Hyperbolic geometry in dimensions 2 and 3

1 Conformal metrics in the plane

Let $\alpha\colon [a,b] \to \mathbb{R}^2$ be a piecewise smooth arc. We can define the *arc length* of α by

$$L(\alpha) = \int_a^b \|\alpha'(t)\| dt$$

where $\|v\|^2 = \langle v, v\rangle$ is the usual Euclidean length of a vector in $\mathbb{R}^n$. The are more general ways of defining the length of such a path, the most common generalization being through a Riemannian metric. Rather, than defining this we will stick to the simple notion of a conformal metric.

Let $U \subset \mathbb{R}^2$ be open. Then a *conformal metric* on U is simply a non-negative function $\rho\colon U \to \mathbb{R}$. In most cases ρ will be smooth and strictly positive although there will be a few places where we will want a more general class of functions. At each point $x \in U$ we use ρ to define new inner product on $\mathbb{R}^2$ by setting $\langle , \rangle_{\rho(x)} = \rho(x)^2\langle , \rangle$. The $\rho(x)$-length of a vector is simply $\|v\|_{\rho(x)} = \rho(x)\|v\|$.

If ρ is constant then we are just rescaling the usual metric. Things are more interesting when ρ varies. We then define the ρ-length of the arc α (assuming its image lies in U) by

$$L_\rho(\alpha) = \int_a^b \|\alpha'(t)\|_{\rho(\alpha(t))} dt.$$

Note that the length of the vector $\alpha'(t)$ depends on the basepoint of the vector $\alpha(t)$. We will usually suppress this dependence on basepoint in our notation. That is instead of writing $\langle , \rangle_{\rho(x)}$ we will write $\langle , \rangle_\rho$, etc. Strictly speaking $\langle , \rangle_\rho$ is not an inner product on $\mathbb{R}^2$ but a map from U to the space of inner products on $\mathbb{R}^2$.

For $i = 1, 2$ let U_i be a domain with a conformal metric ρ_i. Then a diffeomorphism $f\colon U_1 \to U_2$ is an isometry if for every arc α_1 in U_1 we have $L_{\rho_1}(\alpha_1) = L_{\rho_2}(f \circ \alpha_1 = \alpha_2)$. Clearly this condition holds if $\|\alpha_1'(t)\|_{\rho_1} = \|\alpha_2'(t)\|_{\rho_2}$ and if we make some assumptions on the ρ_i (for example the ρ_i are continuous) then it is also necessary. For this equality to hold we need that the derivative at $x \in U_1$ is an isometry from the inner product $\langle , \rangle_{\rho_1(x)}$ to the inner product $\langle , \rangle_{\rho_2(f(x))}$. That is we want

$$\langle v, w\rangle_{\rho_1(x)} = \langle f_*(x)v, f_*(x)w\rangle_{\rho_2(f(x))}.$$

1.1 $\mathbb{R}^2$ v $\mathbb{C}$

This last condition is easiest to characterize if we view $\mathbb{R}^2$ as the complex plane $\mathbb{C}$. There is an obvious map ϕ from $\mathbb{R}^2$ to $\mathbb{C}$ given by $\phi(x,y) = x + iy$. Given an $\mathbb{R}$-linear map $T\colon \mathbb{R}^2 \to \mathbb{R}^2$ there exists $T_x, T_y \in \mathbb{C}$ such that $\phi \circ T(x,y) = T_x x + T_y y$. The real parts of T_x and T_y are the first row of T (in the usual basis) while the imaginary part is the second row. If we also pre-compose with ϕ^{-1} we get a map from $\mathbb{C}$ to $\mathbb{C}$. Here we have $T_z, T_{\overline{z}} \in \mathbb{C}$ such that $\phi \circ T \circ \phi^{-1}(z) = T_z z + T_{\overline{z}}\overline{z}$. A calculation shows that

$$T_z = \frac{1}{2}\left(T_x - iT_y\right) \text{ and } T_{\overline{z}} = \frac{1}{2}\left(T_x + iT_y\right).$$

From here on out we will suppress the map ϕ. The map T is $\mathbb{C}$-linear if and only if $T_{\overline{z}} = 0$.

Given positive real numbers λ_1 and λ_2 we have the inner products $\langle , \rangle_{\lambda_i}$. The $\mathbb{R}$-linear map T is an isometry exactly if $T_{\overline{z}} = 0$ and $|T_z| = \lambda_1/\lambda_2$. We now apply this to the diffeomorphism $f\colon U_1 \to U_2$ where we have conformal metrics ρ_i on U_i. At each point $z \in U_1$ the derivative $f_*(z)$ is an $\mathbb{R}$-linear map.

Let

$$f_z = \frac{\partial f}{\partial z} = \frac{1}{2}\left(\frac{\partial f}{\partial x} - \imath\frac{\partial f}{\partial y}\right)$$

and

$$f_{\overline{z}} = \frac{\partial f}{\partial \overline{z}} = \frac{1}{2}\left(\frac{\partial f}{\partial x} + \imath\frac{\partial f}{\partial y}\right).$$

The $f_z(z)$ is the $\mathbb{C}$-linear part of $f_*(z)$ and $f_{\overline{z}}(z)$ is the $\mathbb{C}$-anti-linear part.

Let $S\colon \mathbb{C} \to \mathbb{C}$ be another $\mathbb{R}$-linear map. For the chain rule it is useful to have a formula for $(S \circ T)_z$ and $(S \circ T)_{\overline{z}}$. We see that

$$\begin{aligned} S \circ T(z) &= S(T_z z + T_{\overline{z}}\overline{z}) \\ &= S_z(T_z z + T_{\overline{z}}\overline{z}) + S_{\overline{z}}(\overline{T_z z + T_{\overline{z}}\overline{z}}) \\ &= (S_z T_z + S_{\overline{z}} T_{\overline{z}})z + (S_z T_{\overline{z}} + S_{\overline{z}} T_z)\overline{z} \end{aligned}$$

and therefore $(S \circ T)_z = S_z T_z + S_{\overline{z}} T_{\overline{z}}$ and $(S \circ T)_{\overline{z}} = S_z T_{\overline{z}} + S_{\overline{z}} T_z$.

1.2 The hyperbolic plane

Let $\mathbb{U} = \{z \in \mathbb{C} | \operatorname{Im} z > 0\}$ be the upper half plane and let $\rho_{\mathbb{U}}(z) = 1/\operatorname{Im} z$ be a conformal metric on $\mathbb{U}$. The pair $(\mathbb{U}, \rho_{\mathbb{U}})$ is a model for the hyperbolic plane, $\mathbb{H}^2$. We begin by studying isometries.

From complex analysis we know that conformal automorphisms of $\mathbb{U}$ are Möbius transformations with real coefficients. This is exactly the group $\operatorname{PSL}_2 \mathbb{R}$. This is also the group of orientation preserving isometries of $\mathbb{H}^2$.

Proposition 1.1 $\mathrm{PSL}_2\,\mathbb{R}$ *is the group of isometries of* $(\mathbb{U}, \rho_{\mathbb{U}})$*.*

Proof. Any orientation preserving isometry $\phi\colon \mathbb{H}^2 \to \mathbb{H}^2$ will be a holomorphic map. By Schwarz reflection, ϕ extends to a conformal automorphism of the Riemann sphere $\widehat{\mathbb{C}}$ and hence must be a Möbius transformation. The coefficients of this Möbius transformation need not be real. However they are only determined up to multiplication by elements of $\mathbb{C}^\times = \mathbb{C}\backslash\{0\}$ and by using that the image of 0, 1 and ∞ are real we can see that the coefficients can be chosen to be real so $\phi \in \mathrm{PSL}_2\,\mathbb{R}$.

Now assume that $\phi(z) = \frac{az+b}{cz+d}$ is in $\mathrm{PSL}_2\,\mathbb{R}$. Note that we can scale the coefficients so that $ad - bc = 1$ (which uniquely determines the coefficients up to sign). With this normalization the derivative of ϕ is $\phi_z(z) = \frac{1}{(cz+d)^2}$. For ϕ to be an isometry we need $|\phi_z(z)|\rho_{\mathbb{U}}(\phi(z)) = \rho_{\mathbb{U}}(z)$. We leave this computation to the reader. $\boxed{1.1}$

1.2.1 Geodesics

A path $\alpha\colon I \to \mathbb{U}$ is a geodesic if for all $s, t \in I$ and all piecewise smooth paths $\beta\colon [a,b] \to \mathbb{U}$ with $\beta(a) = \alpha(s)$ and $\beta(b) = \alpha(t)$ we have

- $L_{\rho_{\mathbb{U}}}(\alpha|_{[s,t]}) = t - s$;
- $L_{\rho_{\mathbb{U}}}(\beta) \geq t - s$

We will first show that vertical lines, suitable parameterized are geodesics.

Lemma 1.2 *Define* $r\colon \mathbb{U} \to \mathbb{U}$ *by* $r(z) = \imath|z|$*. Then for all vectors* v *we have* $\|r_*(z)v\|_{\rho_{\mathbb{U}}} \leq \sin(\arg z)\|v\|_{\rho_{\mathbb{U}}}$.

Proof. To calculate $r_*(z)$ we first take the z and $\overline{z}$ derivatives. We have $r_z(z) = \frac{\imath}{2}\sqrt{\frac{\overline{z}}{z}} = \frac{\imath e^{-\imath \arg z}}{2}$ and $r_{\overline{z}} = \frac{\imath}{2}\sqrt{\frac{z}{\overline{z}}} = \frac{\imath e^{\imath \arg z}}{2}$. We can identify vectors with complex numbers and we then have that

$$\begin{aligned} r_*(z)v &= r_z(z)v + r_{\overline{z}}(z)\bar{v} \\ &= \frac{\imath}{2}\left(e^{-\imath \arg z}v + \overline{e^{-\imath \arg z}v}\right) \\ &= \imath\,\mathrm{Re}(e^{-\imath \arg z}v) \end{aligned}$$

and therefore $\|r_*(z)v\| \leq \|v\|$ and $\|r_*(z)v\|_{\rho_{\mathbb{U}}} \leq (\rho_{\mathbb{U}}(r(z))/\rho_{\mathbb{U}}(z))\|v\|_{\rho_{\mathbb{U}}}$. Since

$$\begin{aligned} \rho_{\mathbb{U}}(r(z))/\rho_{\mathbb{U}}(z) &= \mathrm{Im}\, z/|z| \\ &= |z|\sin(\arg z)/|z| \\ &= \sin(\arg z) \end{aligned}$$

the lemma follows. $\boxed{1.2}$

Theorem 1.3 *The curve $\alpha\colon \mathbb{R} \to \mathbb{U}$ given by $\alpha(t) = \imath e^t$ is a geodesic.*

Proof. Note that $L_{\rho_{\mathbb{U}}}(\alpha|_{[s,t]}) = t - s$ so the first condition holds. Assume that $\beta\colon [a,b] \to \mathbb{U}$ is a piecewise smooth path with $\beta(a) = \imath e^s$ and $\beta(b) = \imath e^t$. If there is some point in the image of β that doesn't lie on the imaginary axis then by Lemma 1.2 $L_{\rho_{\mathbb{U}}}(r \circ \beta) < L_{\rho_{\mathbb{U}}}(\beta)$ so we can assume that the image of β lies on the imaginary axis. Furthermore we can assume that β is injective for it is not then we can replace it with an injective path that is at most as long as the original path. Then β is a just reparameterization of $\alpha|_{[a,b]}$ and will have the same length. $\boxed{\text{1.3}}$

Corollary 1.4 *Any vertical line or semi-circle in $\mathbb{U}$ that is orthogonal to the real axis is the image geodesic.*

Corollary 1.5 *There is a unique geodesic between any two points in $(\mathbb{U}, \rho_{\mathbb{U}})$ and any pair of geodesics will intersect at most once.*

1.2.2 Area

Let (U, ρ) be a conformal metric on a domain $U \subset \mathbb{R}^2 = \mathbb{C}$. If $V \subset U$ is open then $\mathrm{Area}_\rho(V) = \int\int_V \rho^2 dxdy$.

Proposition 1.6 *If $f\colon (U_1, \rho_1) \to (U_2, \rho_2)$ is an isometry and $V_1 \subset U_1$ is open then $\mathrm{Area}_{\rho_1}(V_1) = \mathrm{Area}_{\rho_2}(f(V_1))$.*

1.2.3 Triangles

A geodesic in $(\mathbb{U}, \rho_{\mathbb{U}})$ is uniquely determined by pairs of distinct points in $\mathbb{R} \cup \{\infty\}$. Three distinct points in $\mathbb{R} \cup \{\infty\}$ determine three distinct geodesics and the region they bound is an *ideal triangle*.

Proposition 1.7 *Any two ideal triangles are isometric.*

Proof. Given oriented triples $\{x_0, x_1, x_2\}$ and $\{y_0, y_1, y_2\}$ there is a (unique) $\phi \in \mathrm{PSL}_2\,\mathbb{R}$ such that $\phi(x_i) = y_i$. This ϕ is an isometry between the corresponding ideal triangles. $\boxed{\text{1.7}}$

More generally a triple of points in $\mathbb{U} \cup (\mathbb{R} \cup \infty)$ determines a triangle. Up to isometry each such triangle is determined by its angles.

For each triangle T we can naturally associate an ideal triangle. First induce an orientation on the boundary from the orientation of the triangle. Then extend each side of the triangle in the positive direction to an infinite geodesic ray. These three rays will limit to a triple of distinct points in $\mathbb{R} \cup \infty$ which determines an ideal triangle $I(T)$. Ideal vertices of T will be ideal vertices of $I(T)$. In particular, if T is an ideal triangle the $T = I(T)$.

Proposition 1.8 *A triangle T_θ with two ideal vertices is determined up to isometry by the angle θ of the non-ideal vertex. We also have* $\text{Area}_{\rho_{\mathbb{U}}}(T) = \pi - \theta$.

Proof. We can assume that $I(T_\theta)$ has vertices $\{1, -1, \infty\}$ with 1 and ∞ being vertices of T_θ. The non-ideal vertex lies on the semi-circle of radius one centered at $0 \in \mathbb{C}$. Two of the sides of T_θ will also be sides of $I(T_\theta)$ and the third side will be a vertical geodesic from the the non-ideal vertex to ∞. This non-ideal vertex will have x-coordinate $\cos\alpha$ for some angle α and the angle of the triangle at the non-ideal vertex will be $\theta = \pi - \alpha$. Therefore T_θ is determined by the x-coordinate of its non-ideal vertex which is determined by θ.

The area of T_θ is a double integral. We calculate

$$\begin{aligned}
\text{Area}_{\rho_{\mathbb{U}}}(T_\theta) &= \int_{\cos(\pi-\theta)}^{1} \int_{\sqrt{1-x^2}}^{\infty} \frac{1}{y^2} dy dx \\
&= \int_{\cos(\pi-\theta)}^{1} \frac{1}{\sqrt{1-x^2}} dx \\
&= \pi - \theta.
\end{aligned}$$

The first integral is straightforward. The second integral is also not difficult but perhaps the easiest way to calculate it is to take the derivative of the expression in the second line as a function of θ and see that this derivative is 1. $\boxed{1.8}$

We can now prove the Gauss-Bonnet Theorem for hyperbolic triangles.

Theorem 1.9 *Let T be a hyperbolic triangle with angles θ_0, θ_1 and θ_2. Then* $\text{Area}_{\rho_{\mathbb{U}}}(T) = \pi - \sum \theta_i$.

Proof. The three rays that we used to define $I(T)$ divide $I(T)$ into 4 triangles with disjoint interior. One of these triangles will be T. Three of these triangles will have two ideal vertices. The non-ideal vertices will have angle that is complementary to one of the vertices of T so the area of these triangles will be $\pi - (\pi - \theta_i) = \theta_i$. To get the area of T we take the area of $I(T)$ (which is π) and subtract from it the are of the three other triangles. The theorem follows. $\boxed{1.9}$

Theorem 1.10 *Let T_t be a one parameter family of hyperbolic triangles with angle sum θ_t. Then $\dot{A} = -\dot{\theta}$.*

This is the two-dimensional version of the *Schlafli formula* and it is trivial consequence of the Gauss-Bonnet formula for triangles. In higher dimensions one replaces area with volume and on the right one takes the product of the dihedral angle and volume of co-dimension one faces and sums over all such faces.

