the original (dotted) lamination is one, and the norm is then one. In the approximating lamination each (solid) geodesic also has measure one, and the approximating lamination has norm two. This picture is relevant to discussion in the proof of Lemma 4.6.

Here is a sketch of a more careful way to approximate which avoids this kind of problem. Let 0 be the centre of $\mathbb{D}^2$ and let B_n be the open disk of radius n with centre 0. If $0 \in \ell \in \Lambda$, then we can assume that $\mu(\ell) = 0$. We construct (Λ_n, μ_n) as follows.

4.6.1. All leaves of Λ_n meet B_n.

4.6.2. $\Lambda_n \subset \Lambda$.

4.6.3. For each boundary component V of $\mathbb{D}^2 \setminus \Lambda$, such that ∂V has at least three components meeting B_n, include in Λ_n all components of ∂V meeting B_n.

4.6.4. Include in Λ_n all leaves $\ell \in \Lambda$, such that $\ell \cap B_n \neq \emptyset$ and $\mu(\ell) \geq 1/n$.

4.6.5. By inserting additional elements of Λ into Λ_n, arrange that, for each open geodesic arc $\gamma \subset B_n$ disjoint from Λ_n, we have $\mu(\gamma) < 1/n$.

4.6.6. For each $\ell \in \Lambda_n$, choose $p \in \ell \cap B_n$, and let $[p, q) \subset [p, 0)$ be a half-open geodesic arc, such that (i) q either lies on a g-leaf of Λ_n or $q = 0$ and (ii) $(p, q) \cap \bigcup \Lambda_n = \emptyset$. Define $\mu_n(\ell) = \mu([p, q))$.

For each open geodesic interval α of length 1, it now follows that

$$\mu(\alpha) - 2/n < \mu_n(\alpha) < \mu(\alpha) + 1/n. \tag{4.6.a}$$

The reason the inequality is true is illustrated in Figure 4.6.ii.

It follows that $1 - 2/n \leq \|\mu_n\| \leq 1 + 1/n$. The approximation we seek is $(\Lambda_n, \mu_n/\|\mu_n\|)$. ■

4.7 Normal vectors. We denote the unit tangent bundle of hyperbolic three-space by $T_1(\mathbb{H}^3)$, or by $T_1(\mathbb{D}^3)$ when we use the Poincaré disk model. Since the group of Möbius transformations can be identified with the group of orientation-preserving isometries of hyperbolic 3-space, the Möbius transformations act on $T_1(\mathbb{D}^3)$. We have the real analytic map $\exp_+ : T_1(\mathbb{D}^3) \to \mathbb{S}^2$, which assigns to a unit tangent vector u, based at a point $p \in \mathbb{D}^3$, the end at infinity of the unit speed geodesic through p with tangent vector u.

The next lemma is obvious.

4.8 Lemma. *The map $\exp_+ : T_1(\mathbb{D}^3) \to \mathbb{S}^2$ is equivariant for the group of all isometries of hyperbolic 3-space, and in particular for the group of Möbius transformations.*

We denote by $r : \Omega_0 \to \mathrm{Dome}\,(\Omega_0)$ the nearest point retraction. Given a point $z \in \Omega_0$, we obtain a unit vector $\nu(z) \in T_1(\mathbb{D}^3)$ based at $r(z)$, by taking the tangent vector of the unit speed geodesic from $r(z)$ to z. We denote the image of $\nu : \Omega_0 \to T_1(\mathbb{D}^3)$ by $\mathcal{N}$ and talk of it as the *space of normal vectors to* $\mathrm{Dome}\,(\Omega_0)$. Since $\exp_+$ maps $\mathcal{N}$ to Ω_0, we have proved the following lemma.

4.9 Lemma. *The maps $\exp_+ : \mathcal{N} \to \Omega_0$ and $\nu : \Omega_0 \to \mathcal{N}$ are inverse homeomorphisms.*

The situation of Lemma 4.9 is illustrated in Figure 4.9.i.

The map $\Psi_t : \Omega_0 \to \mathbb{S}^2$ will be defined as the composite

$$\Omega_0 \xrightarrow{\nu} \mathcal{N} \xrightarrow{\mathbf{N}_t} T_1(\mathbb{D}^3) \xrightarrow{\exp_+} \mathbb{S}^2,$$

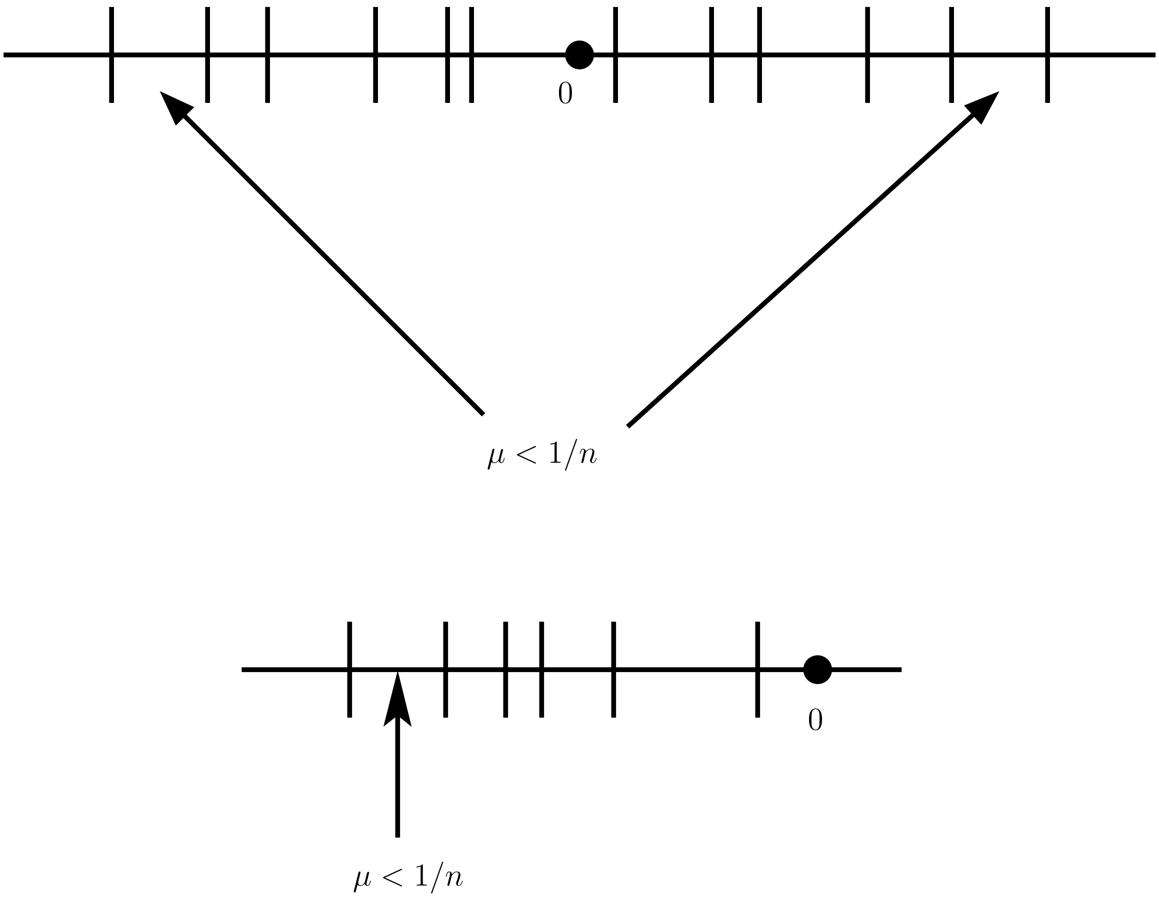

Figure 4.6.ii. These two diagrams illustrate the proof of the two inequalities of 4.6.a. In each case, we have a geodesic segment $\alpha \subset B_n$ of length one. We mark the intersections of α with the geodesics of Λ_n. Each subinterval shown has measure less than $1/n$, and this is indicated by the arrows. The upper diagram shows the situation when α contains 0. We see that the upper diagram leads to the inequalities $\mu_n(\alpha) \leq \mu(\alpha) < \mu_n(\alpha)+2/n$. The other possible configuration, shown in the lower diagram, is that 0 does not lie in α, In this case we have $\mu_n(\alpha) - 1/n < \mu(\alpha) < \mu_n(\alpha) + 1/n$. The inequalities of 4.6.a follow.

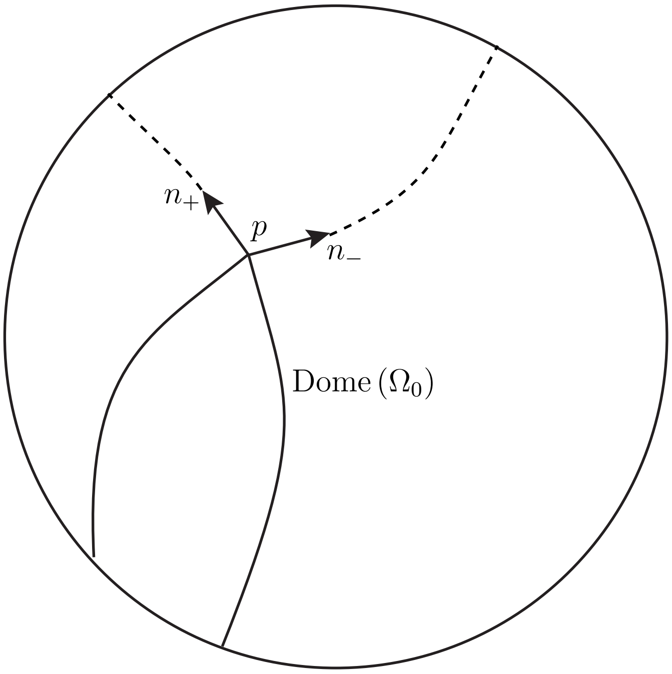

Figure 4.9.i. This picture illustrates Lemma 4.9. It can be thought of as the intersection of the general configuration with a hyperbolic plane orthogonal at a fixed point p to the bending line through p. The set of unit normal vectors (elements of $\mathcal{N}$) at p is, topologically, a closed interval. We show the two endpoints of this interval—the two normal vectors n_+ and n_-. The dotted lines are geodesics leading to the corresponding points in Ω_0.

where the map $\mathbf{N}_t : \mathcal{N} \to T_1(\mathbb{D}^3)$ will be induced by the complex earthquake map $\mathbb{C}E_t$, in a way which we have not yet made explicit.

We call $\mathcal{N}$ the *space of normal vectors* to Dome(Ω_0). Note that a normal vector based at $p \in$ Dome(Ω_0) is orthogonal to one of the support planes at p for the hyperbolic convex hull of $\mathbb{S}^2 \setminus \Omega_0$, and that it points towards Ω_0.

Let N_p denote the set of such normal vectors at p. Most points of Dome(Ω_0) have only one normal vector. But if p lies on a bending line ℓ with positive bending measure $v_0\mu(\ell)$, then N_p corresponds to an arc in the unit tangent bundle $T_1(\mathbb{H}^3)$. The length of the arc is equal to the bending measure. □

4.10 The map $\mathbf{N}_t : \mathcal{N} \to T_1(\mathbb{D}^3)$. We have seen in § 3.5 how, for each $t \in \mathbb{C}$ and each V, a gap or bending line of measure zero, we can determine a Möbius transformation $M_{t,V}$. These transformations are used to define $\mathbf{N}_t$.

