Convex Regions in $\mathbb{C}$ and their Domes

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

17 August 2004

Let Ω be a *euclidean convex* region $\neq \mathbb{C}$. It need not have compact closure. We will study Ω from above by examining its geometric relationship to Dome(Ω). Here Dome(Ω) is the relative boundary component in the upper half 3-space of the hyperbolic convex hull of the closed set $\mathbb{S}^2 \setminus \Omega$. It is constructed by considering maximal disks whose interiors are contained in Ω. Each is bounded by a circle c meeting the boundary at two or more points. From the circle c rises a hemisphere h which divides upper half space in two components. Denote the component which does not abut the disk by H_h. The hyperbolic convex hull is $\cap_h H_h$. The dome is the envelope of the hemispheres. In short, the dome lies over the floor Ω as the dome of a domed stadium lies over the ground under it.

The ambient hyperbolic metric induces a path metric on the dome, which is referred to as the hyperbolic metric on the dome. Thurston proved the Riemann mapping theorem for it, namely there is an isometry between the dome and the unit disk in the respective hyperbolic matrics. It is uniquely determined up to compositions by isometries.

In more detail, the dome is the union of at most countably many open *flat pieces*, which are hyperbolic polygons generally with infinitely many sides, and the complementary closed set. The complementary set is the union of mutually disjoint bending lines. Each bending line is a hyperbolic geodesic with endpoints in the common boundary $\partial\Omega$ of Ω and Dome(Ω).

An isolated bending line ℓ has a well defined *bending angle* θ which we take to be the *exterior* bending angle. Thus $0 < \theta < \pi$ where the extreme $\theta = 0$ means no bending at all, while $\theta = \pi$ means one side of ℓ is folded over the other. In general there is a real valued Borel measure μ on transverse segments τ whose endpoints are in flat pieces or on bending lines with zero atomic measures. It is constructed in a manner akin to the Riamann integral by using finite approximations to the dome.

Its value depends only on the flat pieces or lines containing the end points of τ.

For example, if Ω is a wedge with interior vertex angle less than π, its dome is a half-cone so there are no flat pieces. If Ω is a triangle, its dome is the union of three conical pieces and a flat piece. In fact the dome is a C^1-surface without any isolated bending lines. Similar statements are true for any convex polygon (finite sided).

For an introduction to the general theory see [1].

The basic tool we use is the *nearest point retraction* $r : \Omega \to \mathrm{Dome}(\Omega)$. This is defined as follows. Given $z \in \Omega$ consider the family of spheres in upper halfspace which are tangent to $\mathbb{C}$ at z. As the diameter increases there is a first diameter which the sphere touches the dome, necessarily at a single point which we denote by $r(z)$. The retraction r is continuous, and is a homeomorphism if and only if there are no isolated bending lines. In any case r extends continuously to $\partial\Omega$ and pointwise fixes it.

There are two theorems proved in this paper for euclidean convex regions. The first is Theorem 4.1. It states that r is actually a 2-quasiconformal homeomorphism.

The second, Theorem 6.1, gives a universal bound on the distortion of r at the boundary. At least for convex regions with C^1-boundaries, this result implies that there exists a unique extremal quasiconformal map between the floor and the dome which, like r, pointwise fixes the common boundary. Moreover, it is of Teichmüller type. We also show the maximal dilatation of the extremal mapping cannot exceed 2.

In spite of the transparency of our results for finite sided convex polygons, there are major technical hurdles in extending the results, especially for the second theorem whose proof occupies most of this paper. We were not at all sure what was true when we started our investigation.

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

1 Outline of the paper.

As stated above, Ω will be a proper convex open subset of $\mathbb{C}$ (with a few explicit exceptions). We allow Ω to have non-compact closure. In Theorem 4.1, we show that $r : \Omega \to \mathrm{Dome}\,(\Omega)$, the nearest point retraction, is a quasiconformal homeomorphism with quasiconformal constant at most 2.[1] In Definition 5.6, we recall the definition of substantial boundary dilatation. In Theorem 6.1, we compute the substantial boundary dilatation of r at each point of $\partial\Omega$. This enables us to use the 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.)