Theorem 1.11 *If P is an n-sided polygon then* $\text{Area}_{\rho_{\mathbb{U}}} = \pi(n-2) - \sum \theta_i$.

This follows from dividing the polygon into triangles and applying Theorem 1.9.

Proposition 1.12 *Given non-negative $\theta_0, \theta_1, \theta_2$ with $\sum \theta_i < \pi$ there exists a unique hyperbolic triangle (up to isometry) with the given angles.*

Proof. If $\theta_i = 0$ then the corresponding vertex v_i will be ideal so if two or more of the angles are zero then the proposition follows from Lemma 1.7 and Proposition 1.8. It what follows we will assume that there is at most one ideal vertex.

Fix the vertex v_0 to be the point $\imath \in \mathbb{U}$. Fix a ray r_0 with endpoint v_0 that has ideal endpoint on the positive real axis and such that the angle from r_0 to the imaginary axis in the counter-clockwise is θ_0. For each $y \geq 1$ we similarly fix a ray r_2^y except the endpoint will be $\imath y$ and the angle will be $\pi - \theta_2$. Let $p_0 \in \mathbb{R}^+$ be the ideal endpoint of r_0 and $p_2^y \in \mathbb{R}$ the ideal endpoint of r_2^y. Since $\theta_0 + \theta_2 < \pi$ we have $\theta_0 < \pi - \theta_2$ and therefore $p_2^1 < p_0$. Furthermore if $y < y'$ then $p_2^y < p_2^{y'}$ and $p_2^y \to \infty$ as $y \to \infty$. The map $y \mapsto p_2^y$ is continuous so it follows that this is a homeomorphism from $[1, \infty)$ to $[p_2^1, \infty)$ and there is a unique $y_0 \in (1, \infty)$ such that $p_2^{y_0} = p_0$. For each $y \in (p_2^1, p_2^{y_0}]$ the imaginary axis and the rays r_0 and r_2^y form a triangle T_y. Two of the angles of this triangle will be θ_0 and θ_2. Let $\theta(y)$ be the third angle. This is the angle between r_0 and r_2^y.

At this point one could do some sort of trigonometric computation to see that θ is a bijection from $(1, y_0]$ to $[0, \pi - (\theta_0 + \theta_2))$ which would prove the proposition. The tricky point is to see that as $y \to 1$ that $\theta(y) \to \pi - (\theta_0 + \theta_2)$. Instead of proving this directly we will use Theorem 1.9. Let $A(y)$ be the area of T_y. If $y < y'$ then T_y is proper subset of $T_{y'}$ and $A(y) < A(y')$. By Theorem 1.9 this implies that $y \mapsto \theta(y) = A(y) + \pi - (\theta_0 + \theta_2)$ is injective. We also have $A(y) \to 0$ as $y \to 1$ so $\theta(y) \to \pi - (\theta_0 + \theta_2)$ and therefore θ is the desired bijection. $\boxed{1.12}$

1.3 The disk model

Let $U_1, U_2 \subset \mathbb{C}$ and $f : U_1 \to U_2$ a conformal diffeomorphism. A conformal metric ρ_2 on U_2 can pull backed to a conformal metric on U_1 by $\rho_1(z) = (f^*\rho_2)(z) = |f_z(z)|\rho_2(f(z))$. Then f is an isometry from (U_1, ρ_1) to (U_2, ρ_2).

Let $\Delta = \{z \in \mathbb{C} | |z| < 1\}$ be the open unit disk. Define $f\colon \Delta \to \mathbb{U}$ by $f(z) = \frac{-\imath z+1}{z-\imath}$. We want to calculate the pull-back $\rho_\Delta = f^*\rho_{\mathbb{U}}$. Since $f_z(z) = \frac{-2}{(z-\imath)^2}$ we have $|f_z(z)| = \frac{2}{|z-i|^2}$. We also need to $\operatorname{Im} f(z)$. For this we have

$$
\begin{aligned}
2\imath \operatorname{Im} f(z) &= f(z) - \overline{f(z)} \\
&= \frac{(-\imath z+1)(\overline{z}+\imath) - (\imath\overline{z}+1)(z-\imath)}{|z-\imath|^2} \\
&= \frac{2\imath(1-|z|^2)}{|z-\imath|^2}
\end{aligned}
$$

so

$$\rho_\Delta(z) = \frac{2}{|z-\imath|^2}\frac{|z-\imath|^2}{1-|z|^2} = \frac{2}{1-|z|^2}.$$

The pair (Δ, ρ_Δ) is the Poincare Disk Model for the hyperbolic plane. The disk and the upper half plane are the two models we will generally use for the hyperbolic plane. For doing calculations each has it own uses. We denote $(\mathbb{H}^2, \rho_{\mathbb{H}^2})$ as the hyperbolic plane without reference to a specific model. We also let $d_{\mathbb{H}^2}(z_0, z_1)$ be the minimal length of path between the points $z_0, z_1 \in \mathbb{H}^2$. Note that such a minimal length path exists since there is a geodesic between any two points in $\mathbb{H}^2$. It is easy to check that $(\mathbb{H}^2, d_{\mathbb{H}^2})$ is a metric space.

The (orientation preserving) isometries of the disk model are Möbius transformations that preserve Δ. This group is, of course, isomorphic to $\operatorname{PSL}_2\mathbb{R}$ and the two groups are conjugate subgroups in $\operatorname{PSL}_2(\mathbb{C})$, the group of all Möbius transformations.

1.3.1 Classification of isometries

Given $\phi \in \operatorname{PSL}_2\mathbb{R}$ we define the *translation length* of ϕ by

$$\tau(\phi) = \inf\{d_{\mathbb{H}^2}(x, \phi(x)) | x \in \mathbb{H}^2\}.$$

We use τ to classify isometries:

- ϕ is *hyperbolic* if $\tau(\phi) > 0$;
- ϕ is *elliptic* if $\tau(\phi) = 0$ and there exists an $x \in \mathbb{H}^2$ with $\phi(x) = x$;
- ϕ is *parabolic* if $\tau(\phi) = 0$ and $d_{\mathbb{H}^2}(x, \phi(x)) > 0$ for all $x \in \mathbb{H}^2$.

The trace of an element in $\operatorname{PSL}_2\mathbb{R}$ is well defined up to sign. This classification can also be given in terms of the trace.

Proposition 1.13 *Let $\phi \in \operatorname{PSL}_2\mathbb{R}$. Then ϕ is*

1. *hyperbolic if* $|\operatorname{tr}\phi| > 2$*;*

2. *elliptic if* $|\operatorname{tr}\phi| < 2$*;*

3. *parabolic if* $|\operatorname{tr}\phi| = 2$.

Proof. As a Möbius transformation $\phi(z) = \frac{az+b}{cz+d}$ with $a, b, c, d \in \mathbb{R}$ and $ad - bc = 1$. For the moment assume that $c \neq 0$. To find the fixed points of ϕ we set $\phi(z) = z$ and see that this equation is equivalent to finding zeros of the quadratic polynomial $-cz^2 + (a-d)z + b$. The discriminant of this polynomial is $(a-d)^2 + 4bc = |\operatorname{tr}\phi|^2 - 4$. When $|\operatorname{tr}\phi| < 2$, the discriminant is negative and the quadratic polynomial with real coefficients will have two complex zeros that are conjugate to each other. In particular, ϕ will have exactly one fixed point in $\mathbb{U} = \mathbb{H}^2$ and is therefore elliptic.

If $|\operatorname{tr}\phi| \geq 2$ then the zeros of the polynomial will be real and hence the fixed points of ϕ will lie on the real axis. We can then conjugate ϕ (which won't change the trace) so that one this fixed points is at infinity. We then have $c = 0$ and are left to examine this case. If $|\operatorname{tr}\phi| = 2$, then $a = d = \pm 1$ and the only fixed point of ϕ is at ∞. To calculate $d(z, \phi(z))$ exactly is a bit of work so instead we find an upper bound. Note that $\operatorname{Im} z = \operatorname{Im}\phi(z)$ (since $b \in \mathbb{R}$) and therefore there is a horizontal line of Euclidean length $|b|$ between z and $\phi(z)$. The hyperbolic length of this path is $|b|/\operatorname{Im} z$ so $d(z, \phi(z)) \leq |b|/\operatorname{Im} z$. Therefore as $\operatorname{Im} z \to \infty$, we have $d(z, \phi(z)) \to 0$ and ϕ is parabolic.

We now assume that $|\operatorname{tr}\phi| > 2$. Then ϕ will have one fixed point at ∞ and another in $\mathbb{R}$. we can further conjugate ϕ so that this second fixed point is at zero. Then $b = 0$ and $\phi(z) = \lambda z$ with $\lambda = a/d$. If $\operatorname{Re} z = 0$ then $\operatorname{Re}\phi(z) = 0$ so by Theorem 1.3 we have $d(z, \phi(z)) = |\log\lambda|$. For $z \in \mathbb{U}$ that is not necessarily on the imaginary axis we use the map $r\colon \mathbb{U} \to \mathbb{U}$ from Lemma 1.2. By that lemma $d(z, \phi(z)) \geq d(r(z), r(\phi(z))$ but since r and ϕ commute we have $d(r(z), r(\phi(z))) = d(r(z), \phi(r(z))) = |\log\lambda|$. Therefore $d(z, \phi(z)) \geq |\log\lambda|$ so $\tau(\phi) = |\log\lambda|$ and ϕ is hyperbolic. $\boxed{1.13}$

Up to conjugacy, the trace essentially characterizes elements of $\mathrm{PSL}_2\,\mathbb{R}$. More explicitly if

- ϕ is hyperbolic then it is conjugate to $z \mapsto \lambda z$ with $\lambda \in (1, \infty)$;
- ϕ is elliptic then it is conjugate to $z \mapsto \frac{\cos(\theta)z + \sin(\theta)}{-\sin(\theta)z + \cos(\theta)}$ for some $\theta \in (0, 2\pi)$;
- ϕ is parabolic then it is conjugate to $z \mapsto z + 1$.

A geodesic g is an *axis* for ϕ if it is ϕ-invariant. Note that ϕ will fix the two endpoints of g so if ϕ has an axis it must be hyperbolic. It is easy to check that this is also a sufficient condition. (For example we can check this for $z \mapsto \lambda z$ in the upper half space model.) Furthermore if g is the axis for ϕ the for all $z \in g$ we have $d(z, \phi(z)) = \tau(\phi)$.

1.3.2 The Schwarz-Pick Lemma

Recall the Schwarz Lemma:

Lemma 1.14 *Let $f: \Delta \to \Delta$ be holomorphic with $f(0) = 0$. Then $|f'(z)| \leq 1$ with equality if and only if $f(z) = \lambda z$ with $|\lambda| = 1$.*

Proof. The function $g(z) = f(z)/z$ extends to a holomorphic function on Δ. Apply the maximum principal to g on $\Delta_r = \{z \in \mathbb{C} | |z| \leq r\}$ to see that $|g(z)| \leq 1/r$ for all $r \in (0, 1)$ and therefore $|f'(0)| = |g(0)| \leq 1$. If $|f'(0)| = |g(z)| = 1$ then $|g(z)|$ achieves its maximum and is therefore constant (with absolute value one) so $f(z) = zg(z) = \lambda$ with $|\lambda| = 1$. $\boxed{1.14}$

This very simple lemma is extremely powerful!

Lemma 1.15 *Let $f: \Delta \to \Delta$ be holomorphic. Then for all $z \in \Delta$ and $v \in T_z\Delta$ we have $\|v\|_{\rho_\Delta} \leq \|f_*(z)v\|_{\rho_\Delta}$ with equality if and only f is an isometry (and the restriction of a Möbius transformation).*

Proof. Let $\phi_i: \Delta \to \Delta$ be isometries such that $\phi_1(0) = z$ and $\phi_2(f(z)) = 0$. Then apply Lemma 1.14 to $\phi_2 \circ f \circ \phi_1$. $\boxed{1.15}$

Corollary 1.16 *A conformal automorphism of the disk is the restriction of a Möbius transformation.*

Proof. Let $\phi: \Delta \to \Delta$ be holomorphic, diffeomorphism. Then both ϕ and ϕ^{-1} are (not necessarily strict) contractions of the hyperbolic so ϕ must be an isometry. The corollary then follows from Lemma 1.15. $\boxed{1.16}$

2 Surfaces

Let Σ be a surface. There are various ways to add structure to Σ by fixing an atlas where the transitions maps have some extra property. We will be mostly interested in two cases: hyperbolic structures and complex structures. The *Uniformization Theorem* provides a bijection between these two types of structures but we will get to this later.

A *hyperbolic structure* on Σ is an atlas of charts to $\mathbb{H}^2$ where the transition maps are restrictions of hyperbolic isometries.

A *complex structure* on Σ is an atlas of charts to $\mathbb{C}$ where the transition maps are holomorphic. A surface with a complex structure is usually referred to as a *Riemann surface*.

If we identify $\mathbb{H}^2$ with the unit disk Δ (or the upper half plane $\mathbb{U}$) then hyperbolic isometries are holomorphic so a hyperbolic structure induces a complex structure. The difficult part of the Uniformization Theorem is to show that a complex structure induces a hyperbolic structures.

A parameterized curve on a hyperbolic surface is a *geodesic* if it is a geodesic when restricted to any chart. We will often refer to the image of a geodesics as a geodesic.

A *closed curve* on surface is the continuous image of S^1. Again we won't distinguish between the map and its image. A closed curve is *simple* if the map is also injective. A closed curve is *essential* if it cannot be homotoped to a point, or equivalently, the map on π_1 is non-trivial (and hence, since Σ is orientable, injective). We will often be interested in closed geodesics and simple closed geodesics. Note that a closed geodesic is always essential for if not it would lift to a closed geodesic in $\mathbb{H}^2$.

The hyperbolic metric on Σ always us to measure the length of piecewise smooth curves and the area of open subsets.

2.1 $\pi_1(\Sigma) = \mathbb{Z}$

This simplest example of hyperbolic surface is $\mathbb{H}^2$ itself. One step up is surfaces with $\pi_1(\Sigma) = \mathbb{Z}$. These are constructed by taking the quotient of $\mathbb{H}^2$ by a hyperbolic or parabolic isometry.

Lemma 2.1 *Let Γ_ϕ be the group generated by a hyperbolic $\phi \in \mathrm{Isom}^+(\mathbb{H}^2)$. Then the quotient $\mathbb{H}^2/\Gamma_\phi$ is a hyperbolic annulus and contains a single closed geodesic of length $\tau(\phi)$. Every simple closed curve on the surface is homotopic to a power of the closed geodesic.*

Proof. The axis of ϕ will descend to a closed geodesic on the quotient of length $\tau(\phi)$. Conversely, the pre-image of any closed geodesic on the quotient will be an axis for ϕ. Since ϕ has a unique axis there is a unique closed geodesic on the quotient. The quotient surface also deformation retracts to the closed geodesic and this deformation retract will take every closed curve to a power of the closed geodesic. $\boxed{2.1}$

A *horocycle* in $\mathbb{H}^2$ is a (Euclidean) circle in the disk model that is tangent to the circle at infinity. In the upper half space model horocycles are circles tangent to $\mathbb{R}$ or a horizontal line. There is an isometry between any two horocycles and if ϕ is a parabolic isometry then the horocycles tangent to the fixed point of ϕ are ϕ-invariant. An intrinsic characterization of a horocycle is that it is a curve with geodesic curvature equal to 1. A curve on a hyperbolic surface is a horocycle if it is a horocycle in a chart.

Lemma 2.2 *Let Γ_ϕ be the group generated by a parabolic $\phi \in \mathrm{Isom}^+(\mathbb{H}^2)$. Then the quotient $\mathbb{H}^2/\Gamma_\phi$ is a hyperbolic annulus and for every $L > 0$ there is a simple closed* *horocycle of length L. Any two of these closed horocycles are disjoint but isotopic and every closed curve on the surface is homotopic to the power of a closed horocycle.*

Proof. We can work in the upper half space model and assume that $\phi(z) = z + 1$. Then the horizontal line at height $1/L$ descends to a closed horocycle of length L on the quotient $\mathbb{H}^2/\Gamma_\phi$. As above the quotient surface deformation retracts to any of the closed horocycles so every closed curve is homotopic to a power of a horocycle. $\boxed{2.2}$

2.2 The developing map

As we did for the hyperbolic plane we can use the hyperbolic metric on a surface Σ to make the surface into a metric space. Namely, we let $d_\Sigma(z_0, z_1)$ be the infimum of the length of all paths from z_0 to z_1. We say that Σ is *complete* hyperbolic surface if (Σ, d_Σ) is complete as metric space.

