HYPERBOLIC STRUCTURES ON SURFACES

Javier Aramayona
School of Mathematics, Statistics and Applied Mathematics
National University of Ireland, Galway. IRELAND
Email: Javier.Aramayona@nuigalway.ie

We give a brief introduction to hyperbolic structures on surfaces. Using the concepts of developing map and holonomy, we sketch a proof that every surface equipped with a complete hyperbolic metric is isometric to a quotient of $\mathbb{H}$ by a Fuchsian group. We then define Teichmüller spaces and explain Fenchel-Nielsen coordinates. Finally, we introduce mapping class groups and show that they act properly discontinuously on Teichmüller space.

1. Introduction

This paper is intended as a brief introduction to hyperbolic structures on surfaces, Teichmüller spaces and mapping class groups. It is based on the first half of the course "Hyperbolic structures on surfaces", given by C. Leininger and the author during the programme "Geometry, Topology and Dynamics of Character Varieties" at the Institute for Mathematical Sciences of Singapore in July 2010. It accompanies the article [23], also in this volume, which discusses degenerations of hyperbolic structures.

In order to keep the exposition as concise and self-contained as possible, we have narrowed our attention to three particular strands. First, that a surface S equipped with a complete hyperbolic structure may be identified with a quotient of $\mathbb{H}$ by a torsion-free Fuchsian group, via the developing map. Second, that the Teichmüller space $\mathcal{T}(S)$, that is, the space of complete hyperbolic structures on S, is homeomorphic to some $\mathbb{R}^n$, where n depends only on the topology of S. Finally, that the mapping class group $\mathrm{Mod}(S)$ of S, that is, the group of self-homeomorphisms of S up to homotopy, acts properly discontinuously on $\mathcal{T}(S)$.

The plan of the paper is as follows. In Section 2 we recall some basic facts about plane hyperbolic geometry. In Section 3 we introduce the notion of a hyperbolic structure on a surface, and explain why every complete hyperbolic surface is isometric to $\mathbb{H}/\Gamma$, where Γ is a torsion-free Fuchsian group. In Section 4 we define Teichmüller spaces and describe Fenchel-Nielsen coordinates. Finally, in Section 5 we introduce the mapping class group and prove that it acts properly discontinuously on Teichmüller space.

Acknowledgements. I would like to express my gratitude to the Institute for Mathematical Sciences of Singapore for their hospitality and financial support. The vast majority of this gratitude is targetted at Ser Peow Tan for all his help, support and friendship. I would like to say thanks to Chris Leininger for help and discussions. Thanks also to David Dumas, Jonah Gaster, Add Jaipong, Thomas Koberda, Brice Loustau and Sara Maloni for interesting conversations. Finally, I would like to thank whoever came up with the idea of a "pepper crab", for (s)he deserves a place in heaven.

2. Plane hyperbolic geometry

We refer the reader to [2, 4, 5, 8, 19, 20, 24, 26] for a detailed discussion of the topics presented in this section.

2.1. *Möbius transformations*

Let $\bar{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ be the Riemann sphere. A Möbius transformation is a map $T : \bar{\mathbb{C}} \to \bar{\mathbb{C}}$ of the form

$$T(z) = \frac{az+b}{cz+d},$$

where $a, b, c, d \in \mathbb{C}$ and $ad - bc \neq 0$. Here, $T(\infty) = \frac{a}{c}$ and $T\left(\frac{-c}{d}\right) = \infty$.

Denote by $\text{Möb}(\bar{\mathbb{C}})$ the set of all Möbius transformations. Every element of $\text{Möb}(\bar{\mathbb{C}})$ is a bijection; the inverse of T is

$$T^{-1}(z) = \frac{dz-b}{-cz+a} \in \text{Möb}(\bar{\mathbb{C}})$$

Moreover, the composition of two Möbius transformations is a Möbius transformation, and thus $\text{Möb}(\bar{\mathbb{C}})$ is a group under composition.

To every Möbius transformation

$$T(z) = \frac{az+b}{cz+d}$$

we may associate a matrix of non-zero determinant, namely

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Multiplying the matrix by a non-zero complex number does not change the Möbius transformation it represents, and thus there is a surjective map $\mathrm{SL}(2,\mathbb{C}) \to \mathrm{Möb}(\bar{\mathbb{C}})$. It is easy to verify that this map is in fact a homomorphism, with kernel $\{\pm I\}$. Therefore,

$$\mathrm{Möb}(\bar{\mathbb{C}}) \cong \mathrm{PSL}(2,\mathbb{C}).$$

We will make use of the following important properties of Möbius maps; for a proof, see for instance ([19], Thm. 2.4.1 and 2.11.3).