Let $n \in \mathcal{N}$ be a unit normal vector based at a point $p \in$ Dome(Ω_0). Suppose $p \in \mathbb{C}E_{t_0}(V)$, where V is a gap or a bending line of measure zero. Then we define $\mathbf{N}_t(n) = M_{t,V} \circ M^{-1}_{t_0,V}(n)$.

Now suppose $p \in$ Dome(Ω_0) is a point on a bending line of positive measure. That is, $p \in \mathbb{C}E_{t_0}(\ell)$, where $\ell \in \Lambda$ and $\mu(\ell) > 0$. We orient ℓ so that it makes sense to consider its left and right sides. There are two canonical unit normal vectors based at p, which we denote by n_+ and n_-. The first of these is the limit of unit normal vectors based at points $p_i \in$ Dome(Ω_0) converging to p from the right. The second is the limit with respect to convergence from the left. Since the limit from the right (or left) of support planes has a well-defined limit, it follows that n_+ and n_- are well-defined. The angle between n_- and n_+ is equal to $v_0\mu(\ell) > 0$, where $t_0 = iv_0$. The unit vectors n, n_- and n_+ are coplanar in the tangent 3-space at p.

Let θ be the ratio of the angle between n and n_- to the angle between n_+ and n_-. Then θ plays the same role here that it played in Definition 2.1, and $0 \le \theta \le 1$. Let $M_{t,-}$ be the limit of the Möbius transformations $M_{t,V}$, as V approaches ℓ from the left through gaps or bending lines of measure zero. We define

$$\mathbf{N}_t(n) = M_{t,-} \circ A_{\theta(t-t_0)\mu(\ell)} \circ M^{-1}_{t_0,-},$$

where, for $s \in \mathbb{C}$, A_s is defined to be the Möbius transformation with complex translation length s and oriented axis ℓ lying in the equatorial plane $\mathbb{D}^2 \subset \mathbb{D}^3$.

To summarize, when the complex earthquake $\mathbb{C}E_t$ gives a method of defining $\mathbf{N}_t(n)$ which is unambiguous, then we use the unambiguous formula. Note that the answer is unambiguous exactly at the points where $\mathbb{C}E_t$ is differentiable, and at such points $\mathbf{N}_t$ could be defined using the derivative. When the complex earthquake gives an answer which is not unique, then we interpolate in the only possible reasonable manner.

This completes our description of the continuous map $\mathbf{N}_t : \mathcal{N} \to T_1(\mathbb{D}^3)$. It is defined for all $t \in \mathbb{C}$. □

Ψ_t has now been defined in two different ways when Λ is a finite lamination. The first way was described in § 4.5, and the second way is described just after the statement of Lemma 4.9. The next lemma states that these two ways give us the same answer.

4.11 Lemma. *The composite* $\Omega_0 \xrightarrow{\nu} \mathcal{N} \xrightarrow{\mathbf{N}_t} T_1(\mathbb{D}^3) \xrightarrow{\exp_+} \mathbb{S}^2$ *is equal to* Ψ_t *for finite laminations.*

The proof is an easy induction, using the description of a scaling map in Definition 2.1 and Lemma 2.5.

In order to define the Möbius transformations $M_{t,V}$ associated to a general measured lamination (Λ, μ), as in § 3.5, we approximate by a sequence of finite measured laminations (Λ_n, μ_n). The Möbius transformations $M_{t,V}$ are defined by this limiting process. The limit is well-defined if $V \subset \mathbb{D}^2$ is a gap or a bending line of measure 0. However, if V is a bending line ℓ of positive measure, the limiting values are ambiguous, and the limit value of $M_{t,V}$ can be varied by premultiplication by $A_{t\theta\mu(\ell)}$. Here $0 \le \theta \le 1$ and, for $s \in \mathbb{C}$, A_s is the Möbius transformation with complex translation length s and oriented axis ℓ lying in the equatorial plane $\mathbb{D}^2 \subset \mathbb{D}^3$.

Clearly $A_{t\theta\mu(\ell)}$ depends holomorphically on t. If V is a gap or a bending line of measure zero, then the Möbius transformation $M_{t,V}$ also depends holomorphically on t—this follows from the fact that the limit of a sequence of holomorphic functions which is uniformly convergent on compact subsets is holomorphic.

We have proved the following result.

4.12 Lemma. *For fixed* $n \in \mathcal{N}$, $\exp_+(\mathbf{N}_t(n)) \in \mathbb{S}^2$ *varies holomorphically with* $t \in \mathbb{C}$. *For fixed* $z \in \Omega_0$, $\Psi_t(z) \in \mathbb{S}^2$ *varies holomorphically with* $t \in \mathbb{C}$.

Most of the following results have already been proved.

4.13 Theorem. *Let* (Λ, μ) *be an arbitrary non-negative measured lamination with* $\|\mu\| = 1$. *Fix* $t_0 = iv_0 \in \mathfrak{T}_0$ *and denote by* Ω_0 *the image of* $\mathbb{D}^2$ *under the map* $\Phi_{t_0} : \mathbb{S}^2 \to \mathbb{S}^2$. *We denote by* Ω_t *the image of* $\mathbb{D}^2$ *under* Φ_t. *(A typical image is shown in Figure 4.5.i.) The normalized map (see § 2.6)* $\Psi_t : \Omega_0 \to \mathbb{S}^2$ *has the following properties.*

4.13.1. $\Psi_{t_0} = \mathrm{Id}$.

4.13.2. For each $z \in \Omega_0$, $\Psi(t, z) \in \mathbb{S}^2$ *depends holomorphically on* $t \in \mathbb{C}$.

4.13.3. For each $t \in \mathfrak{T}_\circ$, Ψ_t *can be continuously extended to* $\partial\Omega_0$ *so that Equality 4.5.a is satisfied. In particular,* $\Psi_0 : \partial\Omega_0 \to \mathbb{S}^1$ *and* $\Phi_{t_0} : \mathbb{S}^1 \to \partial\Omega_0$ *are inverse homeomorphisms.*

4.13.4. For $t \in \mathbb{U}^2 \cap \mathfrak{T}_\circ$, $\Psi_t : \Omega_0 \to \mathbb{S}^2$ *is injective, and* $\Psi_t(\Omega_0) = \Phi_t(\mathbb{D}^2) = \Omega_t$.

4.13.5. When $\mathrm{Im}(t) > 0$ *and* $t = u + iv$, $\Psi_t : \Omega_0 \to \mathbb{S}^2$ *is a locally injective* K_t*-quasiregular mapping, where*

$$K_t = \frac{1 + |\kappa(t)|}{1 - |\kappa(t)|}, \quad |\kappa(t)| = \frac{\sqrt{u^2 + (v - v_0)^2}}{\sqrt{u^2 + (v + v_0)^2}}.$$

Proof: All the statements are proved by taking a sequence $((\Lambda_n, \mu_n))_{n\in\mathbb{N}}$ of finite measured laminations of norm one, converging to (Λ, μ). On any compact subset of $\mathbb{C} \times \Omega_0$, $(t, z) \mapsto \left(\exp_+ \circ \mathbf{N}_t \circ \nu\right)(z)$ converges uniformly as n tends to infinity. Note that both domain and range of ν and $\mathbf{N}_t$ vary with n.

To prove 4.13.2, note that if a sequence of holomorphic functions converges uniformly on compact subsets, then its limit is also holomorphic.

The proofs of 4.13.3 and of 4.13.4 follow from Equality 4.5.a, since both $\Phi_t|\mathbb{S}^1$ and Ψ_t are defined through limits of finite measured laminations. The injectivity of Ψ_t for $t \in \mathfrak{T}_\circ$ follows from a degree argument.

One proves 4.13.5 by first showing that it is true for finite laminations, and then taking limits. In particular, the claimed local injectivity follows from the fact that locally we have quasiconformal homeomorphisms with bounded constant. To ensure that the limit is not degenerate, we can normalize (see § 2.6) using a leaf that passes through the small neighbourhood in which we are interested.

Thus, what has been accomplished, through Theorem 4.13 and Theorem 4.4, is the following. We have constructed a holomorphic family of generalized scaling maps $\mathbb{C} \times \Omega_0 \to \mathbb{S}^2$ which for $t \in \mathfrak{T}_\circ$ are related to the boundary values of complex earthquakes $\partial\mathbb{D}^2 \to \partial\mathrm{Dome}(\Omega_t)$ by 4.13.4 and 4.13.3.

5 Sullivan's Theorem: a careful statement

Let $\Omega \subset \mathbb{C} \subset \mathbb{S}^2$ be a simply connected region which is not the whole plane. We recall that Dome$\,(\Omega)$ denotes the boundary of the hyperbolic convex hull of $\mathbb{S}^2 \setminus \Omega$. By a theorem of Thurston, if the dome is given its path metric induced from the hyperbolic metric on $\mathbb{H}^3$, then the dome is isometric to a hyperbolic plane. We begin by stating Sullivan's Theorem in its usual form.

5.1 Theorem. *There is a universal constant $K > 1$ with the following property. Let Ω be a simply connected proper subregion of $\mathbb{C}$. Let Γ be the group of Möbius transformations which preserve Ω. Then there is a Γ-equivariant K-quasiconformal homeomorphism $f_\Omega : \Omega \to \mathrm{Dome}\,(\Omega)$ which extends continuously to the identity map on the common boundary in $\mathbb{S}^2$ of domain and range.*

In interpreting the theorem, there is one exceptional case: Namely when $\Omega \subset \mathbb{S}^2$ is the complement of a circular arc, say the positive real axis $\mathbb{R}_+ \subset \mathbb{C}$, then Dome$\,(\Omega)$ must be interpreted as the two sides of the hyperbolic halfplane rising from $\mathbb{R}_+$. Each point of the halfplane needs to be thought of as giving rise to two points of the dome, one for each side of the halfplane.

This simple example exposes a subtle problem with the stated hypothesis.

Quasiconformal homeomorphisms are normally mappings between open subsets of $\mathbb{S}^2$, and are necessarily orientation preserving. But in the statement of the theorem, the domain and range of f_Ω are in different spaces, and so we have to specify an orientation for each.

If the complement of Ω is a topological arc, it may be possible to "reflect" topologically across the arc, obtaining a solution in the wrong homotopy class. Consider for example the logarithmic spiral with equation $r = e^\theta$ in polar coordinates. The closure A of the spiral is a topological arc in $\mathbb{S}^2$. Set $\Omega = \mathbb{S}^2 \setminus A$. Consider the homeomorphism $g : \mathbb{C} \setminus \{0\} \to \mathbb{C} \setminus \{0\}$ given by $g(r,\theta) = (r, 2\log(r) - \theta)$, using polar coordinates. This extends to a homeomorphism, also denoted by g, of $\mathbb{S}^2$ which fixes the spiral pointwise. In a neighbourhood of a point of the spiral, g interchanges the sides of $\partial\Omega$ locally. Composition with g, combined with changing the orientation of Dome$\,(\Omega)$, gives us, for every standard solution $f_\Omega : \Omega \to \mathrm{Dome}\,(\Omega)$, a strange solution $f_\Omega \circ g : \Omega \to \mathrm{Dome}\,(\Omega)$.