Theorem 2.3 *If Σ is a complete simply connected hyperbolic surface then $\Sigma = \mathbb{H}^2$.*

We prove this by constructing a *developing map* D from Σ to $\mathbb{H}^2$. The construction of the map only requires that Σ is simply connected. We start with a basepoint $p \in \Sigma$ and a chart (U_0, ϕ_0) that contains p. Then on U we set $D|_U = \phi$. To define D at an arbitrary $p' \in \Sigma$ we take an arc $\alpha\colon [0,1] \to \Sigma$ connecting p to p'. We then choose $0 = t_0 < t_1 < \cdots < t_n = 1$ such that $\alpha([t_i, t_{i+1}])$ is contained in a chart (U_i, ϕ_i). The intersection $U_{i-1} \cap U_i$ and the transition map $\phi_{i-1} \circ \phi_i^{-1}$ is the restriction of an isometry $\gamma_i \in \mathrm{Isom}^+(\mathbb{H}^2)$. We then define $D(p') = \gamma_1 \circ \cdots \circ \gamma_n \circ \phi_n(p')$.

Two choices have been map in this definition: the choice of path and the decomposition of the path into subpaths that lie in charts. We need to see that $D(p')$ is independent of these choices. Let α_0 and α_1 be two paths from p to p' and assume that they have be partitioned into subpaths as above and let $D_0(p')$ and $D_1(p')$ be the corresponding points in $\mathbb{H}^2$. We want to show that $D_0(p') = D_1(p')$. Fix a homotopy $\alpha(s,t)$ between the two paths with $\alpha_i(t) = \alpha(i,t)$. (It is exactly here that we use that Σ is simply connected.) The map α is defined on a square which we can partition into rectangles whose α-image is contained in a chart. Such a partition is given by partitioning each of the sides of the rectangle and we label the rectangle whose whose upper right-hand corner is (s_i, t_j), R_{ij}. The rectangles have horizontal and vertical sides. We label the horizontal side with left endpoint (s_i, t_j), h_{ij} and the vertical side with lower endpoint (s_i, t_j), v_{ij}.

We have decomposed the square such that $\alpha(R_{ij})$ is contained in a chart (U_{ij}, ϕ_{ij}). The transition map $\phi_{(i-1)j} \circ \phi_{ij}^{-1}$ is the restriction of an element of $\mathrm{Isom}^+(\mathbb{H}^2)$ which we label γ_{ij}. Similarly $\gamma^{ij} \in \mathrm{Isom}^+(\mathbb{H}^2)$ is the transition map $\phi_{i(j-1)} \circ \phi_{ij}^{-1}$. we make a few observations:

- The charts $U_{ij}, U_{(i-1)j}$ and $U_{(i-1)(j-1)}$ all contain the point $\alpha(t_i, s_j) = p_{ij}$ (and therefore an open neighborhood of p_{ij}) so $\phi_{(i-1)(j-1)} \circ \phi_{ij}^{-1} = (\phi_{(i-1)(j-1)} \circ \phi_{(i-1)j}^{-1}) \circ (\phi_{(i-1)j} \circ \phi_{ij}^{-1})$ on an open neighborhood. Therefore the map $\phi_{(i-1)(j-1)} \circ \phi_{ij}^{-1}$ is the restriction of a hyperbolic isometry that is equal to $\gamma^{(i-1)j} \circ \gamma_{ij}$ on an open neighborhood p_{ij}. Similarly $\phi_{(i-1)(j-1)} \circ \phi_{ij}^{-1}$ is equal to $\gamma_{i(j-1)} \circ \gamma^{ij}$ on an open of p_{ij}. Two elements of $\mathrm{Isom}^+(\mathbb{H}^2)$ that are equal on an open neighborhood are equal so we have $\gamma^{(i-1)j} \circ \gamma_{ij} = \gamma_{i(j-1)} \circ \gamma^{ij}$.

- The maps $\phi_{(i-1)j} \circ \alpha$ (defined on $R_{(i-1)j}$) and $\gamma_{ij} \circ \phi_{ij} \circ \alpha$ (defined on R_{ij}) agree on h_{ij} and therefore together they define a continuous map on $R_{(i-1)j} \cup R_{ij}$. Similarly $\phi_{i(j-1)} \circ \alpha$ and $\gamma^{ij} \circ \phi_{ij} \circ \alpha$ define a continuous map on $R_{i(j-1)} \cup R_{ij}$.

We can define a map $F\colon [0,1]^2 \to \mathbb{H}^2$ by setting F on R_{ij} to be $\gamma^{01} \circ \cdots \circ \gamma^{0j} \circ \gamma_{1j} \circ \cdots \circ \gamma_{ij} \circ \phi_{ij} \circ \alpha$. Note that if $i = 0$ then there are no γ's with lower indices and if $j = 0$ there are no γ's with upper indices. We need to see that this map is continuous. Lexigraphically order the pairs (i,j). Assume that F is continuous on the first $(i, j-1)$ rectangles (assuming $j > 1$). We will show that F is also continuous on R_{ij}. In fact it is clear that F is continuous on each rectangle; the issue is that F may not be well defined on the edges of the rectangles. In particular we need to see that the definition of F on $R_{i(j-1)}$ and R_{ij} agree on v_{ij}. This follows from the second bullet. We also need to know that F on $R_{(i-1)j}$ and R_{ij} agree on h_{ij}. For this we use the first bullet to rewrite the definition of F on R_{ij} as $\gamma_{10} \circ \cdots \circ \gamma_{i0} \circ \gamma^{i1} \circ \cdots \circ \gamma^{ij} \circ \phi_{ij} \circ \alpha$. The continuity along h_{ij} now follows from the second bullet.

We have now finished the hard (or tedious) part of the construction. Note that α is constant along $\{0\} \times [0,1]$ and maps to p and is also constant along $\{1\} \times [0,1]$ and maps to p'. In particular $F(0,1) = F(1,1)$. To finish the proof we need to show $F(i,1) = D_i(p')$ for $i = 0,1$. Note that in the definition of $D_0(p')$ we partitioned the domain of α_0 and then chose a chart for the α_0-image of each sub-interval in the partition. We can assume that the partition of the vertical side of $[0,1] \times [0,1]$ is a refinement of the partition for α_0 and that the charts for the rectangles R_{0j} are the same as the charts for α_0. (We may have to refine the partition of the horizontal side for this to be possible.) It follows that $F(0,1) = D_0(p')$. A similar process shows that $F(1,1) = D_1(p')$. Therefore the map $D\colon \Sigma \to \mathbb{H}^2$ is well defined.

In fact D is a local isometry. For this we just observe that in a neighborhood of p', D is given by $\gamma_1 \circ \cdots \circ \gamma_n \circ \phi_n$. At this point we have only used that Σ is simply connected. In particular, developing maps exist for any simply connected hyperbolic surface. The argument applies in even more generality to general (G,X)-*structures*. (Thurston's book or notes are a good reference for this.)

Where are now left to show that D is a diffeomorphism. For this we will show that D is a covering map; as $\mathbb{H}^2$ is simply connected this will imply that D is a diffeomorphism. We will show that we can lift paths as local homeomorphism with the path lifting property is a covering map. Let $\alpha\colon [0,1] \to \mathbb{H}^2$ be a continuous path and $\tilde{\alpha}(0) \in \Sigma$ with $D(\tilde{\alpha}(0)) = \alpha(0)$. We first show that $\tilde{\alpha}$ can be extended to a lift of α on the entire interval $[0,1]$. Let $I \subset [0,1]$ be the largest connected interval where the lift is defined. Since Σ is complete and D is continuous this interval is closed. (Here we are using that for an sequence t_i converging to the right endpoint of I, $\tilde{\alpha}(t_i)$ is Cauchy.) Since D is a local homeomorphism the interval I is also open. Therefore $I = [0,1]$ and the path can be lifted and D is a covering map.

Corollary 2.4 *The universal cover of a complete hyperbolic surface is* $\mathbb{H}^2$.

Proof. Every cover of a complete hyperbolic surface is complete. (Check this!) Therefore the universal cover is complete and we can apply Theorem 2.3. $\boxed{2.4}$

The group $\mathrm{Isom}^+(\mathbb{H}^2)$ is a Lie group. (For example it can be identified with $\mathrm{PSL}_2\,\mathbb{R}$.) This give $\mathrm{Isom}^+(\mathbb{H}^2)$ a topology. A subgroup is *discrete* if it is a discrete subspace in this topology.

Corollary 2.5 *Let Σ be a complete hyperbolic surface. Then there exist a discrete subgroup $\Gamma \subset \mathrm{Isom}^+(\mathbb{H}^2)$ such that $\Sigma = \mathbb{H}^2/\Gamma$. Every non-trivial element of Γ is hyperbolic or parabolic.*

Proof. By Corollary 2.4, the universal cover of Σ is $\mathbb{H}^2$ so the deck group is a subgroup Γ of $\mathrm{Isom}^+(\mathbb{H}^2)$ and we just need to check that it is discrete. If not there is a sequence $g_i \in \Gamma$ such that $g_i \to g \in \Gamma$ and $g_i \neq g$. Fix a basepoint in $z \in \mathbb{H}^2$. Since Γ is a deck group, the orbit Γz is a discrete subset of $\mathbb{H}^2$ and distinct elements in Γ give distinct translates of z. In particular $g_i(z) \neq g(z)$. But if $g_i \to g$ in $\mathrm{Isom}^+(\mathbb{H}^2)$ then $g_i(z) \to g(z)$, contradicting the discreteness of Γz.

An element of a deck group acts without fixed points so Γ cannot contain elliptics. $\boxed{2.5}$

Corollary 2.6 *Let Σ be a complete hyperbolic surface. Then an essential closed curve γ on Σ is either homotopic to a unique closed geodesic or for all $\epsilon > 0$, γ is homotopic to a curve of length $< \epsilon$.*

Proof. On the level of π_1, the inclusion of γ in Σ is a subgroup of $\pi_1(\Sigma)$ isomorphic to $\pi_1(S^1) = \mathbb{Z}$. By Corollary 2.5 this group is generated by a hyperbolic or parabolic isometry in $\mathrm{Isom}^+(\mathbb{H}^2)$. Let Σ_γ be the corresponding cover. Then γ lifts homeomorphically to Σ_γ. If the subgroup is generated by a hyperbolic element then, by Lemma 2.1, in the cover γ is homotopic to a closed geodesic which will descend to a homotopy to a geodesic in Σ. On the other hand, if γ is homotopic to two geodesics in Σ then there will be two distinct geodesics in Σ_γ, a contradiction. If the group is generated by a parabolic then by Lemma 2.2 in the cover γ is homotopic to a curve of length $< \epsilon$. $\boxed{2.6}$

The *geometric intersection number*, $i(\alpha, \beta)$, of two closed curves α and β is the minimal intersection of all pairs of closed curves in the homotopy class of α and β.

Corollary 2.7 *Let α and β be closed curves that are homotopic to geodesics α^* and $\beta*$. Then $i(\alpha, \beta)$ is the intersection number of α^* and β^*.*

Proof. We lift the picture to the cover Σ_α. The geodesic α^* lifts homeomorphically to the unique closed geodesic in Σ_α. We abuse notation and also label this curve α^*. The key is that the algebraic intersection of a closed curve and a properly embedded arc is a homotopy invariant (as long as the homotopy is through proper arcs). The pre-image (which we denote $\hat{\beta}$) of β^* in Σ_α will be a countable collection of complete geodesics (which are properly embedded arcs). Let I be the number of points of intersection of α^* and β^* in Σ. As α^* lifts homeomorphically to Σ_α the intersection number of α^* with $\hat{\beta}$ will also be I. Any complete geodesic in Σ_α will either be disjoint from the core geodesic α^* or will intersect it exactly once. In particular there will be exactly I components of $\hat{\beta}$ that have algebraic interesction ± 1 with α^* and all other components are disjoint from α^*. The curve β is homotopic to β^* and this homotopy will lift to a homotopy (through proper arcs) between $\hat{\hat{\beta}}$ (the pre-image of β) and $\hat{\beta}$. Therefore for each of the I components of β^* that intersects α^* there is a disctinct component of $\hat{\hat{\beta}}$ that has algebraic intersection ± 1 with α^*. In particular this component will intersect α^* so the total intersection of $\hat{\hat{\beta}}$ with α^* will be at least I. As the intersection of β with α^* in Σ is equal to the intersection of $\hat{\hat{\beta}}$ with α^* in Σ_α we have that the former is at least I. $\boxed{2.7}$

If $\alpha = \beta$ then $i(\alpha, \alpha)$ is the *self-intersection* number of α. That is, it is the minimal number of points of self-intersection over all closed curves homotopic to α. In this case, the above proof shows that this is realized by the geodesic representative of α.

2.3 Ideal quadrilaterals

We start with a very simple problem; classifying ideal quadrilaterals. We will use this to give a parameterization of ideal hexagons which will in turn give a parameterization of right angled hexagons. Right angled hexagons will be our building blocks for hyperbolic structures on surfaces.

A *topological quadrilateral* Q is a closed disk with four distinguished points $V \subset \partial Q$, the *vertices* of Q. The sides of are the arcs in ∂Q between each vertex. An example of a topological quadrilateral is an ideal quadrilateral in $\bar{\mathbb{H}}^2$ where the ideal vertices are included. To define the space of ideal quadrilaterals we fix an oriented topological quadrilateral Q and let $\tilde{\mathcal{IQ}}$ be pairs (R, f) where $R \subset \bar{\mathbb{H}}^2$ is an ideal quadrilateral and $f\colon Q \to R$ is an orientation preserving homeomorphism taking vertices to vertices. This is a huge space! To make it more tractable we say that two pairs (R_0, f_0) and (R_1, f_1) are equivalent if there exists an isometry ϕ taking R_0 to R_1 and $\phi \circ f_0$ is homotopic rel vertices to f_1 and $\mathcal{IQ}$ be the quotient space of *marked* ideal quadrilaterals.

We will parameterize $\mathcal{IQ}$ via the cross ratio. For any collection $z_0, z_1, z_2 \in \widehat{\mathbb{C}}$ of three distinct points three is a unique $\phi \in \mathrm{PSL}_2\,\mathbb{C}$ with $\phi(z_0) = 0$, $\phi(z_1) = 1$ and $\phi(z_2) = \infty$. If we have a fourth point $z_3 \in \widehat{\mathbb{C}}$ then the *cross ratio* is $(z_0 : z_1 : z_2 : z_3) = \phi(z_3)$. The cross ratio has many useful properties:

- For $\gamma \in \mathrm{PSL}_2\,\mathbb{C}$, $(z_0 : z_1 : z_2 : z_3) = (\gamma(z_0) : \gamma(z_1) : \gamma(z_2) : \gamma(z_3))$;
- The points z_0, z_1, z_2, z_3 lie on a circle if and only if $(z_0 : z_1 : z_2 : z_3) \in \mathbb{R}$;
- If z_0, z_1, z_2, z_3 lie on a circle then $(z_0 : z_1 : z_2 : z_3) \in \mathbb{R}^-$ if and only if they are cyclically ordered in the expected way.
- $(z_0 : z_1 : z_2 : z_3) = (z_2 : z_3 : z_0 : z_1)$;
- $(z_0 : z_1 : z_2 : z_3) = \frac{(z_1 - z_2)(z_3 - z_0)}{(z_1 - z_0)(z_3 - z_2)}$.

Proposition 2.8 *The map $\lambda_q\colon \mathcal{IQ} \to \mathbb{R}$ given by*

$$\lambda_q([(R, f)]) = -\log(-(f(v_0) : f(v_1) : f(v_2) : f(v_3)))$$

is well defined and a bijection.