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 § 3. 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 2.1 and finishes with Corollary 3.4. 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 half-cylinder surfaces, with the infinite strip of width one.

Theorem 3.6 shows that the dome of a convex open region is a C^1-submanifold of $\mathbb{H}^3$. Remark 4.4 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 4.1) that the dilatation of the nearest point retraction for an arbitrary convex region has dilatation bounded above by 2.

The Approximation Theorem (Theorem 4.2), proved in [3], 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 5.1 to Remark 5.5).

Next, we state the main result of this chapter (Theorem 6.1), together with the

[1]This proves Thurston's $K = 2$ conjecture for euclidean convex regions. The conjecture is false for certain nonconvex regions—see [2], [3].

one necessary additional definition (Definition 5.6). In § 6.2, Theorem 6.1 is proved in some special situations. In § 6.3, we specify how to move Ω by a similarity, so that its position becomes standard. This is useful when one works in coordinates.

In § 6.6, 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 6.1 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 6.8 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 § 7. 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 7.4, 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, 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.

2 Standard examples.

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

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

The construction of the dome of a wedge is illustrated in Figures 2.0.i and 2.0.ii. The following lemma is proved by a simple calculation.

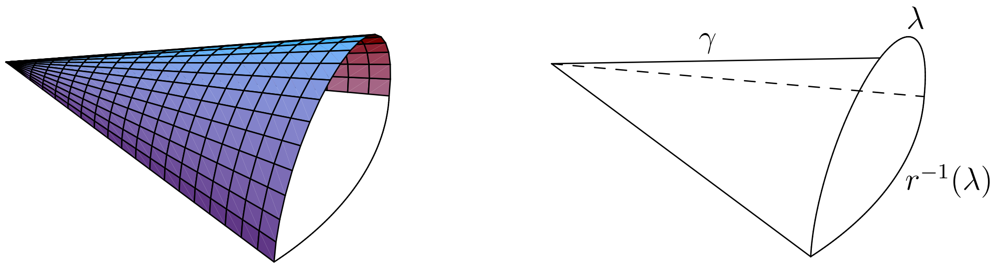

Figure 2.0.i. The dome of a wedge is a cone. In the lefthand picture, the mesh shown consists, on the one hand, of semicircles orthogonal to the plane $\{z = 0\}$ and, on the other hand, rays through the origin. The semicircles are hyperbolic geodesics in $\mathbb{H}^3$. The rays are not geodesics in $\mathbb{H}^3$. 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.

2.1 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

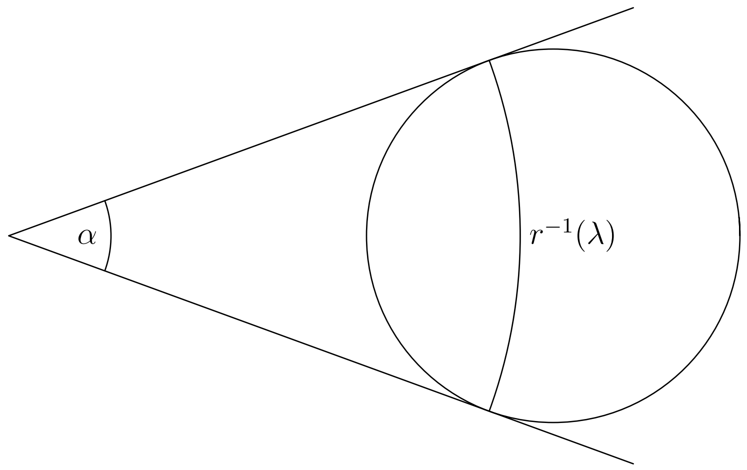

Figure 2.0.ii. This is a picture of the wedge whose dome is shown in Figure 2.0.i. The circle drawn is a circle that is maximal in the wedge. It is the boundary of the hyperbolic plane which contains the geodesic λ and is tangent to the cone surface. This hyperbolic plane is a hemisphere. The diagram 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.

