Complex earthquakes and deformations of the unit disk

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

25 July 2003

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 makes essential use of two tools. The first is a generalization of the method of complex scaling (see [8] for an elementary exposition). The generalization is related to quakebends or complex earthquakes [7], [14]. 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 $\mathrm{Dome}\,(\Omega)$, 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 $f : \Omega \to \mathrm{Dome}\,(\Omega)$ 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 [14]. 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.

The second tool we use in this paper is the result of a close analysis of the nearest point retraction in a neighborhood of each boundary point of a euclidean convex region $\Omega \subset \mathbb{C}$. This leads to the precise computation of boundary dilatations for the family $\mathcal{F}_\Omega$ which is our third main theorem. This difficult analysis occupies nearly half of the paper. In the case of convex regions, at least those with C^1 boundaries, we can conclude that in the class $\mathcal{F}_\Omega$, there is a unique extremal quasiconformal mapping, and it is of Teichmüller type. We also show that the maximal dilatation of the extremal mapping cannot exceed 2.

We want to thank A, Douady, C. Earle, and C. Pommerenke for helpful discussions and communications regarding our work.

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.

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 $0 \le \theta \le 1$, 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 $E_t(z) = E(t, z)$. $\square$

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.

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 $\{z \mapsto e^\lambda.z\}_{\lambda\in\mathbb{R}}$, that fixes each vertex $0, \infty$ and maps W_α onto itself. For $t \in \mathbb{C}$, we

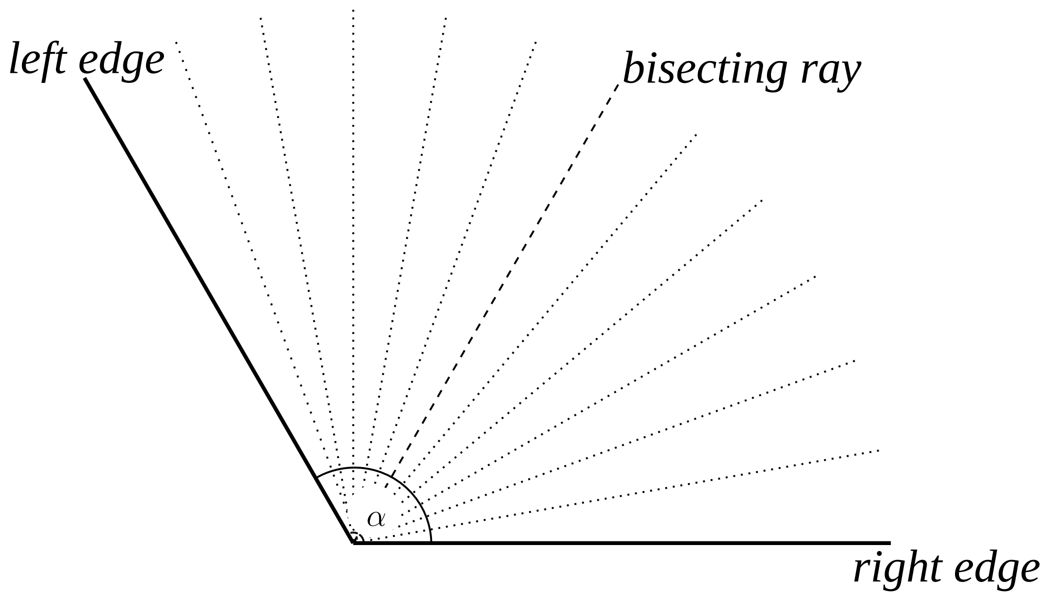

Figure 2.0.i. A standard wedge with angle α. Note the bisecting ray at angle $\alpha/2$. We indicate the foliation of the wedge by rays.

write $A_t : z \in \mathbb{C} \mapsto e^t.z \in \mathbb{C}$. 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

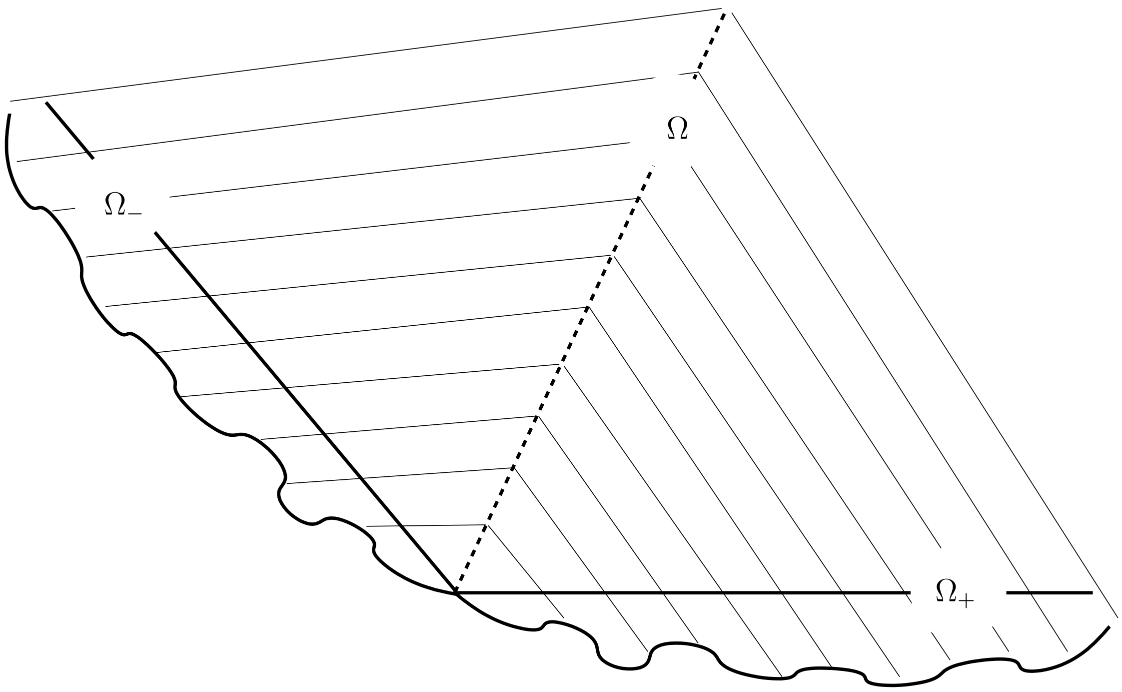

Figure 2.0.ii. This is the same wedge as that shown in Figure 2.0.i. We indicate the subsets Ω, Ω_+ and Ω_- referred to in Definition 2.1, by means of hatching.

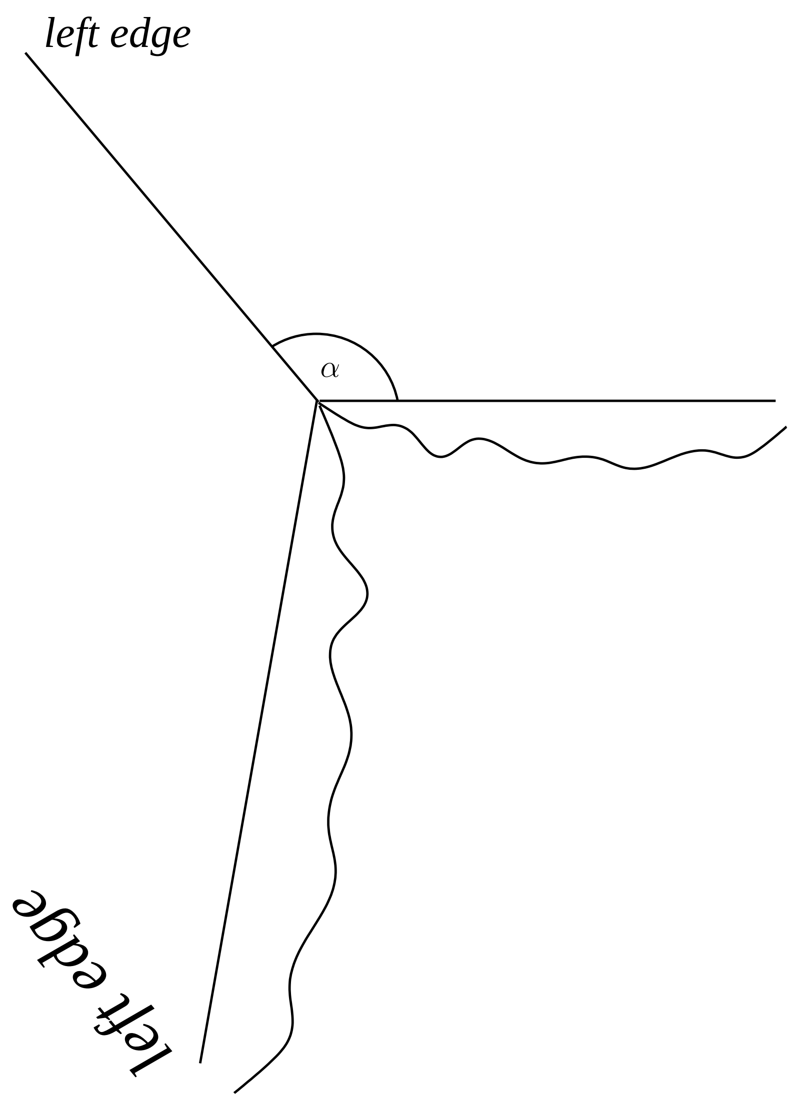

Figure 2.1.i. This picture shows the effect of scaling. The particular scaling parameter t has been chosen so that $\alpha.\mathrm{Re}(t) = \log(2)$ and $\mathrm{Im}(t) = 1$. The effect is to magnify at the lefthand edge by a factor 2 and to rotate by doubling the angle α. To help the reader understand 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.

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.

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.

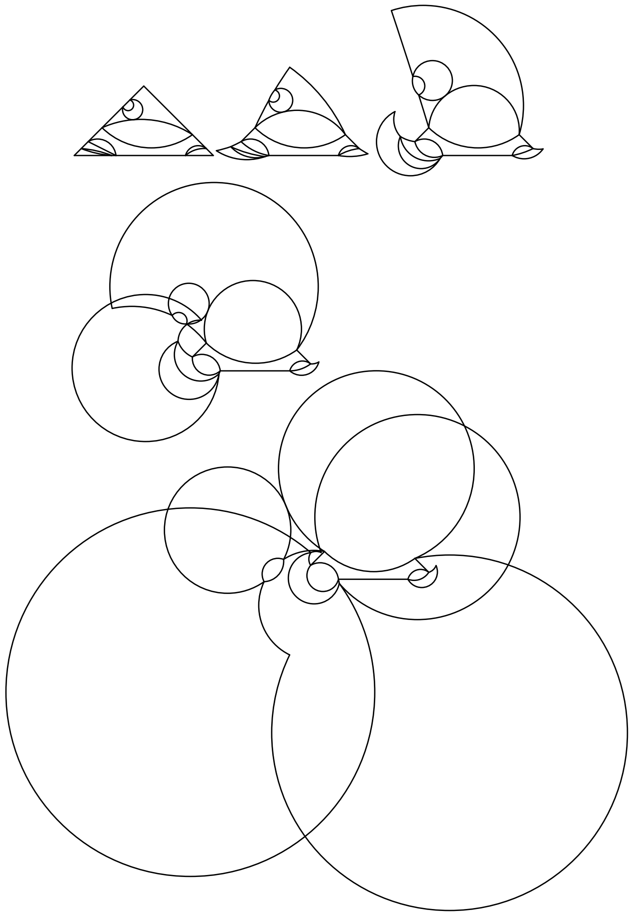

Figure 2.4.i. We start with a number of crescents in a topological disk. In this case, the topological disk is a triangle Δ. We use the notation of Lemma 2.5. The first diagram shows Δ and its crescents. The other four diagrams show the images of the various intervals and arcs in Δ under the scaling 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.

2.6 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.

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. □

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

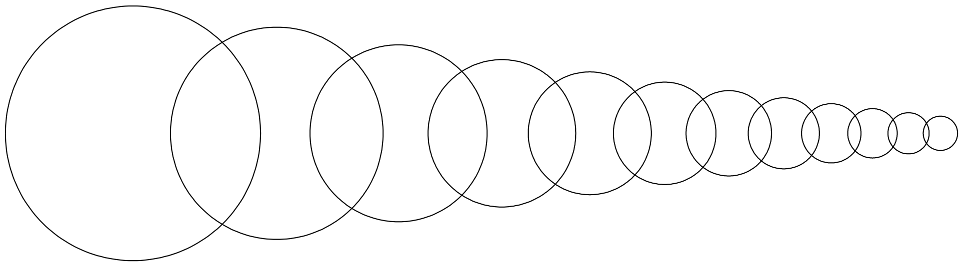

Figure 2.9.i. The picture illustrates Example 2.9. It represents a countable union of round disks, getting smaller and smaller. The union of the interiors is simply connected. The angle between successive circles gets smaller and smaller. In other words, successive disks become more and more nearly tangent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

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.

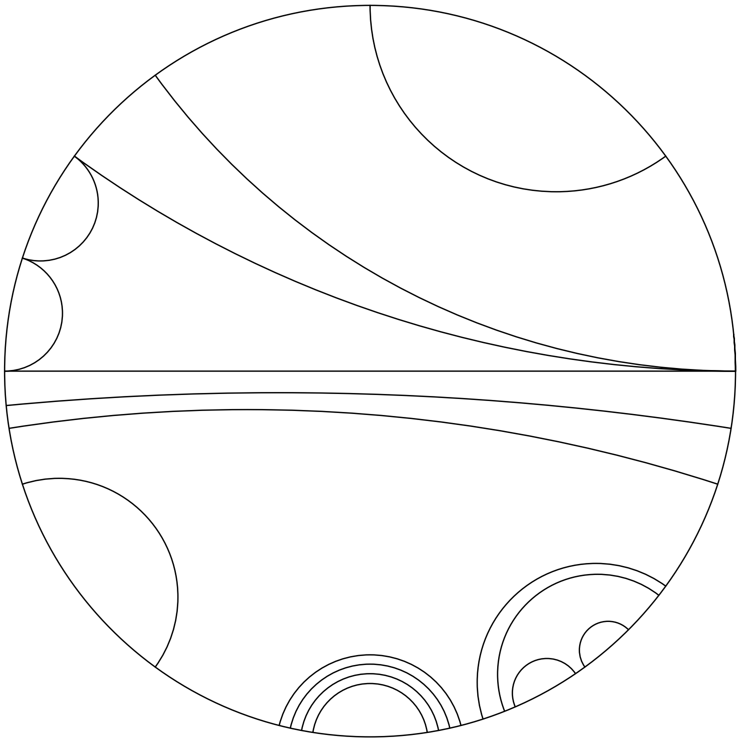

Figure 3.3.i. This is a picture of a finite geodesic lamination. We are using the Poincaré metric on the unit disk. To each of the regions shown is assigned a Möbius transformation. The transformations on different sides of a geodesic differ by a Möbius transformation having that geodesic as an axis.

We 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 $\mathbb{C}E_t(\ell_j)$.

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 $0 \leq \theta \leq 1$. 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 $0 \leq \theta \leq 1$. 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 $\mathfrak{T}_{\mathbf{o}} \times \mathbb{D}^2 \to \mathbb{S}^2$. 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 $\mathfrak{T}_{\mathbf{o}}$ 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. □

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 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

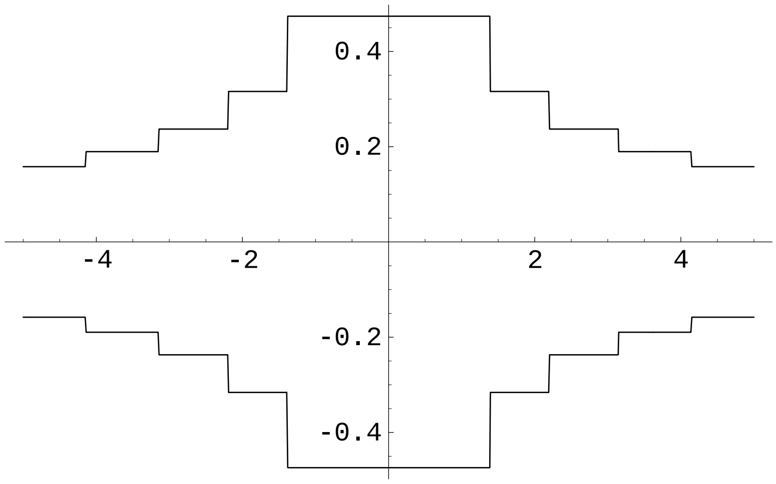

Figure 4.3.i. This is an illustration of $\mathfrak{T}_\circ$. This is the set of values of $t \in \mathbb{C}$ for which we can show that the map Φ_t of Theorem 4.4 is injective. It is the set of points lying between the lower and upper curves shown in the illustration.

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 $t \in \mathfrak{T}_o \cup \mathbb{U}^2$. 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 $\mathbb{U}^2 \cap \mathfrak{T}_o$ 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 $t_0 \in \mathfrak{T}_o$. 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

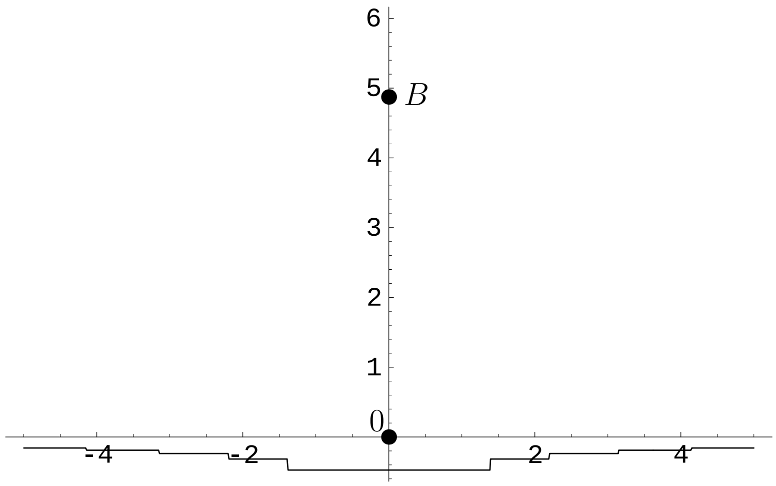

Figure 4.3.ii. Here we illustrate $\mathfrak{T} = \mathfrak{T}_o \cup \mathbb{U}^2$. This is the set of points in $\mathbb{C}$, lying above the curve shown. This is the same as the curve shown in Figure 4.3.i, but drawn to a different vertical scale. As explained in § 4.5, there is a holomorphic map 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 $\mathfrak{T}$.