The function λ_q can be interpreted as a certain hyperbolic length. The diagonal in R connecting $f(v_0)$ to $f(v_2)$ divides R into two ideal triangles. While these two ideal triangles are both isometric to the unique ideal triangle as the sides where they are being glued to form the quadrilateral have infinite length there are an $\mathbb{R}$s worth of possible gluings. There is a natural way to parameterized these gluings. An ideal triangle has a unique inscribed circle which is tangent to each side. We label this point of tangencey as the midpoint of the side. Then $\lambda_1([(R, f)])$ is the signed distance between the two midpoints on the diagonal from $f(v_0)$ to $f(v_2)$. The sign comes from the counter-clockwise orientation on ideal triangles coming from the orientation of the triangles. This signed distance is the *shearing coordinate* for the quadrilateral. We can then restate Proposition 2.8 as:

Proposition 2.9 *The space of of marked ideal quadrilateral is parameterized by a choice (namely a choice of diagonal) of shearing coordinate.*

It is also useful to measure the distance between opposite sides of the ideal quadrilaterals relative to the shearing coordinates. Define functions $\ell_{02}\colon \mathcal{I}Q \to \mathbb{R}^+$ and $\ell_{13}\colon \mathcal{I}Q \to \mathbb{R}^+$ to be the distance between the sides $\{0,1\}$ and $\{2,3\}$ and the sides $\{1,2\}$ and $\{3,0\}$, respectively.

Lemma 2.10 *Let $\lambda_q\colon \mathcal{I}Q \to \mathbb{R}$ be the shearing coordinate associated to the diagonal $\{0,2\}$. Then $\ell_{02} \circ \lambda_q^{-1}\colon \mathbb{R} \to \mathbb{R}^+$ is decreasing and a homeomorphism. The function $\ell_{13} \circ \lambda_q^{-1}$ is an increasing homeomorphism.*

Proof. We work in the upper half plane model. Let g be a vertical geodesic. Then the R-neighborhood of g is bounded by the two rays with same basepoint as g that a make an angle θ with g where θ is an increasing function of R with range from 0 to $\pi/2$. If h a non-vertical geodesic that is disjoint from g then let r be the ray with basepoint the same as g that is tangent to h. Then the distance between g and h is an increasing function of θ. To apply this to our function we arrange the quadrilateral such that the vertex v_2 is at ∞ and v_0 and v_1 are fixed (say v_0 is at 0 and v_1 is at 1). In this picture g will be the vertical geodesic with basepoint v_3 and h will be the geodesic with endpoints 0 and 1. As the shearing coordinate increases the vertex v_3 will increase from $-\infty$ to 0 so the angle θ will decrease 2π to 0. This proves that $\ell_{02} \circ \lambda_q^{-1}$ is a decreasing homeomorphism.

For $\ell_{13} \circ \lambda_q^{-1}$ we fix v_3 (at say -1) and let v_1 vary. The same argument then shows that $\ell_{13} \circ \lambda_q^{-1}$ is an increasing homeomorphism. $\boxed{2.10}$

2.3.1 Hexagons

We can define the space of (marked) ideal hexagons and right-angled hexagons just as we did for ideal quadrilaterals. We first fix a topological oriented hexagon (defined in the obvious way) H and let $\mathcal{I}H$ be equivalence class of marked pairs of ideal hexagons and $\mathcal{R}H$ equivalence classes of marked pairs of right-angled hexagons.

There is a natural bijection from $\mathcal{R}H$ to $\mathcal{I}Q$. If we fix three alternating sides of the a right-angled hexagon we can extend the sides to complete geodesics. These geodesics will be pairwise disjoint and will limit to six distinct endpoints in $\partial\mathbb{H}^2$. These six points in $\partial\mathbb{H}^2$ determine an ideal hexagon. To define the bijection from $\mathcal{R}H$ to $\mathcal{I}Q$ we need to also keep track of markings. This is a bit tedious (and I won't do it here) but it is a good idea to think through the details.

There are a natural collection of functions that we can define on both $\mathcal{R}H$ and $\mathcal{I}Q$. On $\mathcal{R}H$, for each side s_i in H we let $\ell_i\colon \mathcal{R}H \to \mathbb{R}^+$ be the length of the ith side. For ideal hexagons, all sides have infinite length so instead of measuring the length of a side we measure the distance between distinct sides.

Lemma 2.11 *Under the bijection from $\mathcal{R}H$ to $\mathcal{I}Q$, ℓ_i is the distance between the $(i-1)$st and $(i+1)$st (measured mod 6) side.*

Proof. The distance between two sets is the infimum of the distance between points in the two sets. For complete geodesics there are three distinct possibilities:

- The geodesics intersect and the distance between them is zero.
- The geodesics are disjoint in $\mathbb{H}^2$ but have a common endpoint in $\partial\mathbb{H}^2$. In this case the distance is zero.
- There is a geodesic segment meeting both complete geodesics orthogonally. In this case the geodesics are disjoint and have no endpoints in common. Then the distance is the length of the segment.

If we extend the $(i-1)$st and the $(i+1)$st side to complete geodesics then the ith side meets them both orthogonally. The lemma follows. $\boxed{2.11}$

Our goal is the following parameterization of $\mathcal{R}H$:

Theorem 2.12 *The map $\ell\colon \mathcal{R}H \to (\mathbb{R}^+)^3$ given by $\ell = (\ell_0, \ell_2, \ell_4)$ is a bijection.*

By Lemma 2.11 can prove this theorem by parameterizing $\mathcal{I}H$ via the corresponding length functions. This is the approach we will take.

Any three of the ideal vertices of an ideal hexagon determine an ideal triangle. The hexagon can be decomposed as the disjoint union of 4 ideal triangles by taking the triangles with vertices $\{0,2,4\}$, $\{0,1,2\}$, $\{2,3,4\}$ and $\{4,5,0\}$. The triangle is $\{0,2,4\}$ is in the center with side meeting one of the other three triangles. (Picture should be added.) We will define a shear coordinate for each side of the central triangle. The orientation of $\mathbb{H}^2$ induces on orientation on the triangle which in turn induces an orientation on each side. For each side the shear coordinate is the signed distance between the midpoint of the central triangle and the midpoint of the adjacent triangle. This defines a map $\tilde{\lambda}\colon \tilde{\mathcal{I}H} \to \mathbb{R}^3$.

Lemma 2.13 *The map λ is a homeomorphism $\mathbb{R}^3$.*

We will use these shear coordinates to give an alternate parameterization via the distance between certain pairs of sides of the ideal hexagon. Define $\ell\colon \mathcal{I}H \to (\mathbb{R}^+)^3$ by $\ell(H) = (\ell_2(H), \ell_4(H), \ell_6(H))$ (where the ℓ_i are define in Lemma 2.11).

Theorem 2.14 *The map ℓ is a homeomorphism.*

Proof. Using the the homeomorphism $\lambda\colon \mathcal{I}H \to \mathbb{R}^3$ we can identify $\mathcal{I}H$ with $\mathbb{R}^3$. Note that $\mathcal{I}H$ and $(\mathbb{R}^+)^3$ are homeomorphic and simply connected so to show that ℓ is a homeomorphism we only need to show ℓ is a proper, local homeomorphism. The map ℓ is continuous so, by invariance of domain, to show that it is a local homeomorphism we only need to show that it is injective. We will do this by examining the behavior of the functions ℓ_i on lines parallel to the axes of the shear coordinate parameterization of $\mathcal{I}H$. The key is observation is that as in the proof Lemma 2.10 we have for all $(x_1, x_2, x_3) \in \mathbb{R}^3$ the functions

$$
\begin{aligned}
t &\mapsto \ell_2(t, x_2, x_3)\\
t &\mapsto \ell_4(x_1, t, x_3)\\
t &\mapsto \ell_6(x_1, x_2, t)\\
t &\mapsto \ell_6(t, x_2, x_3)\\
t &\mapsto \ell_2(x_1, t, x_3)\\
t &\mapsto \ell_4(x_1, x_2, t)
\end{aligned}
$$

are decreasing homeomorphisms. Let (x_1, x_2, x_3) and (y_1, y_2, y_3) be points in $\mathcal{I}H$. Then the triple $(x_1 - y_1, x_2 - y_2, x_3 - y_3)$ has at least two coordinates with the same sign (where one of the coordinates is allowed to be 0). For example, assume that the first two coordinates are positive. Then $\ell_2(x_1, x_2, x_3) < \ell_2(y_1, y_2, y_3)$. Other cases are similar and it follows that ℓ is injective.

The proof of properness follows a similar strategy. Let $x^i \in \mathcal{I}H \cong \mathbb{R}^3$ be a divergent sequence and assume (after possibly passing to a subsequence) that each coordinate converges in the extended real line. Assume for example that the first coordinate converges to $+\infty$ and the second coordinate either converge in $\mathbb{R}$ or converges to $+\infty$. Then $\ell(x^i) \to 0$. If we replace $+\infty$ with $-\infty$ then $\ell(x^i) \to +\infty$. The other cases are similar and properness follows. $\boxed{2.14}$

2.3.2 Pairs of pants

A *pair of pants* P is a genus zero surface with three boundary components. Up to isotopy, there is a unique collection of disjoint arcs connecting pairs of distinct boundary components. These three arcs decompose P into two topological hexagons. We will use our parameterization of right-angled hexagons to build hyperbolic pairs of pants with geodesic boundary. When doing so we will be able to arbitrarily choose the length of the boundary component.

We need to first discuss hyperbolic surfaces with piecewise geodesic boundary and how we can glue them together. Let Σ be a surface with boundary. A *hyperbolic structure with piecewise geodesic boundary* is a hyperbolic atlas with the boundary components mapped to piecewise geodesics. We want to distinguish a certain collection of *singular points* on the boundary. This will be a discrete set that includes all the points where the boundary is not smooth. The *sides* of the boundary are the smooth geodesic arcs between singular points or the components of the boundary that don't contain a singular point.To each singular point the geodesic boundary forms an angle between 0 and 2π. If the singular component is actually a smooth point of the boundary then this angle will be π. The purpose of allowing smooth points to be singular should be clear in the following gluing theorem.

Theorem 2.15 *Let Σ be an oriented (possibly disconnected) hyperbolic surface with polygonal boundary and let P be the collection of singular points and S the collection of oriented sides. A* gluing *G is a collection of pairs of sides along with an orientation reversing isometry between them such that no side appears in more than one pair. Let Σ_G be corresponding quotient space. Then Σ_G is an oriented surface with boundary. Let P_G be the image of singular points in Σ_G. Assume that*

1. *if $p \in P_G$ is in the interior of Σ_G then the sum of the angles is 2π;*

2. *if $p \in P_G$ is the boundary of Σ_G then the sum of the angle is $< 2\pi$.*

Then Σ_G is a hyperbolic surface with piecewise geodesic boundary.

The proof of this theorem is straightforward but tedious; we'll omit it. Our first application is to find hyperbolic pairs of pants with geodesic boundary with prescribed boundary lengths.

Let R be a hyperbolic structure on P with smooth geodesic boundary. If $f\colon P \to R$ is a orientation preserving homeomorphisms then (R, f) is a *marked hyperbolic structure* on P. We define an equivalence relation on pairs by $(R_0, f_0) \sim (R_1, f_1)$ if there is an orientation preserving isometry $\phi\colon R_0 \to R_1$ with $\phi \circ f_0$ homotopic to f_1. We then let $\mathcal{T}(P)$ be the set of equivalence classes. We will define a topology later.

We can discuss the completeness of hyperbolic surfaces with boundary as before. We need the following extension of Corollary 2.4.

Proposition 2.16 *Let Σ be a compact hyperbolic surface with geodesic boundary. Then $\tilde{\Sigma}$ is subspace of $\mathbb{H}^2$.*

Proof. Let $D\Sigma$ be the doubled surface. Then $D\Sigma$ is compact and hence complete. By Theorem 2.4 $\tilde{D\Sigma}$ is $\mathbb{H}^2$. Let Σ be a component of the pre-image of X in $\tilde{D\Sigma}$. We claim that X is the universal cover of Σ. If X is not simply connected then it contains a closed, essential curve γ. After performing surgery we can assume that γ is simple. In $\mathbb{H}^2$ the curve γ will bound a disk. Since γ doesn't bound a disk in X, the disk must contain a boundary component of ∂X. However, these are all complete geodesics and hence non-compact so we have a contradiction. Therefore X is simply connected and $X = \tilde{\Sigma}$. $\boxed{2.16}$

The assumption of compactness is not necessary. However, then it is a more work to show that the doubled surface is complete. Note that if a component of $\partial\Sigma$ is not compact then one can glue to copies of Σ to itself and obtain non-complete surface.

We can also weaken the assumption of geodesic boundary to the assumption that the boundary is *locally convex*. We may come back to this later.

Lemma 2.17 *Let $f\colon P \to P$ be an orientation preserving homeomorphism that takes each boundary component to itself. Then f is homotopic to the identity.*

Proof. Let $\{a_0, a_1, a_2\}$ be disjoint properly embedded arcs connecting each pair of boundary components. First assume the arcs a_i and $f(a_i)$ are disjoint. Then for each i the complement of a_i and $f(a_i)$ has two components one of which is a disk and the other will contain the third boundary component of P. We can then homotope $f(a_i)$ to a_i through the disk component. Once the map has been homotoped to be the identity on the a_i we can then homotope f to be the identity on the boundary. The complement of the boundary and the three arcs $\{a_0, a_1, a_2\}$ is two disks. On the boundary of each disk the map is the identity. The final step is to homotope the map to be the identity on the two disks.

We are left to show that the map can be homotoped so that the arcs a_i and $f(a_i)$ are disjoint. This is a standard innermost disk argument. $\boxed{2.17}$

Let $\partial P = \{b_0, b_1, b_2\}$. We can then define length functions $\ell_{b_i}\colon \mathcal{T}(P) \to \mathbb{R}^+$ by $\ell_{b_i}([(R, f)])$ to be the length of the boundary component $f(b_i)$ of R. We the define $\ell_{\partial P}\colon \mathcal{T}(P) \to (\mathbb{R}^+)^3$ by $\ell_{\partial P} = (\ell_{b_0}, \ell_{b_1}, \ell_{b_2})$.

Theorem 2.18 *The map $\ell_{\partial P}$ is a bijection.*

Proof. Given $(c_0, c_1, c_2) \in (\mathbb{R}^+)^3$ by Theorem 2.12 there exist unique right angled hexagon with the labeled sides having lengths $(c_0/2, c_1/2, c_2/2)$. Using Theorem 2.15 we double the hexagon along the unlabeled sides to form a hyperbolic pair of pants R with boundary lengths (c_0, c_1, c_2). We then choose a homeomorphism $f\colon P \to R$ that takes b_i to the boundary component of length c_i. This shows that $\ell_{\partial P}$ is surjective.

To prove that the map is injective we need to show that every hyperbolic pair of pants can be built out of right angled hexagons. Let R be a hyperbolic pair of pants with boundary lengths $(c_0, c_1, c_2) \in (\mathbb{R}^+)^3$. We then double R along its boundary to form a genus two surface DR. In DR the arcs a_i become simple closed curves A_i. These curves are essential as they each transversely intersect two of the b_j exactly once. The holonomy of these curves is also hyperbolic as DR is compact. Therefore the A_i are homotopic to simple, mutually disjoint closed geodesics (which we still label A_i). The doubled surface DR has orientation reversing involution that fixes the curves b_i pointwise and is an involution on each A_i. This implies that the A_i meet the b_j orthogonally and that in the original pair of pants R the a_i are homotopic to disjoint geodesic arcs that meet the boundary orthogonally. Cutting open R along these arcs produces two right angled hexagons. H_1 and H_2. These hexagons have three side lengths in common so by uniqueness part of Theorem 2.12 we have that H_1 and H_2 are isometric. The conclusion is that there is an isometry between two hyperbolic pairs of pants with boundary lengths equal.