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 6.1. To fix the notation, we make the following definition.

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

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.

2.3 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 Dome(Ω) 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 Dome(Ω). 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

$$S_1 \xrightarrow{g^{-1}} M \xrightarrow{h^{-1}} \Omega \xrightarrow{r} \mathrm{Dome}(\Omega) \xrightarrow{\iota} S_u. \tag{2.3.a}$$

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

$$\Psi(t, s) = (t, -u/2 + f_\alpha(1/2 + s)), \tag{2.3.b}$$

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.

3 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

$$-2(1+y)\lambda + (1 + 2y + y^2 + z^2)\mu - x\lambda^2(2 + x\mu) = 0. \tag{3.0.c}$$

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

Equation 3.0.c 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'.

3.1 Lemma. *Let f be the function of Equation 2.3.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 § 3.

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 3.0.c 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. ■

3.2 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 2.1, obtaining

$$\tan\left(\frac{f_\alpha(s)\alpha}{\sin\left(\frac{\alpha}{2}\right)}\right) = \frac{-\cos\left(\frac{\alpha}{2}\right) + \cos\left(\frac{\alpha}{2} - s\alpha\right)}{\sin\left(\frac{\alpha}{2}\right)\sin\left(\frac{\alpha}{2} - s\alpha\right)}. \tag{3.2.a}$$

This gives

$$f_\alpha(s) = s + \frac{\alpha}{2\tan(\alpha/2)}s^2 + O(s^3). \tag{3.2.b}$$

□

We recall Definition 2.3.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 2.1 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 2.1 with the next result, we will see that the nearest point retraction does not give an extremal quasiconformal homeomorphism.

3.3 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 2.2). 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 [3, 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

3.4 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 $\pi/2$.*

3.5 The general convex region $\Omega \subset \mathbb{C}$.

3.6 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 $\mathbb{U}^3$.*

3.7 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 4.1 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 3.6, we establish some general properties of the bending lamination for the dome of a convex region.

3.8 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.*

3.8.1. $x \neq \infty$.

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

3.8.3. $\mathrm{Dome}\,(\Omega)$ *has no isolated bending line and each bending line has zero transverse measure.*

3.8.4. The nearest point retraction $r : \Omega \to \mathrm{Dome}\,(\Omega)$ *is a homeomorphism.*

3.8.5. 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 3.8.i.)

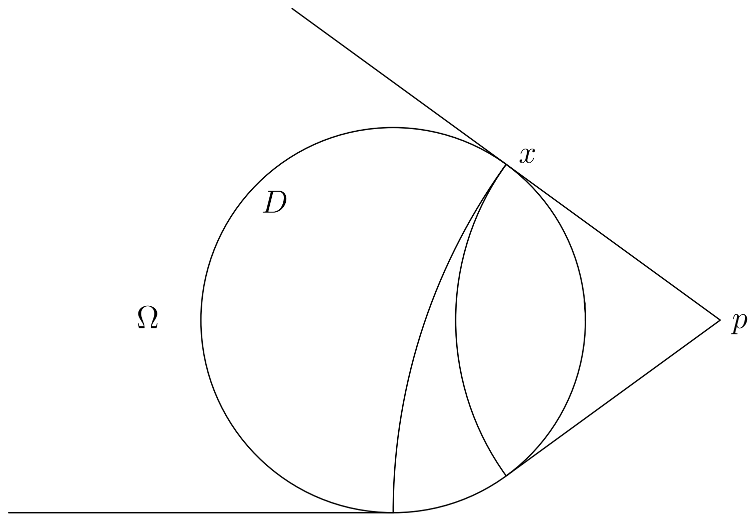

Figure 3.8.i. This illustrates 3.8.5. We show the convex set Ω and a maximal disk $D \subset \Omega$. Two bending lines for $\mathrm{Dome}\,(\Omega)$ have a common endpoint x. The bending lines are not shown. Instead we show their inverse images under the nearest point retraction.