Proposition 2.1: *(1) Every element of* $\mathrm{Möb}(\bar{\mathbb{C}})$ *is conformal.*
(2) Let L *be either a Euclidean circle or a Euclidean line in* $\mathbb{C}$*, and let* $T \in \mathrm{Möb}(\bar{\mathbb{C}})$*. Then* $T(L)$ *is either a Euclidean circle or a Euclidean line in* $\mathbb{C}$*.* $\square$

2.1.1. *Classification in terms of trace and fixed points*

The *trace* of the Möbius transformation $T(z) = \frac{az+b}{cz+d}$ is $\mathrm{tr}(T) := a + d$. Observe that $\mathrm{tr}(T)$ is only defined up to sign; however, $\mathrm{tr}^2(T) = (a + d)^2$ is well-defined and thus yields a function $\mathrm{tr}^2 : \mathrm{PSL}(2,\mathbb{C}) \to \mathbb{C}$ that is continuous with respect to the natural topology on $\mathrm{PSL}(2,\mathbb{C})$, and is constant on each conjugacy class.

If $T \in \mathrm{Möb}(\bar{\mathbb{C}})$ is not the identity, then its fixed points are given by

$$z = \frac{(a-d) \pm \sqrt{\mathrm{tr}^2(T) - 4}}{2c}.$$

Therefore, T has exactly one fixed point if and only if $\mathrm{tr}^2(T) = 4$; otherwise it has two.

If T has exactly one fixed point, then it is called *parabolic*. A parabolic transformation is conjugate in $\mathrm{Möb}(\bar{\mathbb{C}})$ to $z \to z + 1$.

If T has two fixed points then, up to conjugation in $\mathrm{Möb}(\bar{\mathbb{C}})$, $T(z) = \lambda z$ for some $\lambda \in \mathbb{C} \setminus \{0, 1\}$. The number λ is called the *multiplier* of T; note that the multiplier is also a conjugacy invariant, for

$$\mathrm{tr}^2(T) = \lambda + \lambda^{-1} + 2. \tag{2.1}$$

If $|\lambda| = 1$ then T is called *elliptic*; observe that T is elliptic if and only if $\mathrm{tr}^2(T) \in [0, 4)$. Otherwise, T is called *loxodromic*. In the special case that $\lambda \in \mathbb{R}$, T is called *hyperbolic*; observe that T is hyperbolic if and only if $\mathrm{tr}^2(T) > 4$.

2.2. *Models for hyperbolic geometry.*

We will consider two (equivalent) models for plane hyperbolic geometry. The first is the *upper half-plane*

$$\mathbb{H} = \{x + iy \in \mathbb{C} \mid y > 0\},$$

equipped with the Riemannian metric

$$ds^2_{\mathbb{H}} = \frac{dx^2 + dy^2}{y^2}.$$

The second is the *Poincaré disc*, namely the open unit disc

$$\mathbb{D} = \{x + iy \in \mathbb{C} \mid x^2 + y^2 < 1\}$$

in the complex plane, equipped with the Riemannian metric

$$ds^2_{\mathbb{D}} = \frac{4(dx^2 + dy^2)}{(1 - (x^2 + y^2))^2}.$$

2.2.1. *Hyperbolic distance*

Let $\gamma : [a, b] \to \mathbb{H}$ be a piecewise differentiable path. The *hyperbolic length* of γ is defined as

$$l_{\mathbb{H}}(\gamma) = \int_\gamma ds_{\mathbb{H}}.$$

Given $z, w \in \mathbb{H}$, define the *hyperbolic distance* between z and w by

$$d_{\mathbb{H}}(z, w) = \inf\{l_{\mathbb{H}}(\gamma) \mid \gamma \text{ is a piecewise differentiable path from } z \text{ to } w\}$$

Hyperbolic distance in $d_{\mathbb{D}}$ in $\mathbb{D}$ is defined in an analogous way. It is an easy exercise to verify that $(\mathbb{H}, d_{\mathbb{H}})$ and $(\mathbb{D}, d_{\mathbb{D}})$ are metric spaces.