Now assume that $\ell_{\partial P}([(R_0, f_0)]) = \ell_{\partial P}([(R_1, f_1)])$. Then, by the above paragraph, there is an isometry $\phi\colon R_0 \to R_1$. The map $(f_1)^{-1} \circ \phi \circ f_0$ maps each boundary component in ∂P to itself and therefore is homotopic to the identity. It follows that $\phi \circ f_0$ and f_1 are homotopic so $[(R_0, f_0)] = [(R_1, f_1)]$ and $\ell_{\partial P}$ is injective. $\boxed{??}$

2.4 Collars

An extremely useful fact is that simple closed geodesics of collars of definite width that only depends on the length of the curve. Furthermore for disjoint simple closed geodesics we can choose the widths so that the collars of the two geodesics are disjoint.To state the theorem we define $W\colon \mathbb{R}^+ \to \mathbb{R}^+$ such that $W(\ell)$ is the hyperbolic distance, in the upper half space model, between the imaginary axis and the geodesic with endpoints x and $e^{\ell}x$ where $x \in \mathbb{R}\backslash\{0\}$. Note that $W(\ell)$ doesn't depend on the choice x and that this map is an decreasing homeomorphism.

Theorem 2.19 *Let Σ be a complete hyperbolic surface and γ_0 and γ_1 simple closed geodesics of length ℓ_0 and ℓ_1. Then the collars of width $W(\ell_0)$ about γ_0 and γ_1 are embedded and disjoint.*

Proof. If the collar of width $W(\ell_0)$ about γ_0 is not embedded then there is an essential arc connecting γ_0 to itself of length $< W(\ell_0)$. If the collars of widths $W(\ell_0)$ and $W(\ell_1)$ about γ_0 and γ_1 intersect then there is an arc from γ_0 to γ_1 of length $< W(\ell_0)/2 + W(\ell_1)/2 \leq \max\{W(\ell_0), W(\ell_1)\}$. Note that in both cases the arc will lift in the universal cover $\tilde{\Sigma}$ to an arc connecting distinct components of the pre-image $\Gamma \subset \tilde{\Sigma}$ of $\gamma_0 \cup \gamma_1$. We will show that if B is a component of Γ that maps to γ_i then any arc from B to another component of Γ will have length $> W(\ell_i)$. The theorem will then follow.

We work in the upper half space model and arrange that B is the imaginary axis. Since B covers b_i the deck group will contain the element $z \mapsto e^{\ell_i}$. All of the components of Γ are disjoint, complete geodesics in $\mathbb{H}^2$ and this set is invariant under the action of the deck group. The only element of Γ that will be fixed by $z \mapsto e^{\ell_i}z$ is B. If B' is another component then both its endpoints will be in $\mathbb{R}^+$ or $\mathbb{R}^-$ for otherwise B' would intersect B. Assume it is the former and label the left endpoint $x^- \in \mathbb{R}^+$ and right endpoint x^+. Then we must have $x^+ < e^{\ell_i} x^-$ for otherwise B' would intersect its translate. The geodesic from x^- to $e^{\ell_i} x^-$ will have distance $W(\ell_i)$ from B and this will be less than the distance from B'. This completes the proof. $\boxed{2.19}$

Given a simple closed geodesic γ of length ℓ on a complete hyperbolic surface we let $C(\gamma)$ be the collar of width $W(\ell)$. The proposition implies that for any disjoint collection of simple closed geodesics these collars are embedded and disjoint.

Note that the geometry of this collar doesn't depend on the surface Σ and only on the length ℓ of γ. We let $C(\ell)$ be this annulus.

Lemma 2.20 *Each component of $\partial C(\ell)$ has length $\ell\sqrt{\frac{\coth^2(\ell/2)}{4} - \frac{1}{2}} > 1$.*

Proof. The universal cover of a collar of width $2R$ can be identified with the subspace in the upper half plane between the two rays based at the origin that make an angle θ with the imaginary axis where $\cos\theta = \operatorname{sech} R$. From this we can compute that the length of the boundary of the collar of width $2R$ about a geodesic of length ℓ is $\ell \cosh R$ by calculating the length of the path $e^t e^{\imath(\pi/2-\theta)}$ with $t \in [0, \ell]$.

Rather than compute the length of $\partial C(\ell)$ we will compute the length of the boundary of a collar of width $2W(\ell)$. For this we observe that in the universal cover the boundary of the collar of width $2W(\ell)$ will be tangent to a geodesic of distance $W(\ell)$ from the core geodesic. If the lift of the core geodesic is the imaginary axis then the boundary will be the ray with angle θ where $\cos\theta = \tanh(\ell/2)$ so the length of a boundary component of a collar of width $2W(\ell)$ is $\ell\coth(\ell/2)$ so $\ell\cosh(W(\ell)) = \ell\coth(\ell/2)$.

We need to compute $\ell\cosh(W(\ell)/2)$. Standard identities give the result. $\boxed{2.20}$

2.5 Teichmüller space

We are now ready to discuss Teichmüller space! Fixed a compact surface without boundary Σ and assume that the genus of Σ is ≥ 2. Exactly as for a pair of pants we can define a marked hyperbolic structure on Σ with an equivalence relation between pairs. Then the *Teichmüller space* of hyperbolic structures on Σ, denote $\mathcal{T}(\Sigma)$ is the set of equivalence classes. We will see that this space has a natural topology and that in this topology $\mathcal{T}(\Sigma)$ is homeomorphic to $\mathbb{R}^{6g-6}$.

We first need to discuss length functions. Let γ be a closed, homotopically non-trivial curve on Σ. Then for each pair (R, f) the curve $f(\gamma)$ is homotopic to a geodesic on R. The length of this geodesic will only depend on the equivalence of (R, f) so we have a well defined map $\ell_\gamma \colon \mathcal{T}(\Sigma) \to \mathbb{R}^+$ which assigns to each pair the length of γ.

A *pants decomposition* $\mathbf{P}$ of Σ is a maximal collection of essential, homotopicilly distinct, disjoint simple closed curves on Σ. A simple Euler characteristic count show that $\mathbf{P}$ contains $3g-3$ curves. We then have a function $\ell_{\mathbf{P}}\colon \mathcal{T}(\Sigma) \to (\mathbb{R}^+)^{3g-3}$ which takes each pair to its $3g-3$-tuple of lengths. Using Theorem 2.18 it is not hard to see that this map is onto. However, it is very much not injective. Our next goal is to understand this lack of injectivity.

Given a complete hyperbolic surface R and a pants decomposition $\mathbf{P}$ we let $C(\mathbf{P}, R)$ be the standard collar neighborhoods of the geodesic representatives of $\mathbf{P}$ and $V(\mathbf{P}, R)$ the complement of the collars. Note that if R_0 and R_0 are two hyperbolic surfaces and the pants decomposition $\mathbf{P}$ has the same length on both then corresponding collars and complementary pairs of pants are each isometric. However, there still may not be an isometry between the entire surfaces. Furthermore, if the surfaces are marked, it is possible that there are isometries between the surfaces but the isometry may not be in the homotopy class given by the marking.

Lemma 2.21 *Given pairs (R_0, f_0) and (R_1, f_1) with $\ell_{\mathbf{P}}(R_0, f_0) = \ell_{\mathbf{P}}(R_1, f_1)$ there exists a homeomophism $\phi\colon R_0 \to R_1$ such that*

- *$\phi \circ f_0$ is homotopic to f_1;*
- *ϕ restricts to an isometry from $V(\mathbf{P}, R_0)$ to $V(\mathbf{P}, R_1)$.*

Furthermore any such map is equivalent to ϕ on $V(\mathbf{P}, R_0)$.

Proof. We only sketch the proof. The key is Lemma 2.17. Start with an arbitrary map $\psi\colon R_0 \to R_1$ such that $\psi \circ f_0 \sim f_1$. For each component X_0 of $V(\mathbf{P}, R_0)$ there is a corresponding component X_1 of $V(\mathbf{P}, R_1)$ and we lift the map to the covers associated to $\pi_1(X_0)$ and $\pi_1(X_1)$. In this cover we can use Lemma 2.17 to homotopy the lifted map to be an isometry from X_0 to X_1. We can further assume that this homotopy is supported on small neighborhood of X_0 and therefore will descend to a homotopy of ψ that is also supported on a small neighborhood of X_0. In particular, the homotopy won't affect the other components of $V(\mathbf{P}, R_0)$ and we can independently do such a homotopy on each component of $V(\mathbf{P}, R_0)$ such that the resulting map is an isometry on $V(\mathbf{P}, R_0)$.

At this point the map may not be a homeomorphism as we have not controlled its behavior on the collars $C(\mathbf{P}, R_0)$. On the boundary of the collars the map is already well behaved isometry so we just need to perform the homotopy on the interior. Again we can do this by lifting to the corresponding covers. $\boxed{2.21}$

Identify S^1 with the group $\mathbb{R}/\mathbb{Z}$. We to describe a certain homotopy class of maps of the annulus $A = S^1 \times [0, 1]$ to itself. Define $\mathrm{tw}_s\colon A \to A$ by $\mathrm{tw}_s(\theta, x) = (\theta + sx, x)$.

Lemma 2.22 *Let $f\colon A \to A$ be a homotopy equivalence such that f is the identity on $S^1 \times \{0\}$ and a rotation on $S^1 \times \{1\}$. Then there is a unique $s \in \mathbb{R}$ such that f is homotopic to tw_s relative to a homotopy that is stationary on ∂A.*

Furthermore if $f_0, f_1\colon A \to A$ are homotopy equivalences that are the identity on $S^1 \times \{0\}$ and rotations on $S^1 \times \{1\}$ with twist numbers s_i then the twist number for $f_0 \circ f_1$ is $s_0 + s_1$.

Proof. The universal cover $\tilde{A}$ is naturally identified with $\mathbb{R} \times [0,1]$ and there is a unique lift of $\tilde{f}$ of f that is the identity on $\mathbb{R} \times \{0\}$. On $\mathbb{R} \times \{1\}$, $\tilde{f}$ will acts as a translation by some $s \in \mathbb{R}$. We will show that f is homotopic to tw_s by a homotopy that is stationary on the boundary. Let $\tilde{\mathrm{tw}}_s$ be the lift of tw_s that is the identity on the boundary. Then $\tilde{\mathrm{tw}}_s$ is translation of $\mathbb{R} \times \{1\}$ by s. There is then a "straight line" homotopy taking $\tilde{f}$ to $\tilde{\mathrm{tw}}_s$ that is stationary on the boundary and descends to a homotopy taking f to tw_s. Namely define $\tilde{F}$ by $\tilde{F}(\theta, x, \lambda) = \lambda \tilde{f}(\theta, x) + (1-\lambda)\tilde{\mathrm{tw}}_s(\theta, x)$.

For uniqueness we observe that tw_s is homotopic to $\mathrm{tw}_{s'}$ by a homotopy that is stationary on the boundary if and only if $s - s'$.

Since $f_i \sim \mathrm{tw}_{s_i}$ we have $f_0 \circ f_1 \sim \mathrm{tw}_{s_0} \circ \mathrm{tw}_{s_1}$ (where all homotopies are stationary on the boundary). Since $\mathrm{tw}_{s_0} \circ \mathrm{tw}_{s_1} = \mathrm{tw}_{s_0+s_1}$ we have $f_0 \circ f_1 \sim \mathrm{tw}_{s_0+s_1}$. $\boxed{2.22}$

Recall that $C(\ell)$ is the standard collar about a simple closed geodesic of length ℓ. We will implicitly identify each $C(\ell)$ with the annulus A from Lemma 2.22. In particular rotations of A will correspond to isometries of $C(\ell)$. This identification is unique up to isometries of $C(\ell)$ which preserve the boundary components.

We are now ready to describe the fibers of the map $\ell_{\mathbf{P}}$. In what follows we will suppress the map in our pairs (R, f). When we write $\phi\colon R_0 \to R_1$ we will implicitly assume that ϕ is in the homotopy class such that $\phi \circ f_0 \sim f_1$.

Given $\mathbf{L} \in (\mathbb{R}^+)^{3g-3}$ let $F_{\mathbf{L}} = \ell_{\mathbf{P}}^{-1}(\mathbf{L})$. For $R_0, R_1 \in F_{\mathbf{L}}$ let $\phi\colon R_0 \to R_1$ be the map given by Lemma 2.21. If γ_i is a component of $\mathbf{P}$ of length ℓ_i then there is are isometries $\lambda_j\colon C(\ell_i) \to R_j$ for $j = 0, 1$ that take $C(\ell_i)$ to the collars about γ_i in each surface. By Lemma 2.22, there exists a unique s such that tw_s is homotopic, holding the boundary stationary, to $\lambda_1^{-1} \circ \phi \circ \lambda_0$. Define $\mathrm{Tw}_{\gamma_i}\colon F_{\mathbf{L}} \times F_{\mathbf{L}} \to \mathbb{R}$ to be $\mathrm{Tw}_{\gamma_i}(R_0, R_1) = s_i$ and $\mathrm{Tw}_{\mathbf{P}}(R_0, R_1) = (s_1, \dots, s_{3g-3})$. To see that this map doesn't depend on the choice of ϕ we note that if ϕ_0 and ϕ_1 both satisfy the conditions given by Lemma 2.21 then not only do they agree on $V(\mathbf{P}, R_0)$ but there is a homotopy from ϕ_0 to ϕ_1 that is stationary on $V(\mathbf{P}, R_0)$. In particular, the homotopy class of the map ϕ on the collars is well defined.

The following lemma is a direct consequence of Lemma 2.22.

Lemma 2.23 *Given $R_0, R_1, R_2 \in F_{\mathbf{L}}$ we have*

$$\mathrm{Tw}_{\mathbf{P}}(R_0, R_2) = \mathrm{Tw}_{\mathbf{P}}(R_0, R_1) + \mathrm{Tw}_{\mathbf{P}}(R_1, R_2).$$

To show that the map $\mathrm{Tw}_{\mathbf{P}}$ is surjective we need to define a *Dehn twist*. This a purely topological construction. Let Σ be an oriented surface and γ and essential, simple closed curve on Σ. We identify a collar of neighborhood of γ with A and define $D_\gamma\colon \Sigma \to \Sigma$ to be the identity outside of A and tw_1 on A. Note that tw_1 fixes both boundary components of A so that this map is continuous and in fact a homeomorphism.

Any homeomorphism ψ of Σ acts on $\mathcal{T}(\Sigma)$ by taking a pair (R, f) to $(R, f \circ \psi^{-1})$. Again we will often suppress the map f and write $\psi_* R$ for the pair $(R, f \circ \psi^{-1})$. Note that we compose with ψ^{-1} so that $(\psi_0 \circ \psi_1)_* R = (\psi_0)_*(\psi_1)_* R$. Note that the map ψ_* on $\mathcal{T}(\Sigma)$ only depends on the homotopy class of ψ.

Lemma 2.24 *Given $R \in \mathcal{T}(\Sigma)$ and an essential simple closed curve γ on Σ, $\ell_\gamma(R) = \ell_\gamma((D_\gamma)_* R)$. Furthermore if $\gamma_i \in \mathbf{P}$ then $\mathrm{Tw}_{\gamma_i}(R, (D^n_{\gamma_i})_* R) = -n$.*

Proof. The boundary of the annulus A is fixed by D_γ so $D_\gamma(\gamma)$ is freely homotopic to γ and it follows that $\ell_\gamma(R) = \ell_\gamma((D_\gamma)_* R)$.

For the second fact we observe that we can choose A such that $f(A) = C(\gamma_i, R)$ where f is the marking map. In fact we can do this in such a way that $f \circ D_\gamma^{-n} = \phi \circ f$ where ϕ is the identity on the complement of $C(\gamma_i, R)$ and is the twist map tw_{-n} on the collar. The second equality follows. $\boxed{2.24}$

Given $\mathbf{L} \in \mathbb{R}^{3g-3}$ there is a nearly canonical way to build an unmarked hyperbolic structure with pants curves having length $\mathbf{L}$. We start with a collection of hyperbolic pairs of pants that have the correct boundary lengths. On each pair of pants there are three geodesic arcs connecting distinct pairs of boundary components that meet the boundary component orthogonally. On each boundary component of each pair of pants there are the feet of two of these arcs. We pick one for each boundary component. (This choice is why the construction is not canonical.) We then glue the pants together so that each of the chosen feet are identified in the gluing and label the resulting surface $X_{\mathbf{L}}$. In $X_{\mathbf{L}}$ the arcs orthogonal to curves in $\mathbf{P}$ be come a collection of simple closed geodesics which we label $\mathbf{P}^\perp$. Note that the complement of $\mathbf{P} \cup \mathbf{P}^\perp$ in $X_{\mathbf{L}}$ is a collection of right-angled hexagons. In particular they are disks.