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 3.8.5, 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 3.8.5 now follows. The situation is illustrated in Figure 3.8.i. ■

Proof (of Theorem 3.6): 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 3.8 that the Thurston isometry is C^1. ■

4 Quasiconformality of the nearest point retraction

4.1 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 4.1, we will use approximation theorems from [3]. For the main result of this chapter (Theorem 6.1), we will rely heavily on Theorem 4.2, used in the current paper only with the additional (and unnecessary) restriction that Ω and Ω_n are all convex.

4.2 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 4.1, we will not rely on Theorem 4.2, but only on the more easily proved Theorem 4.3. This result, not stated explicitly in [3], is implicitly contained in the proof of [3, 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 4.3 from Theorem 4.2, which is explicitly stated and proved in [3]. Although we need only the convex case, the more general statement given in Theorem 4.3 is true for any Ω and Ω_n.

4.3 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 [3, §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 4.1): 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 3.8.2, 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 4.3.i.

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 $\mathrm{Dome}\,(\Omega)$ 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 2.1, the dilatation lies between 1 and $1 + \cos(\alpha/2) \leq 2$. ■

4.4 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. □

5 Boundary dilatation.

Now we return to the case of a general open convex region in the plane.

5.1 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*

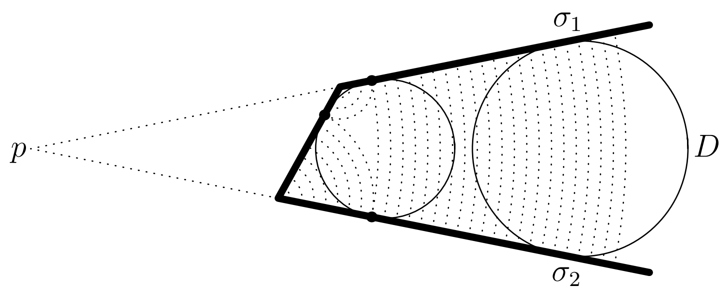

Figure 4.3.i. This illustrates the proof of Theorem 4.1. The boundary of the polygon is indicated by thick lines. The edges σ_1 and σ_2, when extended, meet at p. D denotes the maximal disk D in the proof. In the diagram we show two possible positions for D. The rest of this caption refers only to the 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.

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 5.1.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.

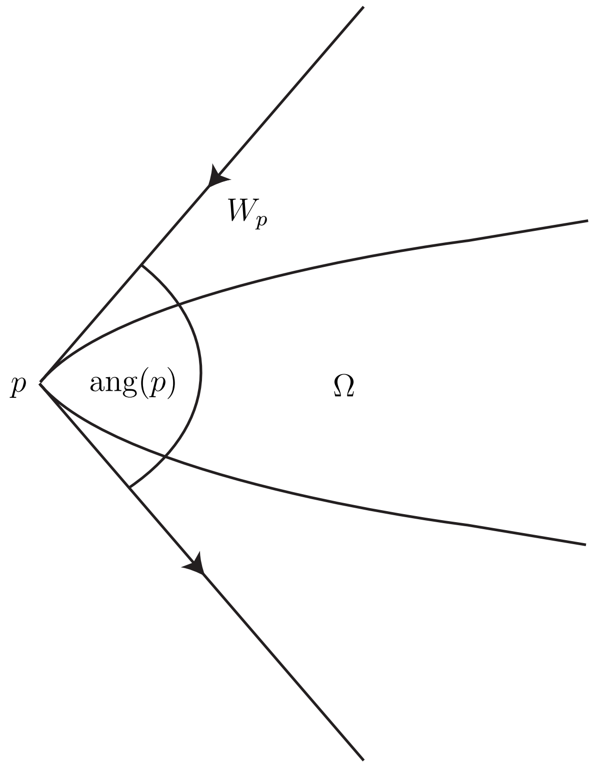

Figure 5.1.i. This illustrates Definition 5.1. W_p is the wedge contained between the two tangents to $\partial\Omega$ at p. We have $\Omega \subset W_p$.

The following lemma is obvious.