2.2.2. *Möbius transformations act by isometries*

Let $\text{Möb}(\mathbb{H})$ (resp. $\text{Möb}(\mathbb{D})$) be the subgroup of $\text{Möb}(\bar{\mathbb{C}})$ consiting of all Möbius transformations that preserve $\mathbb{H}$ (resp. $\mathbb{D}$). As one may easily verify,

$$\text{Möb}(\mathbb{H}) \cong \text{PSL}(2, \mathbb{R}) \quad \text{and} \quad \text{Möb}(\mathbb{D}) \cong \text{PSU}(2, \mathbb{C}).$$

Lemma 2.2: *Every element of* $\text{Möb}(\mathbb{H})$ *(resp.* $\text{Möb}(\mathbb{D})$*) is an isometry of* $\mathbb{H}$ *(resp.* $\mathbb{D}$*).*

Proof: We prove the result for $\mathbb{H}$, as the one for $\mathbb{D}$ is obtained in an analogous way. Let $\gamma : I \to \mathbb{H}$ be a piecewise differentiable path, and let

$$T(z) = \frac{az+b}{cz+d} \in \text{Möb}(\mathbb{H}).$$

Write $w = T(z)$, and observe that

$$\text{Im}(w) = \frac{\text{Im}(z)}{|cz+d|^2},$$

where $\text{Im}(z)$ denotes the imaginary part of z. Then:

$$l_{\mathbb{H}}(T(\gamma)) = \int_{T(\gamma)} \frac{|dw|}{\text{Im}(w)} = \int_{\gamma} \frac{|cz+d|^2}{\text{Im}(z)} \cdot \frac{|dz|}{|cz+d|^2} = l_{\mathbb{H}}(\gamma),$$

as desired. □

2.2.3. *The Cayley transformation*

The *Cayley transformation* is the Möbius map

$$C(z) = \frac{z-i}{z+i} \in \text{Möb}(\bar{\mathbb{C}}).$$

It is easy to check that $C(\mathbb{H}) = \mathbb{D}$. Using a similar calculation to that of Lemma 2.2, we obtain:

Lemma 2.3: *The Cayley transformation $C : \mathbb{H} \to \mathbb{D}$ is an isometry.* □

2.2.4. *Hyperbolic geodesics*

A piecewise differentiable path in $\mathbb{H}$ or $\mathbb{D}$ is said to be *geodesic* if the length of any of its segments realizes the distance between the endpoints. The following gives a full description of geodesics in $\mathbb{H}$ and $\mathbb{D}$.

Proposition 2.4: *(i) The geodesics in $\mathbb{H}$ are either vertical Euclidean lines or Euclidean semicircles perpendicular to $\mathbb{R}$.*
(ii) The geodesics in $\mathbb{D}$ are either diameters of $\mathbb{D}$ or arcs of Euclidean semicircles perpendicular to $\mathbb{S}^1$.

Proof: (i) Let $z, w \in \mathbb{H}$. Suppose first that $z, w \in i\mathbb{R}$; thus, up to relabelling, $z = ip$ and $w = iq$ where $p < q$. Let $\gamma : [a, b] \to \mathbb{H}$ be a piecewise differentiable path from z to w, where $\gamma(t) = (x(t), y(t))$. Then:

$$l_{\mathbb{H}}(\gamma) = \int_a^b \frac{1}{y(t)} \sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2} dt \geq \int_a^b \frac{1}{y(t)} \frac{dy}{dt} dt = \int_p^q \frac{dy}{y} = \log\left(\frac{q}{p}\right),$$

with equality if and only if $\gamma([a,b])$ is the vertical segment from ip to iq.

Now consider arbitrary $z, w \in \mathbb{H}$, and let L be either the vertical Euclidean line through z, w (if z, w have the same real part) or the Euclidean semicircle through z, w and with center in $\mathbb{R}$ (if z, w have different real parts). It is an easy exercise to check that there exists $T \in \text{Möb}(\mathbb{H})$ such that $T(L) = i\mathbb{R}$. The result now follows from the above paragraph and Lemma 2.2.

(ii) The proof for $\mathbb{D}$ is a direct consequence of (i), Lemma 2.3 and Proposition 2.1. $\square$

2.2.5. *The boundary at infinity*

Let $p < 1$. Note that Proposition 2.4 gives that $d_{\mathbb{H}}(i, pi) = -\log(p)$. In particular, $d_{\mathbb{H}}(i, pi) \to \infty$ as $p \to 0$. For this reason, the set $\bar{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ is called the *boundary at infinity* of $\mathbb{H}$. Similarly, $\mathbb{S}^1$ is called the *boundary at infinity* of $\mathbb{D}$.