Now pick an $3g-3$-tuple $t \in (\mathbb{R}/\mathbb{Z})^{3g-3}$. We also orient each curve $\gamma_i \in \mathbf{P}$. We then construct an unmarked surface $X^t_{\mathbf{L}}$ such that the chosen feet at γ_i differ by a twist of size t_i. We then have the following lemma.

Lemma 2.25 *Choose markings $f\colon \Sigma \to X_{\mathbf{L}}$ and $f_t\colon \Sigma \to X^t_{\mathbf{L}}$ so that the $\mathbf{P}$ is taken to the corresponding geodesics. Then $\mathrm{Tw}_{\mathbf{P}}((X_{\mathbf{L}}, f), (X^t_{\mathbf{L}}, f_t))$ is equal to t modulo 1.*

Corollary 2.26 *Fix $R \in F_{\mathbf{L}}$. Given $s \in \mathbb{R}^{3g-3}$ there exists a unique $R' \in F_{\mathbf{L}}$ such $\mathrm{Tw}_{\mathbf{P}}(R, R') = s$.*

Proof. Using Lemma 2.25 we can find an $S \in F_{\mathbf{L}}$ such $\mathrm{Tw}_{\mathbf{P}}(R, S)$ is s modulo 1. We can then us Lemmas 2.23 and 2.24 to apply Dehn twists to S to find R'. $\boxed{2.26}$

.

Given $s \in \mathbb{R}^{3g-3}$ we define a map $\mathrm{Tw}^s_{\mathbf{P}} \colon \mathcal{T}(\Sigma) \to \mathcal{T}(\Sigma)$ as follows. For $R \in F_{\mathbf{L}}$ let $\mathrm{Tw}^s_{\mathbf{P}}(R)$ be the unique point in $F_{\mathbf{L}}$ such that $\mathrm{Tw}_{\mathbf{P}}(R, \mathrm{Tw}^s_{\mathbf{P}}(R)) = s$. This defines an action of $\mathbb{R}^{3g-3}$ on $\mathcal{T}(\Sigma)$. This would natural make $\mathcal{T}(\Sigma)$ a *principal* $\mathbb{R}^{3g-3}$*-bundle* except that we have not yet given $\mathcal{T}(\Sigma)$ at topology. We do that now.

We will do this by defining a map from $(\mathbb{R}^+)^{3g-3}$ to $\mathcal{T}(\Sigma)$ such that the composition with $\ell_{\mathbf{P}}$ is the identity.

Recall that for each $\mathbf{L} \in (\mathbb{R}^+)^{3g-3}$ we have define a hyperbolic surface $X_{\mathbf{L}}$. Fixing a basepoint $L \in (\mathbb{R}^+)^{3g-3}$ for all $\mathbf{L}' \in (\mathbb{R}^+)^{3g-3}$ we have a canonical (up to homotopy) map $\psi_{\mathbf{L},\mathbf{L}'} \colon X_{\mathbf{L}} \to X_{\mathbf{L}'}$. This map is determined by the property that it takes each component of $\mathcal{P}$ in $X_{\mathbf{L}}$ to the corresponding component of $\mathcal{P}$ in $X_{\mathbf{L}'}$ and similarly with $\mathbf{P}^\perp$. Since the complement of $\mathbf{P} \cup \mathbf{P}^\perp$ is disks this map is well defined up to homotopy. These surfaces are still unmarked. To mark them we fix a map $f_{\mathbf{L}} \colon \Sigma \to X_{\mathbf{L}}$ that maps $\mathbf{P}$ in Σ to $\mathbf{P}$ in $X_{\mathbf{L}}$ and define $f_{\mathbf{L}'} = \phi_{\mathbf{L},\mathbf{L}'} \circ f_{\mathbf{L}}$. Let $R_{\mathbf{L}'} = [(X_{\mathbf{L}'}, f_{\mathbf{L}'})]$.

We define $\sigma \colon (\mathbb{R}^+)^{3g-3} \to \mathcal{T}(\Sigma)$ by $\sigma(\mathbf{L}') = R_{\mathbf{L}'}$. We would like to say that σ is continuous but to do so we first need to give $\mathcal{T}(\Sigma)$ a topology! In fact, we will use σ to do this. Namely, we define a bijection from $(\mathbb{R}^+)^{3g-3} \times \mathbb{R}^{3g-3}$ by $(\mathbf{L}', s) \mapsto \mathrm{Tw}^x_{\mathbf{P}}(\sigma(Bl'))$. These are *Fenchel-Nielsen coordinates* for $\mathcal{T}(\Sigma)$ and they give $\mathcal{T}(\Sigma)$ a topology. Note that the $\mathbb{R}^{3g-3}$ action is continuous in these coordinates. However, we also observe that the coordinates depend on the choice $\mathcal{P}$. To see that the topology doesn't depend on $\mathcal{P}$ we need a more canonical way of defining it.

2.6 Quasiconformal maps

Let $f \colon \Omega_0 \to \Omega_1$ be an orientation preserving between domains in $\mathbb{C}$. At each point the derivative of f is a $\mathbb{R}$-linear map between tangent spaces (which are canonically isomorphic to $\mathbb{C}$). We would like to measure how far f is from being conformal. We begin with some linear algebra.

Let $T \colon \mathbb{C} \to \mathbb{C}$ be $\mathbb{R}$-linear. An $\mathbb{R}$-linear map takes round circles to ellipses and we consider the ratio of the outradius to the inradius our measure of how far T is from a conformal map. This the *dilatation* of T. Recall that we can write $Tz = T_z z + T_{\overline{z}}\overline{z}$ where $T_z, T_{\overline{z}} \in \mathbb{C}$. Let $\mu = T_{\overline{z}}/T_z$. This is the *Beltrami differential* of T and it contains much information about the geometry of T. For now the most important thing is that $\frac{1+|\mu|}{1-|\mu|}$ is the dilatation of T. (If $|\mu| = 1$ then T is not invertible. If $|\mu| > 1$ then T is orientation reversing.)

We now apply this to our map f. Let $\mu(z) = \frac{f_{\overline{z}}(z)}{f_z(z)}$. Since f is orientation preserving, $|\mu| < 1$. We say that f is K*-quasiconformal* if $\|\mu\|_\infty = \frac{K-1}{K+1} < 1$ where $\|\mu\|_\infty$ is the sup-norm of μ.

Now let $f\colon X_0 \to X_1$ be a smooth, orientation preserving map between Riemann surfaces. We would like to define the Beltrami coefficient for f. If we apply the previous paragraphs via charts our answer will depend on the choice of charts. To account for this we need to understand how μ changes if we pre or post compose with a conformal map.

This is a problem in linear algebra. Let $T\colon \mathbb{C} \to \mathbb{C}$ be $\mathbb{R}$-linear and $S\colon \mathbb{C} \to \mathbb{C}$ be conformal (or equivalently $\mathbb{C}$-linear). From earlier we have that $(S \circ T)_z = S_z T_z$ and $(S \circ T)_{\overline{z}} = S_z T_{\overline{z}}$ (since $S_{\overline{z}} = 0$) and therefore $(S \circ T)_{\overline{z}}/(S \circ T)_z = T_{\overline{z}}/T_z$.

If we pre-compose with S then things are more complicated. In this case $(T \circ S)_z = T_z S_z$ and $(T \circ S)_{\overline{z}} = T_{\overline{z}} \bar{S}_z$ and $(T \circ S)_{\overline{z}}/(T \circ S)_z = \frac{T_{\overline{z}}}{T_z} \frac{\bar{S}_z}{S_z}$. In this case the Beltrami coefficient of T and $T \circ S$ aren't identical but they have the same norm.

This creates the following problem. We can take the Beltrami differential of f by choosing charts in X_0 and X_1. If we do this the differential won't depend on the choice of chart in X_1 but it will depend on the choice of chart in X_0. Because of this the Beltrami differential is not a function on X_0 but instead it is a section of a certain complex line bundle. We will come back to this later. For now the we will just observe that the absolute value of the Beltrami differential does not depend on the choice of chart on either surface.

More explicitly: Fix a point $z \in X_0$ and choose charts (U_0, ψ_0) and (U_1, ψ_1) such that $z \in U_0$ and $f(z) \in U_1$. Then define $|\mu|(z) = (\psi_1 \circ f \circ (\psi_0)^{-1})_z(\psi(z))/(\psi_1 \circ f \circ (\psi_0)^{-1})_{\overline{z}}(\psi(z))$. The function $|\mu|$ does not depend on the choice of charts. Furthermore if f is smooth (or just has continuous first derivatives) then $|\mu|$ is a continuous function on X_0. The map f is K-quasiconformal if $\|\mu\|_\infty = \frac{K-1}{K+1}$. Since $|\mu(z)|$ is continuous and < 1 if X_0 is compact then f will be K-quasiconformal for some K. We want to minimize K.

We can use this notion to define a metric on $\mathcal{T}(\Sigma)$. Namely, given $R_0, R_1 \in \mathcal{T}(\Sigma)$ define

$$d_{\mathcal{T}}(R_0, R_1) = \log \inf\{K | \text{there exists a } K\text{-quasiconformal map } f\colon R_0 \to R_1\}.$$

We need to check that this is a metric. If we knew there was a map f such that f was K-quasiconformal and $d_{\mathcal{T}}(R_0, R_1) = \log K$ then the proof that $d_{\mathcal{T}}$ is a metric would reduce to understanding composition rules for Beltrami differentials. However, the only time such a map will exist is when $R_0 = R_1$. To find a minimizing map in general we need to expand our definition of quasiconformal map. In particular we need to allow maps that aren't differentiable. There are two approaches, a geometric one and and analytic one. We'll begin with the geometric definition.

2.6.1 The modulus of a quadrilateral

We have already defined a topological quadrilateral. We now want to emphasize it's complex structure. Let Q be a closed (topological) disk in $\mathbb{C}$ with four distinguished points (*vertices*) and cyclically ordered points in ∂Q. Label the points $\{v_0, v_1, v_2, v_3\}$. Then there is a unique rectangle $R = [0, K] \times [0, 1] \subset \mathbb{C}$ such that there exists a conformal map $\phi \colon Q \to R$ with $\phi(v_0) = 0 \in \mathbb{C}$ and all other vertices of Q are taken to vertices of R. (This requires some "classical" but non-elmentary complex analysis.) Then the *modulus* of Q is $m(Q) = K$.

We can now give the *geometric* definition of a K-quasiconformal map. Let $f \colon \Omega_0 \to \Omega_1$ be an orientation preserving homeomorphism. Then f is K-quasiconformal if for all quadrilaterals $Q \subset \Omega_0$ we have

$$m(f(Q))/K \leq m(Q) \leq Km(f(Q)).$$

We emphasize that we are not assuming that f is differentiable. If we let Q' be the quadrilateral Q with vertices rotated by one then $m(Q') = m(Q)^{-1}$ so if $m(f(Q))/K \leq m(Q)$ for all quadrilaterals Q then f is K-quasiconformal.

The first thing to observe is that a smooth K-quasiconformal map as defined earlier is also K-quasiconformal with this new definition.

Lemma 2.27 *Let $f \colon [0, K_0] \times [0, 1] \to [0, K_1] \times [0, 1]$ be a smooth K-quasiconformal map with $K_0 \leq K_1$. Then $K \geq K_1/K_0$ with equality if and only if $f(z) = Kx + \imath y$.*

Proof. We first integrate f along horizontal lines and apply the fundamental theorem of calculus:

$$K_1 = \int_0^{K_0} f_x(x, y)dx$$

for all $y \in [0, 1]$. Using that $f_x = f_z + f_{\overline{z}}$ and taking absolute values we have

$$K_1 \leq \int_0^{K_0} (|f_z| + |f_{\overline{z}}|)dx.$$

Next we integrate in the y-direction:

$$K_1 \leq \int_0^1 \int_0^{K_0} (|f_z| + |f_{\overline{z}}|)dxdy.$$

We can rewrite the inside as

$$\sqrt{\frac{|f_z| + |f_{\overline{z}}|}{|f_z| - |f_{\overline{z}}|}}\sqrt{|f_z|^2 - |f_{\overline{z}}|^2} \quad = \quad \sqrt{\frac{1 + |\mu|}{1 - |\mu|}}\sqrt{|f_z|^2 - |f_{\overline{z}}|^2}$$

The term in the left square root is the dilatation and is bounded by K. The term under the right square root is the Jacobian.

Squaring both sides of the inequality and applying the Cauchy-Schwarz inequality we have

$$K_1^2 \leq \left(\int_0^1 \int_0^{K_0} \frac{1+|\mu|}{1-|\mu|} dxdy\right)\left(\int_0^1 \int_0^{K_0} (|f_z|^2 - |f_{\overline{z}}|^2) dxdy\right).$$

The integral on the left bounded by K times the area of the rectangle $[0, K_0] \times [0, 1]$. This product is KK_0. The right integral is the area of the rectangle $[0, K_1] \times [0, 1]$ so we have

$$K_1^2 \leq (KK_0)(K_1)$$

and rearranging we have $K \geq K_1/K_0$.

For equality we need each inequality to be an equality. For the first inequality we observe that if

$$f_x = |f_z| + |f_{\overline{z}}| = \frac{1}{2}|f_x - \imath f_y| + \frac{1}{2}|f_x + \imath f_y|$$

then f_x is real and non-negative and f_y is imaginary. For such a map the directions of maximal and minimal stretch will be the x and y-axes. Since f is orientation preserving we must have that $-\imath f_y$ is real and positive. The dilatation is the $\max\{\imath f_x/f_y, \imath f_y/f_x\}$ and the Jacobian is $-\imath f_x f_y$. For the Cauchy-Schwarz inequality to be an equality both functions need to be constant. This implies that f_x and f_y are constant and f is an affine map. $\boxed{2.27}$

Corollary 2.28 *Let $f\colon \Omega_0 \to \Omega_1$ is a diffeomorphism and let $K = \frac{1+\|\mu\|_\infty}{1-\|\mu\|_\infty}$. Then f is K-quasiconformal (in the geometric definition).*

Proof. Let Q be a quadrilateral in Ω_0. We can pre- and post-compose f with conformal maps such that the composition is a map between rectangles as in Lemma 2.27. Note that the norm of the Beltrami differential of the composed map doesn't change. Then $m(Q) = K_0$ and $m(f(Q)) = K_1$. We need to show that $m(f(Q))/K \leq m(Q)$. If $K_1 \leq K_0$ this follows automatically. If not we apply Lemma 2.27. $\boxed{2.28}$

However, there are quasiconformal maps under the geometric definition that are not smooth. We would like to weaken the analytic definition to allow this possibility but we need to be careful. For example we could look at maps f that are differentiable almost everywhere. Then the Beltrami differential μ would be defined almost everywhere and we could ask if $\|\mu\|_\infty < 1$. The following example shows that this definition is to weak.

Let $C \subset [0, 1]$ be the usual Cantor set. Then the Lesbegue measure of C is zero. Fix $K > 1$ and let ν be a measure on C without atoms such that $\nu(C) = K - 1$. Let μ be the sum of ν and Lesbegue measure and define

$$f_0(x) = \int_0^x d\mu.$$

Then f_0 is a homeomorphism from $[0,1]$ to $[0,K]$ and $(f_0)'(x) = 1$ on $[0,1]\backslash C$. We then define a homeomorphism $f\colon [0,1]\times[0,1] \to [0,K]\times[0,1]$ by $f(x,y) = (f_0(x), y)$. The Lesbesgue measure of $C \times [0,1]$ is still zero so f is differentiable almost everywhere. In fact, $f_{\overline{z}} = 0$ almost everywhere. This example shows that between any to quadrilaterals there is a homeomorphism that is conformal almost everywhere. Such a map would not be quasiconformal by the geometric definition.

The problem in the example is that the map f_0 is not *absolutely continuous*. Absolutely continuous functions are differentiable almost everywhere and obey the fundamental theorem of calculus.