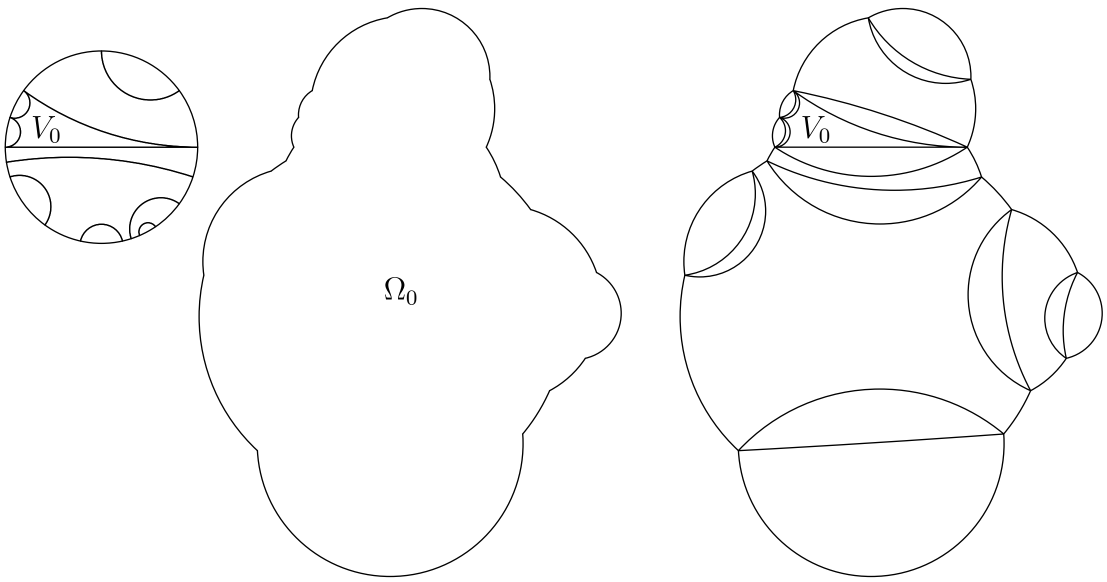

Figure 4.5.i. The lefthand picture shows the unit disk, thought of as a hyperbolic plane with the Poincaré metric, together with a measured geodesic lamination. The measure of each geodesic is proportional to the angle of the corresponding crescent in the righthand picture. The middle picture shows Ω_0, 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.

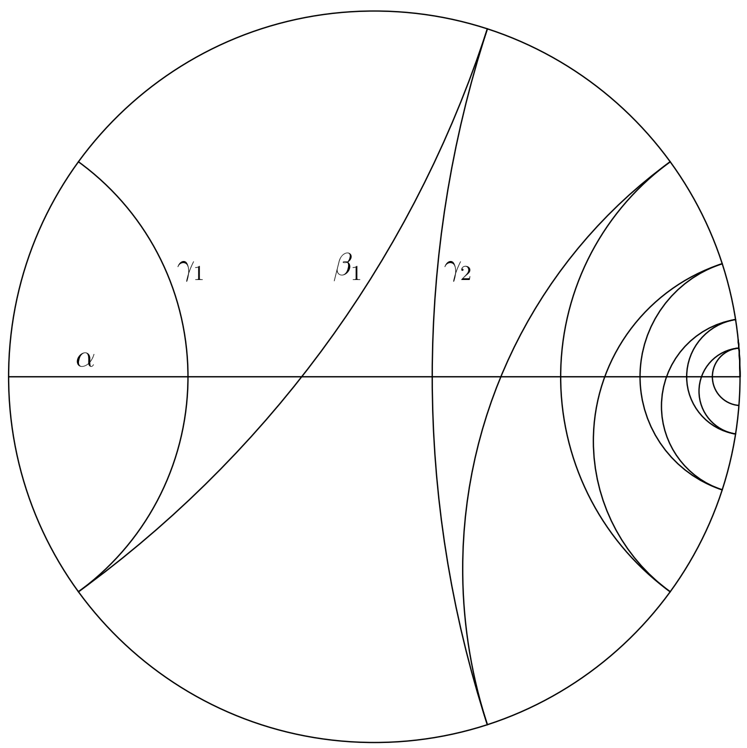

Figure 4.5.ii. For a finite measured lamination of $\mathbb{D}^2$, it is always possible to fit disjoint crescents, with vertices at the endpoints of geodesics in the lamination, and crescent angles proportional to the measure of the corresponding leaf. The definition of Φ_t—in our case carried out using 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 $\mathbb{D}^2$.

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 $v_0\mu(\ell_i)$.

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 $t \in \mathfrak{T}_\mathbf{o}$ (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 $t \in \mathfrak{T}_\mathbf{o}$. 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. $\square$

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.

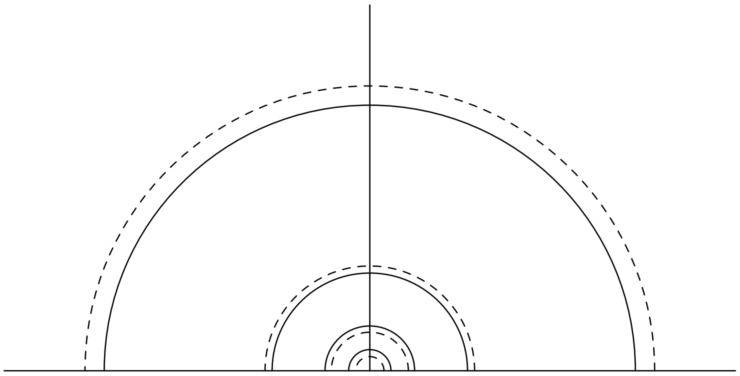

Figure 4.6.i. We show a picture in the upper halfplane with the hyperbolic metric. The dotted semicircles are geodesics a hyperbolic distance one apart. The solid semicircles are approximating geodesics. The measure of each leaf in the original (dotted) lamination is one, and the norm is then one. In 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.*

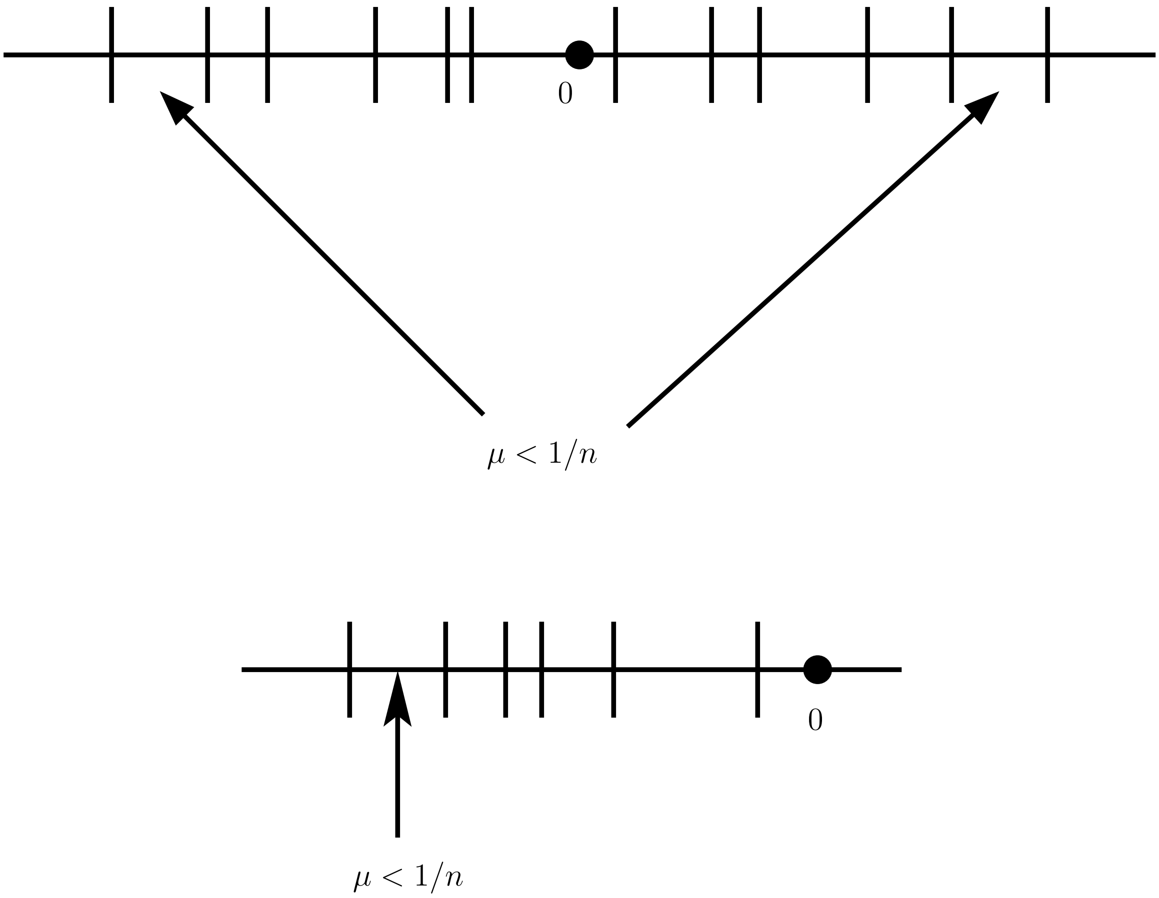

Figure 4.6.ii. These two diagrams illustrate the proof of the two inequalities of 4.6.a. In each case, we have a geodesic segment $\alpha \subset B_n$ of length one. We mark the intersections of α with the geodesics of Λ_n. Each subinterval shown has measure less than $1/n$, and this is indicated by the arrows. The upper diagram shows the situation when α contains 0. We see that the upper diagram leads to the inequalities $\mu_n(\alpha) \le \mu(\alpha) < \mu_n(\alpha) + 2/n$. 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.

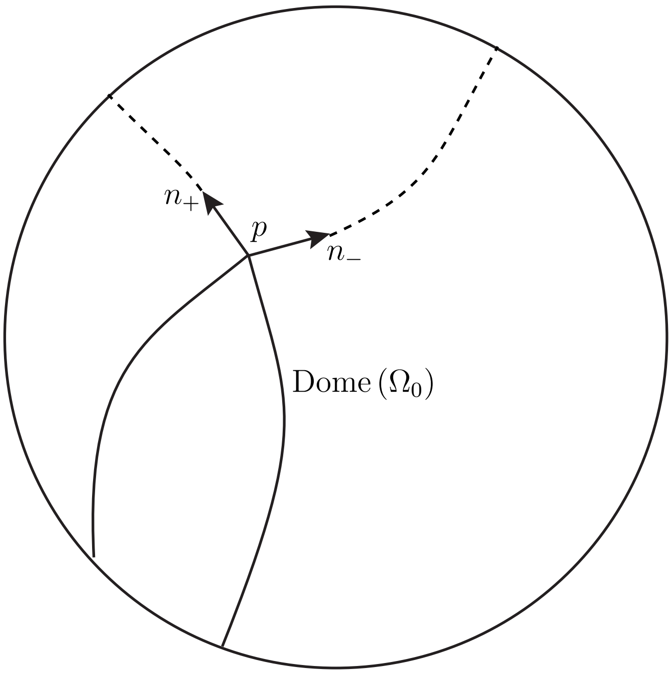

Figure 4.9.i. This picture illustrates Lemma 4.9. It can be thought of as the intersection of the general configuration with a hyperbolic plane orthogonal at a fixed point p to the bending line through p. The set of unit normal vectors (elements of $\mathcal{N}$) at p is, topologically, a closed interval. 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.

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

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

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 $\mathrm{Dome}\,(\Omega_0)$. Note that a normal vector based at $p \in \mathrm{Dome}\,(\Omega_0)$ 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 $\mathrm{Dome}\,(\Omega_0)$ 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 $p \in \mathrm{Dome}\,(\Omega_0)$. 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 $p \in \mathrm{Dome}\,(\Omega_0)$ 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 $p_i \in \mathrm{Dome}\,(\Omega_0)$ 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 $t_0 = iv_0 \in \mathfrak{T}_{\mathsf{o}}$ 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 $t \in \mathfrak{T}_{\mathsf{o}}$, 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 $t \in \mathbb{U}^2 \cap \mathfrak{T}_{\mathsf{o}}$, $\Psi_t : \Omega_0 \to \mathbb{S}^2$ 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 $t \in \mathfrak{T}_{\mathsf{o}}$ 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 $t \in \mathfrak{T}_\mathfrak{o}$ 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 $\mathrm{Dome}\,(\Omega)$ 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 $\mathrm{Dome}\,(\Omega)$ 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 $\mathrm{Dome}\,(\Omega)$, 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. □

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 $\overline{\mathbb{D}}^2$. ■

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).

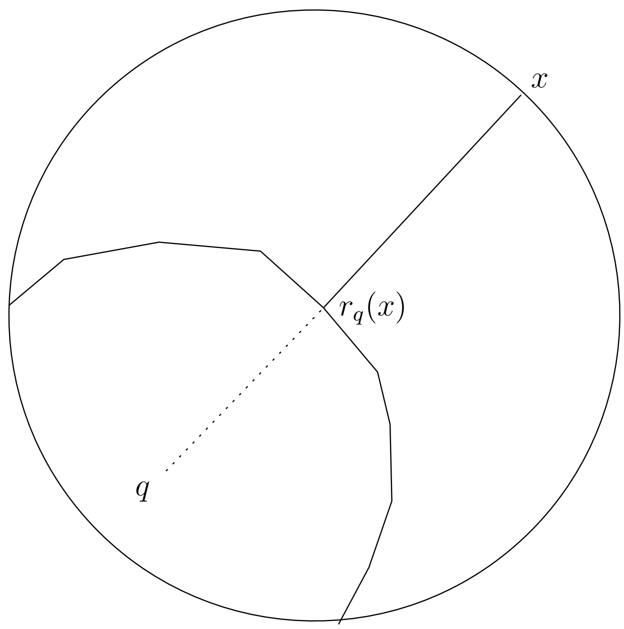

Figure 5.6.i. This illustration is in the projective model, where hyperbolic convex sets are the same as euclidean convex sets. This makes it easier to draw hyperbolic convex sets. We show the centre of projection q. In the projective model, hyperbolic radial projection from q is the same as euclidean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

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 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 $\mathrm{Dome}\,(\Omega)$. 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 $\mathrm{Dome}\,(\Omega)$ 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. ∎

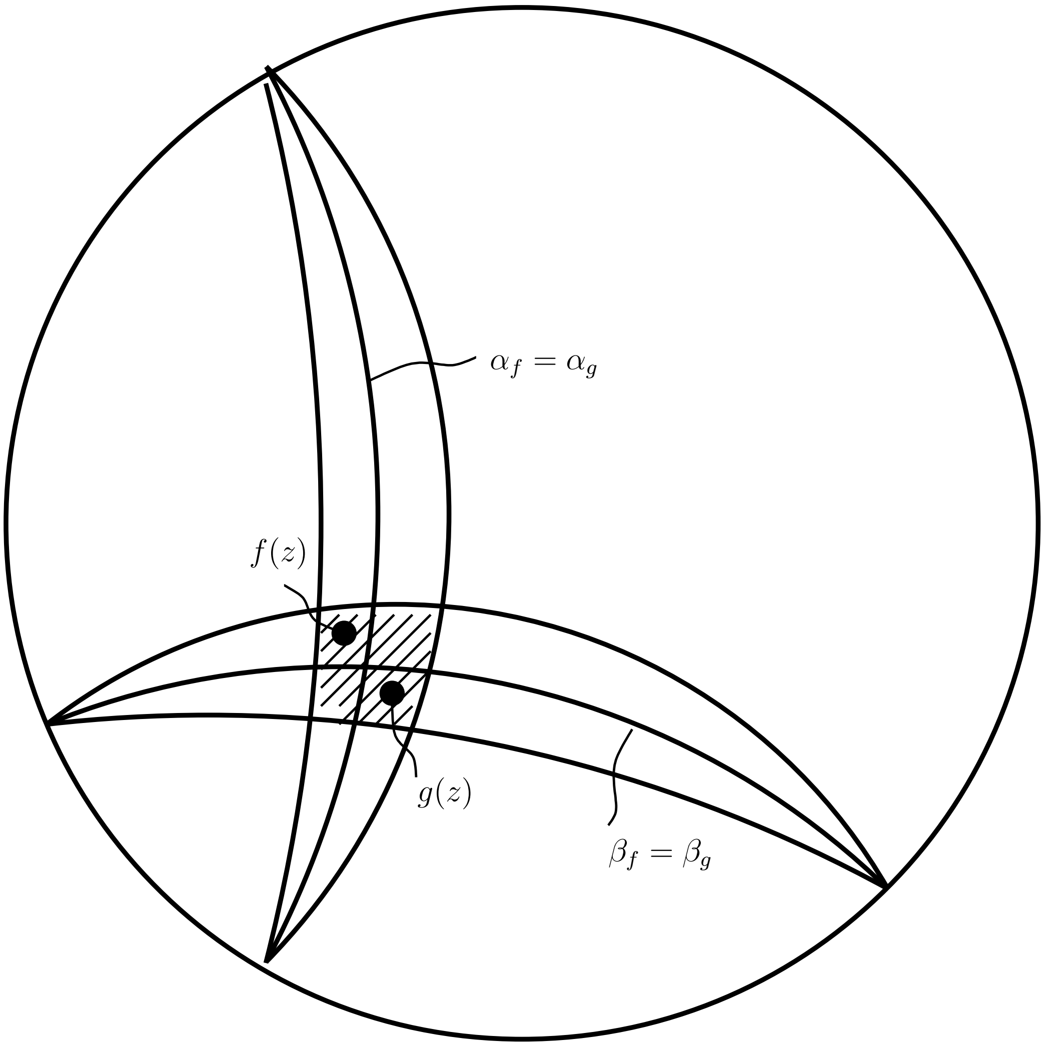

Figure 5.10.i. This picture illustrates the proof of Lemma 5.10. We illustrate the geodesics $\alpha_f = \alpha_g$ and $\beta_f = \beta_g$ and their ϵ-neighbourhoods.

6 Sullivan's Theorem: a new proof

This section is devoted to a new proof of Sullivan's Theorem (Theorem 5.1). In [18] 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 $\mathfrak{T}_\circ \subset \mathbb{C}$ be the open subset defined in Definition 4.3. Given $t \in \mathfrak{T}_\circ$, 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

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

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 $t_0 \in \mathfrak{T}_{\mathsf{o}}$ 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

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

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

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 $\mathfrak{T} = \mathfrak{T}_{\mathsf{o}} \cup \{t : \mathrm{Im}(t) > 0\}$, 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 $\{t \in \mathfrak{T}_{\mathsf{o}} : \mathrm{Im}(t) > 0\}$. 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 $\mathfrak{T}_\mathfrak{o}$ 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 $\iota : \mathbb{D}^2 \to$ Dome$\,(\Omega_{ic})$, which fixes the prime ends corresponding to -1, 1 and i.

Furthermore, g_{ic} and the nearest point retraction $r : \Omega_{ic} \to$ 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 $G(ic) \in \mathfrak{T}$ (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

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

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

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

□

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*

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

So $K \le 13.88$.