This raises the concern that there might perhaps be an infinite number of distinct homotopy classes of solutions to Sullivan's Theorem for some regions where the boundary has bad local connectedness properties. Here we are talking of homotopy classes of maps $f : \overline{\Omega} \to \overline{\mathrm{Dome}\,(\Omega)}$, where the homotopy restricted to $\partial\Omega$ is always the identity map. We will consider some foundational questions concerning solutions to Sullivan's Theorem, with the objective of showing that this concern is unfounded.

We first recall some standard definitions.

5.2 Definition. The topological definition and basic properties of prime ends are given in [2, Theorem 4.3]. Fix a basepoint $z_0 \in \Omega$. A *crosscut* is a closed arc in $\overline{\Omega}$, with the property that only its endpoints lie in $\partial\Omega$. A *fundamental sequence* in Ω is (not a sequence but) an infinite countable set A of points of Ω, with the property that, for each $\epsilon > 0$ there exists a crosscut of diameter less than ϵ, separating all except a finite number of points of A from z_0. Two such sets are said to be *equivalent* if their union is also a fundamental sequence. This is an equivalence relation. A *prime end* is an equivalence class of fundamental sequences. The *impression* of the prime end is the intersection of the closures of the regions cut off by smaller and smaller crosscuts, and containing all except a finite number of elements of A. There is no loss of generality in restricting to crosscuts whose endpoints do not lie in the impression. $\square$

5.3 Lemma. *The nearest point retraction $r : \Omega \to \mathrm{Dome}\,(\Omega)$ extends continuously to a homeomorphism of the spaces of prime ends.*

The definition of prime ends given above can be equally well applied to the dome, using the euclidean diameter in $\mathbb{D}^3$.

Proof: The nearest point retraction is a quasi-isometry. Any quasi-isometry from $\mathbb{D}^2$ to itself induces a homeomorphism of $\partial\mathbb{D}^2$. A point of $\mathbb{S}^1 = \partial\mathbb{D}^2$ is a prime end. If the quasi-isometry is continuous, then with its extension to the boundary it is continuous on the closure $\overline{\mathbb{D}}^2$. $\blacksquare$

5.4 Theorem. *Let Ω be an open simply connected subset of $\mathbb{S}^2$ and let f be a homeomorphism of $\overline{\Omega}$ with itself such that the restriction of f to $\partial\Omega$ is the identity. Then*

5.4.1. EITHER $f|\Omega$ extends continuously to the space of prime ends, and, on the space of prime ends is the identity map

5.4.2. OR f reverses orientation and $\partial\Omega$ is a topological arc or a point. The case of a point cannot arise in this paper, because we do not allow $\Omega = \mathbb{C}$.

Proof: Extend f to $\mathbb{S}^2$ by defining it to be the identity in the complement of $\overline{\Omega}$. The fact that f extends to the space of prime ends follows from the uniform continuity of f. Moreover the extension to the space of prime ends is continuous. Let ω be a prime end of Ω, and let p be a point in the impression of ω. Let U be a small round open disk with centre p, and let V be a smaller concentric open disk with centre p such that $fV \subset U$.

To this situation, we can apply [6, Theorem 4.1], which we now state.

5.5 Theorem. *Let M be a connected n-manifold. Let X be a closed subset and let U be an open subset of M. Let $f : U \to M$ be an embedding, such that $f|X \cap U = \mathrm{Id}$. Let V be an open connected subset of M, such that $V \cap X \neq \emptyset$ and $V \cup fV$ is orientable, and is contained in U. We suppose that f preserves orientations on $X \cap U$. Then, for each $y \in V \setminus X$, y and fy lie in the same component of $U \setminus X$.*

We apply this theorem with $n = 2$, $M = \mathbb{S}^2$, $X = \partial\Omega$, and with V, U and f as above. Since U and V are arbitrarily small, it is easy to deduce from Theorem 5.5 that if f preserves orientation, then f must fix each prime end.

So we now assume that f reverses orientation. Since f is the identity on $\mathbb{S}^2 \setminus \overline{\Omega}$, this subset must be empty. We can now apply [6, Theorem 2.5], which we now state in a form applicable to our situation.

5.6 Theorem. *Let $f : \mathbb{S}^2 \to \mathbb{S}^2$ be an orientation reversing homeomorphism, such that f fixes $\partial\Omega$ pointwise and sends Ω to itself, as above. Then $\partial\Omega$ is an interval or a point.*

This completes the proof of Theorem 5.4. ∎

In proving Sullivan's Theorem (Theorem 5.1), we are looking for a particular homeomorphism $f_\Omega : \Omega \to \mathrm{Dome}\,(\Omega)$. We don't yet know f_Ω, but we want at least to determine some elementary topological properties that it must have in order to satisfy Theorem 5.1. We will do this by applying Theorem 5.4. However, that result applies to homeomorphisms from a region to itself. This leads us to seek a standard homeomorphism $\sigma : \Omega \to \mathrm{Dome}\,(\Omega)$, which would enable us to apply Theorem 5.4 to $\sigma^{-1} f_\Omega$.

If the convex hull of $\mathbb{S}^2 \setminus \Omega$ has no interior, it is contained in a hyperbolic plane. Then we know that Ω is a slit plane. In this case there is a unique solution for f_Ω which minimizes the maximal dilatation K among all quasiconformal maps with the same boundary values. The extremal map can be found explicitly (it is the appropriate map introduced in §2).

So, from now on, we assume that the convex hull has non-empty interior. Given a point q in the interior of the convex hull, let $r_q : \Omega \to \mathrm{Dome}\,(\Omega)$ be the radial projection along geodesics through q. Then r_q is a homeomorphism which extends continuously to a homeomorphism, again denoted by $r_q : \overline{\Omega} \to \overline{\mathrm{Dome}\,(\Omega)}$, which is the identity on $\partial\Omega$. We illustrate r_q in Figure 5.6.i.

Clearly $r_q(z)$ is continuous in the pair of variables (q, z). Therefore, varying q causes r_q to change by an isotopy. It follows that the effect on the prime ends of any two such radial projections is the same. This observation allows us to prove the following result.

5.7 Lemma. *The nearest point retraction* $r : \Omega \to \mathrm{Dome}\,(\Omega)$ *and any radial map* $r_q : \Omega \to \mathrm{Dome}\,(\Omega)$ *have the same effect on prime ends: Each extends continuously to a homeomorphism between the spaces of prime ends and their values on the prime*

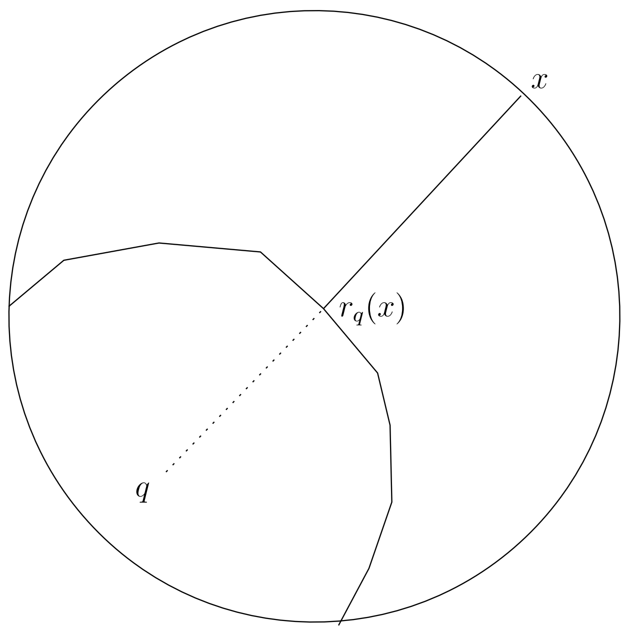

Figure 5.6.i. This illustration is in the projective model, where hyperbolic convex sets are the same as euclidean convex sets. This makes it easier to draw hyperbolic convex sets. We show the centre of projection q. In the projective model, hyperbolic radial projection from q is the same as euclidean radial projection. We indicate the effect of r_q on $x \in \Omega$.

ends are identical.

Proof: If ℓ is a bending line, then $r^{-1}(\ell)$ is a crescent which is foliated by circular arcs. The bisector of this crescent is a circular arc. Once a positive direction for ℓ has been specified, the positive endpoint of ℓ determines a prime end p_1 for Dome (Ω). We orient the bisector so that r sends the orientation of the bisector to the orientation of ℓ. The bisector defines a prime end p_2 for Ω and r sends p_2 to p_1.

The bisector together with ℓ determine a hyperbolic plane in $\mathbb{U}^3$ containing ℓ and orthogonal to the bisector. We choose q to lie in this plane. Then $r_q(p_2) = p_1$. By the observation made before the statement of this lemma, we obtain the same equality for any other radial projection.

A similar argument works for a boundary point of a flat. But boundary points of flats together with endpoints of bending lines are dense in $\partial\Omega$, and dense in the space of prime ends. It follows that r and radial projection have the same effect on all prime ends. From Lemma 5.3 we see that the extensions of r and r_q are homeomorphisms giving rise to the same homeomorphism between the spaces of prime ends. ∎

We orient Dome (Ω) so that radial projection from Ω preserves orientation.

5.8 Lemma. *Suppose $\Omega \subset \mathbb{S}^2$ is a simply connected region and $f : \mathrm{Dome}\,(\Omega) \to \Omega$ is any homeomorphism that continuously extends to $\partial\Omega$, where it is the identity. Assume that $\partial\Omega$ is not a topological arc or a point. Then f is orientation preserving and it has the same effect on prime ends as the nearest point retraction and as radial projection.*

Proof: Composing with the inverse of radial projection, we obtain a homeomorphism from $\overline{\Omega}$ to itself, which is the identity on $\partial\Omega$. To this homeomorphism apply Theorem 5.4. This shows that f has the same effect on prime ends as the nearest point retraction. ∎

To prove Sullivan's theorem, we will work consistently with the following analytic characterization of the homeomorphisms of interest.

5.9 Theorem. *Suppose $f : \Omega \to \mathrm{Dome}\,(\Omega)$ is a quasiconformal or anti-quasiconformal homeomorphism. Then the following two conditions are equivalent:*

5.9.1. f extends to the identity on $\partial\Omega$ and f preserves orientation.

5.9.2. $d(f(z), r(z))$ *is uniformly bounded for* $z \in \Omega$. *Here* r *is the nearest point retraction and* d *is the path metric in the dome.*

If Ω *is a Jordan region, then the above conditions are automatically true.*

Proof: In one direction, the result is immediate: Suppose $d(f(z), r(z))$ is uniformly bounded. Since r extends continuously to the identity on $\partial\Omega$, the same is true for f. Moreover, f and r, both extended, are homotopic by a homotopy which is the identity on $\partial\Omega$, and therefore along with r, f preserves orientation.

The proof in the reverse direction requires the following well-known result. For completeness, we will outline a proof.

5.10 Lemma. *Let* $f, g : \mathbb{D}^2 \to \mathbb{D}^2$ *be two quasi-isometries with respect to the hyperbolic metric, which have the same boundary values. Then there is a constant* c*, such that, for each* $z \in \mathbb{D}^2$, $d(f(z), g(z)) < c$. *The constant* c *depends only on the constants of quasi-isometry.*

Proof: The following proof is illustrated in Figure 5.10.i.