5.2 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 $(p_i)_{i\in\mathbb{N}}$.*

5.3 Definition. We extend Definition 5.1 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 5.3.i.

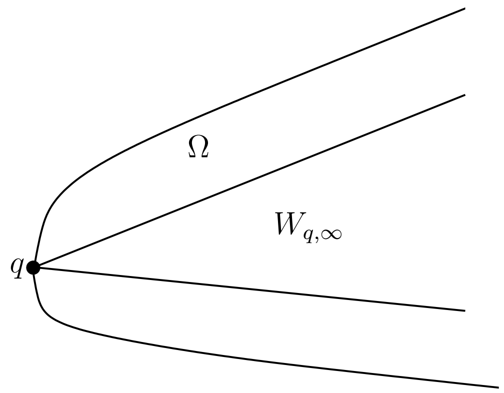

Figure 5.3.i. This illustrates Definition 5.3. We have $q \in \partial\Omega$. $W_{\infty,q} \subset \Omega$ is the wedge defined in 5.3

$W_{\infty,q} \subset \Omega$ 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 $0 \le \mathrm{ang}(\infty) \le \pi$, where 0 corresponds to the degenerate case.

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

5.4 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. □

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

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.)

5.6 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 [4, 5] that the maximum is attained for some $p \in \partial U$. □

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 [7] for an exposition of this topic.

6 The main theorem.

6.1 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 5.3).*

The proof will occupy the remainder of this paper.

6.2 Proof for special cases. 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 3.3 is much stronger than Theorem 6.1. So, for the remainder of the proof of Theorem 6.1, we assume that Ω has at most one boundary point in $\mathbb{S}^2$ equal to ∞.

Next we think about the case when $\operatorname{ang}(p) = \pi$. We need to prove that the substantial boundary dilatation at p is 1. If $p = \infty$ and $\operatorname{ang}(p) = \pi$, then, as already pointed out in Remark 5.5, Ω is a halfplane and that the nearest point retraction is conformal. This reduces consideration of $\operatorname{ang}(p) = \pi$ 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 4.1. 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 6.1, 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 6.1, we can confine our attention to points $p \in \partial\Omega$, finite or infinite, such that $0 \le \operatorname{ang}(p) < \pi$. 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.

6.3 Standard position for Ω. We recall from § 6.2 that we may assume that $\operatorname{ang}(p) < \pi$ 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 5.3). 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 6.3.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.

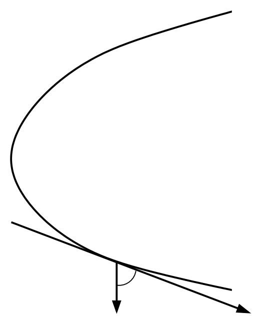

Figure 6.3.i. This illustrates the first paragraph of § 6.3.

We recall that p is the corner point defined in Theorem 6.1. 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 $\operatorname{ang}(p) > 0$. Following Definition 5.1, 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 $\log : \Omega \to \mathbb{C}$.

6.4 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 5.4, 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.

6.5 Definition. As in Definition 2.3, 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.

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 $\partial_-\Omega$. □

6.6 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 5.6). 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.)

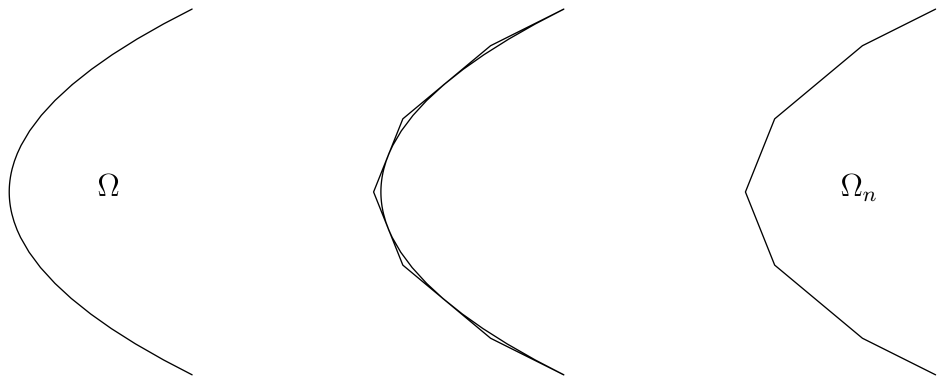

Figure 6.6.i. The lefthand figure is a convex curve, enclosing a convex region Ω. The point p is at infinity on the right. The middle picture shows Ω together with a number of support lines, as described in § 6.6. Using the support lines only, we obtain the polygonal approximation Ω_n, as shown in the righthand picture.