Proof (of Theorem 6.11): From Step 6.2.5 we have the holomorphic map $G : \mathfrak{X} \to \mathfrak{T}$. 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 $\mathfrak{X} \subset \mathbb{C}$ is simply the hyperbolic metric. In $\mathfrak{T}$ on the other hand, it is a famous theorem of [17], extended to the general case in [12] that the Kobayashi metric is the Teichmüller metric. In particular, in the respective metrics,

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

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. □

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

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

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 $X = \{z : |\mathrm{Im}(z)| \le \pi/\mu(\lambda)\}$. 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 $\mathrm{Im}(t)\,(\mathrm{Re}(t)_*(\mu))$.*

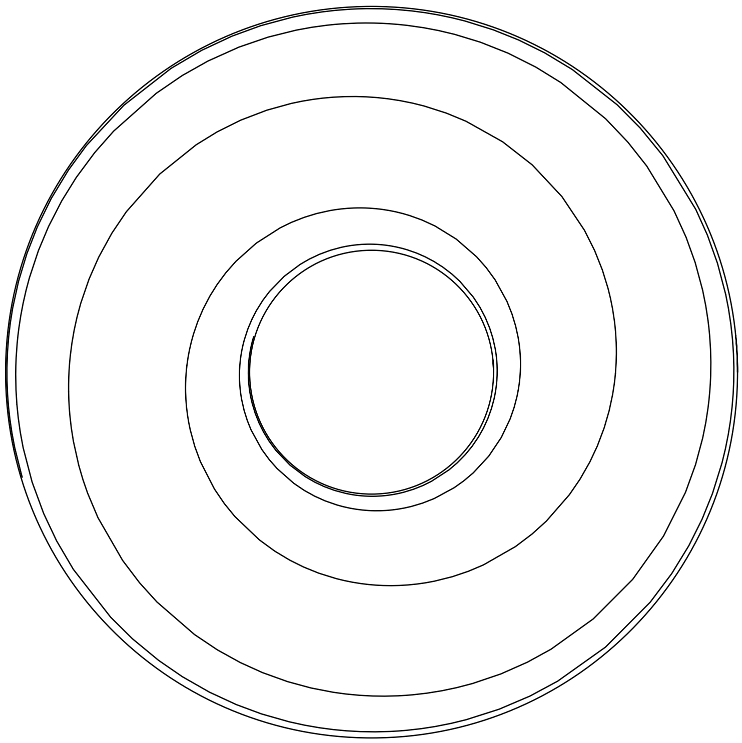

□

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

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

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 $\partial U = P_{\Lambda,t\mu} = \mathrm{Dome}\,(\Omega_t)$. ■

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 $P_{\Lambda,t\mu} = \mathrm{Dome}\,(\Omega_t)$. 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 $\left(\mathbb{C}E_{i\mathrm{Im}(t_\infty)}(x_i)\right)_{i\in\mathbb{N}}$ is a sequence converging to some point $x_\infty \in \mathbb{D}^3$, and prove a

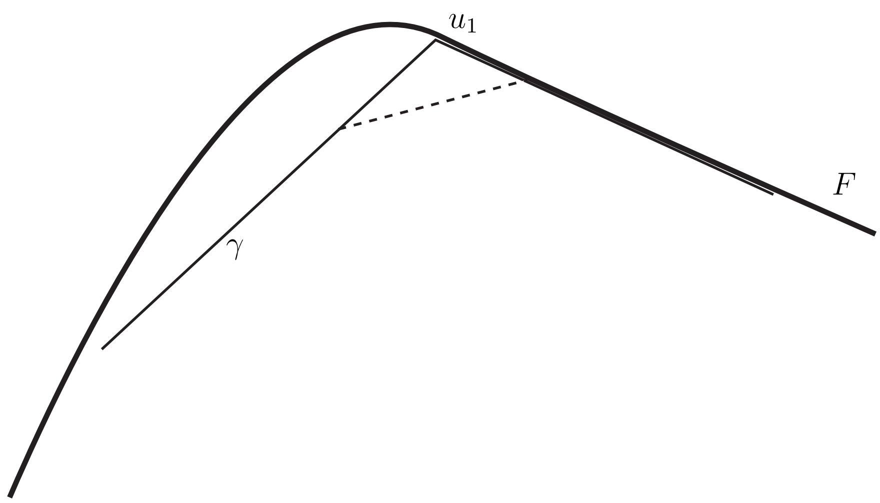

Figure 7.4.i. This illustrates the proof of Hadamard's Theorem, which arises in the proof of Lemma 7.4. F here represents a flat. If u_1 is an endpoint of $\gamma \cap F$, but not an endpoint of γ, then we can shorten γ near u_1, giving a contradiction. The dotted line shows a typical shortcut.

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.

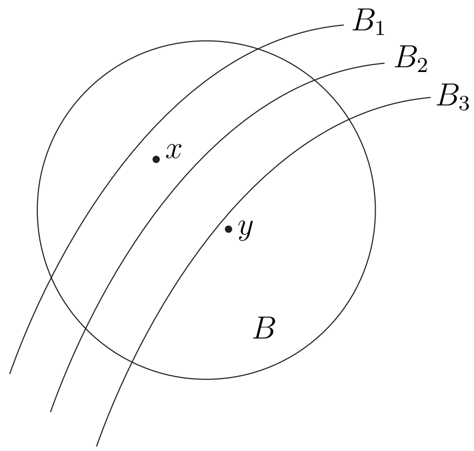

Figure 7.5.i. This illustrates part of the proof of Proposition 7.5 and we follow that notation. B is a 3-dimensional disk, with three sheets passing through it, labelled B_1, B_2 and B_3. The points x and y lie on the convex sides of B_1 and B_3 respectively. But the geodesic joining x to y must then 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 $r : \Omega_0 \to \mathrm{Dome}\,(\Omega_0)$ onto a flat or bending line of $\mathrm{Dome}\,(\Omega_0)$.

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 $r_t : \Omega_t \to \mathrm{Dome}\,(\Omega_t)$ to a flat or bending line. It follows that $\mathrm{Dome}\,(\Omega_t)$ 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 $U_i \subset \mathrm{int}(U_{i+1})$, 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 $t \in \mathbb{C}$, $\Psi_t|\overline{V} = \mathrm{Id}\,|\overline{V}$. 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. ∎

8 Convex regions of the plane

In this chapter, Ω will be a proper convex open subset of $\mathbb{C}$, unless otherwise stated. We allow Ω to have non-compact closure. In Theorem 8.12, we show that $r : \Omega \to \mathrm{Dome}(\Omega)$, the nearest point retraction, is a quasiconformal homeomorphism with quasiconformal constant at most 2. This proves the $K = 2$ conjecture for euclidean convex regions. We recall that we have produced (non-convex) examples where the conjecture is false (see [10] and [9]). In Definition 8.21, we recall the definition of substantial boundary dilatation. In Theorem 8.22, we compute the substantial boundary dilatation of r at each point of $\partial\Omega$. This enables us to use the

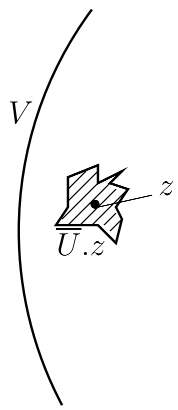

Figure 7.8.i. This illustrates the proof of the disk theorem, 7.8. V is a gap or c-leaf which is kept fixed, $z \in \Omega_0$ is a point which is near V and not near $\partial\Omega_0$. $\overline{U}.z$ is compact, near V, and contains z.

method of Chris Bishop to show that the Makarov Integral Means Spectrum for a euclidean convex region covers a larger range of exponents than that predicted by the Brennan Conjecture. In fact, the application to Brennan does not need the substantial boundary dilatation, because the derivative of the Riemann mapping of a convex domain to the unit disk is uniformly bounded. (We thank C. Pommerenke for pointing this out to us.)

We now outline the chapter. Our main objective is to study the substantial boundary dilatation at a point $p \in \partial\Omega$. So we are only interested in the situation near p, and it is reasonable to "zoom in". This method is discussed more formally in § 8.4. We zoom in as follows. We take a small round circle C, centred at p. This makes sense also if $p = \infty \in \mathbb{S}^2$. Let $A = \Omega \cap C$. We apply a similarity to $\mathbb{C}$, so that A becomes an arc of unit length, with the midpoint of A at the origin, and tangent direction at the midpoint vertical. Then, as C converges to p, Ω converges either to a wedge or to an infinite strip of width one.

We therefore study first, as important special cases of convex regions, wedges and strips, together with their domes. This study starts with Lemma 8.1 and finishes with Corollary 8.8. We also discuss some standard homeomorphisms between them. These standard isomorphisms provide a model that we try to follow in more complicated situations. In particular, there are canonical conformal isomorphisms of wedges, of half-cone surfaces and of helf-cylinder surfaces, with the infinite strip of width one.

Theorem 8.9 shows that the dome of a convex open region is a C^1-submanifold of $\mathbb{H}^3$. Remark 8.15 shows that a convex polygon with a finite number of sides breaks up canonically into a finite number of real analytic pieces, each of which corresponds, under the nearest point retraction, to some conical or hemispherical piece in the dome. From this we deduce (Theorem 8.12) that the dilatation of the nearest point retraction for an arbitrary convex region has dilatation bounded above by 2.

The Approximation Theorem (Theorem 8.13), proved in [10], is an essential tool in arguing that the finite sided case implies the general case. We prove the well-known fact that there is a well-defined angle at each point of the boundary of a general convex region in the plane (from Definition 8.16 to Remark 8.20).

Next, we state the main result of this chapter (Theorem 8.22), together with the one necessary additional definition (Definition 8.21). In § 8.23, Theorem 8.22 is proved in some special situations. In § 8.24, we specify how to move Ω by a similarity, so that its position becomes standard. This is useful when one works in coordinates.

In § 8.27, we present a rather explicit approximation of Ω by a region Ω_n, which, if $p = \infty$, is a locally finite convex polygon. If p is finite, $\partial\Omega_n$ is locally finite except in a neighbourhood of p. It turns out that the proof of Theorem 8.22 is easier when Ω is replaced by Ω_n. After proving the result for Ω_n, we can deduce it for Ω by letting n tend to infinity. In order to make Ω and Ω_n more similar to a strip, we apply log, getting subspaces M and M_n of $\mathbb{C}$ respectively. (In fact, sometimes it is convenient to apply $-\log$ and sometimes the identity map, instead of log. We use the notation h as a shorthand for whichever of the three possibilities is appropriate.)

We provide M_n with a basepoint $\omega_{x,n}$, whose horizontal coordinate x is large and positive, and which lies in the centre of M_n. Theorem 8.29 shows that the pair $(M_n, \omega_{x,n})$ can be mapped by a similarity $z \mapsto az + b$ where $a > 0$ to a very close approximation to the pair $(S_1, 0)$.

In our analysis it is important to distinguish two distinct types of bending line near p. This is done in § 8.31. Let $\sigma \subset \partial\Omega$ be a connected neighbourhood of p. A bending line near p has its endpoints in the same component of $\sigma \setminus \{p\}$ or in different components. In the former case, the bending line is called a *fringe bending lines*. We prove results about their shape and position. The other type of bending line is called an *arch*.

In Lemma 8.35, we prove that the nearest point retraction has dilatation near one, provided we are sufficiently close to $p \in \partial\Omega$ and sufficiently near the boundary, as measured in M. The map $\psi : S_1 \to S_u$ is obtained from the nearest point retraction $r : \Omega \to \mathrm{Dome}\,(\Omega)$ by composing with conformal isomorphisms. It therefore has the same dilatation as r. So properties required of ψ can be deduced from analogous properties of r.

In § 8.36, we define foliations of Ω, Ω_n and their domes. The leaves of the foliations of the domes are geodesics, and the nearest point retraction induces a map between foliations on the floor and the dome. More precisely, the foliations are defined only near p. These foliations will help us to compute partial derivatives and hence dilatations.

Having completed out overview of this section, we start the real work with a discussion of some standard examples: the wedge, the halfplane (the wedge with angle π) and the infinite strip, and some special homeomorphisms $\Omega \to \mathrm{Dome}\,(\Omega)$ in these standard cases. The particular homeomorphisms that interest us most are the nearest point retraction and the extremal quasiconformal homeomorphism which extends to the identity on $\partial\Omega$. This extremal happens to be unique for these special cases. We use the wedge

$$U_\alpha = \{z : -\alpha/2 < \arg(z) < \alpha/2\}\,.$$

This is the same as W_α (Definition 2.1), apart from a rotation. The construction of the dome of a wedge is illustrated in Figures 8.0.i and 8.0.ii. The following lemma is proved by a simple calculation.

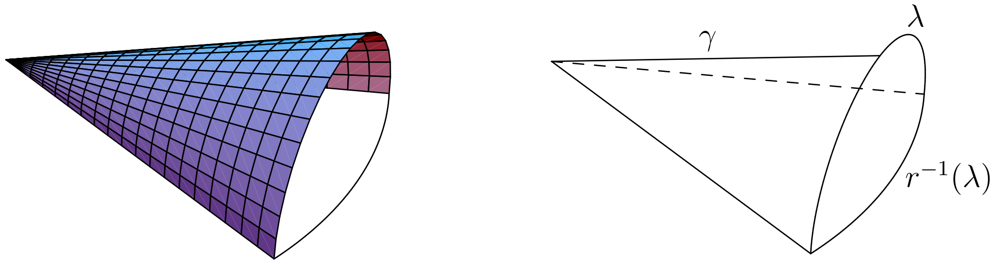

Figure 8.0.i. The dome of a wedge is a cone. In the lefthand picture, the mesh shown consists, on the one hand, of semicircles orthogonal to the plane $\{z = 0\}$ and, on the other hand, rays through the origin. The semicircles are hyperbolic geodesics in $\mathbb{H}^3$. The rays are not geodesics in $\mathbb{H}^3$. However, the ray γ, running along the highest points of the semicircles, is a geodesic for the induced Riemannian metric on the dome. One of the semicircles is labelled λ, in agreement with the surrounding text, and its inverse image $r^{-1}(\lambda)$ under the nearest point retraction r is also labelled.

8.1 Lemma. $\mathrm{Dome}\,(U_\alpha)$ *is a half-cone in* $\mathbb{U}^3$. *The nearest point retraction* $r : U_\alpha \to \mathrm{Dome}\,(U_\alpha)$ *is a real analytic diffeomorphism. Let*

$$u = s(\cos(\alpha/2), \sin(\alpha/2)\cos(\theta), \sin(\alpha/2)\sin(\theta)) \in \mathrm{Dome}\,(U_\alpha)\,,$$

where $s > 0$ *and* $0 < \theta < \pi$. *Then the dilatation at* u *of the inverse of the nearest point retraction is* $1 + \sin(\theta)\cos(\alpha/2)$.

In the case of the infinite strip, which we assume has width 2, the nearest point retraction is again a real analytic diffeomorphism. We take $u = (s, \cos(\theta), \sin(\theta))$. *The dilatation at* u *of the inverse of the nearest point retraction is* $1 + \sin(\theta)$.

The dilatation of the nearest point retraction $r : \Omega \to \mathrm{Dome}\,(\Omega)$, *when* Ω *is a convex wedge or an infinite strip, is at most 2.*

Proof: Using hyperbolic isometries of $\mathbb{U}^3$ which preserve the domain and its dome, we can restrict to a single bending line λ and its inverse image under the nearest point retraction r. Let u and v be the endpoints of λ. Then $r^{-1}(\lambda)$ is an arc of circle centred at the vertex of the wedge. The maximal disk $D \subset U_\alpha$, with $u, v \in \partial D$, has its centre c_0 at the intersection of the perpendiculars to ∂U_α at u and v. We may assume that D has radius 1. We can compute the coordinates of c_0 in terms of α. Let $p_0 = c_0 - (0,0,1)$. Then radial projection from p_0 sends $r^{-1}(\lambda)$ to λ and, on this set is equal to r. (To see this, recall that the nearest point retraction $D \to \mathrm{Dome}\,(D)$ is given by radial projection from p_0.)

The details of the computation are left to the reader. ■

Infinite strips play a fundamental role in our proof of Theorem 8.22. To fix the notation, we make the following definition.

8.2 Definition. For each $u > 0$, we set $S_u = \{z : -u/2 < \mathrm{Im}(z) < u/2\}$, the infinite horizontal strip of height u. □

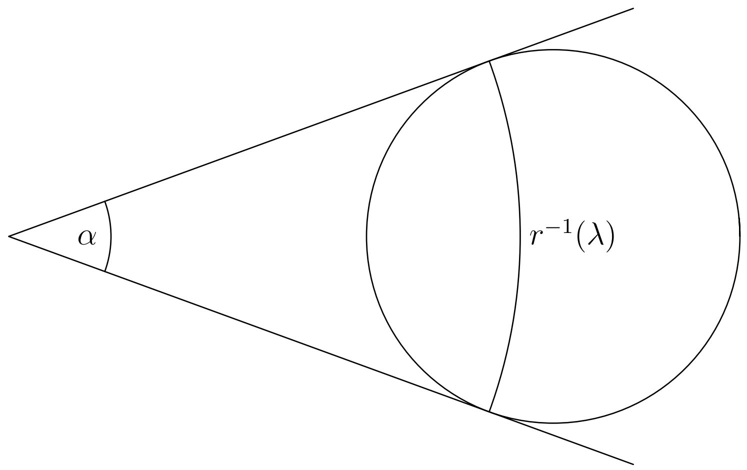

Figure 8.0.ii. This is a picture of the wedge whose dome is shown in Figure 8.0.i. The circle drawn is a circle that is maximal in the wedge. It is the boundary of the hyperbolic plane which contains the geodesic λ and is tangent to the cone surface. This hyperbolic plane is a hemisphere. The diagram has been rescaled so that the corresponding sphere has radius 1. We choose coordinates so that the wedge vertex is the origin. Then $S = (\cot(\alpha/2), 0, -1)$ the lowest point of that corresponding sphere. The nearest point retraction $r : r^{-1}(\lambda) \to \lambda$ is given by radial projection from S.

When constructing quasiconformal homeomorphisms $\Omega \to \mathrm{Dome}\,(\Omega)$, it is much more convenient to construct instead quasiconformal homeomorphisms from one infinite strip to another. The reason this is more convenient, is that we can then work with standard concepts involving two real variables, such as partial derivatives—much easier than trying to construct and work with the corresponding concepts on a convex hull boundary. To justify this transition to maps between strips, we need to construct conformal isomorphisms between Ω and an infinite strip, and also between $\mathrm{Dome}\,(\Omega)$ and another infinite strip.

The construction for general convex Ω will need in depth investigation. Here we construct these conformal isomorphisms only for the special cases of a wedge or infinite strip. This will give us a useful introduction to the properties applying to general convex subsets Ω.

8.3 Definition. If $\alpha > 0$ and we are concentrating on the boundary point $p = 0 \in \partial\Omega \subset \mathbb{S}^2$ of the wedge $\Omega = U_\alpha \subset \mathbb{C}$, we set $h = -\log : U_\alpha \to S_\alpha$. Here S_α is the infinite horizontal strip of height α. If $\alpha > 0$ and we are concentrating on the boundary point $p = \infty \in \partial\Omega \subset \mathbb{S}^2$ of the wedge $\Omega = U_\alpha \subset \mathbb{C}$, we set $h = \log : U_\alpha \to S_\alpha$. If $\alpha = 0$, so that $\Omega = S_v$ is an infinite strip for some $v > 0$, we set $h : S_v \to S_v$ equal to the identity, and take $p = +\infty \subset \partial S_v \subset \mathbb{S}^2$ to be the boundary point of interest. We set M to be the image of h, so that $M = S_\alpha$ or $M = S_v$, and $h : \Omega \to M$. In each case, h maps $p \in \partial\Omega$ to $+\infty \in \partial M \subset \mathbb{S}^2$.