Given a point $z \in \mathbb{D}^2$, we take two orthogonal geodesics α and β through z. Since f and g are quasi-isometries, they have well-defined, equal, extensions to $\mathbb{S}^1 = \partial\mathbb{D}^2$. We continue to denote the extended maps from the closed disk to itself by f and g respectively. We denote by $\alpha_f = \alpha_g$ the geodesic which connects the image under f of the endpoints of α and similarly for $\beta_f = \beta_g$. The angle between the geodesics α_f and β_f is bounded away from 0 by a positive constant, depending only on the constants of quasi-isometry. For each $\epsilon > 0$, the intersection of the ϵ-neighbourhoods of α_f and β_f has compact closure. We can choose ϵ, depending only on the quasi-isometry constants, so that the ϵ-neigbourhood of α_f contains $f(\alpha) \cup g(\alpha)$ and the ϵ-neighbourhood of β_f contains $f(\beta) \cup g(\beta)$. The diameter c of the intersection of these two neighbourhoods can be bounded in terms of the constants of quasi-isometry. Both $f(z)$ and $g(z)$ lie in this bounded set.

We can now complete the proof of Theorem 5.9. Let $f : \Omega \to \mathrm{Dome}\,(\Omega)$ be a quasiconformal homeomorphism which extends to the identity on $\partial\Omega$ and which preserves orientations. Then by Lemma 5.8, f and r have the same effect on prime ends. By [9, Theorem 5.1], every K-quasiconformal homeomorphism $h : \mathbb{D}^2 \to \mathbb{D}^2$ is a $(K, a(K))$-quasi-isometry. Also r is a quasi-isometry. We apply Lemma 5.10 to complete the proof. ∎

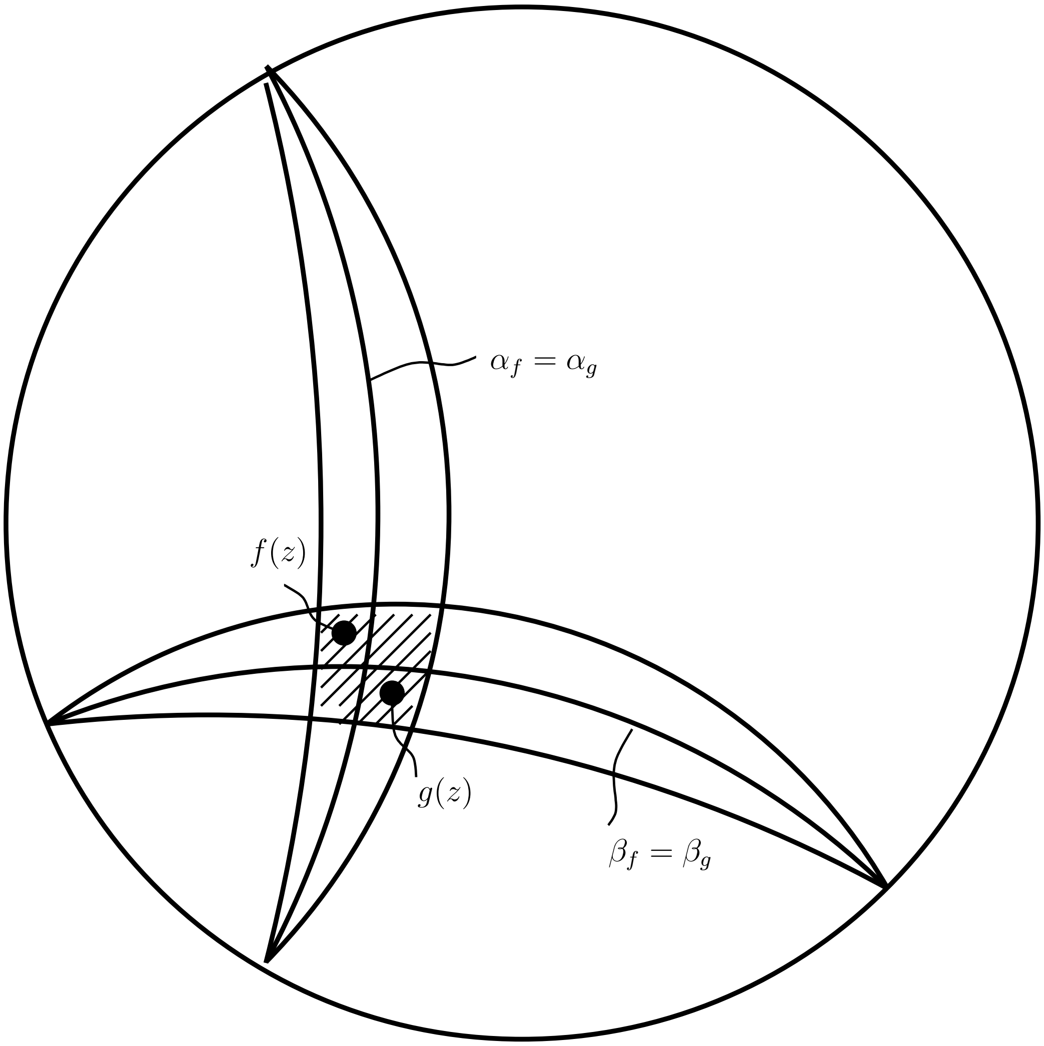

Figure 5.10.i. This picture illustrates the proof of Lemma 5.10. We illustrate the geodesics $\alpha_f = \alpha_g$ and $\beta_f = \beta_g$ and their ϵ-neighbourhoods.

6 Sullivan's Theorem: a new proof

This section is devoted to a new proof of Sullivan's Theorem (Theorem 5.1). In [13] we find a half-page sketch proof. A complete proof, with an estimate $K < 82.8$, appeared in [7] and a simpler proof, without any estimates was suggested by Bishop and presented in [8]. Our original proof, involving difficult estimates of a differential geometric nature, was quite complicated. The proof we now give is much simpler and seems to us to be the most natural, involving exactly those elements of the situation which have to be involved. It also allows for much improved estimates.

6.1 Teichmüller space. We will use universal Teichmüller space $\mathcal{T}$. This is the space of all quasisymmetric homeomorphisms of $\mathbb{S}^1$, modulo the action of the group of Möbius transformations by composition on the left. This quotient can be equally be taken as the set of quasisymmetric homeomorphisms of $\mathbb{S}^1$ which fix three points. As usual, we will take these points to be -1, 1 and i. This space can be given a suitable topology and complex structure—see [1, Chapter 6]. The basic result we need is that the map from the space of Beltrami differentials, that is the open unit ball in $L^\infty(\mathbb{D}^2)$, to $\mathcal{T}$, given by solving the Beltrami equation, is holomorphic.

Let $g : \mathbb{D}^2 \to \mathbb{S}^2$ be a locally injective quasiregular map; we recall that this means that $g = h \circ f$ where f is a quasiconformal homeomorphism and h is locally injective and holomorphic on the image of f. The locally injective property of g can be used to transfer the complex structure on $\mathbb{S}^2$ to $\mathbb{D}^2$ using g. We obtain a complex structure C_g on $\mathbb{D}^2$. We then have a quasiconformal homeomorphism $\hat{g} : \mathbb{D}^2 \to C_g$.

Let $R : C_g \to \mathbb{D}^2$ be a Riemann mapping. $R \circ \hat{g} : \mathbb{D}^2 \to \mathbb{D}^2$ is a quasiconformal homeomorphism. By choosing R appropriately, we can assume that the three points $-1, 1$ and i are fixed. The map $\hat{g}$ has a homeomorphic extension to $\partial\mathbb{D}^2$. Its restriction to $\partial\mathbb{D}^2$ is quasisymmetric. We will denote this normalized boundary function by $qs(g)$.

In this way the locally injective quasiregular g determines a point $qs(g) \in \mathcal{T}$. If g is itself a normalized quasiconformal map of $\mathbb{D}^2$ onto itself, then $qs(g)$ is given by the boundary values of g. □

6.2 The program. Return to our geodesic lamination $\Lambda \subset \mathbb{D}^2$ with a transverse non-negative measure μ satisfying $\|\mu\| = 1$. Our program is as follows.

6.2.1. Let $\mathfrak{T}_\mathrm{o} \subset \mathbb{C}$ be the open subset defined in Definition 4.3. Given $t \in \mathfrak{T}_\mathrm{o}$, according to Theorem 4.4 there is a quasiconformal map

$$f_t = \Phi_{(\Lambda,t\mu)} : \mathbb{D}^2 \to \Omega_t \subset \mathbb{S}^2.$$

The map extends to a homeomorphism of the closed disk $\overline{\mathbb{D}^2}$ onto $\overline{\Omega_t}$, which is therefore a quasidisk. We will make use the map

$$F : t \in \mathfrak{T}_\circ \mapsto qs(f_t) \in \mathcal{T}.$$

6.2.2. From Theorem 4.13, for each point in the upper halfplane $\{t : \mathrm{Im}(t) > 0\}$, there is a locally injective quasiregular map $\Psi_t = \Psi(t, \cdot) : \Omega_0 = \Omega_{t_0} \to \mathbb{S}^2$. The map Ψ depends on a choice of $t_0 \in \mathfrak{T}_\circ$ lying on the positive imaginary t-axis.

Consider the locally injective quasiregular map

$$g_t = \Psi_t \circ \Phi_{t_0} : \mathbb{D}^2 \to \Omega_0 \to \mathbb{S}^2,$$

and the associated map

$$G : \mathbb{U}^2 \to \mathcal{T} \text{ given by } G(t) = qs(g_t).$$

6.2.3. We prove below that $F = G$ on their common domain

$$\{t \in \mathfrak{T}_\circ : \mathrm{Im}(t) > 0\}.$$

6.2.4. We prove below that F is holomorphic for t in its full domain $\mathfrak{T}_\circ$.

6.2.5. It is harder to prove that G is holomorphic for $t \in \{\mathrm{Im}(t) > 0\}$. We will prove this using some more abstract methods.

6.2.6. Putting the above material together means that G is the holomorphic extension of F to $\mathfrak{T} = \mathfrak{T}_\circ \cup \{t : \mathrm{Im}(t) > 0\}$, as in Definition 4.3. We will finally show how the existence of G allows us to deduce Sullivan's Theorem. □

6.3 Proof that $F = G$ on $\{t \in \mathfrak{T}_\circ : \mathrm{Im}(t) > 0\}$. From Theorem 4.13, especially (4.13.4), we see that f_t and g_t, though not equal, have the same boundary values and quasidisk image Ω_t. Therefore $qs(f_t) = qs(g_t)$. Thus F and G are identical on their common domain, as announced in Step 6.2.3. □

6.4 Proof that F is holomorphic. This is essentially a consequence of the following theorem.

6.5 Theorem. *[3]. Let $Q : (t, z) \in \mathbb{D}_r \times \mathbb{D}_s \mapsto Q(t, z) \in \mathbb{S}^2$ be a continuous map, where $\mathbb{D}_r, \mathbb{D}_s$ are open disks about the origin with radii r, s, respectively. Assume that*

- *For each $t \in \mathbb{D}_r$, the map $z \mapsto Q(t,z)$ is injective, where $z \in \mathbb{D}_s$.*
- *For each $z \in \mathbb{D}_s$ the map $t \mapsto Q(t,z)$ is holomorphic, where $t \in \mathbb{D}_r$.*
- *$z \mapsto Q(0,z)$ is a quasiconformal map of $\mathbb{D}_s$.*

Then for each $t \in \mathbb{D}_r$, the map $z \mapsto Q(t,z)$ is quasiconformal on $\mathbb{D}_s$. Moreover the complex dilatation with respect to z, namely

$$\kappa(t,z) = \frac{\partial Q/\partial \overline{z}}{\partial Q/\partial z},$$

has the following property. For fixed $z \in \mathbb{D}_s$, the function $t \mapsto \kappa(t,z) \in L^\infty(\mathbb{D}_s)$ is holomorphic, $t \in \mathbb{D}_r$.