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.

6.6.1. $\partial\beta_0 \subset X_n$.

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

6.6.3. X_n is locally finite.

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

6.6.5. 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.

6.6.6. 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 $i \ge 1$, the geodesic $x_i x_{i+1}$ is a bending line in $\mathrm{Dome}\,(\Omega)$ joining $\partial_-\Omega$ to $\partial_+\Omega$. By 3.8.3, such a chain must satisfy $k \le 3$.

6.6.7. 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 $0 < \mathrm{ang}(p) < \pi$, $h = \log$ if $p = \infty$ and $0 < \mathrm{ang}(p) < \pi$, and $h = \mathrm{Id}$ if $p = \infty$ and $\mathrm{ang}(p) = 0$.

6.7 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 § 6.5. 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 ∂_-M_n and ∂_+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 ∂_-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 $\sigma_{x,n}$.

6.8 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 6.8 continues to be true in this greater generality.

Proof: We apply Caratheodory's Kernel Theorem (see [6, 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 § 6.5 and § 6.7 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.

6.9 Corollary. *We use the notation of Theorem 6.8. 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 $(n(j))_{j\in\mathbb{N}}$.*

7 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 7.0.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 5.2, 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 $\mathrm{ang}(p) = 0$. (Recall that $\mathrm{ang}(p) < \pi$, 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 § 7, joining $\partial_-\Omega$ (or $\partial_+\Omega$) to itself, to the euclidean size of a nearby almost vertical interval joining ∂_-M to ∂_+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.

7.1 Definition. We call a bending line like λ in § 7, which joins $\partial_-\Omega$ (or $\partial_+\Omega$) to itself, a *fringe*. We call a bending line like λ_0 in § 7, 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.

7.2 Lemma. *Let $\lambda \subset \operatorname{Dome}(\Omega)$ 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 $\operatorname{Dome}(\Omega)$. By Lemma 3.8, 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

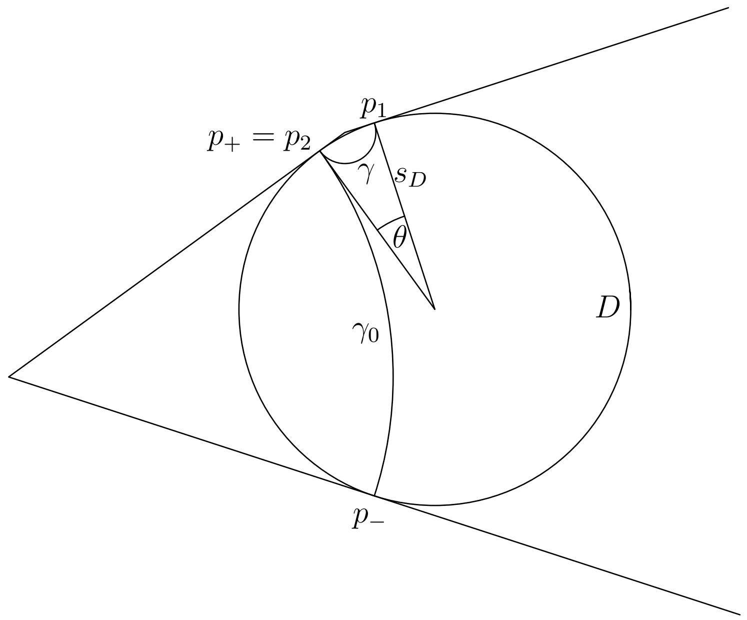

Figure 7.0.i. This picture illustrates the discussion in § 7. We have taken a case where, among all possible shapes of Ω and among all possible configurations within Ω, D happens to meet $\partial_-\Omega$: this makes s_D as large as possible consistent with our other assumptions. We have $\gamma = r^{-1}(\lambda)$ 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 $s_0 = s_D$.