Theorem 2.29 *Let $f\colon [a,b] \to \mathbb{R}$ be absolutely continuous. Then f is differentiable almost everywhere and*

$$\int_a^b f'(x)dx = f(b) - f(a).$$

We say that a function $f\colon \Omega_0 \to \Omega_1$ is *absolutely continuous on lines* (or ACL) if f restricted to almost every horizontal or vertical line is absolutely continuous. By Fubini's theorem this implies that f_x and f_y (and hence f_z and $f_{\overline{z}}$) exist almost everywhere. We can now give the analytic definition of K-quasiconformal. The homeomorphism $f\colon \Omega_0 \to \Omega_1$ is *analytically* K-quasiconformal if

1. f is ACL;
2. $|f_{\overline{z}}| \leq k|f_{\overline{z}}|$ almost everywhere with $K = \frac{1+k}{1-k}$.

We have the following important theorem:

Theorem 2.30 *The analytic and geometric definitions of K-quasiconformality are equivalent.*

We can now see that Lemma 2.27 holds for general K-quasiconformal maps. We use the analytic definition.

There are two places that the proof of Lemma 2.27 fails in our bad example. The first is the very first equality

$$K_1 = \int_0^{K_0} f_x(x,y)dx$$

which is the fundamental theorem of calculus. For an ACL function this inequality holds for almost every y. This is enough to carry through the rest of the proof until we get to the area integral

$$\int_0^1 \int_0^{K_0} (|f_z|^2 - |f_{\overline{z}}|^2) dx dy$$

that appears in the Cauchy-Schwarz inequality. This integral should be the area of the second rectangle but in our example it will be the area of the first integral. However, for quasiconformal maps we have the following result (which we will not prove).

Proposition 2.31 *Let $f\colon \Omega_0 \to \Omega_1$ be a quasiconformal map. Then*

$$\int_{\Omega_0} (|f_z|^2 - |f_{\overline{z}}|^2) dx dy$$

is the Euclidean area of Ω_1.

Assuming this we have:

Lemma 2.32 *Lemma 2.27 holds for analytically K-quasiconformal maps.*

Note that this implies that the analytic definition implies the geometric definition.

2.7 Riemann surfaces from polygons

Let P be a finite collection of disjoint polygons in the $\mathbb{C}$ with side pairings given by pure translations or rotations by π. Then the quotient space is a Riemann surface X. Let $V \subset X$ be the image of the vertices of P in X. Outside of V, X has a Euclidean structure. At each point in V the metric is singular. It will be a cone point with cone angle πn for some $n \in \mathbb{Z}^+$. We assume that $n \geq 2$. We will say more about Euclidean cone-structures below.

Given $K \geq 1$ we let $f_K\colon \mathbb{C} \to \mathbb{C}$ be the affine map $f_K(x + \imath y) = Kx + \imath y$ and let $P_K = f_K(P)$. The conjugation of pure translations and rotations by π by f_K will again be pure translations or π-rotations so the side pairing for P conjugate to side pairing for P_K which determines a new Riemann surface X_K. The map f_K descends to a map from X to X_K.

Theorem 2.33 *Let $f\colon X \to X_K$ be a K'-quasiconformal map that is homotopic to f_K. Then $K \leq K'$ with equality if and only if $f = f_K$.*

Proof. After some setup we will see the proof is almost the same as the proof of Lemma 2.27. The one place where we will need a new idea is in the proof of the first inequality; after that the proof is exactly the same.

We start with the setup. We can normalize the area of the polygons so that the total area of P is one. Then the area of P_K will be K. Next we observe that the map f can be thought of as a map between the polygons P and P_K. This map may not be continuous as f may not map sides of P (as arcs in X) to sides of P_K. However, this discontinuity will occur on a set of measure zero. As a map on $P \subset \mathbb{C}$ to $P_K \subset \mathbb{C}$ we can take the derivative f_x, f_z and $f_{\overline{z}}$. Exactly as in the proof of Lemma 2.27 we have

$$\begin{aligned}\left(\int\int_P |f_x| dxdy\right)^2 &\leq \int\int_P \frac{1+|\mu|}{1-|\mu|} dxdy \int\int_P (|f_z|^2 - |f_{\overline{z}}|^2) dxdy \\ &\leq (\text{area}(P)K') \,\text{area}(P_K) = KK'.\end{aligned}$$

To finish the proof we need to show that

$$\text{area}(P_K) \leq \int\int_P |f_x| dxdy.$$

This will require a new approach.

A Euclidean structure on a surface is an atlas of charts to $\mathbb{R}^2 = \mathbb{C}$ with transition maps Euclidean isometries. Note that this is also defines a conformal structure on the surface. A Euclidean structure is very restrictive; the only compact, oriented surface that supports one is the torus. To allow more general surfaces we need to allow *cone points*. We begin by defining a cone point. Let $\mathbb{C}_\infty$ be the universal cover $\mathbb{C}^\times = \mathbb{C}\backslash\{0\}$. Complex multiplication gives $\mathbb{C}^\times$ a group structure and this lifts to a group structure on $\mathbb{C}_\infty$. On $\mathbb{C}^\times$, $\arg z$ is only define modulo 2π while in $\mathbb{C}_\infty$ $\arg z$ is a real number. On $\mathbb{C}^\times$ multiplication by $e^{\imath\theta}$ is determined by θ modulo 2π while on $\mathbb{C}_\infty$ it is determined by θ in $\mathbb{R}$. The Euclidean metric on $\mathbb{C}$ pulls back to a metric on $\mathbb{C}_\infty$ and multiplication by $e^{\imath\theta}$ is an isometry. We let $\mathbb{C}_\theta$ be the metric completion of $\mathbb{C}_\infty/\langle e^{\imath\theta}\rangle$. The metric completion will contain one more point than $\mathbb{C}_\infty/\langle e^{\imath\theta}\rangle$; label this point p_θ.

A *Euclidean cone metric* on a surface Σ with a discrete set of points V is a Euclidean structure on $\Sigma\backslash V$ such that each $v \in V$ has a neighborhood isometric to a neighborhood of p_θ for some $\theta \in \mathbb{R}^+$. Then θ is the *cone angle* at v. The surface we constructed at the beginning of this section is an example of a Euclidean cone metric. Note that it has extra structure as the transition maps are pure translations or π-rotations.

Lemma 2.34 *Assume that X has a Euclidean cone metric with an atlas such all transition maps are pure translations or π-rotations. Then all cone angles are multiples of π.*

Theorem 2.35 *Let P be a polygon in $\mathbb{R}^2$ with angles $\{\theta_1, \dots, \theta_k\}$. Then $\sum(\pi - \theta_i) = 2\pi$.*

We will be especially interested in cone metrics where all cone angles are $\geq 2\pi$. In many ways these surfaces behave like hyperbolic surfaces.

Given a piecewise smooth path $\gamma\colon I \to X$ we can measure its length in the usual way on the non-singular part of the surface and we let this be the length of the path and denote it $L_X(\gamma)$. We can the define a metric d_X on the surface, again in the usual way. We say that γ is a *geodesic* if every $t \in I$ has a neighborhood $I_t \subset I$ such that for all $a, b \in I_t$ we have $d_X(\gamma(a), \gamma(b)) = |a - b|$.

Note that if we have two Euclidean rays in $\mathbb{C}_\theta$ that are based at p_θ they make two angles, one in the clockwise direction and one in the counter clockwise direction.

Lemma 2.36 *A path γ on a Euclidean cone metric is a geodesic if and only if in each Euclidean chart it is a straight line and the two angles at each cone point are $\geq \pi$.*

We need to generalize this theorem to Euclidean polygons with cone points.

Theorem 2.37 *Let P be a polygon with cone points of angles $\{\alpha_1, \dots, \alpha_k\}$ and angles $\{\theta_1, \dots, \theta_j\}$. Then $\sum(2\pi - \alpha_i) + \sum(\pi - \theta_i) = 2\pi$.*

Proof. Triangulate P such that all the cone points and then double to form a triangulated sphere S. Let V, E and T be the number of vertices, edges and triangles in the triangulation. Then $E = 3T/2$ and $V - E + T = 2$. Together this implies that $2V = 4 + T$. Let β_i^j be the angles at the jth vertex. Note that the total sum of all the angles is πT. Therefore

$$\begin{aligned}
\sum_j \left(2\pi - \sum_i \beta_i^j\right) &= 2\pi V - \pi T \\
&= \pi(4 + T - T) \\
&= 4\pi.
\end{aligned}$$

There are three types of vertices: vertices corresponding to cone points of P, vertices on the interior of P that are not cone points and vertices on the boundary of P. For each cone vertex of P there will be two vertices in S. For each interior vertex the sum $2\pi - \sum_i \beta_i^j$ is zero. Each boundary vertex of P will appear only once in S. Therefore we have

$$\sum(2\pi - \alpha_i) + \sum(\pi - \theta_i) = \frac{1}{2}\sum_j \left(2\pi - \sum_i \beta_i^j\right) = 2\pi.$$

$\boxed{2.37}$

Theorem 2.38 *Let X be a complete Euclidean cone metric with all cone angles $\geq 2\pi$. Then every arc on X is homotopic (rel endpoints) to a unique geodesic.*

Proof. We can assume that X is simply connected. If not we replace X with its universal cover which will also be complete. Every arc downstairs will lift to an arc in the universal cover and if we can homotop it to a geodesic in the cover we can do so downstairs also. Furthermore in the universal cover there is a unique homotopy class of arcs between any two points.

We first show that geodesics exist. This is a standard argument that needs to be slightly modified to account for the cone points. In general, on a complete Riemannian manifold, a point and a unit tangent vector determines a unique complete (defined for all $\mathbb{R}$) geodesic. On a cone-manifold this uniqueness fails when the geodesic hits a cone point.

We begin with the standard argument. For every $z \in X$ there is an $\epsilon > 0$ such that the ϵ-neighborhood of z is standard ball either in $\mathbb{C}$ or if z is a smooth point or an ϵ-neighborhood of p_θ in $\mathbb{C}_\theta$ if z is a cone point of angle θ. In particular there is a unique geodesic between every point in this ball and z. Now let z_0 and z_1 be points in X. Let S be the sphere bounding the ϵ-ball centered at z_0. Note that if $d_X(z_0, z_1) \leq \epsilon$ then we are done. If not there is a unique $z \in S$ that is closest to z_1. Let γ be the geodesic ray based at z_0 that goes through z and extend this ray as far as possible until it either hits z_1, in which case we stop, or a cone point. Every time that the ray hits a cone point we repeat the construction replacing z_0 with the cone point.

We claim that this path is a geodesic and that it will eventually hit z_1. To see this let $I \subset [0, d(z_0, z_1)]$ be the largest interval, with left endpoint 0, such that for $t \in I$, $d(\gamma(t), z_1) = d(z_0, z_1) - t$. We first observe that I contains $[0, \epsilon]$ and, since the distance function is continuous, I is closed. We need to show that I is also open. For this we pick a $t \in I$ and construct a geodesic ray β, as we constructed γ, with z_0 replaced with $\gamma(t)$. As for γ, for s in a neighborhood of 0 we have $d(\beta(s), z_1) = d(\beta(0), z_1) - s$. If we can show that $\gamma(s) = \beta(s-t)$ then we have $d(\gamma(s), z_1) = d(z_0, z_1) - s$ for s in a neighborhood of t and I is open. If $\gamma(t)$ is a cone point then by construction $\gamma(s) = \beta(s - t)$. When $\gamma(t)$ is not a cone point we need to show that the angle between γ and β at $\gamma(t)$ is π. If this is not the case then we can shorten the path $\gamma|_{[0,t]} \cup \beta$ to see that for $s < t$ near t, $d(\gamma(s), z_1) < d(z_0, z_1) - s$, contradicting $s \in I$.

The more interesting part of the proof is uniqueness. Let γ_0 and γ_1 be two geodesic connecting z_0 and z_1. After a surgery argument we can assume that they intersect only at z_0 and z_1. Since they are homotopic rel endpoints the will then bound a polygon P with cone points. Except for possibly the angles at z_0 and z_1 (which we label θ_0 and θ_1) all other angles will be $\geq \pi$. The cone angles in the interior of P are all $\geq 2\pi$ so in the formula from Theorem 2.37 all terms are ≤ 0 except for possible terms with θ_0 and θ_1 so we have $(\pi - \theta_0) + (\pi - \theta_1) \leq 2\pi$. Since $\theta_0, \theta_1 \geq 0$ this implies that $\theta_0 = \theta_1 = 0$.

Therefore γ_0 and γ_1 agree in a neighborhood of z_0 and z_1, a contradiction. $\boxed{2.38}$

We now return to our surfaces X and X_K that we defined at the beginning of the section. A *horizontal geodesic* is a geodesic that is parallel to the x-axis in every chart. (This is well defined since transition maps are pure translations or π-rotations which take lines parallel to the x-axis to lines parallel to the x-axis.)

Lemma 2.39 *Let $f\colon X \to X_K$ be a homeomorphism homotopic to f_K. Then there exists an $M = M(f)$ such that for all horizontal geodesics γ on X we have*

$$L_{X_K}(f \circ \gamma) \geq K L_X(\gamma) - M.$$

Proof. We first observe that for horizontal geodesics $L_{X_K}(f_K \circ \gamma) = K L_X(\gamma)$.

We work in the universal cover where there is a unique homotopy class of arc between any two points.

The homotopy from f to f_K lifts to the universal covers $\tilde{X}$ and $\tilde{X}_K$. In particular there are lifts $\tilde{f}$ and $\tilde{f}_K$ that are equivariantly homotopic. Therefore the function $\tilde{z} \mapsto d_{\tilde{X}_K}(\tilde{f}(\tilde{z}), \tilde{f}_K(\tilde{z}))$ is a continuous equivariant function $\tilde{X}$ which descends to a function on the compact surface X. This implies that the function is bounded by some constant, say $M/2$.

Let $\tilde{\gamma}$ be a lift of γ to $\tilde{X}$. Then $\tilde{\gamma}$ is the unique geodesic between its endpoints which we label z_0 and z_1 so $L_X(\gamma) = L_{\tilde{X}}(\tilde{\gamma}) = d_{\tilde{X}}(z_0, z_1)$. We also have $L_{X_K}(f_K \circ \gamma) = L_{\tilde{X}_K}(\tilde{f}_K \circ \gamma) = d_{\tilde{X}}(\tilde{f}_K(z_0), \tilde{f}_K(z_1))$. We can now apply the triangle inequality to see that

$$d_{\tilde{X}_K}(\tilde{f}_K(z_0), \tilde{f}_K(z_1)) \leq d_{\tilde{X}_K}(\tilde{f}_K(z_0), \tilde{f}(z_0)) + d_{\tilde{X}_K}(\tilde{f}(z_0), \tilde{f}(z_1)) + d_{\tilde{X}_K}(\tilde{f}(z_1), \tilde{f}_K(z_1)).$$

The first and last term on the right are bounded by $M/2$ and $d_{\tilde{X}_K}(\tilde{f}(z_0), \tilde{f}(z_1)) \leq L_{\tilde{X}_K}(f \circ \tilde{\gamma}) = L_{X_K}(f \circ \gamma)$. Combining this last sentence with the previous inequality gives the lemma. $\boxed{2.39}$

We can now complete the proof of Theorem 2.33. Let $\hat{X}$ be the set of horizontal unit tangent vectors on X. Then $\hat{X}$ is a double cover of the non-singular part of X. (In fact $\hat{X}$ could just be two copies of X.) There is a natural flow $F_t\colon \hat{X} \to \hat{X}$. We define F_t as follows. Each $v \in \hat{X}$ is a horizontal unit tangent vector at a point $z \in X$. There is a unique horizontal geodesic γ on X with $\gamma(0) = z$ and $\gamma'(0) = v$. Generically, γ will be defined for all $\mathbb{R}$ but there will be a measure zero set where the horizontal geodesic limits in the forward or backward (or both) directions to a cone point. Ignoring this set of measure zero we define $F_t(v)$ to be the tangent vector of γ at $\gamma(t)$.

The Euclidean structure on X lifts to a Euclidean structure on $\hat{X}$. We have the following lemma:

Lemma 2.40 *The flow F_t is area preserving.*

Proof. Let $R \subset \hat{X}$ be a rectangle with sides horizontal and vertical geodesics and assume the interior of R doesn't contain cone points. In $\hat{X}$ all cone points will have cone angles that are multiples of 2π (instead of π as in X). If the cone angle is $2\pi n$ then there will be n horizontal geodesics that flow forward into the cone point. If we flow backwards from the cone points for time t we will form horizontal geodesics of length t (or less if the backwards flow hits a cone point). If there are k cone points of cone angle $2\pi n_i$ then total number of geodesic segments will be $\sum n_i$ and the total length will be (at most $t \sum n_i$). Therefore the intersection of these segments with R will have finitely many components: There can be at must $\sum n_i$ segments that don't cross the entire rectangle. The number for segments that do bound the entire rectangle is bounded by $t \sum n_i$ divided by the horizontal length of the rectangle.