Actually this is a slight strengthening of [3] in that we are not assuming that $Q(0,z) = z,\ z \in \mathbb{D}_s$. We will need the stronger form. However the stronger form as stated follows immediately from the expressions in [1] for the complex dilatation of the composition of two quasiconformal homeomorphisms.

We apply Theorem 6.5 to the holomorphic motion resulting from Theorem 4.4. This gives a holomorphic map from $\mathfrak{T}_o$ to the unit ball in $L^\infty(\mathbb{D}^2)$, by taking the complex dilatation of $\Phi_t = \Phi_{(\Lambda,t\mu)} : \mathbb{D}^2 \to \mathbb{S}^2$. The quotient map from the unit ball in $L^\infty(\mathbb{D}^2)$ to $\mathfrak{T}$ is also holomorphic. This shows that F is holomorphic. □

Now we come to Step 6.2.5. We first prove a local version of what we need.

6.6 Lemma. *Let $U \subset \{t : \mathrm{Im}(t) > 0\}$ be a relatively compact, open neighborhood of the basepoint $t_0 = iv_0$. Given $z_1 \in \mathbb{D}^2$, there exists a neighborhood $z_1 \subset N_1 \subset \mathbb{D}^2$ with the following two properties: For each $t \in U$, the map $z \mapsto f_t(z)$ is injective for $z \in N_1$. For each $z \in N_1$, the map $t \mapsto f_t(z)$ is holomorphic for $t \in U$.*

Proof: Recall from (6.2) that $f_t = \Phi_t$. We will determine N_1 by determining its image $N_0 = f_{t_0}(N_1) \subset \Omega_{t_0} = \Omega_0$. Recall that Ω_0 has a partial foliation or lamination by c-leaves. Each c-leaf is a circular arc contained in the inverse image of a bending line under the nearest point retraction $r_0 : \Omega_0 \to \mathrm{Dome}\,(\Omega_0)$. This foliation has a transverse measure induced from the bending measure $v_0\mu$ of $\mathrm{Dome}\,(\Omega_0)$.

A given point $z_0 \in \Omega_0$ lies either on a c-leaf, or it lies in a gap. For each $t \in U$, there is a Möbius transformation A_t such that $A_t \circ \Psi_t : \Omega_0 \to \mathbb{S}^2$ fixes the gap or c-leaf in which z_0 lies; this is just a renormalization of Ψ.

If z_0 lies in a gap, we take N_0 to be that gap.

Suppose instead that z_0 lies on a c-leaf. Choose a neighborhood N_0 so small such that its transverse measure with respect to $t_0\mu$ is small (as N_0 shrinks to z_0, this measure tends to 0). Here we can assume that the leaves of the foliation and lamination of Ω_0 meet N_0 in a parallel family. Since $\overline{U}$ is compact, we can take N_0 small enough so that it also has small transverse $t\mu$-measure for $t \in U$.

Consequently, for each $t \in U$, $A_t \circ \Psi_t : \mathbb{D}^2 \to \mathbb{S}^2$ restricts to an embedding of N_0 in $\mathbb{S}^2$. The same is then true of Ψ_t itself. Now apply Theorem 6.5 to complete the proof. ∎

We next show how to combine the small neighbourhoods of Lemma 6.6. First we introduce some terminology.

Let $\{H_i\}_{i\in\mathbb{N}}$ be a countable collection of Banach spaces. Define the *L^∞-product* of this collection to be the space H of all bounded sequences $\{v = (v_1, v_2, \ldots)\}$ with $v_i \in H_i$. Associated with H are the projections $\pi_i : H \to H_i$. Define the norm of $v \in H$ to be $\|v\| = \sup_i\{\|v_i\|\}$.

6.7 Lemma. *Suppose we are given a region $U \subset \mathbb{C}$ and a map $h : U \to H$ with bounded image such that, for each $i \in \mathbb{N}$, the composite $h_i = \pi_i \circ h : U \to H_i$ is holomorphic. Then h is holomorphic.*

Proof: First we have to prove that h is continuous. Choose $t \in U$ and a circle $\gamma \in U$ around t. By the Cauchy Integral Formula, for each index i,

$$h_i(t) = \frac{1}{2\pi i}\int_\gamma \frac{h_i(z)}{z-t}\,dz.$$

Since h has bounded image, $\{h_i\}$ is an equicontinuous family. It follows that h is continuous.

In particular h can be integrated so that the expression

$$h(t) - \frac{1}{2\pi i}\int_\gamma \frac{h(z)}{z-t}\,dz$$

is meaningful. In fact, it is equal to zero, since, for each i, its projection to H_i is zero. The fact that h is analytic follows immediately. ∎

6.8 Corollary. *Let X be a measure space which is the union of measurable subspaces $X = \bigcup_{i\in\mathbb{N}} X_i$. Suppose we are given a region $U \subset \mathbb{C}$ and a map $h : U \to L^\infty(X)$ with bounded image such that, for each i, the composition $U \to L^\infty(X) \to L^\infty(X_i)$ is holomorphic. Then h is holomorphic.*

Proof: From Lemma 6.7, the associated map from U to the L^∞ product of the $L^\infty(X_i)$ is holomorphic. The result follows since $L^\infty(X)$ is isometric to a subspace of the L^∞ product. ■

6.9 Proposition. *The map $G : \{t : \mathrm{Im}(t) > 0\} \to \mathcal{T}$, defined by $G(t) = qs(g_t)$, is holomorphic.*

Proof: Choose a relatively compact open set $U \subset \{t : \mathrm{Im}(t) > 0\}$. Let $\{N_i\}_{i \in \mathbb{N}}$ be an open covering of $\mathbb{D}^2$, such that, for each index i and each $t \in U$, (i) the map $f_t : N_i \to \mathbb{S}^2$ is an embedding of N_i, and (ii) for each $z \in N_i$, the map is holomorphic for $t \in U$. The existence of $\{N_i\}$ is proved in Lemma 6.6.

For each $t \in U$ and each index i, the complex dilatation of f_t at $z \in \mathbb{D}^2$, namely $\kappa(t, z) \in L^\infty(\mathbb{D}^2)$, satisfies $|\kappa(t, z)| < 1$. According to Corollary 6.8, the map $\mu(t, z) : U \to L^\infty(\mathbb{D}^2)$ is a holomorphic function of $t \in U$. As pointed out in (6.1) the map from the unit ball of $L^\infty(\mathbb{D}^2)$ to $\mathcal{T}$, given by solving the Beltrami differential equation, is holomorphic.

Consequently the map $G : \{t : \mathrm{Im}(t) > 0\} \to \mathcal{T}$ (see (6.2.2)) is holomorphic. ■

This completes our discussion of Step 6.2.5.

We are finally ready to complete our proof of Sullivan's Theorem, as promised in (6.2.6).

Start with a simply connected region $\Omega \subset \mathbb{C}$, with $\Omega \neq \mathbb{C}$. Denote the bending measure of Dome$\,(\Omega)$ by $c\mu$, where $c > 0$ and $\|\mu\| = 1$.

The map $g_{ic} : \mathbb{D}^2 \to \mathbb{S}^2$ defined in (6.2.2) is known to be quasiregular. We know from (4.13.4) that $\Psi_{ic} : \Omega_0 \to \Omega_{ic}$ is a homeomorphism. From (4.13.5) we know that it is quasiconformal. In fact then, g_{ic} is a quasiconformal mapping of $\mathbb{D}^2$ onto Ω_{ic}. We have normalized consistently so that -1, 1 and i are fixed. We also have an isometry $\iota : \mathbb{D}^2 \to \mathrm{Dome}\,(\Omega_{ic})$, which fixes the prime ends corresponding to -1, 1 and i.

Furthermore, g_{ic} and the nearest point retraction $r : \Omega_{ic} \to \mathrm{Dome}\,(\Omega_{ic})$ induce inverse bijections on those prime ends which correspond to endpoints of bending lines or boundary points of flats. Therefore the homeomorphism of prime ends induced by g_{ic} is inverse to the homeomorphism of prime ends induced by r. From the results of Section 5, it further follows that

$$\iota \circ g_{ic}^{-1} : \Omega_{ic} \to \mathrm{Dome}\,(\Omega_{ic})$$

is a quasiconformal homeomorphism which extends continuously to the identity on $\partial\Omega_{ic}$.

Applying for example Theorem 5.9, we see that any quasiconformal extension $\hat{g}$ of the quasisymmetric map $G(ic) = qs(g_{ic})$ to $\mathbb{D}^2$ gives rise to a map $\iota^{-1} \circ \hat{g} \circ R : \Omega_{ic} \to \mathrm{Dome}\,(\Omega_{ic})$ of the kind demanded in Sullivan's Theorem. Here $R : \Omega_{ic} \to \mathbb{D}^2$ is a normalized Riemann map. Finding the quasiconformal homeomorphism in Theorem 5.1 is therefore the same as finding a good representative for $G(ic) \in \mathcal{T}$ (see Step 6.2.2 for notation).

The Teichmüller distance between two quasisymmetric maps $f, g : \mathbb{S}^1 \to \mathbb{S}^1$, representing points of Teichmüller space and fixing -1, 1 and i, is defined as

$$d_{\mathcal{T}}(f, g) = \log \inf K(\hat{f}^{-1} \circ \hat{g}),$$

where the infimum is taken over the maximal dilation K of all quasiconformal extensions $\hat{f}, \hat{g} : \mathbb{D}^2 \to \mathbb{D}^2$.

6.10 Definition. For each proper simply connected open subset $\Omega \subset \mathbb{C}$, let $c(\Omega)$ be the norm of the associated bending measure. Let $c_1 = \sup_\Omega c(\Omega)$, as Ω varies over all such subsets. In [5] it is shown that

$$c_1 \leq 2\pi - 2\sin^{-1}\left(\frac{1}{\cosh 1}\right) \approx 4.8731. \tag{6.10.a}$$

□

6.11 Theorem. *Let K be Sullivan's constant, the infimum of the quasiconformal constants of homeomorphisms between Ω and* Dome (Ω), *pointwise fixed on the boundary. Let c_1 be as in Definition 6.10. Then*

$$\log(K) \leq d_{\mathcal{T}}(G(ic_1), \mathrm{Id}) \leq d_{\mathfrak{T}}(ic_1, 0) \leq d_{\mathfrak{T}}(4.8732i, 0) \approx 2.63.$$

So $K \leq 13.88$.