2.2.6. *The full isometry group*

Denote by $\text{Isom}^+(\mathbb{H})$ (resp. $\text{Isom}^+(\mathbb{D})$) the group of orientation-preserving isometries of $\mathbb{H}$ (resp. $\mathbb{D}$). We have:

Proposition 2.5: $\text{Isom}^+(\mathbb{H}) \cong \text{PSL}(2, \mathbb{R})$ *and* $\text{Isom}^+(\mathbb{D}) \cong \text{PSU}(2, \mathbb{C})$.

Proof: Again, we prove the result only for $\mathbb{H}$. In view of Lemma 2.2, we must show that every element of $\text{Isom}^+(\mathbb{H})$ is a Möbius transformation. Let $F \in \text{Isom}^+(\mathbb{H})$. Composing with an element of $\text{PSL}(2, \mathbb{R})$ if necessary, we may assume that F fixes two distinct points $z_1, z_2 \in i\mathbb{R}$. Choose a point $w \notin i\mathbb{R}$, noting that $T(w)$ lies on the hyperbolic circle C_i of centre z_i and radius $d_{\mathbb{H}}(z_i, w)$, for $i = 1, 2$.

Now, it is not difficult to verify (see [24], Ch. 2) that every hyperbolic circle is also a Euclidean circle. Therefore, C_1 and C_2 intersect at two points: one of them is w, and the other one is on the other side of $i\mathbb{R}$ from w. Since F is orientation-preserving, we get that $F(w) = w$. Therefore, F is an isometry fixing three points, and hence the identity. $\square$

2.2.7. *Dynamics of elements of* $\text{Isom}^+(\mathbb{H})$

Recall the classification of Möbius transformations into *elliptic, parabolic* and *loxodromic*. Since $\text{Isom}^+(\mathbb{H}) \cong \text{PSL}(2, \mathbb{R})$, loxodromic isometries of $\mathbb{H}$

are all *hyperbolic*. We now make a few comments on the dynamics of the different types of isometries:

(i) If $T \in \mathrm{Isom}^+(\mathbb{H})$ is parabolic, then it has exactly one fixed point on $\bar{\mathbb{R}}$. Thus, up to conjugation in $\mathrm{Isom}^+(\mathbb{H})$, $T(z) = z + a$ for some $a \in \mathbb{R}$; observe that T leaves invariant every Euclidean line of the form $y =$ constant.

(ii) If $T \in \mathrm{Isom}^+(\mathbb{H})$ is hyperbolic, then it has two fixed points on $\bar{\mathbb{R}}$; the geodesic between them is called the *axis* of T. Up to conjugation, $T(z) = \lambda z$ for some $\lambda \in \mathbb{R}$. The map T acts on its axis as a hyperbolic translation, with translation distance $l = \log \lambda$; observe that (2.1) gives that $\mathrm{tr}^2(T) = 4\cosh^2(l/2)$.

(iii) Finally, if $T \in \mathrm{Isom}^+(\mathbb{H})$ is elliptic, then it has exactly one fixed point in $\mathbb{H}$. Up to conjugation, $T(z) = \frac{\cos(t)z+\sin(t)}{\sin(t)z+\cos(t)}$, for some $t \in \mathbb{R}$.

We refer the interested reader to ([24], Chapter 3) for pictures showing the dynamics of the different types of elements of $\mathrm{Isom}^+(\mathbb{H})$.

2.3. *Fuchsian groups and fundamental domains*

2.3.1. *Fuchsian groups*

Let Γ be a group acting by homeomorphisms on a metric space $\mathbb{X}$. We say that Γ acts *properly discontinously* on $\mathbb{X}$ if, for all compact subsets $K \subset \mathbb{X}$, the set

$$\{\gamma \in \Gamma \mid \gamma(K) \cap K \neq \emptyset\}$$

is finite. We say that Γ acts *freely* on $\mathbb{X}$ if every non-identity element of Γ acts without fixed points.

We will be interested mainly in the case where $\mathbb{X} = \mathbb{H}$ and Γ is subgroup of $\mathrm{PSL}(2,\mathbb{R})$ that is *discrete* with respect to the natural topology inherited from $\mathrm{PSL}(2,\mathbb{R})$. Discrete subgroups of $\mathrm{PSL}(2,\mathbb{R})$ are called *Fuchsian groups*. We will need the following well-known result; for a proof see ([5], Prop. B.1.6), for instance.