$p_1, p_2 \in \partial D$. 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.

7.3 Definition. In parallel with Definition 2.3.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 $+\infty$. □

7.4 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 § 6.6 how Ω_n is constructed.

We now discuss the three possibilities for u.

7.4.1. First, u may lie in a flat of $\mathrm{Dome}\,(\Omega_n)$. In this case the dilatation of r_n is one.

7.4.2. 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 7.1). We apply Lemma 2.1 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 3.1 or Remark 3.2. Here our region is not a genuine halfcone or halfcylinder, but only an approximation to such. However, Theorem 6.8 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.

7.4.3. 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 2.1 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 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 $\overline{\mathrm{Dome}(\Omega_n)}$. 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 8.0.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.

8.1 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 $d(b_1, b_2)$. ∎

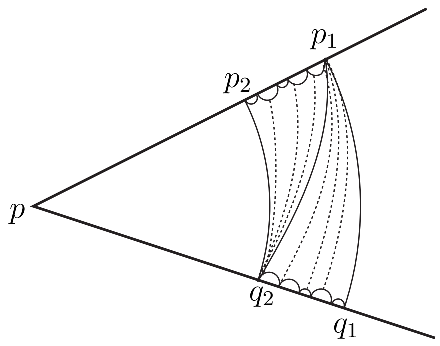

Figure 8.0.i. This illustrates § 8, where the bending lamination is extended to a foliation of a neighbourhood of p by geodesics. The picture should really be drawn on $\mathrm{Dome}\,(\Omega)$, but we have instead drawn the inverse image in Ω under the nearest point retraction r, which is easier to 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.

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

8.2 Lemma. *Let $Y_{m,n} = r_n\left(h^{-1}\left(g_n^{-1}\left([m,\infty) \times (0,1)\right)\right)\right)$. 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 $Y_{m,n} \subset$ Dome (Ω_n) is bounded by ϵ_m. $\Gamma_n \cap Y_{m,n}$, parametrized by hyperbolic pathlength, is a $(1,\epsilon_m)$-quasigeodesic.*

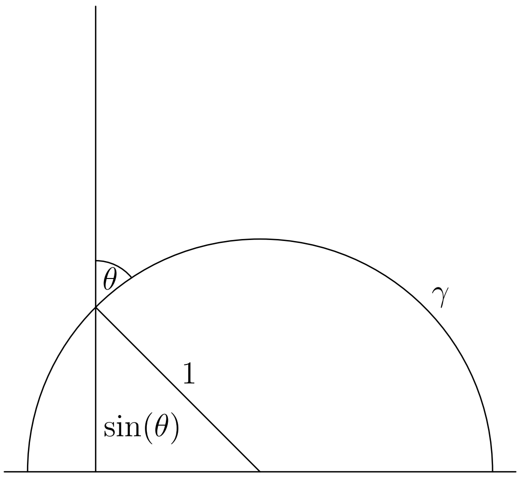

Figure 8.1.i. This diagram illustrates the proof of Lemma 8.1. We show the geodesic γ of that proof as a semicircle of radius 1 in the upper halfplane. We work out the relationship between the signed distance t measured along γ, and the angle that a vertical geodesic makes with γ. Since the point of 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 8.0.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 6.9 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 2.3). We recall the definition of $g_n : M_n \to S_1$ from § 6.7 and of $\psi_n : S_1 \to S_u$ from Definition 7.3 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.

8.3 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 6.9 and Lemma 8.2. 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. ∎

8.4 Example. It is tempting to think that a central line in M_n could be added to Proposition 8.3 so as to give three asymptotic curves, rather than two. This can indeed be done if $\text{ang}(p) > 0$. But if $p = \infty$ and $\text{ang}(p) = 0$, 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. □

8.5 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 $\text{ang}(p) > 0$ 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 6.8, 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 8.2 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 § 7—occur within a distance $\delta/4$ of ∂S_1 if m is large enough.)