Proof (of Theorem 6.11): From Step 6.2.5 we have the holomorphic map $G : \mathfrak{T} \to \mathcal{T}$. The Kobayashi metric can be defined on any complex manifold. It is a generalization of the hyperbolic metric on Riemann surfaces; as in the Riemann surface case, it has the important property that distances get *reduced* under a holomorphic map from one manifold to another.

In the case at hand, the Kobayashi metric in the simply connected region $\mathfrak{T} \subset \mathbb{C}$ is simply the hyperbolic metric. In $\mathcal{T}$ on the other hand, it is a famous theorem of [12], extended to the general case in [10] that the Kobayashi metric is the Teichmüller metric. In particular, in the respective metrics,

$$d_{\mathcal{T}}(G(ic_1), \mathrm{Id}) \leq d_{\mathfrak{T}}(ic_1, 0).$$

Now $d_{\mathcal{T}}(G(ic_1), \mathrm{Id}) = \inf \log(K)$, where the "inf" (or "min") is taken, for fixed (Λ, μ), over all quasiconformal maps sending $\Omega_{ic_1} \to \mathrm{Dome}\,(\Omega_{ic_1})$ which are the identity on the common boundary, as discussed in Section 5. Therefore the best constant K for Sullivan's theorem for fixed (Λ, μ) is the *supremum* —or equivalently *maximum*—of the minima computed for all admissible $\{c\}$.

The rightmost inequality in Theorem 6.11 is obtained by bringing in Bridgeman's estimate (6.10.a), which is independent of the choice of μ with $\|\mu\| = 1$. ■

6.12 The equivariant situation. Now let Γ be the group of (orientation preserving) Möbius transformations which preserve Ω.

The group Γ acts as well by isometries on $\mathrm{Dome}\,(\Omega)$ and acts injectively on the common prime ends of Ω and $\mathrm{Dome}\,(\Omega)$. To satisfy Theorem 5.1, we now look for Γ-equivariant quasiconformal homeomorphisms $f_\Omega : \Omega \to \mathrm{Dome}\,(\Omega)$.

Let $K(f_\Omega)$ be the maximum dilatation of f_Ω. We define $K_{eq}(\Omega)$ to be the infimum, or minimum, of values of $K(f_\Omega)$ as f_Ω varies over Γ-equivariant quasiconformal homeomorphisms. We define K_{eq} to be the supremum or maximum over all Ω of $K_{eq}(\Omega)$. We next discuss the value of K_{eq}. □

First we recall the general theory. Let Γ be a group of Möbius transformations preserving $\mathbb{D}^2$. It has a right action on $\mathcal{T}$ as follows. If $f \in \mathcal{T}$ is a normalized quasisymmetric function and $\gamma \in \Gamma$, the right action of γ sends f to $[f \circ \gamma]$, the result of normalizing $f \circ \gamma$ to fix $-1, 1, i$, as usual.

There is a corresponding right action by $\gamma \in \Gamma$ on the complex dilatations $\{\kappa(g)\} \subset L^\infty_1(\mathbb{D}^2)$, where g is a quasiconformal map $\mathbb{D}^2 \to \mathbb{D}^2$, and L^∞_1 denotes the unit ball in L^∞. Namely, $\kappa(g) \mapsto \kappa(g \circ \gamma)$ where

$$\kappa(f \circ \gamma) = \frac{\overline{\gamma'(z)}}{\gamma'(z)} \cdot \kappa(g) \circ \gamma.$$

If, for all $\gamma \in \Gamma$, $\kappa(g \circ \gamma) = \kappa(g)$, then $\kappa(g)$ is called a *Beltrami differential* for Γ. For $\kappa(g)$ to be a Beltrami differential means that there is an isomorphism $\varphi : \Gamma \to \Gamma'$ to a new group Γ' of Möbius transformations, such that $g \circ \gamma = \varphi(\gamma) \circ g$. Such a g is called Γ-equivariant.

The Teichmüller space $\mathcal{T}(\Gamma) \subset \mathcal{T}$ for Γ is the set of quasisymmetric maps which fix -1, 1 and i, as usual, and which extend to Γ-equivariant quasiconformal homeomorphisms of the disk. Such an extension is possible if and only if the quasisymmetric map itself is Γ-equivariant.

Using the fact that all of the constructions developed for the proof of Theorem 6.11 are equivariant, we obtain the following strengthening of Theorem 6.11, with the same numerical upper bound.

6.13 Theorem. *Let K_{eq} be Sullivan's constant for equivariant quasiconformal homeomorphisms between Ω and* Dome(Ω)*, pointwise fixed on the boundary. Let c_1 be as in Definition 6.10. Then*

$$\log(K_{eq}) \leq d_{\mathfrak{T}}(G(ic_1), \mathrm{Id}) \leq d_{\mathfrak{T}}(ic_1, 0) \leq d_{\mathfrak{T}}(4.8732i, 0) \approx 2.63.$$

So $K_{eq} \leq 13.88$.

6.14 Remark. We would like to thank Toby Driscoll for his Matlab package SC, which was used to compute the number 13.88 appearing in the statement of Theorems 6.11 and 6.13. This number is substantially better than 82.8, the upper bound for K_{eq} found by Epstein and Marden in [7]. It is worse than the bound 7.8 found by Chris Bishop for K in [4]. The method given here looks only at the transverse measure of open geodesic intervals of length one. There should be an improvement on the method which takes into account simultaneously the transverse measures of all intervals of all lengths. Such an adaptation of our method should give a result better than 13.88 for K_{eq}. Improvements over Bridgeman's bound for c_1 (and analogues of his result for intervals of lengths other than one) could result in a substantially improved bound for K_{eq}.

Note that the approach here will not give the best result for K_{eq}, no matter how far one pushes the details of the computation, because we are looking at the distance between two points in exactly one slice of universal Teichmüller space. The actual Teichmüller distance is given by the minimum of the hyperbolic distance as one varies over all slices of complex dimension one passing through the two points—Royden's theorem proving this is discussed in the proof we are about to start. $\square$

7 The Disk Theorem

Let (Λ, μ) be a geodesic lamination on $\mathbb{D}^2$, endowed with the Poincaré metric, and let μ be a non-negative transverse measure, with support equal to Λ. By Definition 3.6, for each $t \in \mathbb{C}$, we have the map $\mathbb{C}E_t : \mathbb{D}^2 \to P_{\Lambda,t\mu} \subset \mathbb{D}^3$, whose image is a pleated surface $P_{\Lambda,t\mu}$ in hyperbolic 3-space. $\mathbb{C}E_t$ is not necessarily continuous; it factors as

$$\mathbb{D}^2 \xrightarrow{\mathbb{C}E_{\mathrm{Re}(t)}} P_{\Lambda,\mathrm{Re}(t)\mu} = \mathbb{D}^2 \xrightarrow{\mathbb{C}E_{i\mathrm{Im}(t)}} P_{\Lambda,t\mu} \subset \mathbb{D}^3. \tag{7.0.a}$$

We recall that this factorization was introduced in § 3.8, together with the following notation. We set Λ^* to be the image lamination and μ^* to be the image transverse measure on Λ^* under the earthquake $\mathrm{Re}(t)\mu$. Then $\mathbb{C}E_{i\mathrm{Im}(t)}$ is pure bending, using the bending measure $\mathrm{Im}(t)\mu^*$ on Λ^*. Also $P_{\Lambda,t\mu} = P_{\Lambda^*,i\mathrm{Im}(t)\mu^*}$.

7.1 Definition. Let $X = X_{\Lambda,\mu} \subset \mathbb{C}$ be the set of t with the following properties:

7.1.1. For each geodesic $\lambda \in \Lambda$, we have $|\mathrm{Im}(t)|\mu(\lambda) < \pi$. (If Λ contains only one geodesic, we allow equality.)

7.1.2. The pure bending factor $\mathbb{C}E_{i\mathrm{Im}(t)} : \mathbb{D}^2 \to \mathbb{D}^3$ of $\mathbb{C}E_t$ in 7.0.a is a proper embedding—see Example 7.2.

7.2 Example. To see that Condition 7.1.2 is not vacuous, we construct an example of a non-proper C^∞-embedding $\mathbb{R} \to \mathbb{D}^2$ which maps each unit tangent vector to $\mathbb{R}$ to a unit tangent vector for the hyperbolic plane structure on $\mathbb{D}^2$. For example, we may assume that the image is a spiral in an annulus, with limit at each end one of the circles bounding the annulus. Such a spiral has a certain geodesic curvature, which we can think of as the bending measure. We take the inverse image of this spiral under the (hyperbolic) orthogonal projection $\mathbb{D}^3 \to \mathbb{D}^2$ to obtain a pleated surface which is not properly embedded.

Let X_+ be the subset of X where the imaginary part of t is positive, and X_- the subset of X where the imaginary part of t is negative. □

Note that $X = \mathbb{R} \cup X_+ \cup X_-$. Complex conjugation interchanges X_+ and X_-. It follows that understanding X is equivalent to understanding X_+. Moreover, the real line is in the interior of X—this follows readily from the discussion in Section 4.

7.3 Remark. If Λ consists of a single geodesic λ, then $X = \{z : |\mathrm{Im}(z)| \leq \pi/\mu(\lambda)\}$. Laminations consisting of exactly one geodesic are easy to deal with, and the relevant arguments will be left to the reader. From now on, we assume that Λ contains at least two geodesics. □

7.4 Lemma. *X is the set of $t \in \mathbb{C}$ such that $P_{\Lambda,t\mu} = \mathrm{Dome}(\Omega_t)$ for some simply connected open subset $\Omega_t \subset \mathbb{S}^2$, and such that the dome has bending measure $\mathrm{Im}(t)\,(\mathrm{Re}(t)_*(\mu))$.*

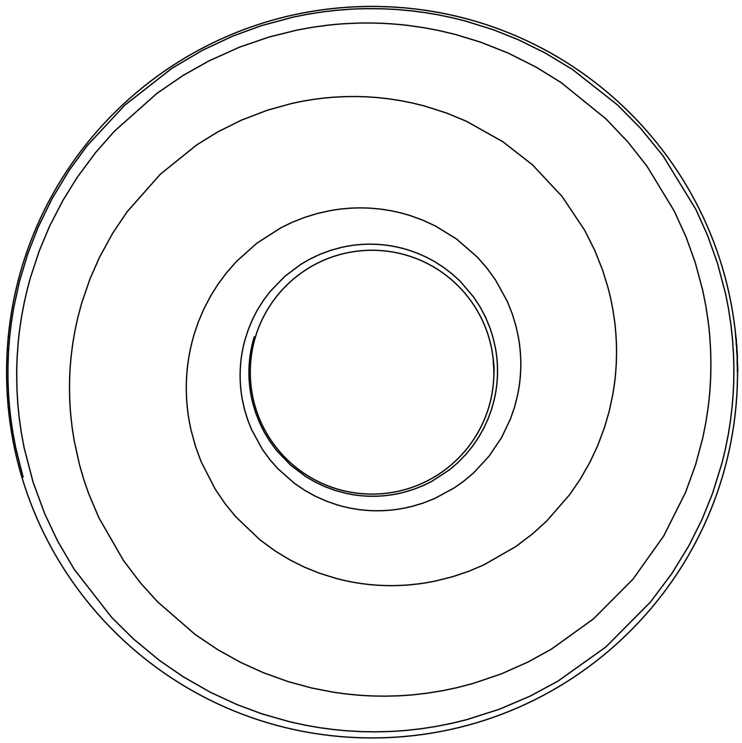

□

Figure 7.2.i. This illustrates Example 7.2. It is a spiral converging to the two boundary edges of an annulus.