Proposition 2.6: *Let Γ be a subgroup of $\mathrm{PSL}(2,\mathbb{R})$. The following conditions are equivalent:*

(1) Γ acts freely and properly discontinuously on $\mathbb{H}$.
(2) $\mathbb{X}/\Gamma$ is Hausdorff and the projection $\mathbb{X} \to \mathbb{X}/\Gamma$ is a covering map.
(3) Γ is a torsion-free Fuchsian group. $\square$

2.3.2. *Fundamental domains*

We now introduce the concept of fundamental domain for the action of a group on a metric space.

Definition 2.7: (Fundamental domain) Let Γ be a group acting properly discontinously by homeomorphisms on a metric space $\mathbb{X}$. A *fundamental domain* for the action of Γ on $\mathbb{X}$ is a closed subset $C \subset \mathbb{X}$ such that:

(1) The interior $\mathrm{int}(C)$ of C is not empty.
(2) If $T \neq \mathrm{id}$ then $T(\mathrm{int}(C)) \cap \mathrm{int}(C) = \emptyset$.
(3) The Γ-translates of C *tessellate* $\mathbb{X}$; that is, $\bigcup_{T\in\Gamma} T(C) = \mathbb{X}$.

As it turns out, every Fuchsian group admits a particularly nice type of fundamental domain, known as the *Dirichlet domain*, which we now describe.

Definition 2.8: (Dirichlet domain) Let Γ be a group acting properly discontinously by isometries on a metric space $\mathbb{X}$, and let $z_0 \in X$ be a point not fixed by any non-trivial element of Γ. The *Dirichlet domain* of Γ centered at z_0 is

$$\mathcal{D}_\Gamma(z_0) = \{x \in \mathbb{X} \mid d(x, z_0) \leq d(x, T(z_0)), \forall T \in \Gamma\}.$$

Proposition 2.9: *Let Γ be a Fuchsian group, and let z_0 be a point not fixed by any non-trivial element of Γ. Then $\mathcal{D}_\Gamma(z_0)$ is a convex fundamental domain for the action of Γ on $\mathbb{H}$.*

Proof: First, $\mathcal{D}_\Gamma(z_0)$ is closed and convex since it is the intersection of closed half-planes of $\mathbb{H}$. Moreover, $z_0 \in \mathrm{int}(\mathcal{D}_\Gamma(z_0))$ since Γ is discrete.

We now claim that the Γ-translates of $\mathcal{D}_\Gamma(z_0)$ tessellate $\mathbb{H}$. Let $z \in \mathbb{H}$. Since Γ is discrete, there exists $T \in \Gamma$ such that

$$d(z, T(z_0)) = \min_{S\in\Gamma}\{d(z, S(z_0))\}.$$

Thus $T^{-1}(z) \in \mathcal{D}_\Gamma(z_0)$, and therefore $z \in T(\mathcal{D}_\Gamma(z_0))$, as desired.

Finally, suppose for contradiction that there exists $z \in \mathrm{int}(\mathcal{D}_\Gamma(z_0))$ and $T \neq \mathrm{Id}$ such that $T(z) \in \mathrm{int}(\mathcal{D}_\Gamma(z_0))$. In particular,

$$d(z, z_0) < d(z, T^{-1}(z_0)) = d(T(z), z_0)$$

and

$$d(T(z), z_0) < d(T(z), T(z_0)) = d(z, z_0),$$

which is impossible. □

Observe that the fact that Γ acts properly discontinously on $\mathbb{H}$ implies that $\mathcal{D}_\Gamma(z_0)$ is *locally finite*, that is, for every compact set $K \subset \mathbb{H}$, there are only finitely many Γ-translates of $\mathcal{D}_\Gamma(z_0)$ that intersect K; see ([20], Thm. 3.5.1) for details.

Example 2.10: (a) Let $T \in \text{Möb}(\mathbb{H})$ be a parabolic isometry so, up to conjugation, $T(z) = z + a$ for some $a \in \mathbb{R} \setminus \{0\}$. Let $\Gamma = \langle T \rangle$ and $z_0 \in \mathbb{H}$. Then $\mathcal{D}_\Gamma(z_0) = \{z \in \mathbb{C} \mid \text{Re}(z_0) - \frac{a}{2} \leq \text{Re}(z) \leq \text{Re}(z_0) + \frac{a}{2}\}$.