In order to tie the choice of conformal isomorphism with an infinite strip more tightly to the geometry, and in order to make stronger the parallel with later more general constructions, we proceed as follows. We fix a bending line β_0 in $\mathrm{Dome}\,(\Omega)$, and set x_0 to be the euclidean midpoint of this semicircular arc in $\mathbb{U}^3 \subset \mathbb{R}^3$. We choose $z_0 \in \Omega$ so that $r(z_0) = x_0$. Let $g : M \to S_1$ be the translation plus change of scale sending $h(z_0) \in M$ to $0 \in S_1$.

Next we specify a conformal isomorphism of $\mathrm{Dome}\,(\Omega)$ with an infinite strip. We set γ to be the locus of the euclidean midpoints of the bending lines. Since Ω is a wedge or an infinite strip, this curve is a hyperbolic geodesic with respect to the hyperbolic path metric on $\mathrm{Dome}\,(\Omega)$. We set $u = u(\alpha) = \pi \sin(\alpha/2)/\alpha$. Let $\iota : \mathrm{Dome}\,(\Omega) \to S_u$ be the isometry of hyperbolic metrics which sends p to $+\infty$ and sends $x_0 \in \mathrm{Dome}\,(\Omega)$ to $0 \in S_u$. Then $\iota(\gamma)$ is the unique horizontal hyperbolic geodesic in S_u.

We define $\Psi : S_1 \to S_u$ as the composite

$$S_1 \xrightarrow{g^{-1}} M \xrightarrow{h^{-1}} \Omega \xrightarrow{r} \mathrm{Dome}\,(\Omega) \xrightarrow{\iota} S_u. \qquad (8.3.a)$$

We call Ψ the *model*. We will use it for purposes of comparison with other maps, similarly defined, but using a general convex region $\Omega \subset \mathbb{C}$ with boundary point $p \in \partial\Omega \subset \mathbb{S}^2$, instead of a wedge or an infinite strip. □

Note that, in the case of a wedge of angle α or an infinite strip

$$\Psi(t, s) = (t, -u/2 + f_\alpha(1/2 + s)), \tag{8.3.b}$$

where $s \in (-1/2, 1/2)$ and $t \in \mathbb{R}$. The important function $f_\alpha : (0, 1) \to (0, u)$ is defined by this equation.

We will need to prove a number of properties of f_α. This can be done either by direct calculation or by using the method of "zooming in". The second method will give less information, though still sufficient for nearly all our purposes in this paper. In view of its general importance, and its ease of use, we now explain this method in the present context.

8.4 Parametrized families of surfaces. The equation

$$y^2 + z^2 = \lambda^2 x^2.$$

defines a cone surface in $\mathbb{R}^3$.

We consider a sequence of points $(u_n)_{n\in\mathbb{N}}$ in the plane $z = 0$, converging to a point on the cone (not to $(0, 0, 0)$). We apply a change of scale and translation, so that u_n becomes the origin, and the distance of u_n from the cone is 1. This transforms the equation of the cone to a form where it depends real analytically on a certain parameter. Setting this parameter equal to zero, we obtain the equation of a vertical plane.

In psychological terms, imagine yourself moving towards the cone surface, while you shrink in size at the same rate. The cone surface towers above you, looking more and more vertical as you approach. In the limit the surface becomes a vertical halfplane.

Here is a version of the process in formulas. Let $u(\mu) = (1, -\lambda + \mu, 0)$ be our varying point u. We define the similarity $A_\mu : \mathbb{R}^3 \to \mathbb{R}^3$ by $A_\mu(v) = \mu v + u(\mu)$. Qualitatively, A_μ has the effect we require. With respect to the new coordinates given by A_μ, the cone surface has the equation

$$-2(1 + y)\lambda + (1 + 2y + y^2 + z^2)\mu - x\lambda^2(2 + x\mu) = 0. \tag{8.4.a}$$

Setting $\mu = 0$, we obtain the equation $1 + y + \lambda x = 0$ of a vertical plane.

Equation 8.4.a makes rigorous the statement that a family of cone surfaces converges analytically to a vertical plane. We make all objects of interest, for example Riemannian metrics, nearest point retractions, etc., depend analytically on μ. The analytic continuation of the nearest point retraction as μ changes sign is no longer a nearest point retraction. However, the map itself continues to be defined and meaningful, and it depends analytically on μ.

In a similar way, we can see that various objects vary analytically in the angle α as it moves through 0 to negative values or through π. When working in this way, we have to be careful to rethink the meaning of the objects, using analysis rather than particular accidents of the English language chosen to describe the objects. For example, there is absolutely no reason why the description "nearest point retraction" should continue to be a valid description when analytic continuation takes one beyond the original geometric framework in which the nearest point retraction was originally defined.

This method is called 'rescaling' or 'zooming in'. $\square$

8.5 Lemma. *Let f be the function of Equation 8.3.b. There is an $\epsilon > 0$, such that $f_\alpha(s)$ is the restriction of an analytic function of α and s, defined if $|s| < \epsilon$ and $|\alpha| < \pi$. (In fact, we can take $\epsilon = 1/2$, but we do not need this.) Also $f_\alpha(s) = s + O(s^2)$.*

Proof: We have explicit formulas for the conformal isomorphisms $h^{-1} \circ g^{-1} : S_1 \to U_\alpha$ and $\iota : \mathrm{Dome}\,(U_\alpha) \to S_u$. We are left with proving appropriate analyticity for the nearest point retraction $r : U_\alpha \to \mathrm{Dome}\,(U_\alpha)$. Analyticity in s and α can be proved using the considerations of § 8.4.

To see that $f(s) = s + O(s^2)$, we use the formula for the nearest point retraction of a halfplane onto a vertical halfplane. We can then apply Equation 8.4.a to see that a similar formula must hold for a cone.

Alternatively, note that the nearest point retraction approximately preserves euclidean distances from ∂U_α, for points very near ∂U_α and far from the vertex of U_α. $\blacksquare$

8.6 Remark. Although detailed information about $f_\alpha = f$ is not needed in this paper, we provide it in case readers are interested, and also as a check on the above reasoning. We calculate, along the lines indicated in Lemma 8.1, obtaining

$$\tan\left(\frac{f_\alpha(s)\alpha}{\sin\left(\frac{\alpha}{2}\right)}\right) = \frac{-\cos\left(\frac{\alpha}{2}\right) + \cos\left(\frac{\alpha}{2} - s\alpha\right)}{\sin\left(\frac{\alpha}{2}\right)\sin\left(\frac{\alpha}{2} - s\alpha\right)}. \tag{8.6.a}$$

This gives

$$f_\alpha(s) = s + \frac{\alpha}{2\tan(\alpha/2)} s^2 + O(s^3). \tag{8.6.b}$$

□

We recall Definition 8.3.a, where the model Ψ is defined. Let $t \in \mathbb{R}$ and $s \in (0,1)$. The dilatation of Ψ at $(t,s) \in S_1$ is $f'_\alpha(s)$. Since the other maps involved in the definition of Ψ are conformal, this is also the dilatation of the inverse of r, the nearest point retraction, at the appropriate point. This is recorded in Lemma 8.1 as being equal to $1+\sin(\theta)\cos(\alpha/2)$, where θ is the angular coordinate, on a certain vertical semicircle, of the point corresponding to $(s,t) \in S_1$. We see that $\theta = f_\alpha(s)\alpha/\sin(\alpha/2)$ and so

$$f'_\alpha(s) = 1+\sin(\theta)\cos(\alpha/2) = 1+\sin\left(\frac{f_\alpha(s)\alpha}{\sin(\alpha/2)}\right)\cos(\alpha/2) = 1+O(s).$$

By comparing Lemma 8.1 with the next result, we will see that the nearest point retraction does not give an extremal quasiconformal homeomorphism.

8.7 Lemma. *With respect to the euclidean pathmetric,* Dome (U_α) *is isometric to a wedge of angle $\beta = \pi\sin(\alpha/2)$. There is a unique extremal quasiconformal homeomorphism between this wedge and the wedge U_α, corresponding to the identity map on ∂U_α, namely the homeomorphism that preserves euclidean distance from the vertex and scales linearly the angle subtended at the vertex. The dilatation is constant and is equal to β/α.*

Proof: The logarithm function maps the domain U_α conformally to S_α (for the meaning of S_α, see Definition 8.2). Dome (U_α), which, with the euclidean pathmetric is euclidean isometric to a wedge of angle β, is mapped by the logarithm to S_β. We have a quasiconformal homeomorphism $q : S_\alpha \to S_\beta$, given by $q(x,y) = (x, \beta y/\alpha)$. The restriction of q to ∂S_α is the obvious map corresponding to the identity on ∂U_α.

The corresponding quasiconformal homeomorphism, still denoted by q, but now a map $q : S_1 \to S_{\beta/\alpha}$, is given by the same formula $q(x,y) = (x, \beta y/\alpha)$. It is better to divide by α because the new maps are applicable also when $\alpha = 0$, in which case U_α can be interpreted as an infinite strip, and β/α as $\pi/2$.

To prove our lemma, we apply a result due to Strebel and, independently, to Beurling–Ahlfors, which states that, given an affine map $A : \mathbb{R}^2 \to \mathbb{R}^2$, such that $A(S_1) = S_u$, there is a unique extremal quasiconformal map $q : S_1 \to S_u$ among homeomorphisms $f : S_1 \to S_u$ such that $f|\partial S_1 = A|\partial S_1$, namely $q = A|S_1$. An easy proof, if one requires to know only the extremality, which is all that is required in this paper, and not the uniqueness, is provided in [10, Section 3]. ■

The above proof also works for $\alpha = \pi$, when the nearest point retraction r is the relevant extremal quasiconformal mapping, and the dilatation is 1. We have also proved the following result

8.8 Corollary. *With respect to the euclidean pathmetric,* Dome (S_1) *is isometric to the strip* $S_{\pi/2}$. *There is a unique extremal quasiconformal homeomorphism* $q : S_1 \to S_{\pi/2}$ *corresponding to the identity map on* ∂S_1*, given by* $q(x,y) = (x, \pi y/2)$. *The dilatation is constant and is equal to* $\pi/2$.

We now consider the case of a general convex open subset $\Omega \subset \mathbb{C}$.

8.9 Theorem. *The Thurston isometry* $\iota : \mathbb{H}^2 \to$ Dome (Ω) *is a proper* C^1*-map with derivative everywhere of rank 2.* Dome (Ω) *is a proper* C^1*-submanifold of* $\mathbb{U}^3$.

8.10 Remark. Although the dome is differentiable, the nearest point retraction $r : \Omega \to$ Dome (Ω) may not be differentiable. In fact, it is unusual to have a convex region Ω for which the nearest point retraction is everywhere differentiable. Since quasiconformal maps are differentiable almost everywhere, one can however deduce from Theorem 8.12 that r is differentiable almost everywhere. A related point is that the dome of a convex region is usually not a C^2-submanifold of $\mathbb{U}^3$. □

Before proving Theorem 8.9, we establish some general properties of the bending lamination for the dome of a convex region.

8.11 Lemma. *Let* $x \in \partial\Omega$ *be the endpoint of at least one bending line. (The boundary of* Ω *is taken in* $\mathbb{S}^2$*.) Then the following must be true.*

8.11.1. $x \neq \infty$.

8.11.2. $\partial\Omega$ *has a unique tangent at* x.

8.11.3. Dome (Ω) *has no isolated bending line and each bending line has zero transverse measure.*

8.11.4. The nearest point retraction $r : \Omega \to$ Dome (Ω) *is a homeomorphism.*

8.11.5. There are at most two bending lines with x *as endpoint. If* x *is the endpoint of two distinct bending lines, then there is a flat such that each of the two bending lines is a boundary component of the flat. (See Figure 8.11.i.)*

Proof: First we deal with the case $x = \infty$. Since Ω is euclidean convex, a maximal disk in Ω can have ∞ in its closure only if it is equal to Ω. But then there are no bending lines and x cannot exist. For the remainder of the proof, we assume that x is finite.

If a bending line has positive transverse measure, then $\partial\Omega$ has a corner at the endpoint x, and the bending is in the wrong direction, preventing Ω from being convex. The first three conclusions follow.

In order to prove 8.11.5, suppose x is an endpoint of a bending line ℓ. Let P be a support plane for Ω which contains ℓ. There is a corresponding maximal disk $D \subset \Omega$, such that $x \in \partial D = \partial P$. Let y be the other endpoint of ℓ. Then $y \in \partial D$.

Since the tangent line to $\partial\Omega$ at x is unique, there is exactly one maximal disk $D \subset \Omega$ such that $x \in \partial D$. It follows that any bending line with x as an endpoint must lie in P. Moreover, $P \cap \mathrm{Dome}\,(\Omega)$ is the convex hull of $\partial\Omega \cap \partial D$. The statement 8.11.5 now follows. The situation is illustrated in Figure 8.11.i. ∎

Proof (of Theorem 8.9): Let $u \in \mathrm{Dome}\,(\Omega)$. If u lies in a flat, then there is clearly a smooth chart. If u lies on a bending line, we see from Lemma 8.11 that the Thurston isometry is C^1. ∎

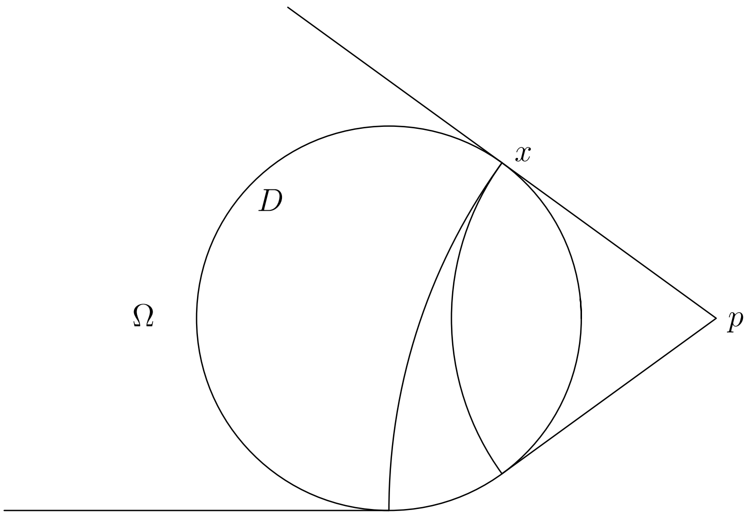

Figure 8.11.i. This illustrates 8.11.5. We show the convex set Ω and a maximal disk $D \subset \Omega$. Two bending lines for $\mathrm{Dome}\,(\Omega)$ have a common endpoint x. The bending lines are not shown. Instead we show their inverse images under the nearest point retraction.

Here is another general result.

8.12 Theorem. *Let Ω be a proper open convex subset of $\mathbb{C}$, possibly with non-compact closure, and let $r : \Omega \to \mathrm{Dome}\,(\Omega)$ be the nearest point retraction. Then r is a 2-quasiconformal homeomorphism.*

To prove Theorem 8.12, we will use approximation theorems from [10]. For the main result of this chapter (Theorem 8.22), we will rely heavily on Theorem 8.13, used in the current paper only with the additional (and unnecessary) restriction that Ω and Ω_n are all convex.

8.13 Theorem (The Approximation Theorem). *Let $K \geq 1$, and let $(\Omega_n)_{n\in\mathbb{N}} \subset \mathbb{C}$ be a sequence of simply connected proper open subsets converging to a simply connected proper open subset $\Omega \subset \mathbb{C}$. (The convergence is Hausdorff convergence of the compact complements in $\mathbb{S}^2$.) Given any sequence*

$$\left(\phi_n : \Omega_n \to \mathrm{Dome}\,(\Omega_n) \subset \overline{\mathbb{D}^3}\right)_{n\in\mathbb{N}}$$

of K-quasiconformal homeomorphisms which extend continuously to the identity on the common boundary, there is a subsequence, converging uniformly on compact subsets, to a K-quasiconformal $\phi : \Omega \to \mathrm{Dome}\,(\Omega) \subset \overline{\mathbb{D}^3}$ which extends continuously to the identity on the common boundary.

In order to prove Theorem 8.12, we will not rely on Theorem 8.13, but only on the more easily proved Theorem 8.14. This result, not stated explicitly in [10], is implicitly contained in the proof of [10, Lemma 7.5]. Alternatively, if we assume Ω and Ω_n convex, which is all that we need in the present paper, then it is easy to deduce Theorem 8.14 from Theorem 8.13, which is explicitly stated and proved in [10]. Although we need only the convex case, the more general statement given in Theorem 8.14 is true for any Ω and Ω_n.

8.14 Theorem. *Let $(\Omega_k)_{k\in\mathbb{N}}$ be a sequence of open simply-connected regions in $\mathbb{S}^2$ converging to the simply-connected region Ω, in the sense of Hausdorff convergence of their compact complements in $\mathbb{S}^2$. Let $r_k : \Omega_k \to \mathrm{Dome}\,(\Omega_k)$ and $r : \Omega \to \mathrm{Dome}\,(\Omega)$ be the nearest point retractions. Let $\iota_k : \mathrm{Dome}\,(\Omega_k) \to \mathbb{H}^2$ and $\iota : \mathrm{Dome}\,(\Omega) \to \mathbb{H}^2$ be the Thurston isometries, normalized consistently in some sensible way (as explained in [10, §7.8]). Then the sequence $(\iota_k \circ r_k)_{k\in\mathbb{N}}$ converges to $\iota \circ r$ uniformly on compact subsets of Ω.*

Proof (Theorem 8.12): By the above approximation theorems, we need only prove that r is 2-quasiconformal when Ω is a finite sided convex euclidean polygon, possibly with non-compact closure. Let V be a flat or a bending line in $\mathrm{Dome}(\Omega)$. If V is a flat, then r is conformal on $r^{-1}(V)$. Otherwise V is a geodesic joining two of the edges of the polygon Ω. Note that V cannot end at a vertex of the polygon. We denote these two edges of the polygon by σ_1 and σ_2. The union of bending lines with one endpoint in σ_1 and the other endpoint in σ_2 is a closed subset of $\mathrm{Dome}(\Omega)$. By 8.11.2, none of the four points which are endpoints of the intervals σ_1 or σ_2 can be the endpoints of any bending line. See Figure 8.14.i.

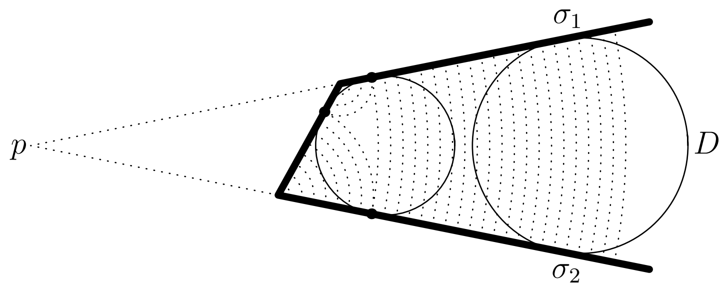

Figure 8.14.i. This illustrates the proof of Theorem 8.12. The boundary of the polygon is indicated by thick lines. The edges σ_1 and σ_2, when extended, meet at p. D denotes the maximal disk D in the proof. In the diagram we show two possible positions for D. The rest of this caption refers only to the extreme position of D towards the left. The geodesic V is the last bending line joining σ_1 to σ_2. The hyperbolic plane P with $\partial P = \partial D$ is divided into two halves by V. The left half of P meets $\mathrm{Dome}(\Omega)$ and the right half does not meet it. This is the basis of the general reasoning which shows that we can move D a little to the right by scaling with centre p, while retaining the property that it is a maximal disk of the polygon Ω, tangent both to σ_1 and to σ_2.