Proof: If $P_{\Lambda,t\mu} = \text{Dome}\,(\Omega_t)$ for some simply connected open subset $\Omega_t \subset \mathbb{S}^2$, and the dome has bending measure $\text{Im}(t)\,(\text{Re}(t)_*(\mu))$, then it is clear that $t \in X$. We need to prove that, if $t \in X$, then $P_{\Lambda,t\mu} = \text{Dome}\,(\Omega_t)$ for some simply connected open subset $\Omega_t \subset \mathbb{S}^2$, and that the dome has bending measure $\text{Im}(t)\,(\text{Re}(t)_*(\mu))$.

If $t \in \mathbb{R}$, then $P_{\Lambda,t\mu}$ is a hyperbolic plane with boundary a round circle. So it is the dome of the disk bounded by the circle, and the desired result follows. For the remainder of this proof, we may assume, by symmetry, that $t \in X_+$.

We may assume without loss of generality that $t = iy$, with $y > 0$ (to prove this, change Λ to Λ^* and μ to μ^*). Then $\mathbb{C}E_{iy} : \mathbb{D}^2 \to \mathbb{D}^3$ is a proper embedding, separating $\mathbb{D}^3$ into two components U and V. We choose the names so that $\partial U = \partial V$ is locally convex from the point of view of U and locally concave from the point of view of V.

We repeat a standard argument, due to Hadamard, to prove that $\overline{U}$ is convex. Part of the argument is illustrated in Figure 7.4.i. Given two points in $\overline{U}$, the Ascoli Theorem gives us a shortest path γ in $\overline{U}$, parametrized by pathlength and joining the two points. Inside U, γ must be a geodesic. Suppose $u \in \gamma \cap \partial U$. Then $u \in F$, where F is either a g-leaf of Λ or the closure of a flat. Clearly, $\gamma \cap F$ is either a point or a closed interval. Let u_1 and u_2 be the endpoints in the case of a closed interval, and let $u_1 = u_2 = \gamma \cap F$ otherwise. Then u_1 and u_2 must be the endpoints of γ. For otherwise, local convexity shows that one can shorten γ, as shown in Figure 7.4.i. This shows that γ is a geodesic in $\mathbb{H}^3$. So $\overline{U}$ is convex.

Let $\Omega_t \subset \mathbb{S}^2$ be the complement in $\mathbb{S}^2$ of the closure of V in $\overline{\mathbb{D}^3}$ and let $u \in U$. Radial projection from u gives a homeomorphism between ∂U and Ω_t. This shows that Ω_t is homeomorphic to an open disk. Since each flat or g-leaf of Λ is the convex hull of its ideal boundary points, we see that ∂U is contained in the convex hull of $\mathbb{S}^2 \setminus \Omega_t$. It follows that $\overline{U}$ is the hyperbolic convex hull of $\mathbb{S}^2 \setminus \Omega_t$, and that $\partial U = P_{\Lambda,t\mu} = \text{Dome}\,(\Omega_t)$. ∎

7.5 Proposition. *X is a closed subset of $\mathbb{C}$.*

Proof: We prove the equivalent statement that X_+ is a closed subset of $\mathbb{U}^2$. Let $(t_i)_{i\in\mathbb{N}}$ be a sequence of complex numbers converging to $t_\infty \in \mathbb{U}^2$, and suppose that, for each $i < \infty$, $t_i \in X_+$. We must show that $t_\infty \in X_+$.

By Lemma 7.4, we have, for each $t \in X_+$, the open simply connected subset $\Omega_t \subset \mathbb{S}^2$, such that $P_{\Lambda,t\mu} = \text{Dome}\,(\Omega_t)$. In order to avoid double subscripts, we write Ω_m instead of Ω_{t_m}. Set $A_m = \mathbb{S}^2 \setminus \Omega_m$ and C_m equal to the hyperbolic convex hull of A_m.

First we show that, for each $\lambda \in \Lambda$, $\mathrm{Im}(t_\infty)\mu(\lambda) < \pi$. We will suppose that $\mathrm{Im}(t_\infty)\mu(\lambda) = \pi$ and prove a contradiction. By Remark 7.3, there is a short open geodesic interval A which is disjoint from λ, such that $\mu(A) > 0$. Let $K \subset \mathbb{D}^2$ be a closed hyperbolic disk containing both A and a portion of λ. For each $t \in \mathbb{C}$, each flat or g-leaf V of Λ, and for $0 \leq \theta \leq 1$, there is a Möbius transformation $M_{t,V,\theta}$, as discussed in § 3.5. By Lemma 3.7, the dependence on t is holomorphic. The parameter θ only comes into play if V is a g-leaf with positive measure.

Let $m \in \mathbb{N}$ be sufficiently large, so that, for each flat or g-leaf V meeting K, the Möbius transformation $M_{t_m,V,\theta}$ is extremely close to the Möbius transformation $M_{t_\infty,V,\theta}$. It follows that, for each flat or g-leaf V of Λ that meets K, the image $\mathbb{C}E_{t_m}(V) \subset P_{\Lambda,t_m\mu} = \mathrm{Dome}\,(\Omega_m)$ is very near to $\mathbb{C}E_{t_\infty}(V) \subset P_{\Lambda,t_\infty\mu}$.

It follows that the bending line $\mathbb{C}E_{t_m}(\lambda)$ has a bending angle which is arbitrarily near to π. By the geometry of convex sets, we see that $\mu(A)$ is arbitrarily small and therefore $\mu(A) = 0$. But this is a contradiction, showing that $\mathrm{Im}(t_\infty)\mu(\lambda) < \pi$.

Next we show by contradiction that $\mathbb{C}E_{i\mathrm{Im}(t_\infty)} : \mathbb{D}^2 \to \mathbb{D}^3$ is a proper map. We assume that we have a sequence $(x_i)_{i\in\mathbb{N}}$ in $\mathbb{D}^2$, converging to a point ∂D, such that $\big(\mathbb{C}E_{i\mathrm{Im}(t_\infty)}(x_i)\big)_{i\in\mathbb{N}}$ is a sequence converging to some point $x_\infty \in \mathbb{D}^3$, and prove a

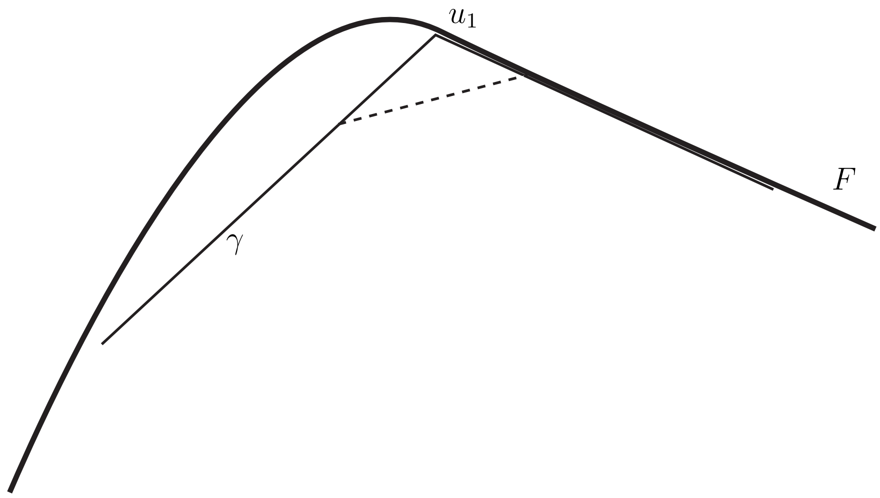

Figure 7.4.i. This illustrates the proof of Hadamard's Theorem, which arises in the proof of Lemma 7.4. F here represents a flat. If u_1 is an endpoint of $\gamma \cap F$, but not an endpoint of γ, then we can shorten γ near u_1, giving a contradiction. The dotted line shows a typical shortcut.

contradiction. We choose three points from the sequence: $y_j = x_{i_j}$, for $j = 1, 2, 3$, such that, firstly, they are at least a hyperbolic distance 100 from each other in $\mathbb{D}^2$, and, secondly, they have images which are a hyperbolic distance less than 10^{-4} from x_∞. For $j = 1, 2, 3$, let $B_j \subset \mathbb{D}^2$ be a closed disk of radius 1, centred at y_j. Now choose n large enough, so that the embedding $\mathbb{C}E_{i\mathrm{Im}(t_n)}$ is a very close approximation to $\mathbb{C}E_{i\mathrm{Im}(t_\infty)}$ on $B_1 \cup B_2 \cup B_3$. Let B be a ball in $\mathbb{D}^3$, centred at x_∞, of radius 10^{-2}. Then $\mathbb{C}E_{i\mathrm{Im}(t_n)}(B_1 \cup B_2 \cup B_3)$ separates B into components in such a way that at least two distinct components are on the convex side of $\mathbb{C}E_{i\mathrm{Im}(t_n)}(\mathbb{D}^2)$. This is illustrated in Figure 7.5.i. But these components are joined by a geodesic interval in $\mathbb{D}^3$, and this interval does not meet the surface $\mathbb{C}E_{i\mathrm{Im}(t_n)}(\mathbb{D}^2)$, so we have a contradiction. This completes the proof that $\mathbb{C}E_{i\mathrm{Im}(t_\infty)} : \mathbb{D}^2 \to \mathbb{D}^3$ is a proper map.

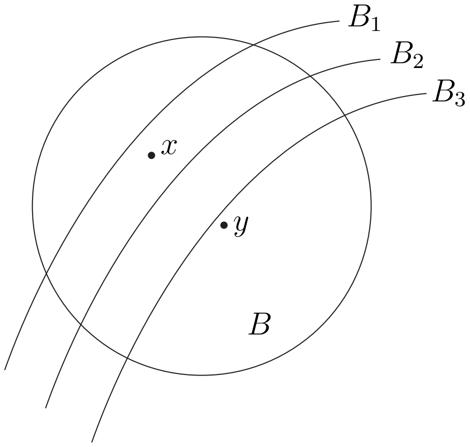

Figure 7.5.i. This illustrates part of the proof of Proposition 7.5 and we follow that notation. B is a 3-dimensional disk, with three sheets passing through it, labelled B_1, B_2 and B_3. The points x and y lie on the convex sides of B_1 and B_3 respectively. But the geodesic joining x to y must then cross B_2, which is impossible.

Our next claim is that, given a compact subset $K \subset \mathbb{D}^2$, there is an $\epsilon > 0$, with the property that, for each disk $D \subset K$ of radius ϵ, the restriction $\mathbb{C}E_{i\mathrm{Im}(t_\infty)}|D$ is an embedding. For suppose this is not true. Then we can find sequences $(u_m)_{m\in\mathbb{N}}$ and $(v_m)_{m\in\mathbb{N}}$ in K with the same limit $w \in K$, such that, for each $m \in \mathbb{N}$, $u_m \neq v_m$ and $\mathbb{C}E_{i\mathrm{Im}(t_\infty)}(u_m) = \mathbb{C}E_{i\mathrm{Im}(t_\infty)}(v_m)$. So w is contained in a g-leaf $\lambda \in \Lambda$, and not in a flat. We have already shown that $\mathrm{Im}(t_\infty)\mu(\lambda) < \pi$. We take a very short transverse open interval A which meets λ, such that $\mathrm{Im}(t_\infty)\mu(A) < \pi$. This contradicts the existence of u_m and v_m as above.