(b) Let $\Gamma = \text{SL}(2, \mathbb{Z})$ and let $z_0 = 2i$, which is not fixed by any element of Γ. Then $\mathcal{D}_\Gamma(2i) = \{z \in \mathbb{H} \mid -1/2 \leq \text{Re}(z) \leq 1/2, |z| \geq 1\}$.

(b) Consider the Euclidean isometries $A : (x, y) \to (x + 1, y)$ and $B : (x, y) \to (x, y + b)$, where $b > 0$. Let $G \cong \mathbb{Z}^2$ be the group generated by A and B. Then the Dirichlet domain for the action of Γ on Euclidean plane $\mathbb{E}^2$ is generically a hexagon; in the special case when $b = 1$, it is a square.

2.3.3. *The action of a group on a Dirichlet domain*

Let Γ be a discrete group acting properly discontinuously by isometries on $\mathbb{X} = \mathbb{H}$ or $\mathbb{E}$, and let D be a Dirichlet domain for the action of Γ. We will assume, for simplicity, that Γ acts freely on $\mathbb{X}$, so that $\mathbb{X}/\Gamma$ is a smooth surface, and that D is compact, so that D is a finite sided polygon. In the case when Γ is a Fuchsian group, then the fact that D is compact implies that Γ has no parabolic elements; see ([20], Thm. 4.2.1). The group Γ identifies the sides of D in pairs and, in fact, Γ is generated by the (finite) collection of all side pairings; see ([20], Thm. 3.5.4). Each Γ-orbit of vertices of D is called a *cycle*, and the sum of the internal angles at the vertices of a cycle is always equal to 2π; see ([20], Thm. 3.5.3).

A converse to this situation is described in *Poincaré's Polygon Theorem*, which we now state; for a proof, see ([24], Ch. 7). Again, $\mathbb{X} = \mathbb{H}$ or $\mathbb{E}$.

Theorem 2.11: *(Poincaré) Let $P \subset \mathbb{X}$ be a compact polygon whose sides are identified in pairs by isometries of $\mathbb{X}$, and let Γ be the group generated by those isometries. Suppose that, for each Γ-orbit of vertices of P, the internal angles at the vertices in that orbit add up to 2π. Then Γ is a discrete group acting freely and properly discontinously on $\mathbb{X}$; moreover, D is a fundamental domain for the action.*

3. Hyperbolic structures on surfaces

We refer the reader to the texts [5, 10, 11, 24, 26] for a more detailed discussion on the material presented in this section.

3.1. *Definition and examples*

Let $\mathbb{X}$ denote the hyperbolic plane $\mathbb{H}$, the Euclidean plane $\mathbb{E}$ or the unit sphere $\mathbb{S}^2$ in $\mathbb{R}^3$.

Definition 3.1: (Geometric structure) Let S be a topological surface. A *geometric structure* on S consists of an open cover $\{U_i\}_{i\in I}$ of S and a collection $\{\phi_i\}_{i\in I}$ of maps, with $\phi_i : U_i \to \mathbb{X}$, such that

(1) ϕ_i is a homeomorphism onto its image, for each $i \in I$, and
(2) if $U_i \cap U_j \neq \emptyset$, the restriction of the transition map

$$\phi_i \circ \phi_j^{-1} : \phi_j(U_i \cap U_j) \to \phi_i(U_i \cap U_j)$$

to each connected component of $\phi_j(U_i \cap U_j)$ is an orientation-preserving isometry of $\mathbb{X}$.

In the case where $\mathbb{X} = \mathbb{H}$ (resp. $\mathbb{X} = \mathbb{E}$ or $\mathbb{X} = \mathbb{S}^2$), we say that the surface S is equipped with a *hyperbolic structure* (resp. *Euclidean or spherical structure*). In the definition above, each pair (U_i, ϕ_i) is called a *chart*. The set of all charts is called an *atlas* of S; note that every atlas is contained in a unique maximal atlas. Finally, observe that a surface equipped with a geometric structure supports a natural path-metric, obtained by deeming each chart map to be an isometry.