We extend those components of the intersection with R that don't cross the rectangle so that they do. This now splits the rectangle into finitely many rectangles with disjoint interior whose areas sum to the area of R. For each of these smaller rectangles F_t is an isometry, and hence area preserving, on the interior. This implies that $\text{area}(F_t(R)) = \text{area}(R)$ and the lemma follows. $\boxed{2.40}$

We are interested in the integral

$$\int \int_X |f_x| dxdy.$$

Let $\hat{f}$ be the lift of the function $|f_x|$ to $\hat{X}$. Then

$$\int \int_{\hat{X}} \hat{f} dxdy = 2 \int \int_X |f_x| dxdy.$$

We will bound below the integral on the left. Since F_t is area preserving for all $t \in \mathbb{R}$ we have

$$\int \int_X \hat{f} \circ F_t dxdy = \int \int_X \hat{f} dxdy.$$

If we integrate the left hand integral with respect to t from $-D$ to D we have

$$\int_{-D}^{D} \left(\int \int_X \hat{f} \circ F_t dxdy \right) dt = 2D \int \int_X \hat{f} dxdy.$$

We will change the order of integration so that we are integrating with respect to t first. Before we do this we observe that if $\gamma \colon [a, b] \to X$ is a horizontal geodesic with v the tangent vector at $\gamma(0)$ then

$$\begin{aligned} L_{X_K}(\gamma \circ f) &= \int_a^b |f_x \circ \gamma(t)| dt \\ &= \int_a^b \hat{f} \circ F_t(v) dt. \end{aligned}$$

Since $L_X(\gamma) = b - a$ for almost every $v \in \hat{X}$ we have

$$2DK - M \leq \int_{-D}^{D} f \circ F_t(v) dt.$$

We now put everything together

$$\begin{aligned} 4D \int\int_X |f_x| dxdy &= \int_{-D}^{D} \left(\int\int_X \hat{f} \circ F_t dxdy \right) dt \\ &= \int\int_{\hat{X}} \left(\int_{-D}^{D} f \circ F_t(v) dt \right) dxdy \\ &\geq \int\int_{\hat{X}} (2DK - M) dxdy \\ &= 2(2DK - M). \end{aligned}$$

Dividing both sides by D and taking the limit as $D \to \infty$ we have

$$\int\int_X |f_x| dxdy \geq K.$$

This is exactly the inequality we need to complete the proof of Theorem 2.33. $\boxed{2.33}$

2.7.1 (p, q)-differentials

Let $\mathcal{I}$ be the index set of an atlas $\mathcal{A}$ for a Riemann surface X. A (p, q)-differential is a collection of functions $\Phi = \{\phi_\alpha\}_{\alpha \in \mathcal{I}}$ where ϕ_α are functions on $\psi_\alpha(U_\alpha) \subset \mathbb{C}$ and on $\psi_\beta(U_\alpha \cap U_\beta)$ we have $\phi_\alpha \circ \psi_{\alpha\beta}(z)(\psi_{\alpha\beta})_z^p(z)\overline{(\psi_{\alpha\beta})_z^q(z)} = \phi_\beta(z)$. (Or we could write $\phi(z)dz^p d\bar{z}^q$ in z-coordinates becomes $\phi(z(w))z'(w)^p \overline{z'(w)}^q dw^p d\bar{w}^q$ in w-coordinates where we view z as a holomorphic function in w.)

Proposition 2.41 *Let $\Phi = \{\phi_\alpha\}_{\alpha \in \mathcal{I}}$ be a (p, q)-differential and $\Theta = \{\theta_\alpha\}_{\alpha \in \mathcal{I}}$ a (r, s)-differential. Then*

1. *$\Phi\Theta = \{\phi_\alpha \theta_\alpha\}_{\alpha \in \mathcal{I}}$ is a $(p + r, q + s)$-differential;*

2. $\Phi/\Theta = \{\phi_\alpha/\theta_\alpha\}_{\alpha\in\mathcal{I}}$ *is a* $(p-r, q-r)$*-differential;*

3. $\bar{\Phi} = \{\bar{\phi}_\alpha\}_{\alpha\in\mathcal{I}}$ *is a* (q,p)*-differential;*

4. *if* $p = 0$ *then* $\Phi_z = \{(\phi_\alpha)_z\}_{\alpha\in\mathcal{I}}$ *is a* $(1,q)$*-differential;*

5. *if* $q = 0$ *then* $\Phi_{\overline{z}} = \{(\phi_\alpha)_{\overline{z}}\}_{\alpha\in\mathcal{I}}$ *is a* $(p,1)$*-differential.*

The (p,q)-differentials form a vector space. By changing the class of functions that we allow the ϕ_α to be we can get different spaces. For example in the case of $(2,0)$-differentials we will restrict to holomorphic functions. This space of *holomorphic quadratic differentials* is extremely important in Teichmüller theory.

There are some important relationships between vector fields and differential forms and with (p,q)-differentials.

Lemma 2.42

- *A vector field is a* $(-1,0)$*-differential.*
- *A* 1*-form is a* $(1,0)$*-differential.*
- *A* 2*-form is a* $(1,1)$*-differential.*
- *A Beltrami differential is* $(-1,1)$*-differential.*

2.7.2 Quadratic differentials

Let Φ be a holomorphic $(1,0)$-form. (Φ is an *Abelian differential.*) Given a piecewise smooth path $\gamma\colon [a,b] \to X$ we can take the integral $\int_\gamma \Phi$. In local coordinates this is just a contour integral. In particular if the image of a γ is contained in a chart (U,ψ) and $\phi\colon \psi(U) \to \mathbb{C}$ is the function representing Φ in the chart then

$$\int_\gamma \Phi = \int_{\psi\circ\gamma} \phi(z)dz.$$

In general we calculate $\int_\gamma \Phi$ by breaking the arc into sub-arcs that lie in charts.

Lemma 2.43 *Let* γ_0 *and* γ_1 *be arcs in* X *that are homotopic rel endpoints. Then* $\int_{\gamma_0} \Phi = \int_{\gamma_1} \Phi$.

Proof. This is essentially Cauchy's theorem. The homotopy from γ_0 to γ_1 is a map on a square S into X where two of the parallel sides represent to the two arcs. The map is constant on the other two sides of the square since the homotopy is rel endpoints. If we orient ∂S so that it traverses the γ_0-side in the positive direction and the γ_1-side in the negative direction we have

$$\int_{\partial S} \Phi = \int_{\gamma_0} \Phi - \int_{\gamma_1} \Phi$$

since the integral on the two sides where the map is constant is zero.

We now break the square into rectangles with disjoint interior so that each rectangle is contained in a chart. By Cauchy's theorem, the integral around the perimeter of each rectangle is zero. The integral $\int_{\partial S} \Phi$ is the sum of the integrals around the perimeter (suitably oriented) of each rectangle. The lemma follows. $\boxed{2.43}$

On simply connected neighborhood $U \subset X$, all arcs with the same endpoints are homotopic rel endpoints. After fixing a basepoint we can use this to define a function on U. If Φ is holomorphic then this function will be holomorphic and we can understand its local behavior through the zeros of its derivative. Note that the zeros (and their order) of Φ (or any (p,q)-differential) are well defined independent of chart. If Φ is holomorphic and non-constant (and X is connected) then the zeros are isolated.

Lemma 2.44 *Let $U \subset X$ be simply connected and open and $z_0 \in U$ a basepoint. For $z \in U$ let γ_z be an arc in U from z_0 to z. Then*

$$z \mapsto \int_{\gamma_z} \Phi$$

is a holomorphic function. The derivative of the function has a zero of order n at z if and only if Φ has a zero of order n at z.

The key point is that where Φ is non-zero we can use it define an atlas where the translations maps are pure translations of the Euclidean metric. For each $z \in X$ where Φ is non-zero we choose a a neighborhood U_x that is simply connected such that the function ψ_z from Lemma 2.44 is injective on U_z. Let $\mathcal{A}_z = \{(U_z, \psi_z)\}$ where z varies over the points in X where Φ is non-zero.

Proposition 2.45 *Let Φ be a (non-constant) holomorphic quadratical differential on X and let $V \subset X$ be the zeros of Φ. Then there is a conformal atlas $\mathcal{A}$ for $X \backslash V$ such that for each chart $(U_\alpha, \psi_\alpha) \in \mathcal{A}$ we have $\phi_\alpha \equiv 1$. The transition maps for $\mathcal{A}$ are restrictions of pure translations or π-rotations.*

Proof. Given a point $z_0 \in X \backslash V$ choose a chart (U_β, ψ_β) that contains z_0. We can assume that U_β is simply connected and disjoint from V. (If not we just shrink U_β.) Let ϕ_β be the function representing Φ on (U_β, ψ_β). Then define $\psi_\alpha \colon U_\beta \to \mathbb{C}$ as follows. Fix a branch of the square root on $\psi_\beta(U_\beta)$ so that $\sqrt{\phi_\beta}$ is well defined and for each $z \in U_\beta$ choose an arc γ in U connecting z_0 to z_1. Then let

$$\psi_\alpha(z) = \int_{\psi_\beta \circ \gamma} \sqrt{\phi_\beta(w)} dw.$$

By Lemma 2.43, $\psi(z)$ doesn't depend on the choice of γ and will be a holomorphic function with non-zero derivative at z_0. Therefore we can choose a neighborhood $U_\alpha \subset U_\beta$ of z_0 where ψ_α is injective so (U_α, ψ_α) is a chart. The derivative of $\psi_\alpha \circ \psi_\beta^{-1}$ at $z \in \phi_\beta(U)$ is $\sqrt{\phi_\beta(z)}$ so we have

$$\phi_\alpha(\psi_\alpha \circ \psi_\beta^{-1}(z)) \left(\sqrt{\phi_\beta(z)}\right)^2 = \phi_\beta(z)$$

and therefore $\phi_\alpha \equiv 1$ and we have constructed the required atlas.

For any atlas that supports a quadratic differential where all functions are $\equiv 1$ the square of the derivative of the transition map will be $\equiv 1$ so the derivative itself is $\equiv \pm 1$. This implies that the transition maps are restrictions of pure translations or π-rotations. $\boxed{2.45}$

This defines a Euclidean metric on $X \backslash V$. We need to show that this a cone-metric.

Proposition 2.46 *Let $z_0 \in V$ be a zero of Φ of order n. Then the Euclidean metric on $X \backslash V$ extends to a cone-point of angle $(n+2)\pi$ at z_0.*

Proof. Choose a chart (U, ψ) with $z_0 \in U$, $\psi(z_0) = 0$ and Φ in (U, ψ) given by the function $\phi(z) = \left(\frac{n}{2} + 1\right)^2 z^n$. We can assume that $\phi(U)$ is an open disk centered at 0 of radius ϵ which we label Δ_ϵ. Let $\Delta_\epsilon^\times = \Delta_\epsilon \backslash \{0\}$ and let $\pi \colon \widetilde{\Delta_\epsilon^\times} \to \Delta_\epsilon^\times$ be the universal cover. Note that $\widetilde{\Delta_\epsilon^\times}$ is homeomorphic to $\mathbb{R}^2$ but it is natural to give $\widetilde{\Delta_\epsilon^\times}$ coordinates $(0, \epsilon^{\frac{\alpha}{2\pi}}) \times \mathbb{R}$ so that $\pi(r, \theta) = r^{\frac{2\pi}{\alpha}} e^{\imath 2\pi\theta/\alpha}$ where $\alpha = (n+2)\pi$. In these coordinates the deck transformations are generated by the map $(r, \theta) \mapsto (r, \theta + \alpha)$.

The surface $\widetilde{\Delta_\epsilon^\times}$ has a conformal structure lifted from the conformal structure on $\Delta_\epsilon^\times$. For this lifted conformal structure the covering map will always be locally conformal (holomoprhic and locally injective). We can also lift the quadratic differential Φ to a quadratic differential $\tilde{\Phi}$ on $\widetilde{\Delta_\epsilon^\times}$. We would like to find an atlas for $\widetilde{\Delta_\epsilon^\times}$ where the quadratic differential is identically one. To do so we define a map $\tilde{\psi}(r, \theta) = re^{\imath\theta}$. This map is locally conformal. To see this we observe that for any simple connected neighborhood in $\Delta_\epsilon^\times$ and any inverse of π on this neighborhood $\tilde{\psi} \circ \pi^{-1}$ is holomorphic. In particular it is a branch of the map $z \mapsto z^{\frac{\alpha}{2\pi}}$.

We can use the map $\tilde{\psi}$ to define charts on $\widetilde{\Delta_\epsilon^\times}$ by restricting $\tilde{\psi}$ to neighborhoods in $\widetilde{\Delta_\epsilon^\times}$ where $\tilde{\psi}$ is injective. On such charts we will show that $\tilde{\Phi}$ is identically one. For charts given by the covering map π, $\tilde{\Phi}$ is $\left(\frac{n}{2} + 1\right)^2 z^n$. The transition map is given by $z \mapsto z^{\frac{\alpha}{2\pi}} = z^{\frac{n}{2}+1}$ so the derivative of the transition map is $\left(\frac{n}{2} + 1\right) z^{\frac{n}{2}}$. (This expression is well defined up to sign.) Therefore $\tilde{\Phi}$ in a chart coming from the map $\tilde{\psi}$ is a function

$\tilde{\phi}$ with

$$\tilde{\phi}(z)\left(\left(\frac{n}{2}+1\right)z^{\frac{n}{2}}\right)^2 = \left(\frac{n}{2}+1\right)^2 z^n$$

which implies that $\tilde{\phi}(z) = 1$.

The holomorphic quadratic differential $\tilde{\Phi}$ on $\widetilde{\Delta_\epsilon^\times}$ defines a Euclidean metric on $\widetilde{\Delta_\epsilon^\times}$. The map π (which we can view as a map to $X\backslash V$) will be a local isometry to the Euclidean metric on $X\backslash V$ given by Φ. The map $\tilde{\psi}$ is a local isometry to the Euclidean metric on $\mathbb{C}^\times \subset \mathbb{C}$.

$\widetilde{\Delta_\epsilon^\times}$ is naturally a subspace of $\mathbb{C}_\infty$ where we give $\mathbb{C}_\infty$ coordinates $\mathbb{R}^+ \times \mathbb{R}$. The quotient map $\mathbb{C}_\infty \mapsto \mathbb{C}_\infty/\langle e^{\imath\alpha}\rangle$ can be written as a composition of π and a injective map from $\Delta_\epsilon^\times$ to $\mathbb{C}_\alpha$. This last map must be an isometry (as the other two maps are) completing the proof of the proposition. $\boxed{2.46}$

Let Φ be a non-constant holomorphic quadratic differential on X with $V \subset X$ the set of zeros. Recall that X is an atlas on the topological surface Σ where the transition maps are conformal. By Proposition 2.45 the quadratic differential Φ defines an atlas $\mathcal{A}$ on $\Sigma\backslash V$ that is a conformal atlas for $X\backslash V$. We define a new atlas $\mathcal{A}_K$ on $\Sigma\backslash V$ by taking each chart (U_α, ψ_α) and post-composing ψ_α with the map $x + \imath y \mapsto Kx + \imath y$ to form a new chart $(U_\alpha, \psi_\alpha^K)$. The transition maps for this new atlas will be restrictions of translations and π-rotations. In particular this defines a conformal structure on $\Sigma\backslash V$. We also have a map $f_K \colon X\backslash V \to X_K$ given by taking the identity map on $\Sigma\backslash V$. Of course this maps extends to all of Σ and we can use this to extend X_K to a conformal structure on all of Σ.

Theorem 2.47 *Let Φ be a holomorphic quadratic differential on X. Let $f \colon X \to X_K$ be K'-quasiconformal and assume that $f \sim f_K$. Then $K' \geq K$ with equality if and only if $f = f_K$.*