Suppose that, when extended, σ_1 and σ_2 meet at a point p. (If the sides are parallel, they meet at infinity. We leave to the reader the obvious interpretation of the following remarks in this case.) Let $D \subset \Omega$ be the maximal disk such that the endpoints of V lie on ∂D. Let P be the hyperbolic plane with $\partial P = \partial D$. Then $P \cap \mathrm{Dome}(\Omega)$ is either equal to V or to an ideal hyperbolic polygon, one of whose sides is V. Therefore at least one side of V in P does not meet $\mathrm{Dome}(\Omega)$. It follows that scaling by positive real numbers either in $(1-\epsilon, 1]$ or by positive real numbers in $[1, 1+\epsilon)$, with p as fixed point, will move D through maximal disks in Ω to which both σ_1 and σ_2 are tangent.

So the subset of Dome(Ω) consisting of bending lines with endpoints on σ_1 and σ_2 has non-empty interior. The computation of the dilatation of r can therefore be reduced to the computation for such a subset. This can be carried out by extending σ_1 and σ_2 until they meet, bounding a wedge of angle α with vertex p. By Lemma 8.1, the dilatation lies between 1 and $1 + \cos(\alpha/2) \leq 2$. ∎

8.15 Remark. We have shown above that, for a finite sided convex region, the dome consists of a finite number of pieces, each of which is either flat or part of a cone surface. □

Now we take Ω to be a general open convex region in the plane.

8.16 Definition. For each point $p \in \partial\Omega$, we have a minimal wedge $W_p \supset \Omega$, with vertex at p. We set $\mathrm{ang}(p)$ equal to the angle of W_p. We define the *directed tangent lines* at p to Ω to be the two lines containing the two edges of W_p. We direct the two tangent lines so that they point in a counterclockwise direction around Ω. See the picture in Figure 8.16.i. □

The angle between the two directed tangent lines is equal to $\pi - \mathrm{ang}(p)$. If $\mathrm{ang}(p) = \pi$, then the two directed tangent lines are equal.

The following lemma is obvious.

8.17 Lemma. *Let $(p_i)_{i\in\mathbb{N}}$ be a sequence in $\partial\Omega$ converging to p from one side of p. For each i, let τ_i be either of the two directed tangent lines at p_i. Then $(\tau_i)_{i\in\mathbb{N}}$ converges to τ, the directed tangent line at p on the same side of p as $(p_i)_{i\in\mathbb{N}}$.*

8.18 Definition. We extend Definition 8.16 to the case $p = \infty$ as follows. We fix a point $q \in \partial\Omega$ and define $W_{\infty,q} \subset \Omega$ to be the closure of the maximal wedge with vertices at q and ∞ contained in Ω. We define $\mathrm{ang}(\infty)$ to be the angle of $W_{\infty,q}$. It is easy to see that this angle is independent of q. See the picture in Figure 8.18.i.

$W_{\infty,q} \subset \Omega$ may be degenerate, collapsing to a single line. Whether it is degenerate or not, we obtain two directed tangent lines to $\partial\Omega$ at $\infty \in \mathbb{S}^2$. We have $0 \leq \mathrm{ang}(\infty) \leq \pi$, where 0 corresponds to the degenerate case.

We define $0 \leq \mathrm{ang}_0 \leq \pi$ to be the minimum of $\mathrm{ang}(p)$, as p varies over $\partial\Omega \subset \mathbb{S}^2$. □

8.19 Example. An example where $\mathrm{ang}(\infty) = 0$ is given by the set

$$\left\{(x,y) : y^2 < x\right\}.$$

This is one component of the complement of a parabola. □

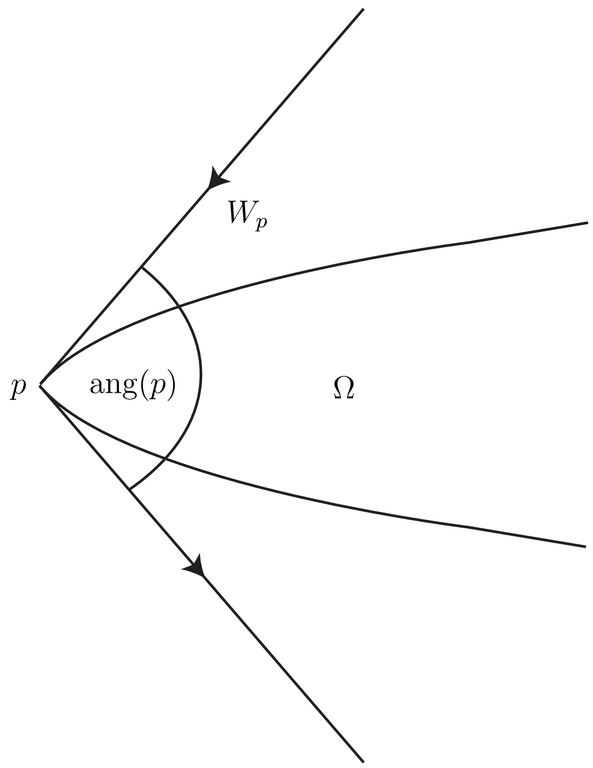

Figure 8.16.i. This illustrates Definition 8.16. W_p is the wedge contained between the two tangents to $\partial\Omega$ at p. We have $\Omega \subset W_p$.

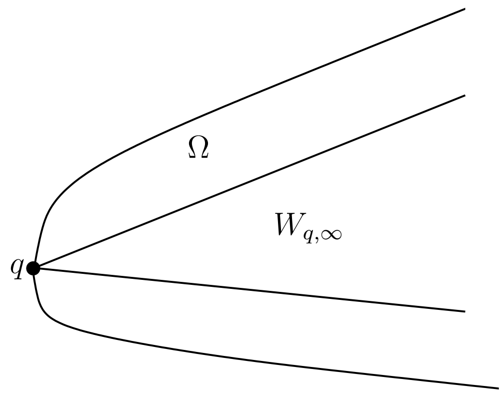

Figure 8.18.i. This illustrates Definition 8.18. We have $q \in \partial\Omega$. $W_{\infty,q} \subset \Omega$ is the wedge defined in 8.18

8.20 Remark. If $p = \infty$ and $\operatorname{ang}(p) = \pi$, then it is easy to see that Ω is a halfplane.$\square$

Let U and V be proper open simply connected subsets of the plane. Let $h_U : U \to \mathbb{D}^2$ and $h_V : V \to \mathbb{D}^2$ be Riemann mappings. Two quasiconformal homeomorphisms $f_1, f_2 : U \to V$, are said to have *the same boundary values* if $h_V f_1 h_U^{-1}, h_V f_2 h_U^{-1} : \mathbb{D}^2 \to \mathbb{D}^2$ have the same boundary values. (Recall that a quasiconformal homeomorphism between two open unit disks can always be continuously extended to the closed disks, so the "boundary values" of these quasiconformal homeomorphisms makes sense.)

8.21 Definition. Given a quasiconformal homeomorphism $f : U \to V$ and a point $p \in \partial U$, the notion of *substantial boundary dilatation* $H_p(f)$ of f at p is defined as follows. Denote the disk of radius s about p by $B_s(p)$. Suppose we have a quasiconformal map g of U with the same boundary values as f in the sense just described (write $g \sim f$). We denote the maximal dilatation of g restricted to $B_s(p) \cap U$ by $K_s(g)$. We make the following sequence of definitions:

$$\widetilde{H_p}(g) = \lim_{s \searrow 0} K_s(g), \quad H_p(f) = \inf_{g \sim f} \widetilde{H_p}(g), \quad H(f) = \max_{p \in \partial U} H_p(f).$$

The number $H(f)$ is called the *boundary dilatation* of f. It is known [11, 13] that the maximum is attained for some $p \in \partial U$. $\square$

Automatically $H(f) \leq K_0$ where K_0 is the maximal dilatation of an extremal quasiconformal map with the boundary values of f. Typically, there is strict inequality. See [16] for an exposition of this topic.

The following is the main result of this chapter.

8.22 Theorem. *Let $r : \Omega \to \operatorname{Dome}(\Omega)$ be the nearest point retraction. If $\operatorname{ang}(p) > 0$, then the substantial boundary dilatation of r at p is given by*

$$H_p(r) = \pi \sin(\operatorname{ang}(p)/2)/\operatorname{ang}(p).$$

If $\operatorname{ang}(p) = \pi$, then $H_p(r) = 1$. If $\operatorname{ang}(p) = 0$, then $H_p(r) = \pi/2$. The boundary dilatation is obtained by substituting ang_0 for $\operatorname{ang}(p)$ in these formulas (see Definition 8.18).

8.23 Special cases of Theorem 8.22. Suppose the boundary in $\mathbb{S}^2$ of the convex open subset $\Omega \subset \mathbb{C}$ includes ∞. There are at most two ways of going to ∞ in Ω, and, if there are two, then Ω must be an infinite strip. In this case, Lemma 8.7 is much stronger than Theorem 8.22. So, for the remainder of the proof of Theorem 8.22, we assume that Ω has at most one boundary point in $\mathbb{S}^2$ equal to ∞.

Next we think about the case when $\text{ang}(p) = \pi$. We need to prove that the substantial boundary dilatation at p is 1. If $p = \infty$ and $\text{ang}(p) = \pi$, then, as already pointed out in Remark 8.20, Ω is a halfplane and that the nearest point retraction is conformal. This reduces consideration of $\text{ang}(p) = \pi$ to the case when p is finite.

Given $u \in \Omega$ near p, we need to show that the dilatation of r near u is not much bigger than 1. We approximate Ω by a finite sided convex polygon, with internal angles near p very close to π. We complete the proof in this special case following the method of Theorem 8.12. The point is that here we are able to use the nearest point retraction, because its dilatation is very near one. Use of the nearest point retraction as the relevant homeomorphism satisfying Theorem 8.22, rather than some specially constructed homeomorphism, is appropriate only because we are near a point with angle π.

Our conclusions from these few remarks are that, when proving Theorem 8.22, we can confine our attention to points $p \in \partial\Omega$, finite or infinite, such that $0 \leq \text{ang}(p) < \pi$. Moreover, if $p = \infty$, we may assume it is the only boundary point at infinity.

Given a convex open region in the plane, its closure in $\mathbb{S}^2$ is a closed topological disk, except for an infinite strip. But this is excluded by the previous paragraph. So we assume from now on that the closure of Ω in $\mathbb{S}^2$ is a closed topological disk. □

8.24 Standard position for Ω. We recall from § 8.23 that we may assume that $\text{ang}(p) < \pi$ and that Ω has at most one infinite end. We will find a similarity that puts Ω into a convenient position for applying the log function.

First we look at the case $p = \infty$. For each finite $q \in \partial\Omega$, we have a closed wedge (or single ray) $W_{\infty,q} \subset \Omega$ (see Definition 8.18). As q varies, the edges of $W_{\infty,q}$ vary in either one or two parallel families of lines.

It follows from continuity (see Figure 8.24.i) that we can find a point q, with the property that the bisector of $W_{\infty,q}$ is orthogonal to a support line for Ω at q. We apply a similarity to make the support line vertical with Ω on its right, and move q to 0.

We recall that p is the corner point defined in Theorem 8.22. If $p = \infty$, we have defined q and described a normalization, with $q = 0$. This normalization fixes a branch of $\log : \Omega \to \mathbb{C}$.

Now we consider the case where p is finite. Then $\text{ang}(p) > 0$. Following Definition 8.16, we set W_0 equal to the minimal wedge containing Ω with vertex at

p. We may assume that $p = 0$ and that W_0 is a wedge with rays in the range $(-\operatorname{ang}(p)/2, \operatorname{ang}(p)/2)$. Once again, this fixes a branch of $\log : \Omega \to \mathbb{C}$. □

8.25 Definition.

- If $p = 0$, we set $h = -\log$.
- If $p = \infty$ and $\operatorname{ang}(\infty) > 0$, we set $h = \log$.
- If $p = \infty$ and $\operatorname{ang}(\infty) = 0$, we set h equal to the identity map.

In each case, we define $M = h(\Omega)$. After extending h continuously to the closure of Ω in $\mathbb{S}^2$, p is sent to a point $+\infty \in \partial M$. In each case $M = h(\Omega)$ is somewhat like a horizontal strip. The word "somewhat" does need to be taken seriously, as we see with Example 8.19, when $p = \infty$, $\operatorname{ang}(p) = 0$ and $M = \Omega$. The notation reminds us that $+\infty$ is thought of as lying at infinity on the right of M. □

Since Ω is convex, $\partial\Omega$ has a well-defined tangent except at a countable set of points, at each of which there are exactly two distinct tangent directions. It follows that ∂M has a well-defined tangent except at a countable set of points, where there are exactly two distinct tangents. Near the boundary point $+\infty$, the tangents to ∂M are almost horizontal.

8.26 Definition. As in Definition 8.3, we fix a bending line β_0 whose image under ι joins the two components of ∂S_u. We fix x_0 to be the euclidean midpoint of the semicircular arc $\beta_0 \subset \operatorname{Dome}(\Omega)$, and set $z_0 = r^{-1}(x_0) \in \Omega$. Without loss of generality (changing the choice of β_0), we may assume that all tangent lines to ∂M to the right of $h(z_0)$ have slope less than 0.001 in absolute value.

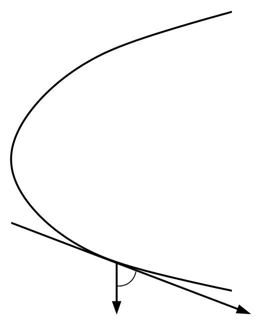

Figure 8.24.i. This illustrates the first paragraph of § 8.24.

To the right of $h(z_0)$, ∂M breaks up into two components: the upper component $\partial_+ M$ and the lower component $\partial_- M$. The corresponding parts of $\partial\Omega$ are denoted by $\partial_+\Omega$ and $\partial_-\Omega$. □

8.27 Approximating Ω. Some proofs turn out to be easier when Ω has a special form, namely when $\partial\Omega \setminus \{p\}$ is the union of a locally finite family of straight edges. For this reason, we will work with an approximation Ω_n to Ω, which has such a locally finite family of edges.

One slightly subtle point is that we make our proofs and estimates using Ω_n, but with conclusions which are independent of n. To see why we find this necessary, recall that we are computing the substantial boundary dilatation of the nearest point retraction $r : \Omega \to \mathrm{Dome}(\Omega)$ at p (see Definition 8.21). We need to show that we can find a very small neighbourhood N of p and $\tilde{r} : \Omega \to \mathrm{Dome}(\Omega)$, with the same boundary values as r and such that $\tilde{r}|N$ has dilatation not much larger than $\pi\sin(\mathrm{ang}(p)/2)/\mathrm{ang}(p)$. We will show that N can be chosen independently of n, so that, if Ω is replaced by Ω_n and r by r_n, there is an $\tilde{r}_n$ with the same boundary values as r_n, such that $\tilde{r}_n|N$ has dilatation uniformly near $\pi\sin(\mathrm{ang}(p)/2)/\mathrm{ang}(p)$. (That is, the closeness of the dilatation does not depend on n, but only on N.) As usual, interchanging of two limits—in this case, convergence of N to p and convergence of n to infinity—depends on some kind of uniform behaviour, or independence of the exact value of n.

We define Ω_n as the intersection of certain supporting open halfplanes of Ω. We specify Ω_n by specifying $X_n \subset \partial\Omega$, the set of points at which we take supporting halfplanes. If $x \in X_n$ is a point with two tangent lines, we use both supporting halfplanes. (Or, if one prefers, one can avoid the countable set of points where $\partial\Omega$ has two distinct tangent lines.)

Suppose we have defined subsets

$$X_i \subset (\partial\beta_0 \cup \partial_+\Omega \cup \partial_-\Omega)$$

for $1 \le i \le n-1$, so that

$$\partial\beta_0 \subset X_1 \subset \ldots \subset X_{n-1} \subset \partial\Omega \setminus \{p\}.$$

We define X_n so that the following conditions apply.

8.27.1. $\partial\beta_0 \subset X_n$.

8.27.2. X_n is closed in $\partial\Omega \setminus \{p\}$.

8.27.3. X_n is locally finite.

8.27.4. $X_{n-1} \subset X_n$.

8.27.5. Each component of $(\partial_-\Omega \cup \partial_+\Omega) \setminus X_{n-1}$ containing the endpoint of a bending line joining $\partial_-\Omega$ to $\partial_+\Omega$, also contains a point of X_n that is the endpoint of such a bending line.

8.27.6. If $x \in X_n$ is an endpoint of a bending line joining $\partial_-\Omega$ to $\partial_+\Omega$, then the other endpoint is also in X_n. To see that this requirement is consistent with the local finiteness requirement, consider a chain $x_1, \ldots, x_k \in \partial\Omega$ such that, for each $i \geq 1$, the geodesic $x_i x_{i+1}$ is a bending line in Dome(Ω) joining $\partial_-\Omega$ to $\partial_+\Omega$. By 8.11.3, such a chain must satisfy $k \leq 3$.

8.27.7. Each point of $\partial_+ M \cup \partial_- M$ is within a distance 2^{-n} of the image of X_n in ∂M.

Each point $x \in X_n$ gives rise to either one or two halfplanes containing Ω. Their intersection is a convex open region $\Omega_n \supset \Omega$ for each n, and $(\Omega_n)_{n\in\mathbb{N}}$ converges to Ω. We have $p \in \partial\Omega_n$ We set $M_n = h(\Omega_n)$, where $h = -\log$ if p is finite and $0 < \operatorname{ang}(p) < \pi$, $h = \log$ if $p = \infty$ and $0 < \operatorname{ang}(p) < \pi$, and $h = \mathrm{Id}$ if $p = \infty$ and $\operatorname{ang}(p) = 0$. □

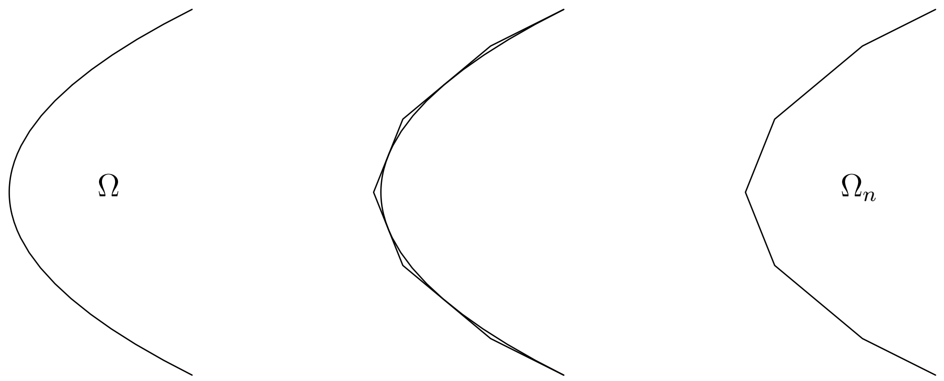

Figure 8.27.i. The lefthand figure is a convex curve, enclosing a convex region Ω. The point p is at infinity on the right. The middle picture shows Ω together with a number of support lines, as described in § 8.27. Using the support lines only, we obtain the polygonal approximation Ω_n, as shown in the righthand picture.