Remark 3.2: (Geometric structure on a covering space) Suppose that S is equipped with a geometric structure $\{(U_i, \phi_i)\}_{i\in I}$ and let $\pi : \tilde{S} \to S$ be a covering map; without loss of generality, we assume that U_i is evenly covered for all $i \in I$. Then $\tilde{S}$ comes equipped with a natural geometric structure, where the open sets are the preimages of the U_i under π and the chart maps are the restrictions of $\phi_i \circ \pi$ to each of these sets.

A geometric structure on a covering space also induces a geometric structure on the quotient space; see Examples 3.4 and 3.5 below. From now on, we will focus our attention mainly on the case $\mathbb{X} = \mathbb{H}$.

Remark 3.3: (Hyperbolic structure on a surface with boundary) If S has boundary, we define a *hyperbolic structure with geodesic boundary* on S by requiring that each of the sets U_i in Definition 3.1 be an open set of a closed half-plane in $\mathbb{H}$.

Example 3.4: (Hyperbolic structure on a surface of genus $g \geq 2$). Let $g \geq 2$ and P be a regular hyperbolic $4g$-gon in $\mathbb{D}$ with internal angles $\pi/2g$. To see that such polygon exists, consider $4g$ equispaced geodesic rays in $\mathbb{D}$ emanating from the origin O, as in Figure 1. Consider the hyperbolic polygon P_t whose vertices are the points of intersection between these rays and the hyperbolic circle of center O and hyperbolic radius $t > 0$. As t increases, the internal angle of P_t decreases from the Euclidean value $(4g-2)\pi$, down to 0. By continuity, there is a value of t for which the internal angle is equal to $\pi/2g$.

Suppose that the sides of P are identified in pairs by elements of $\mathrm{PSL}(2,\mathbb{R})$ according to the labelling outlined in Figure 1. We see that the hypotheses of Poincaré's Polygon Theorem are satisfied, and thus the group Γ generated by the side pairings is a Fuchsian group acting freely on $\mathbb{H}$; observe there is only one Γ-orbit of vertices. The quotient space $\bar{P} = \mathbb{H}/\Gamma$ is homeomorphic to a closed surface of genus g. We define a hyperbolic structure on $\bar{P}$ by a specifying a chart around each point in $\bar{P}$; such charts are schematically shown in Figure 1 for $g = 2$, depending on whether a lift of the point is in the interior of P, on one of the sides of P, or is a vertex of P. Observe that, since the angle around any vertex of P is $\pi/2g$, then the angle around the corresponding point in $\bar{P}$ is 2π and thus the chart is well-defined. Finally, note the natural path-metric on $\bar{P}$ is complete.

In fact, the previous example is a special case of a more general situation, as we now explain.

Example 3.5: (Quotient of $\mathbb{H}$ by a Fuchsian group) Let Γ be a Fuchsian group acting freely on $\mathbb{H}$. Let $S = \mathbb{H}/\Gamma$ and let $\pi : \mathbb{H} \to S$ be the natural covering map. We endow $S = \mathbb{H}/\Gamma$ with a hyperbolic structure by specifying a chart around each point $p \in S$, as follows. Let U_p be an evenly covered open neighbourhood of p, and let f_p be a homeomorphism identifying U_p with any of the open sets in $\mathbb{H}$ covering U_p. The collection $\{(U_p, f_p)\}_{p\in S}$ gives a hyperbolic structure on S; again, the natural path-metric on the surface is complete.

It is easy to see that the hyperbolic structures on $\mathbb{H}/\Gamma$ given in Examples 3.4 and 3.5 are in fact equivalent.

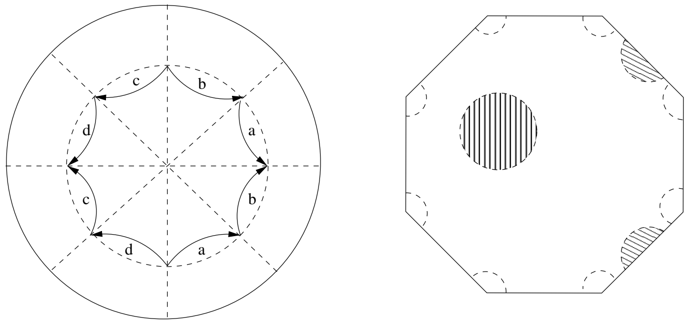

Fig. 1. The left figure shows a regular hyperbolic octagon; for exactly one value of the radius of the shaded circle, the internal angles will be $\pi/4$. The sides are identified by isometries $a, b, c, d \in \mathrm{PSL}(2, \mathbb{R})$ according to the labelling shown. The quotient surface $\mathbb{H}/\Gamma$, where $\Gamma = \langle a, b, c, d \rangle$, is homeomorphic to a closed surface of genus 2, and is equipped with a complete hyperbolic structure. The schematics of the charts around a point p are shown in the right figure, depending on whether a lift of p lies in the interior of the polygon, or in the interior of a side, or is a vertex.