9 A key proposition.

Now we come to one of the main intermediate results required in the proof of Theorem 6.1.

9.1 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 2.3 and Equation 2.3.b.) The convergence is uniform on the compact subset and with respect to n. (See the statement of Theorem 6.8 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 $z_{n,m} = h^{-1}\left(g_n^{-1}(m)\right) =$

$\exp(-g_n^{-1}(m)) \in \Omega_n$. 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 9.1 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 2.1), given by $z \mapsto \exp(-\,\mathrm{ang}(p)z)$. To see this, note first that from Corollary 6.9, Theorem 6.8 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 4.3, namely it is a nearest point retraction, followed by an isometry $\mathrm{Dome}\,(\Omega_{n,m}) \to S_u$. In order to apply Theorem 4.3, 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 8.3, 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 8.3, this geodesic is sent to a curve in S_u which is asymptotic to $y = 0$. Theorem 4.3 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 9.1 now follows if p is finite.

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

When $p = \infty$ and $\operatorname{ang}(p) = 0$, 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. § 6.7. We define $\Omega_{n,m} = \sigma_{n,m}^{-1}(\Omega_n)$. As in the case when $p = 0$ and $\operatorname{ang}(p) > 0$, we factorize the map $S_1 \to S_u$ in the statement of Proposition 9.1 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 8.3 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 4.3, this converges to the model map.

This proves Proposition 9.1 when $p = \infty$ and $\operatorname{ang}(p) = 0$. ■

9.2 The derivative of ψ_n. The next thing we have to do is to estimate the derivative of the map $\psi_n : S_1 \to S_u$ (see Definition 7.3). Since the derivative of translation is the identity, this is the same as estimating the derivative of the composite of Proposition 9.1. 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.

9.3 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:*

9.3.1. 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.)

9.3.2. In 9.3.1, 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 § 8.5, 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 4.1), 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 9.3.1.

To prove 9.3.2, we examine further the composite $\tau_{-m'} \circ \psi_n \circ \tau_m : S_1 \to S_u$ introduced in Proposition 9.1. 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 6.8, 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 2.3.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 7.4 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 9.3.2. ∎

9.4 A quasiconformal map $S_1 \to S_u$. 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 6.1. 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 $x = M_0$.

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

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

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

9.4.4. 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 9.5, 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 9.3, 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 9.4.4. 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 $1/k \leq X_V(\rho(a), \rho(c); \rho(b), \rho(x_0)) \leq k$ for all choices of a, b, c in ∂U giving cross-ratio -1 as already explained.

In order to fill in the gap at 9.4.4, 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 $\left(M_0 + 1, \frac{1}{2} - \delta\right) \in S_1$.

9.5 Lemma. *There exists $k \geq 1$, $\delta_0 > 0$ and $m_0 \in \mathbb{Z}$, such that, for all n, for all $0 < \delta \leq \delta_0$ and for all $M_0 \geq m_0$, 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.

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

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

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

9.5.4. 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.

9.5.5. 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 $M_0 = M_i$.

9.5.6. 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.

9.5.7. 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.

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

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

A theorem of Rado (see [6, Theorem 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 9.3, 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 9.5 is complete. ∎

10 Completion of Theorem 6.1

.

Proof: By Lemma 9.5, 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 9.5 could be used to give an explicit of k in Lemma 9.5, and therefore of k_1, in terms of $\operatorname{ang}(p)$.)

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 $\Omega_n \to \operatorname{Dome}(\Omega_n)$ which extend continuously to the identity map on $\partial\Omega_n$. By Theorem 4.2, ρ_δ corresponds to a quasiconformal homeomorphism $\Omega \to \operatorname{Dome}(\Omega)$ 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 6.1. ∎

10.1 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. $\square$

References

[1] 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.

[2] 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`.

[3] 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`.

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

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

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

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

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