8.28 Construction of g and g_n. We next choose a conformal isomorphism $g : M \to S_1$, such that $g \circ h$ sends the $p \in \partial\Omega$ to the point $+\infty \in \partial S_1$. We require that $g(h(z_0)) = 0$, where z_0 is defined in § 8.26. This completely determines g.

Similarly, we set $M_n = h(\Omega_n)$ and define $g_n : M_n \to S_1$ to be the unique conformal isomorphism such that $g_n(h(z_0)) = 0$ and $g_n \circ h$ sends p to $+\infty$. □

We use the notation $\partial_- M_n$ and $\partial_+ M_n$, with its obvious meaning. For large $x > 0$, set $v_{x,n} \subset M_n$ to be the vertical interval with x-coordinate equal to x, joining $\partial_- M_n$ to $\partial_+ M_n$. Let $\sigma_{x,n} : \mathbb{C} \to \mathbb{C}$ be the affine map $z \mapsto az + b$, sending the interval $(-i/2, i/2)$ to $v_{x,n}$. Here $a > 0$ is the length of $v_{x,n}$ and $\mathrm{Re}(b) = x$. The imaginary part of b depends on the situation: unless $p = \infty$ and $\mathrm{ang}(p) = 0$, the imaginary part of b converges to zero as x tends to infinity. If p is finite, then the length of $v_{x,n}$ increases monotonically as x tends to infinity, with limit $\mathrm{ang}(p)$. If $p = \infty$ and $0 < \mathrm{ang}(p) < \pi$, then its length decreases monotonically, with limit $\mathrm{ang}(p)$.

Let $P_{x,n}$ be the inverse image of M_n under $\sigma_{x,n}$.

8.29 Theorem. *Let $(x(j))_{j\in\mathbb{N}}$ be a sequence of real numbers tending to $+\infty$, and let $(n(j))_{j\in\mathbb{N}}$ be an arbitrary sequence of positive integers. For each j, let τ_j be the unique horizontal translation such that the composite, ρ_j given by*

$$\rho_j : P_{x(j),n(j)} \xrightarrow{\sigma_{x(j),n(j)}} M_{n(j)} \xrightarrow{g_{n(j)}} S_1 \xrightarrow{\tau_j} S_1,$$

maps 0 to a point in S_1 on the y-axis. Given a compact subset C of the infinite open strip S_1 and an $\epsilon > 0$, there is an j_0 such that, if $j \geq j_0$, then $|\rho_j - \mathrm{Id}\,| < \epsilon$ on C. Moreover, j_0 can be chosen independently of the sequence $(n(j))_{j\in\mathbb{N}}$.

We can allow one or more of the $n(j)$ above to equal ∞. In this case, we replace $\Omega_{n(j)}$ by Ω, and correspondingly for the other terms, such as M, g, v and σ. Theorem 8.29 continues to be true in this greater generality.

Proof: We apply Caratheodory's Kernel Theorem (see [15, Theorem 1.8]) to the fixed point 0 in the varying region $P_{x(j),n(j)}$. We see that the conformal mapping $\phi_j : S_1 \to P_{x(j),n(j)}$, fixing 0, and with positive derivative at 0, converges to the identity map $S_1 \to S_1$, uniformly on compact sets. An argument by contradiction shows that the rate of convergence is independent of the choice of the sequence $(n(j))_{j\in\mathbb{N}}$.

The Poincaré metric in S_1 is $\pi|dz|/\sin(\pi y)$. Therefore the metric in $P_{x(j),n(j)}$ is near this on any compact subset $C \subset S_1$, for large j. We deduce that any geodesic for the hyperbolic metric on $P_{x(j),n(j)}$ for large j, that travels a substantial horizontal distance, must follow close to the central geodesic of S_1, the line $y = 0$, for some (lesser) substantial distance. If one wanted to, one could get numerical estimates through one's knowledge of the shapes of rightangled hyperbolic triangles. It follows from § 8.26 and § 8.28 that the inverse image under $\rho_j : P_{x(j),n(j)} \to S_1$ of the central line of S_1 is a real analytic path that approximates the central line closely in the C^∞-topology in any compact subset $C \subset S_1$, provided that j is large enough.

The claimed convergence follows. The independence of j_0 on the sequence $(n(j))_{j\in\mathbb{N}}$ follows using an argument by contradiction. ∎

We have also proved the following result.

8.30 Corollary. *We use the notation of Theorem 8.29. The geodesic in $P_{x(j),n(j)}$ from $\left(\sigma_{x(j),n(j)}\right)^{-1}(h(z_0))$ to $+\infty$ approximates the line $y = 0$ more and more closely in the C^∞-topology in the region $x \geq 0$, as j tends to infinity. This convergence is uniform in the sense that it is independent of the sequence $(n(j))_{j\in\mathbb{N}}$.*

8.31 Fringe bending lines. In order to get precise results about the dilatation of the nearest point retraction, we need to study bending lines for Ω and Ω_n in more detail. As in Figure 8.31.i, we consider the position of a bending line $\lambda \subset \mathrm{Dome}\,(\Omega)$, such that both endpoints p_1 and p_2 of λ lie in $\partial_+\Omega$ (or equivalently, both in $\partial_-\Omega$). We take λ very near to p. Then there is a unique maximal disk $D \subset \Omega$, such that the $p_1, p_2 \in \partial D$. Set s_D to be the radius of D and θ to be the change in the angle of the directed tangent line between the two endpoints of λ. By Lemma 8.17, θ is arbitrarily small. Then $d(p_1, p_2) \approx \theta.s_D$. So $\gamma = r^{-1}(\lambda)$ is a circular arc, which is almost a semicircle, of radius approximately equal to $\theta.s_D/2$.

Let $D_0 \subset \Omega$ be a maximal disk as near as possible to p_1, such that ∂D_0 meets both $\partial_-\Omega$ and $\partial_+\Omega$. Then $\mathrm{Dome}\,(\Omega)$ contains a bending line λ_0, with endpoints $p_- \in \partial_-\Omega$ and $p_+ \in \partial_+\Omega$. Let $\gamma_0 = r^{-1}(\lambda_0) \subset \Omega$. Let s_0 be the radius of D_0, and let α be the angle between the tangent lines at p_- and p_+. Then s_D/s_0 is bounded above (approximately) by 1, and α is very near to $\mathrm{ang}(p)$.

The length of the circular arc γ_0 is $s_0\alpha/\tan(\alpha/2)$ or $2s_0$ if $\alpha = 0$. Let y_1 be the length of the projection of $h(\gamma)$ onto the vertical coordinate of M, and let y_2 be the length of the vertical interval in M through $h(p_1)$. So y_1 is the angle subtended by γ at p (if $\mathrm{ang}(p) > 0$). The ratio y_1/y_2 is bounded above (approximately) by

$$\frac{s_D\theta/2}{s_0\alpha/\tan(\alpha/2)} \leq \frac{\theta\tan(\alpha/2)}{2\alpha} \approx \frac{\theta\tan(\mathrm{ang}(p)/2)}{2\,\mathrm{ang}(p)},$$

or by $\theta/4$ if $\operatorname{ang}(p) = 0$. (Recall that $\operatorname{ang}(p) < \pi$, so there is no danger that tan might become infinite.) Since θ tends to zero, this becomes arbitrarily small as we move towards p.

In M, as we move towards $+\infty$, the ratio of the euclidean size of fringe bending lines like λ in § 8.31, joining $\partial_-\Omega$ (or $\partial_+\Omega$) to itself, to the euclidean size of a nearby almost vertical interval joining $\partial_- M$ to $\partial_+ M$, converges to zero. In other words, if we zoom at p, the shape looks more and more like a wedge of angle $\operatorname{ang}(p)$, and the fringe bending lines become less and less noticeable.

8.32 Definition. We call a bending line like λ in § 8.31, which joins $\partial_-\Omega$ (or $\partial_+\Omega$) to itself, a *fringe*. We call a bending line like λ_0 in § 8.31, which joins $\partial_-\Omega$ to $\partial_+\Omega$ an *arch*. These definitions are only relevant near p. These terms are also used for bending lines of $\operatorname{Dome}(\Omega_n)$. □

The following lemma is an immediate consequence of the properties of log and of the geometry.

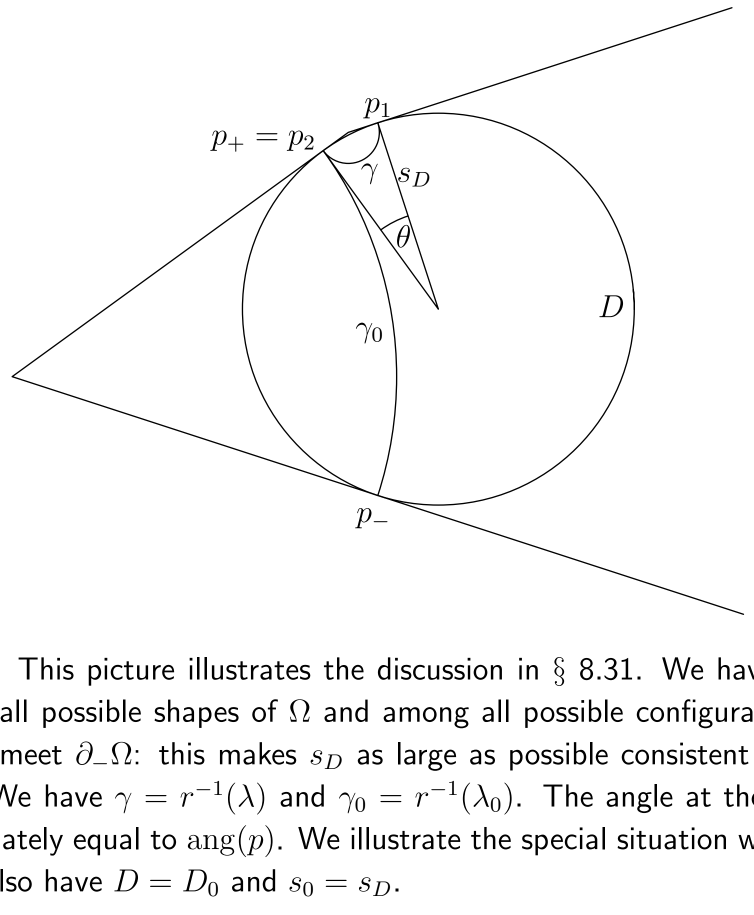

Figure 8.31.i. This picture illustrates the discussion in § 8.31. We have taken a case where, among all possible shapes of Ω and among all possible configurations within Ω, D happens to meet $\partial_-\Omega$: this makes s_D as large as possible consistent with our other assumptions. We have $\gamma = r^{-1}(\lambda)$ and $\gamma_0 = r^{-1}(\lambda_0)$. The angle at the vertex on the left is approximately equal to $\operatorname{ang}(p)$. We illustrate the special situation where p_2 and p_+ coincide. We also have $D = D_0$ and $s_0 = s_D$.

8.33 Lemma. *Let $\lambda \subset \mathrm{Dome}\,(\Omega)$ be an arch. Then $h(r^{-1}(\lambda))$ is an analytic curve in M. The angles which this curve makes with the vertical converge to zero as we get closer to p.*

Let λ be a bending line of $\mathrm{Dome}\,(\Omega)$. By Lemma 8.11, the endpoints p_1 and p_2 of λ have unique tangent lines to Ω. There is a unique maximal disk $D \subset \Omega$ with $p_1, p_2 \in \partial D$. Let $P \subset \mathbb{U}^3$ be the hyperbolic plane such that $\partial P = \partial D$. We define $P_\infty = P \cap \mathrm{Dome}\,(\Omega)$. Then P_∞ is a convex polygon in the hyperbolic plane P, possibly with an infinite number of sides. We have $\lambda \subset P_\infty$.

Now suppose that $p_1, p_2 \in X_n$, with tangent lines σ_1 and σ_2 to Ω. Then D is also a maximal disk in Ω_n. We define $P_n = P \cap \mathrm{Dome}\,(\Omega_n)$. Now P_n is either equal to λ or to a finite sided convex hyperbolic polygon with ideal vertices. In either case, p_1 and p_2 are ideal vertices of P_n and of P_∞. Each ideal vertex of P_n is in

$$\partial D \cap \overline{\Omega} \cap \partial\Omega_n \subset \partial D \cap \partial\Omega.$$

It follows that $P_n \subset P_\infty$. Either $\lambda = P_\infty$ or λ is an edge of ∂P_∞. We deduce that λ is also a bending line of $\mathrm{Dome}\,(\Omega_n)$. This will also hold for all larger values of n.

Among the bending lines of $\mathrm{Dome}\,(\Omega)$ which join $\partial_-\Omega$ to $\partial_+\Omega$, we therefore see certain bending lines of $\mathrm{Dome}\,(\Omega_n)$, more and more densely as n tends to infinity.

8.34 Definition. In parallel with Definition 8.3.a, we define $\psi_n : S_1 \to S_u$ as the composite

$$\psi_n : S_1 \xrightarrow{g_n^{-1}} M_n \xrightarrow{h^{-1}} \Omega_n \xrightarrow{r} \mathrm{Dome}\,(\Omega_n) \xrightarrow{\iota_n} S_u,$$

where again $u = \pi \sin(\mathrm{ang}(p)/2)/\,\mathrm{ang}(p)$.

Here $g_n : M_n \to S_1$ is the unique conformal isomorphism sending $h(z_0)$ to 0 and $+\infty$ to $+\infty$. Also $\iota_n : \mathrm{Dome}\,(\Omega_n) \to S_u$ is the unique conformal isomorphism sending x_0 to 0 and p to $+\infty$. □

8.35 Lemma. *There is a constant $c > 0$ and an $m_0 \in \mathbb{N}$, both independent of n, ($1 \le n < \infty$), such that the dilatation of $\psi_n|[m_0, \infty) \times (0, \delta]$ is bounded above by $1 + c\delta$. By symmetry of the hypotheses, the same is true for $\psi_n|[m_0, \infty) \times [1 - \delta, 1)$.*

Proof: Suppose the first coordinate t of $x = (t, s)$ is large, and that $0 < s \le \delta$. In order to prove the claimed bound on the dilatation near x, we need only prove the same bound for the dilatation of each nearest point retraction $r_n : \Omega_n \to \mathrm{Dome}\,(\Omega_n)$ in a neighbourhood of the point $u = h^{-1}\left(g^{-1}(x)\right) \in \Omega_n$. Recall from § 8.27 how Ω_n is constructed.

We now discuss the three possibilities for u.

8.35.1. First, u may lie in a flat of Dome (Ω_n). In this case the dilatation of r_n is one.

8.35.2. Second u may lie in a conical piece of Dome (Ω_n), corresponding to a wedge of angle α, with α very near to $\mathrm{ang}(p)$. (If $\mathrm{ang}(p) = 0$, we interpret "cone" to include "cylinder".) In this case, the bending line is an arch (see Definition 8.32). We apply Lemma 8.1 and see that the dilatation is $1 + \cos(\alpha/2)\sin(\theta)$, where $0 < \theta < \pi$ is the angular position of u in its bending line. For the model halfcone or halfcylinder, θ and δ can be explicitly computed in terms of each other, using Lemma 8.5 or Remark 8.6. Here our region is not a genuine halfcone or halfcylinder, but only an approximation to such. However, Theorem 8.29 shows that δ here is a very close approximation to δ in the model. We deduce that $\theta = O(\delta)$, and so the dilatation is $1 + O(\delta)$, as required.

8.35.3. Third, u may lie in a conical piece of Dome (Ω_n), corresponding to a wedge of angle α, with α very near to π. In this case, the angle θ in the formula of Lemma 8.1 bears no relationship with δ, a vertical measurement in S_1. By taking m_0 large enough, we can ensure that α is arbitrarily near to π. Note that m_0 can be chosen independently of n. Then $\cos(\alpha/2)$ is arbitrarily small, in particular smaller than δ. So $1 + \cos(\alpha/2)\sin(\theta) < 1 + \delta$, as required. ■

8.36 The foliation. The bending lamination and its inverse image under r provide convenient geodesics on which we can compute r precisely. However, we also need to carry out computations in the flats where the bending lamination does not exist. It is therefore helpful to extend the bending lamination to cover some of the flat regions.

We extend the bending lamination of Dome (Ω_n) to a foliation of those flat pieces F whose boundary components include two bending lines each joining $\partial_+\Omega$ to $\partial_-\Omega$. We denote these geodesics by p_1q_1 and p_2q_2, where, for $i = 1, 2$, $p_i \in \partial_+\Omega$ and $q_i \in \partial_-\Omega$. We choose the notation so that p_2q_2 separates p_1q_1 from the boundary point $p \in \partial\Omega$ in $\overline{\text{Dome}\,(\Omega_n)}$. Let $F \setminus p_1q_2 = F_1 \cup F_2$, where p_1q_1 is a side of F_1 and p_2q_2 is a side of F_2. We foliate F_1 by geodesics ending at p_1 and F_2 by geodesics ending at q_2. The situation is shown in Figure 8.36.i.

We denote the foliation by $\mathcal{F}_n$. Note that this is a foliation of a neighbourhood in Dome (Ω_n) of the boundary point $p \in \partial$Dome (Ω_n), not of the whole of Dome (Ω_n).

As we move towards p, the images in M_n under h_n of these leaves become more and more nearly vertical. □

The following lemma is well-known.

8.37 Lemma. *Let $\mathcal{F}$ be any foliation of a region $U \subset \mathbb{H}^2$ by geodesics, such that the complete geodesics are disjoint in $\mathbb{H}^2$. Let $\gamma : \mathbb{R} \to \mathbb{H}^2$ be any geodesic, parametrized by hyperbolic pathlength. Let $\theta(t)$ be the angle γ makes to $\mathcal{F}$ at $\gamma(t) \in U$. Then θ is a real-valued lipschitz function of a real variable, with lipschitz constant 1.*

Proof: First note that $\theta'(t) = \sin(\theta(t))$ for the foliation of $\mathbb{U}^2$ by vertical geodesics. This proves the result for any foliation by geodesics with a single common endpoint in $\partial\mathbb{H}^2$. Now let β_1 and β_2 be any two disjoint geodesics. We construct a foliation of $\mathbb{H}^2$ such that β_i is a leaf for $i = 1, 2$. Denote by β_3 one of the two diagonal geodesics joining an endpoint of β_1 to the diagonally opposite endpoint of β_2. Then β_3 separates $\mathbb{H}^2$ into two components. The component containing β_1 is foliated by all geodesics through the common endpoint of β_1 and β_3, and the other component is foliated by all geodesics through the common endpoint of β_2 and β_3. The known result for θ' can be applied to each component.