We now prove that the assumptions $u_1, u_2 \in \mathbb{D}^2$, $u_1 \neq u_2$ and $\mathbb{C}E_{i\mathrm{Im}(t_\infty)}(u_1) = \mathbb{C}E_{i\mathrm{Im}(t_\infty)}(u_2)$ lead to a contradiction. For $i = 1, 2$, let $D_i \subset \mathbb{D}^2$ be a hyperbolic disk centred at u_i, and suppose $D_1 \cap D_2 = \emptyset$. For m sufficiently large, $\mathbb{C}E_{i\mathrm{Im}(t_m)}|D_1 \cup D_2$ is a very close approximation to $\mathbb{C}E_{i\mathrm{Im}(t_\infty)}|D_1 \cup D_2$. The images of D_1 and D_2 under $\mathbb{C}E_{i\mathrm{Im}(t_m)}$ must then be convex pieces of surface, facing each other, with the hyperbolic convex hull C_m (in the notation introduced near the beginning of this proof) between them.

If we let m tend to infinity, we see that the images of D_1 and D_2 must flatten up against each other. It follows that D_1 and D_2 must lie inside flats F_1 and F_2. Convexity also shows that the images of F_1 and F_2 must flatten up against each other, matching exactly, as m tends to infinity. We then get a contradiction by examining the consequences of non-trivial convexity in the vicinity of boundary geodesics of F_1 and F_2. It follows that $t_\infty \in X_+$. ∎

We use the notation introduced in Section 4. We recall that $\Phi_{t_0} : \mathbb{D}^2 \to \Omega_0$ is a quasiconformal homeomorphism onto a quasidisk. The map $(t, z) \mapsto \Psi_t(z)$, gives a function $\mathbb{C} \times \Omega_0 \to \mathbb{S}^2$, which has the properties specified in Theorem 4.13. In particular, it is holomorphic as a function of t. Recall also that $\Psi_{t_0} = \mathrm{Id} : \Omega_0 \to \Omega_0$.

For each flat or g-leaf $V \subset \mathbb{D}^2$ of Λ, for each $t \in \mathbb{C}$, and for each θ with $0 \leq \theta \leq 1$, we have a Möbius transformation $M_{t,V,\theta}$.

7.6 Lemma. *The image of Ψ_t is equal to $\bigcup_{V,\theta} M_{t,V,\theta}(\mathbb{D}^2)$.*

Proof: Note that

$$\Psi_t(\mathbb{D}^2) = \bigcup_V \Psi_t(V) = \bigcup_V M_{t,V,\theta}(V) \subset \bigcup_V M_{t,V,\theta}(\mathbb{D}^2).$$

So we need only show that $M_{t,V,\theta}(\mathbb{D}^2) \subset \Psi_t(\Omega_0)$. This is easy when Λ is a finite lamination. In general, we approximate (Λ, μ) by a finite lamination. Then $M_{t,V,\theta}$ is closely approximated by the similar Möbius transformation defined for the finite lamination. The result follows by taking the limit. ∎

7.7 Lemma. *X_+ is the set of $t \in \mathbb{U}^2$ such that $\Psi_t : \Omega_0 \to \mathbb{S}^2$ is injective.*

Proof: If $t \in X_+$, Lemma 7.4 shows that $\Psi_t : \Omega_0 \to \Omega_t$ is injective.

Conversely, suppose that $\Psi_t : \Omega_0 \to \Omega_t$ is a bijection. Since it is continuous, it is a homeomorphism. We must show that $\mathrm{Dome}\,(\Omega_t) = \mathbb{C}E_t(\mathbb{D}^2)$. Let $V \subset \Omega_0$ be a gap or a c-leaf mapped homeomorphically by the nearest point retraction $r : \Omega_0 \to \text{Dome}(\Omega_0)$ onto a flat or bending line of Dome (Ω_0).

By Lemma 7.6, $\Omega_t = \bigcup_V M_{t,V,\theta}(\mathbb{D}^2)$. Now, $M_{t,V,\theta}(\mathbb{D}^2)$ is a maximal disk in Ω_t: to see this, note that $M_{t,V,\theta}(V)$ will have ideal boundary points inside any larger disk, and so the larger disk cannot be a subset of Ω_t. It follows that $M_{t,V,\theta}(V)$ is a gap or c-leaf sent homeomorphically by $r_t : \Omega_t \to \text{Dome}(\Omega_t)$ to a flat or bending line. It follows that Dome (Ω_t) is the union of subsets of the form $M_{t,V,\theta}(V)$. Therefore $t \in X_+$. ■

Let $X = X_{\Lambda,\mu}$ be the set of Definition 7.1.

7.8 Theorem (The disk theorem). *7.8.1. $\mathbb{C} \setminus X$ has no bounded components.*

7.8.2. If X' is a component of X, then $\mathbb{C} \setminus X'$ has no bounded component.

7.8.3. Every component of the interior of X is simply connected.

Proof: Suppose by contradiction that B is a bounded component of $\mathbb{C} \setminus X$. Since X is closed, B is open. Given any compact subset of B, we can enlarge it to a connected compact submanifold with boundary in B. Exactly one of the boundary components of the submanifold bounds a disk in $\mathbb{C}$ which contains the submanifold. We can therefore find a sequence of closed disks $U_i \subset \text{int}(U_{i+1})$, whose union U is an open topological disk. Then $B \subset U \subset \mathbb{C}$ and $\partial U \subset X \cap \overline{B}$. It follows that $\overline{U} = U \cup \partial U$ is compact.

Note that $B \cap \mathbb{R} \subset B \cap X = \emptyset$. Therefore the above construction takes place entirely in X_+ or entirely in X_-. Without loss of generality, we will work in X_+. Fix $t \in B$. By Lemma 7.7, $\Psi_t : \Omega_0 \to \mathbb{S}^2$ is not injective. Therefore there exists $z_0 \in \Omega_0$ and a gap or c-leaf V, such that $z_0 \notin V$ and $\Psi_t(z_0) \in \Psi_t(V)$.

We normalize (compare § 2.6) so that, for each $t \in \mathbb{C}$, $\Psi_t|\overline{V} = \text{Id}\,|\overline{V}$. Then $z_0 \notin \overline{V}$.

For $t \in \mathbb{C}$ and $z \in \Omega_0$, we write $t.z = \Psi_t(z)$. Given $P \subset \mathbb{C}$ and $Q \subset \Omega_0$, $P.Q$ will denote the set of all points $\Psi_t(z)$ with $t \in P$ and $z \in Q$.

If $z \in \overline{V}$, we have $\overline{U}.z = z$. Now Ψ_t is locally injective, and, if variables in both $\mathbb{C}$ and Ω_0 vary over compact subsets, we can choose a fixed radius for the disk in Ω_0 on which Ψ_t is injective. So, by continuity, if $z \in \Omega_0$ is near enough to V (while staying more than some definite positive spherical distance from $\partial\Omega_0$), then $\overline{U}.z$ is disjoint from $\overline{V}$.

By the maximum modulus theorem, for all $z \in \Omega_0$, the topological boundary of $\overline{U}.z$ is contained in $\partial U.z$. Now, for each $z \in \Omega_0 \setminus \overline{V}$,

$$\partial U.z \cap \overline{V} \subset X.z \cap \overline{V} = \emptyset.$$

Moreover, for z near $\overline{V}$ as above, $\overline{U}.z \cap \overline{V} = \emptyset$. The situation is shown in Figure 7.8.i. By continuity, it therefore follows that $\overline{U}.z \cap \overline{V} = \emptyset$, for all $z \in \Omega_0 \setminus \overline{V}$. But this contradicts our choice of z_0, and we have proved that there is no bounded component of $\mathbb{C} \setminus X$.

It follows immediately that, if X' is a component of X, then $\mathbb{C} \setminus X'$ has no bounded component.

Now consider a component W of the interior of X. If W is not simply connected, then there is a Jordan curve in W that bounds a disk in $\mathbb{C}$, but not a disk in W. But the disk would then contain a bounded component of $\mathbb{C} \setminus X$, and this is impossible.■

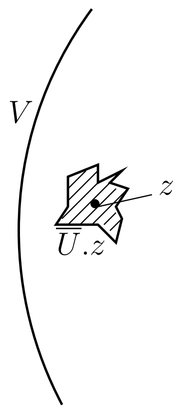

Figure 7.8.i. This illustrates the proof of the disk theorem, 7.8. V is a gap or c-leaf which is kept fixed, $z \in \Omega_0$ is a point which is near V and not near $\partial\Omega_0$. $\overline{U}.z$ is compact, near V, and contains z.

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.J. Bishop. An explicit constant for Sullivan's convex hull theorem. *Contemporary Math.*, to appear. Proceedings of the 2002 Ahlfors-Bers Colloquium.

[5] M. Bridgeman. Bounds on the average bending of the convex hull of a kleinian group. *Preprint*, 2002.

[6] D.B.A. Epstein. Pointwise periodic homeomorphisms. *P.L.M.S.*, 62:pp.415–460, 1981.

[7] D.B.A. Epstein and A. Marden. Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces. In *Analytical and Geometric Aspects of Hyperbolic Space: LMS Lecture Notes Series 111*, pages 113–253. Cambridge University Press, 1987.

[8] D.B.A. Epstein, A. Marden, and V. Markovic. Complex angle scaling. In C. Series, editor, *Proceedings in Hyperbolic Geometry*. Cambridge University Press, to appear.

[9] D.B.A. Epstein, A. Marden, and V. Markovic. Quasiconformal homeomorphisms and the convex hull boundary. *Annals of Math*, to appear. available at `http://www.maths.warwick.ac.uk/~dbae` or at `http://www.maths.warwick.ac.uk/~markovic`.

[10] F. Gardiner and N. Lakic. *Quasiconformal Teichmüller theory*, volume 76. American Mathematical Society, Providence, RI, 2000. ISBN: 0-8218-1983-6.

[11] C.T. McMullen. Complex earthquakes and Teichmüller theory. *Jour. A.M.S.*, 11:283–320, 1998.

[12] H.L. Royden. Automorphisms and isometries of Teichmüller space. In *Advances in the Theory of Riemann Surfaces: Annals of Math. Studies* **66**. Princeton University Press, 1971.

[13] D.P. Sullivan. Travaux de Thurston sur les groupes quasi-fuchsiens et les variétés hyperboliques de dimension 3 fibrés sur S^1. In *Springer Lecture Notes 842*. Springer-Verlag, 1981.

D.B.A. Epstein, Mathematics Institute, University of Warwick.
`dbae@maths.warwick.ac.uk`
A. Marden, Mathematics Department, University of Minnesota.
`am@math.umn.edu`
V. Markovic, Mathematics Institute, University of Warwick.
`markovic@maths.warwick.ac.uk`