For $i = 1, 2$, let $\gamma \cap \beta_i = \{b_i\}$. By the Mean Value Theorem, the difference in the angle at which γ meets β_1 and β_2 is bounded by the hyperbolic distance $d(b_1, b_2)$. ∎

We transfer the foliation $\mathcal{F}_n$ on (part of) Dome (Ω_n), to the various spaces S_1, S_u,

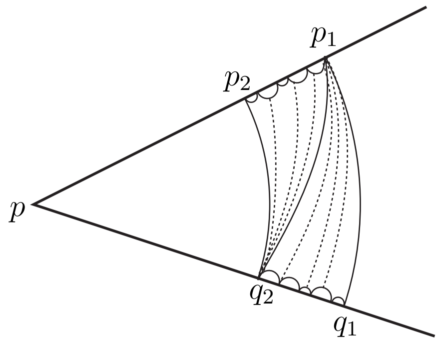

Figure 8.36.i. This illustrates § 8.36, where the bending lamination is extended to a foliation of a neighbourhood of p by geodesics. The picture should really be drawn on Dome (Ω), but we have instead drawn the inverse image in Ω under the nearest point retraction r, which is easier to do. Notice the small fringe geodesics. For convenience, we have drawn the boundary of Ω in two straight pieces, but we must have some deviation from straightness if there is to be any fringe. The solid curves are inverse images of bending lines. The dotted curves are inverse images of the new geodesics introduced in order to create a foliation.

M_n and Ω_n by using the maps

$$S_1 \xleftarrow{g_n} M_n \xleftarrow{h} \Omega_n \xrightarrow{r_n} \mathrm{Dome}\,(\Omega_n) \xrightarrow{\iota_n} S_u.$$

By abuse of notation, we will denote the corresponding foliations on any of these spaces by $\mathcal{F}_n$.

Let Γ_n be the curve in $\mathrm{Dome}\,(\Omega_n)$ which starts at x_0, the midpoint of the bending line β_0, and ends at p. It follows the locus of highest points on the various leaves of $\mathcal{F}_n$, using the height in upper halfspace. On a conical piece, Γ_n follows a geodesic path. We wish to understand how Γ_n behaves on a flat piece. Using the definition of $\mathcal{F}_n$ from § 8.36, we see that it is sufficient to deal with the case that all the geodesics in $\mathcal{F}_n$ lying in this particular flat have a single boundary point in common.

8.38 Lemma. *Let $Y_{m,n} = r_n\,(h^{-1}\,(g_n^{-1}\,([m,\infty) \times (0,1))))$. There is a sequence $(\epsilon_m)_{m\in\mathbb{N}}$ of positive numbers, converging to zero, such that, for each m, the total hyperbolic length of Γ_n within flat pieces meeting $Y_{m,n} \subset \mathrm{Dome}\,(\Omega_n)$ is bounded by ϵ_m. $\Gamma_n \cap Y_{m,n}$, parametrized by hyperbolic pathlength, is a $(1,\epsilon_m)$-quasigeodesic.*

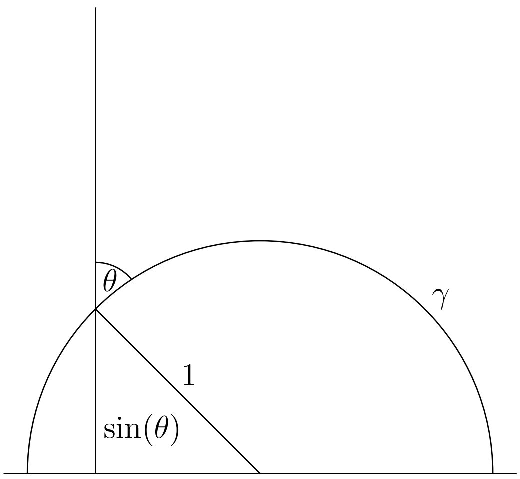

Figure 8.37.i. This diagram illustrates the proof of Lemma 8.37. We show the geodesic γ of that proof as a semicircle of radius 1 in the upper halfplane. We work out the relationship between the signed distance t measured along γ, and the angle that a vertical geodesic makes with γ. Since the point of interest is at height $\sin(\theta)$, we see that $dt = d\theta/\sin(\theta)$.

Proof: Let F be a flat piece crossed by Γ_n. By the construction (see Figure 8.36.i) of $\mathcal{F}_n$, we may assume that all the leaves of $\mathcal{F}_n$ in F have a boundary point in common. F lies in a hyperbolic plane P. We choose coordinates so that ∂P is the unit circle and such that the common boundary point of the leaves is $1 \in \partial P$. Suppose that the arch is the geodesic in upper halfspace joining $(1,0,0)$ to $(\cos(\theta), \sin(\theta), 0)$. The highest point on the geodesic is

$$u(\theta) = ((1+\cos(\theta))/2, \sin(\theta)/2, \sin(\theta/2)),$$

for $0 < \theta < 2\pi$. We have

$$u'(\theta) = (-\sin(\theta)/2, \cos(\theta)/2, \cos(\theta/2)/2).$$

The hyperbolic length of $u'(\theta)$ is strictly smaller than $2/\sin(\theta/2)$. If m is large, then $\theta \approx \pi - \operatorname{ang}(p)$. So $\|u'(\theta)\| \leq 2/\cos(\operatorname{ang}(p)/2)$.

The change in θ over $\partial F \cap \partial P$ is equal to the change in the angle that the tangent to Ω_n makes with the horizontal. The sum of such changes therefore tends to zero as m tends to infinity.

To see that Γ_n, parametrized by pathlength, is a $(1, \epsilon_m)$-quasigeodesic, let $t_1 < t_2$ be points in the domain of Γ_n, considered as a curve in $Y_{m,n} \subset \operatorname{Dome}(\Omega_n)$. Then $d(\Gamma_n(t_1), \Gamma_n(t_2)) \leq t_2 - t_1$. Let λ be the geodesic from $\Gamma_n(t_1)$ to $\Gamma_n(t_2)$. λ has to cross the same conical pieces as Γ_n. The length of λ is bounded below by the sum of the lengths of λ lying in conical pieces. This is no smaller than the sum of the lengths of Γ_n in the same conical pieces. In fact, unless λ and Γ_n coincide on such a conical piece, λ is definitely longer. This in turn is greater than $t_2 - t_1 - \epsilon_m$. In symbols, this means

$$t_2 - t_1 - \epsilon_m \leq d(\Gamma_n(t_1), \Gamma_n(t_2)) \leq t_2 - t_1$$

as required. ∎

The following consequence of Corollary 8.30 helps to show that our intuitive picture of the situation corresponds to reality. Let $u = \pi \sin(\operatorname{ang}(p)/2)/\operatorname{ang}(p)$ and let $\iota_n : \operatorname{Dome}(\Omega_n) \to S_u$ be the conformal isomorphism taking $x_0 \in \operatorname{Dome}(\Omega_n)$ to $0 \in S_u$ and the boundary point p to $+\infty$ (the definition is given above in the simpler situation of Definition 8.3). We recall the definition of $g_n : M_n \to S_1$ from § 8.28 and of $\psi_n : S_1 \to S_u$ from Definition 8.34 as the composite

$$S_1 \xleftarrow{g_n} M_n \xleftarrow{h} \Omega_n \xrightarrow{r_n} \operatorname{Dome}(\Omega_n) \xrightarrow{\iota_n} S_u.$$

In S_1 and in S_u, we have the central horizontal geodesic $y = 0$. This can be transferred to the other spaces using the homeomorphisms g_n, h, r_n and ι_n. In this way, we get two curves in each of the five spaces.

8.39 Proposition. *In each of the five spaces, the two curves are close to each other in the following sense. Firstly, we look at the two curves in S_1 through a window whose shape is $[m-1, m+1] \times (0,1) \subset S_1$. In the other four spaces, the shape is the subset which corresponds under the given homeomorphisms. Secondly, we use Hausdorff metrics with respect to the underlying hyperbolic metric. In S_1 and S_u, this is equivalent to using the Hausdorff metric based on the euclidean metric (because the curves lie near the centre of the strip). Thirdly, the tangent directions of nearby points on the two curves are near each other. Fourthly, the errors in how well the two curves approximate each other, both for position and direction, tends to zero as m tends to infinity.*

Proof: The statement about Hausdorff metrics follows from Corollary 8.30 and Lemma 8.38. In order to see that the statements about directions are true, note that the nearest point retraction r_n preserves $\mathcal{F}_n$ by definition. The direction orthogonal to $\mathcal{F}_n$ is also preserved by r_n. This is true on flat pieces where r_n is conformal, and on conical pieces because it is locally exactly the same as in the model situation. The other maps g_n, h and ι_n are conformal, so they have the same property. It follows that a curve which is almost orthogonal to $\mathcal{F}_n$ in one of the five spaces has that same property when transferred to the other spaces. ∎

8.40 Example. It is tempting to think that a central line in M_n could be added to Proposition 8.39 so as to give three asymptotic curves, rather than two. This can indeed be done if $\operatorname{ang}(p) > 0$. But if $p = \infty$ and $\operatorname{ang}(p) = 0$, there are counter-examples. For example, consider the subset of $\mathbb{R}^2$ lying between the positive x-axis and the upper half of the parabola $y^2 = x$. This is a convex region Ω, equal to M, and there is no central line corresponding to $p = \infty$. A hyperbolic geodesic in Ω limiting at ∞ cannot be asymptotic to a straight line. □

8.41 The leaves of $\mathcal{F}_n$ near p. We now discuss the shape of a leaf λ of $\mathcal{F}_n$ near the boundary point of interest (p or $+\infty$). We start in Ω_n, normalizing coordinates so that λ has length one. If $\operatorname{ang}(p) > 0$ and p is finite, then a leaf λ of $\mathcal{F}_n$ in Ω_n is an arc of circle, provided λ is near enough to p. As λ moves towards p, we continually rescale, normalizing with λ of length one, the centre of the circle containing the arc at zero, and the midpoint of λ real and positive.

As we do this, p tends towards the centre of the normalized circle containing λ as a subarc of length one. If $\operatorname{ang}(p) > 0$ and $p = \infty$, we proceed in the same way, but now the centre of the circle has no particular relationship to p. In each case, the radius of the circle whose circumference has λ as a subarc converges to $1/\operatorname{ang}(p)$. If $\operatorname{ang}(p) = 0$, λ converges to an interval of length one and, after normalization, p tends to infinity.

It follows that, on M_n near the boundary point $+\infty$, the leaves of $\mathcal{F}_n$ are nearly vertical curves (at any rate, away from the boundary curves), the tangent directions tending to the vertical as m tends to infinity, uniformly in n and in the position along the leaf. By Theorem 8.29, it follows that, on $[m, \infty) \times [\delta/2, 1 - \delta/2] \subset S_1$, $\mathcal{F}_n$ is nearly vertical, the error converging to zero as m tends to infinity, uniformly in n.

Recall from Lemma 8.38 that the isometry ι_n sends the hyperbolic geodesic Γ_n to the central line of S_u. On compact subsets of S_u, tangent vectors to $\mathcal{F}_n$ therefore converge to the vertical. (All fringe effects—see § 8.31—occur within a distance $\delta/4$ of ∂S_1 if m is large enough.) □

Now we come to one of the main intermediate results required in the proof of Theorem 8.22.

8.42 Proposition. *Let $u = \pi \sin(\operatorname{ang}(p)/2)/\operatorname{ang}(p)$ and let $\tau_z : \mathbb{C} \to \mathbb{C}$ be translation by z. For all $m, n \in \mathbb{N}$, consider the composite*

$$S_1 \xrightarrow{\tau_m} S_1 \xrightarrow{\psi_n} S_u \xrightarrow{\tau_{-m'}} S_u,$$

where $m' \in \mathbb{R}$ is chosen so that the composite sends 0 to a point on the y-axis. As m tends to infinity, this composite converges on compact subsets of S_1 to the model map Ψ—that is to the map $(t, s) \mapsto (t, -\frac{u}{2} + f_\alpha(s + \frac{1}{2}))$. (See Definition 8.3 and Equation 8.3.b.) The convergence is uniform on the compact subset and with respect to n. (See the statement of Theorem 8.29 for the meaning of uniform convergence with respect to n.)

Proof: The proof is based on rescaling. It seems clearest to separate the argument for different cases, even though the proofs of the different cases run along similar lines.

We start with p finite and $0 < \operatorname{ang}(p) < \pi$. We take $p = 0$. The two tangents at p are the rays $\arg(z) = -\operatorname{ang}(p)/2$ and $\arg(z) = \operatorname{ang}(p)/2$. Let $z_{n,m} = h^{-1}\left(g_n^{-1}(m)\right) = \exp\left(-g_n^{-1}(m)\right) \in \Omega_n$. For m large, $|z_{n,m}|$ is very small. Let $\Omega_{n,m}$ be the result of scaling Ω_n by a factor $|z_{n,m}|^{-1}$. In $\Omega_{n,m}$, our point of interest $z_{n,m}$ is sent into the unit circle.

The composite in the statement of Proposition 8.42 is factored into two other composites. Each of these is uniformly convergent, and we will be able to compute the limits. The first composite is

$$S_1 \xrightarrow{\tau_m} S_1 \xrightarrow{g_n^{-1}} M_n \xrightarrow{h^{-1}} \Omega_n \xrightarrow{1/|z_{n,m}|} \Omega_{n,m}$$

and the second is

$$\Omega_{n,m} \xrightarrow{r_{n,m}} \mathrm{Dome}\,(\Omega_{n,m}) \xrightarrow{|z_{n,m}|} \mathrm{Dome}\,(\Omega_n) \xrightarrow{\iota_n} S_u \xrightarrow{\tau_{-m'}} S_u.$$

Here the map between domes is the hyperbolic isometry of $\mathbb{H}^3$ which is equal to multiplication by the real number $|z_{n,m}|$ on $\partial\mathbb{H}^3 = \mathbb{C}$.

We claim that the sequence of first composites converges to the model map $S_1 \to U_{\mathrm{ang}(p)}$ (using the notation of Lemma 8.1), given by $z \mapsto \exp(-\mathrm{ang}(p)z)$. To see this, note first that from Corollary 8.30, Theorem 8.29 and the fact that $h^{-1}(z) = \exp(-z)$, we can deduce that the image of $y = 0$ in S_1 is a curve in Ω_n which ends at p and is asymptotic to the bisector of the angle between the two tangents at p. It follows that $z_{n,m}/|z_{n,m}|$ converges to 1 as m tends to infinity, and the convergence is uniform with respect to changes in n. Furthermore, the euclidean distance of $z_{n,m}/|z_{n,m}|$ to $\partial\Omega_{n,m}$ converges to $\sin(\mathrm{ang}(p)/2)$. The Koebe One Quarter Theorem now implies that the sequence of derivatives of the first composite at 0 is uniformly bounded away from 0 and ∞.

The second composite has the form demanded by Theorem 8.14, namely it is a nearest point retraction, followed by an isometry $\mathrm{Dome}\,(\Omega_{n,m}) \to S_u$. In order to apply Theorem 8.14, we also need to check normalizations. By construction, $z_{n,m}/|z_{n,m}|$ is sent to a point in S_u on the y-axis. By Proposition 8.39, this point is near $y = 0$. It therefore converges to $0 \in S_u$ as m tends to infinity. Consider the geodesic in $\Omega_{n,m}$ from $z_{n,m}/|z_{n,m}|$ to 0. By Proposition 8.39, this geodesic is sent to a curve in S_u which is asymptotic to $y = 0$. Theorem 8.14 now tells us that the limit is the map

$$U_{\mathrm{ang}(p)} \xrightarrow{r} \mathrm{Dome}\left(U_{\mathrm{ang}(p)}\right) \xrightarrow{\iota} S_u,$$

where ι is the hyperbolic isometry which (i) preserves orientation, (ii) sends $r(1)$ to 0 and (iii) sends the geodesic along the top of the cone to the central horizontal line of S_u.

Since the uniform limit of a composition is the composition of the uniform limit, Proposition 8.42 now follows if p is finite.

When $p = \infty$ and $\mathrm{ang}(p) > 0$, the same proof works, though now $z_{n,m}$ tends to infinity.

When $p = \infty$ and $\mathrm{ang}(p) = 0$, the proof is very slightly different. We define $z_{n,m} = x_{n,m} + iy_{n,m} = g_n^{-1}(m)$. Then $x_{n,m}$ tends to $+\infty$ as m tends to infinity, uniformly with respect to n. Also $y_{n,m}/xn, m$ tends to zero. Let $\sigma_{n,m}(z) = a(m,n)z + z_{n,m}$, where $a(m,n)$ is the length of the vertical interval from $\partial_- M$ to $\partial_+ M$ with x-coordinate equal to m. § 8.28. We define $\Omega_{n,m} = \sigma_{n,m}^{-1}(\Omega_n)$. As in the case when $p = 0$ and $\mathrm{ang}(p) > 0$, we factorize the map $S_1 \to S_u$ in the statement of Proposition 8.42 as

$$S_1 \xrightarrow{\tau_m} S_1 \xrightarrow{g_n^{-1}} \Omega_n \xrightarrow{\sigma_{n,m}^{-1}} \Omega_{n,m}$$

with

$$\Omega_{n,m} \xrightarrow{r} \mathrm{Dome}\,(\Omega_{n,m}) \xrightarrow{\sigma_{n,m}} \mathrm{Dome}\,(\Omega_n) \xrightarrow{\iota_n} S_u \xrightarrow{\tau_{-m'}} S_u.$$

The first composite sends 0 to 0. It sends the horizontal line $y = 0$ to a curve in $\Omega_{n,m}$ which, near $0 \in \Omega_{n,m}$, is nearly horizontal. It follows that the limit is the identity map. The second composite sends $0 \in \Omega_{n,m}$ to a point on the y-axis. By Proposition 8.39 this point converges to $0 \in S_u$. Also the second composite sends a horizontal interval of length two and centre 0 in $\Omega_{n,m}$ to a nearly horizontal curve passing near $0 \in S_u$. From Theorem 8.14, this converges to the model map.

This proves Proposition 8.42 when $p = \infty$ and $\mathrm{ang}(p) = 0$. ∎

The next thing we have to do is to estimate the derivative of the map $\psi_n : S_1 \to S_u$ (see Definition 8.34). Since the derivative of translation is the identity, this is the same as estimating the derivative of the composite of Proposition 8.42. In particular, we want to estimate, for small $\delta > 0$, the derivative of $\psi_n|[m-1, m+1] \times \{1/2 - \delta\}$. The estimate for the $\psi_n|[m-1, m+1] \times \{-1/2 + \delta\}$ will be the same, by symmetry of the hypotheses. For this discussion, we will use, in an essential way, the approximations Ω_n to Ω. (Previous uses of the approximation were not essential.)

The nearest point retraction $r_n : \Omega_n \to \mathrm{Dome}\,(\Omega_n)$ is real analytic in the interior of each conical, cylindrical or flat (that is, hemispherical) piece. On a bending line marking the transition between such pieces, r_n is not differentiable. The restriction of r_n to any leaf of $\mathcal{F}_n$ is real analytic.

8.43 Proposition. *There is an m_0 such that the restriction of ψ_n to the closed interval $\ell = [m-1, m+1] \times \{1/2 - \delta\}$ has the following properties:*

8.43.1. The tangent vector to the piecewise analytic curve $\psi_n(\ell)$ (which has two possible values at the endpoint of an analytic piece) is nearly horizontal. The convergence to the horizontal, as m tends to infinity, is uniform with respect to n. (The positive integer n can vary as m varies.)

8.43.2. In 8.43.1, the length of the tangent vector is $1 + O(\delta)$. The constants in the O-notation are independent of n and of $m \geq m_0$.

Proof: By 8.41, the leaves of $\mathcal{F}_n$ in S_1 and S_u are nearly vertical if m is large. The factors g_n, h and ι_n preserve angles and the foliation $\mathcal{F}_n$. The nearest point retraction r_n does not preserve angles. But it does preserve the foliation $\mathcal{F}_n$ and the direction orthogonal to the leaves of this foliation. Since r_n has dilatation at most 2 (by Theorem 8.12), very small angles which are almost horizontal or almost vertical in S_1 are sent by ψ_n to angles which are small and almost horizontal or almost vertical respectively. In particular the curve $\psi_n\left([m-1, m+1] \times \{1/2 - \delta\}\right)$ is almost horizontal in S_u. This proves 8.43.1.

To prove 8.43.2, we examine further the composite $\tau_{-m'} \circ \psi_n \circ \tau_m : S_1 \to S_u$ introduced in Proposition 8.42. We describe the effect of the composite on a leaf of $\mathcal{F}_n$ which passes through some fixed compact subset $C \subset S_1$. We parametrize the leaf in S_1 by euclidean pathlength, and we examine the effect of composing the parametrization with the above composite.

Since we are assuming m large, the leaf is almost vertical in C. In each of the spaces we look at the leaf will follow, more and more closely on a compact subset, the model direction. We will leave details of this to the reader, and concentrate on explaining more subtle features.

We compose the pathlength parametrization of a leaf in S_1 with the map

$$S_1 \xrightarrow{\tau_m} S_1 \xrightarrow{g_n^{-1}} M_n.$$

By Theorem 8.29, on the compact subset $C \subset S_1$, the parametrization of the image curve in M_n is C^1-near to some parametrization proportional to pathlength.

The next map to apply is $h^{-1} : M_n \to \Omega_n$ which is $z \mapsto \exp(\pm z)$ or the identity. The image in Ω_n of our leaf is an arc of circle whose centre is arbitrarily near 0 if $\mathrm{ang}(p) > 0$, and is an interval or arc of circle with very large radius if $\mathrm{ang}(p) = 0$. In the case of an arc of circle, the parametrization is C^1-near to a parametrization proportional to angle subtended at the centre. For both arcs of circle and intervals in Ω_n, the parametrization is C^1-near a parametrization proportional to pathlength. The constant of proportionality is the pathlength of the arc or interval.

If we stick to a single leaf of $\mathcal{F}_n$ in Ω_n, the effect of the nearest point retraction can be computed exactly, using the map f_α (see Equation 8.3.b). Let P be the hyperbolic plane tangent to Dome(Ω_n) and containing the corresponding leaf in Dome(Ω_n). Then the subscript α of f_α is the angle between the two tangents to P at the two ends of the leaf. So α converges to $\mathrm{ang}(p)$ as m tends to infinity, and f_α is C_∞ approximated by $f_{\mathrm{ang}(p)}$. The derivative of ψ_n along a leaf of $\mathcal{F}_n$ is therefore approximately $f'_{\mathrm{ang}(p)}(\delta) = 1 + O(\delta)$ at a point of the curve $\psi_n\left([m-1, m+1] \times \{-1/2+\delta\}\right)$.

If we are in a flat piece of Dome (Ω_n), the nearest point retraction is conformal, and the directional derivative has the same magnitude in each direction. In particular, the length of the derivative vector with respect of $\tau_{-m'} \circ \psi_n \circ \tau_m$ along $\tau_{-m'}\psi_n\tau_m|[-1,+1] \times \{1/2-\delta\}$ with respect to the parameter $t \in [-1,1]$ is $1 + O(\delta)$.

If we are in a conical piece, we can argue directly from the shape to prove that the derivative has magnitude arbitrarily near 1. However, we prefer to follow the style of reasoning just used, which gives a weaker result. Namely, Lemma 8.35 shows that the nearest point retraction is $(1+O(\delta))$-quasiconformal along the curve of interest. It follows that the directional derivative in any direction must have magnitude $1+O(\delta)$. This completes the proof of 8.43.2. ■

Using the previous results, we now define a quasiconformal map $\rho_{n,\delta} : S_1 \to S_u$ which will give rise to the sequence we require for the proof of Theorem 8.22. The definition depends on a certain positive integer M_0. We will prove that M_0 can be chosen large enough so that all the properties required of $\rho_{n,\delta}$ are true. We choose the leftmost leaf λ of $\mathcal{F}_n$ which passes through the strip $\mathbb{R} \times [-\frac{1}{2}+\delta, \frac{1}{2}-\delta]$ so that its intersection with the strip is entirely (weakly) to the right of the vertical line $x = M_0$.

8.43.1. If $s \leq -\frac{1}{2}+\delta$ or for $s \geq \frac{1}{2}-\delta$, then $\rho_{n,\delta}(t,s) = \psi_n(t,s)$.

8.43.2. If $-\frac{1}{2}+\delta \leq s \leq \frac{1}{2}-\delta$ and (t,s) is to the left of λ, then $\rho_{n,\delta}(t,s) = \psi_n(t,s)$.

8.43.3. Let $t \geq M_0+1$. On the vertical interval $\{t\} \times [-\frac{1}{2}+\delta, \leq \frac{1}{2}-\delta] \subset S_1$, $\rho_{n,\delta}(t,s)$ is a linear interpolation of the values already defined at the two endpoints.

8.43.4. We have not yet defined $\rho_{n,\delta}$ on an almost rectangular region whose four sides are given by λ, $y = \frac{1}{2}-\delta$, $y = -\frac{1}{2}+\delta$, and $x = M_0+1$. Using Lemma 8.44, we will define $\rho_{n,\delta}$ so that it is K-quasiconformal, for some fixed K, independent of n, δ and M_0. Here $\delta > 0$ must be chosen small enough and M_0 large enough, so that the basic topological features of the picture of the map $\rho_{n,\delta}$ are as we expect them to be. For example M_0 must be chosen greater than the m_0 of Proposition 8.43, and δ must be chosen small enough so that the image of the (almost) rectangular gap really is an (almost) rectangle.

We start by filling the gap at 8.43.4. Recall that the cross-ratio of a quadruple of points in $\mathbb{CP}^1$ is a point in $\mathbb{CP}^1$, defined provided that at least three of the four points are distinct. We fix conventions by defining the cross-ratio $X(z, 1; 0, \infty) = z$. Given a Jordan region, $U \subset \mathbb{S}^2$ and points $x_1, x_2, x_3, x_4 \in \partial U$, at least three of which are distinct, we define the cross-ratio $X_U(x_1, x_2; x_3, x_4)$ by choosing a Riemann mapping $f : U \to \mathbb{H}^2$, extending f continuously to $f : \overline{U} \to \overline{\mathbb{H}^2}$, and then taking the usual cross-ratio $X(fx_1, fx_2; fx_3, fx_4)$. The result is independent of the choice of Riemann mapping.

Given two Jordan regions U and V, and a homeomorphism $\rho : \partial U \to \partial V$, ρ is said to be *quasisymmetric* if it does not distort cross-ratios too much. More precisely, the standard theory (due to Ahlfors), says that we can fix a basepoint $x_0 \in U$ and confine our attention to counter-clockwise quadruples $(a, b, c, x_0) \in U^3$ with $X_U(a, c; b, x_0) = -1$. Here a, b, c, x_0 are distinct points in ∂U. We say that ρ is k-quasisymmetric if the image cross-ratio satisfies $1/k \le X_V(\rho(a), \rho(c); \rho(b), \rho(x_0)) \le k$ for all choices of a, b, c in ∂U giving cross-ratio -1 as already explained.

In order to fill in the gap at 8.43.4, we apply these considerations to $\rho_{n,\delta}$. The basepoint for computation of cross-ratios will be the upper righthand corner of the rectangular region—explicitly $\left(M_0 + 1, \frac{1}{2} - \delta\right) \in S_1$.

8.44 Lemma. *There exists $k \ge 1$, $\delta_0 > 0$ and $m_0 \in \mathbb{Z}$, such that, for all n, for all $0 < \delta \le \delta_0$ and for all $M_0 \ge m_0$, the restriction of $\rho_{n,\delta}$ to the boundary of the rectangular gap is k-quasisymmetric.*

Proof: Suppose this is false. Then we can construct the following data.

8.44.1. A sequence $(\delta_i)_{i\in\mathbb{N}}$ tending monotonically to zero.

8.44.2. A sequence $(M_i)_{i\in\mathbb{N}}$ tending monotonically to infinity.

8.44.3. A sequence $(n_i)_{i\in\mathbb{N}}$ which does not necessarily converge.

8.44.4. The corresponding maps

$$T_i = \tau_{-m'} \circ \rho_{n,\delta} \circ \tau_m : S_1 \to S_u,$$

with $m = M_i$, defined except on the interiors of the rectangular gaps.

8.44.5. For each i, an almost rectangular domain gap $U = U_i$ and an almost rectangular image gap $V = V_i$, corresponding to $n = n_i$, $\delta = \delta_i$ and $M_0 = M_i$.

8.44.6. Convergent sequences $(a_i)_{i\in\mathbb{N}}$, $(b_i)_{i\in\mathbb{N}}$ and $(c_i)_{i\in\mathbb{N}}$, with $a_i, b_i, c_i \in \partial U_i$, with limits a_∞, b_∞ and c_∞ respectively, and cross-ratio $X_U(a_i, c_i; b_i, x_0) = -1$.

The data has the following properties.

8.44.7. The cross-ratio $X_V(T_i(a_i), T_i(c_i); T_i(b_i), T_i(x_0))$ converges either to infinity or to zero.

8.44.8. $(U_i)_{i\in\mathbb{N}}$ converges to the rectangle $U_\infty = [0,1] \times [-\frac{1}{2}, \frac{1}{2}]$.

8.44.9. $(V_i)_{i\in\mathbb{N}}$ converges to the rectangle $V_\infty = [0,1] \times [-\frac{u}{2}, \frac{u}{2}]$.

A theorem of Rado (see [15, Theoremm 2.11]) shows that the Riemann maps for the regions U_i (normalized for example by sending three of the four corners to three fixed points in the boundary of the disk) converge to the Riemann map for the rectangle U_∞ and the Riemann maps for the regions V_i converge to the Riemann map for the rectangle V_∞. Rado's theorem uses the fact that the boundaries of the regions are parametrized, and the parametrizations converge uniformly to the parametrization of the boundary of the limit.

Note that the boundary maps $T_i : \partial U_i \to \partial V_i$ together with their inverses $T(i)^{-1} : \partial V_i \to \partial U_i$ are piecewise differentiable maps, with uniformly bounded derivatives if we parametrize by pathlength. We have four euclidean distances in S_1, namely $|x_0 - a_i|$, $|a_i - b_i|$, $|b_i - c_i|$ and $|c_i - x_0|$. If the smallest of the four distances is bounded below away from zero, then the uniform boundedness of derivatives shows that the same is true for the images under T_i. But this is impossible, because then the image cross-ratio would tend neither to zero nor to infinity. We therefore assume that, independently of i, the distance occupying the same position (1,2, 3 or 4) is the smallest of the four, and converges to zero.

Let (v_i, w_i) denote that pair in the four element set

$$\{(x_0, a_i), (a_i, b_i), (b_i, c_i), (c_i, x_0)\}$$

which has smallest distance apart. We define a similarity $\sigma_{U,i}$ with multiplicative factor $1/|v_i - w_i|$, such that $\sigma_{U,i}(v_i)$ converges to some point in $\mathbb{C}$. Now look at the image of each of the four "corners" of U_i under $\sigma_{U,i}$. We may assume that each of these images converges, either to a finite point, or to infinity. At most one of the four corners has a finite limit, because the distance between two corners, after applying $\sigma_{U,i}$ is very large. If this happens, choose the constant term in the similarity $\sigma_{U,i}$ so that it maps that corner to 0. If no corner gives a finite limit, choose $\sigma_{U,i}$ so that $\sigma_{U,i}(v_i) = 0$ for each i.

We define the similarity $\sigma_{V,i}$ with multiplicative constant $1/|v_i - w_i|$, and so that, for each i, $\hat{T}_i = \sigma_{V,i} \circ T_i \circ \sigma_{U,i}^{-1}$ sends 0 to 0. We set $\hat{U}_i = \sigma_{U,i}(U_i)$ and $\hat{V}_i = \sigma_{V,i}(V_i)$. So $\hat{U}_i$ and $\hat{V}_i$ are very large almost rectangles.

If one of the corners of $\hat{U}_i$ and $\hat{V}_i$ is at 0 for each i, then $\hat{U}_i$ and $\hat{V}_i$ converge to one of the coordinate quadrants of $\mathbb{C}$. Otherwise, $\hat{U}_i$ and $\hat{V}_i$ converge to halfplanes (with either vertical or horizontal boundary). Riemann maps with image $\mathbb{H}^2$ for these limits $\hat{U}_\infty$ and $\hat{V}_\infty$ are given by rotation about 0 through a multiple of $\pi/2$, possibly followed by squaring $z \mapsto z^2$.

The derivatives of $\hat{T}_i$ are the same as those of T_i. From Proposition 8.43, we see that the limit $\hat{T}_\infty$ of the $\hat{T}_i$ is the identity on horizontal boundaries. On a lefthand vertical boundary it is equal to multiplication by $f'_{\mathrm{ang}(p)}(s)$, where $s \in [0,1]$ corresponds to the limit v_∞ of $(v_i)_{i\in\mathbb{N}}$. On a righthand vertical boundary, $\hat{T}_\infty$ is multiplication by $u = \pi \sin(\mathrm{ang}(p)/2)/\mathrm{ang}(p)$.

Now $\hat{v}_\infty$ and $\hat{w}_\infty$ are finite, by construction, corresponding to two of the four sequences $(\sigma_{U,i}(a_i))_{i\in\mathbb{N}}$, $(\sigma_{U,i}(b_i))_{i\in\mathbb{N}}$, $(\sigma_{U,i}(c_i))_{i\in\mathbb{N}}$, $(\sigma_{U,i}(x_0))_{i\in\mathbb{N}}$. Recall that we have chosen $|u_i - v_i|$ to be the smallest of four choices. Therefore the other two limits are distinct from $\hat{v}_\infty$ and $\hat{w}_\infty$, so that three of the four limit points are distinct and their cross-ratio is well-defined. By definition, the cross-ratio $X_{U_i}(a_i, c_i; b_i, x_0) = -1$. By Rado's theorem, cited above, the cross-ratio of the limits $X_{U_\infty}(\hat{a}_\infty, \hat{c}_\infty; \hat{b}_\infty, \hat{x}_\infty) = -1$. It follows that the four limits are distinct points of the boundary of the quadrant or halfplane $\hat{U}_\infty$. It further follows that the four image points in $\partial\hat{V}_\infty$ are distinct. So their cross-ratio is neither 0 nor ∞. But this is a contradiction, and so the proof of Lemma 8.44 is complete. ∎

Proof (completion of Theorem 8.22): By Lemma 8.44, we can fill the gap in the definition of $\rho_{n,\delta}$ with a k_1 quasiconformal homeomorphism, where $k_1 > 1$ is a uniform constant. (The method of proof of Lemma 8.44 could be used to give an explicit of k in Lemma 8.44, and therefore of k_1, in terms of $\mathrm{ang}(p)$.)

We first prove that $\rho_{n,\delta} : S_1 \to S_u$ is a homeomorphism (for small δ and large m_0). This follows because it is injective and sends the boundary to the boundary. Therefore, by topology, it is surjective and a homeomorphism.

In order to show that it is quasiconformal, with an appropriate constant, we need to compute the derivative on $(m_0+1, \infty) \times [-\frac{1}{2}+\delta, \frac{1}{2}-\delta]$. For an appropriate choice of $\lambda(s)$, the 2×2 derivative matrix at (x, s) is equal to

$$\left[\lambda(s)\frac{\partial\psi_n}{\partial x}(p_1) + (1-\lambda(s))\frac{\partial\psi_n}{\partial x}(p_2) \quad \frac{\psi_n\left(x, \frac{1}{2}-\delta\right) - \psi_n\left(x, -\frac{1}{2}+\delta\right)}{1-2\delta}\right],$$

where $p_1 = (x, \frac{1}{2} - \delta)$ and $p_2 = (x, \frac{1}{2} + \delta)$, and each entry is a column vector of length two. For large enough m_0, this matrix can be estimated as

$$\begin{bmatrix} 1 + O(\delta) & O(\delta) \\ O(\delta) & u + O(\delta) \end{bmatrix},$$

where the constants in the $O(\delta)$ terms do not depend on n. By taking δ sufficiently small, we can make the dilatation arbitrarily near u.

Now $\rho_{n,\delta} : S_1 \to S_u$ is a sequence of quasiconformal homeomorphism with a uniform quasiconformal constant. These all send 0 to 0 and $+\infty$ to $+\infty$, and so there is a limit $\rho_\delta : S_1 \to S_u$ with the same quasiconformal constant. The $\rho_{n,\delta}$ correspond to quasiconformal homeomorphisms $\Omega_n \to \mathrm{Dome}\,(\Omega_n)$ which extend continuously to the identity map on $\partial\Omega_n$. By Theorem 8.13, ρ_δ corresponds to a quasiconformal homeomorphism $\Omega \to \mathrm{Dome}\,(\Omega)$ which extends continuously to the identity on $\partial\Omega$. If we let δ tend to zero, we obtain the type of sequence necessary in order to find the substantial boundary dilatation. This completes the proof of Theorem 8.22. ■

8.45 Remark. Our results about convex regions of the plane can be made invariant under Möbius transformations, by suitably generalizing our class of regions. We consider a class of regions more restrictive than regulated regions. (Our results on substantial boundary dilatation are false for general regulated regions.) Let $\Omega \subset \mathbb{S}^2$ be a simply connected open region, not the complement of a point. For each boundary point p, we suppose that p has a neighbourhood, such that $\partial\Omega$ near p is locally the graph of a function f_p. We assume that the second derivative of f_p in the distributional sense is a constant plus a positive Borel measure. In this case we say that Ω is *infinitesimally convex*. We thank Adrien Douady for making this condition explicit for us. The condition is invariant under Möbius transformation.

Given an infinitesimally convex set, it is possible to change coordinates by a Möbius transformation, so that a neighbourhood of p becomes convex. Our proof will then apply. □

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] D.B.A. Epstein and V. Markovic. The logarithmic spiral: a counter-example to the $k = 2$ conjecture. available at `http://www.maths.warwick.ac.uk/~dbae` or at `http://www.maths.warwick.ac.uk/~markovic`.

[11] R. Fehlman. Ueber extremale quasikonforme abbildungen. *Comment. Math. Helv.*, 56:558–580, 1981.

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

[13] N. Lakic. Substantial boundary points for plane domains and Gardiner's conjecture. *Ann. Acad. Sci. Fenn.*, 25:285–306, 2000.

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

[15] Ch. Pommerenke. *Boundary Behaviour of Conformal Maps.* Springer-Verlag, 1992.

[16] E. Reich. Extremal quasiconformal mappings of the disk. In R. Kühnau, editor, *Geometric Function Theory, Vol.1, Handbook of Complex Analysis.* Elsevier, 2001.

[17] 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.

[18] 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`