Remark 3.6: (Euclidean structure on a torus) By applying the same reasoning as above, we obtain a Euclidean structure on a surface of genus $g = 1$ by identifying opposite sides of a rectangle in the Euclidean plane $\mathbb{E}$. More generally, the quotient of $\mathbb{E}$ by a discrete group of Euclidean isometries is naturally equipped with a Euclidean structure.

Remark 3.7: (Geometric structures on closed surfaces) Observe that a surface equipped with a geometric structure has constant Gaussian curvature. Therefore, a closed surface of genus $g \geq 1$ admits a hyperbolic (resp. Euclidean) structure if and only if $g \geq 2$ (resp. $g = 1$), as follows from Example 3.4, Remark 3.6 and the Gauss-Bonnet theorem.

3.2. *The Cartan-Hadamard Theorem. Developing map and holonomy*

The next result, a special case of the celebrated *Cartan-Hadamard Theorem*, asserts that Example 3.5 is the only way of obtaining a surface equipped with a hyperbolic structure, provided we restrict our attention to *complete* structures. We refer the reader to [3, 9, 12] for more general versions of the Cartan-Hadamard Theorem, and to ([8], Ch. 6) for a discussion on *incomplete* hyperbolic structures on surfaces.

Theorem 3.8: *(Cartan-Hadamard) Let X be a connected surface equipped with a hyperbolic structure, and suppose that the natural path-metric on*

X is complete. Then X is isometric to $\mathbb{H}/\Gamma$, where Γ is a Fuchsian group acting freely on $\mathbb{H}$.

The rest of this section is devoted to give a sketch of the proof of Theorem 3.8. The strategy is as follows. First, we will construct an isometry

$$\mathrm{Dev} : \tilde{\mathrm{X}} \to \mathbb{H},$$

called the *developing map*; here $\tilde{X}$ denotes the universal cover of X. The map Dev will induce an isomorphism

$$\mathrm{Hol} : \pi_1(\mathrm{X}) \to \Gamma,$$

where Γ is a torsion-free Fuchsian group; the map Hol is called the *holonomy representation* of $\pi_1(X)$. Once all this has been established, it will easily follow that X is isometric to $\mathbb{H}/\Gamma$.

Next, we explain some of the details, and refer the reader to [5, 11, 24] for a more thorough discussion; we remark that one obtains a Cartan-Hadamard Theorem for Euclidean surfaces using the same ideas as below, with the obvious modifications.

3.2.1. *The developing map*

Let $\{(U_i, \phi_i)\}_{i\in I}$ be an atlas defining the hyperbolic structure on X. Fix, once and for all, a basepoint $p \in U_0$. Let $\tilde{X}$ be the universal cover of X, namely the set of homotopy classes of paths in X that start at p; recall that $\tilde{X}$ has a natural hyperbolic structure coming from that of X, by Remark 3.2.

Let $[\gamma] \in \tilde{X}$ and choose a representative $\gamma : [0, 1] \to X$ of $[\gamma]$. We cover $\gamma([0, 1])$ with a finite collection $\{(U_i, \phi_i)\}_{i=0}^n$ of charts as shown in Figure 2; in particular $U_i \cap U_{i+1}$ is connected. We define the map Dev successively, as follows. First, set

$$\mathrm{Dev}|_{(\mathrm{U_0}\cap\gamma)} = \phi_0|_{(\mathrm{U_0}\cap\gamma)}.$$

Since $U_0 \cap U_1$, is connected, the definition of hyperbolic structure yields that $\phi_0 \circ \phi_1^{-1}|_{(U_0\cap U_1)} = T_1 \in \mathrm{PSL}(2, \mathbb{R})$. Set

$$\mathrm{Dev}|_{(\mathrm{U_1}\cap\gamma)} = \mathrm{T_1} \circ \phi_1|_{(\mathrm{U_1}\cap\gamma)},$$

and note Dev is now defined on $(U_0 \cup U_1) \cap \gamma$. Repeating this process, we obtain maps $T_2, \ldots, T_n \in \mathrm{PSL}(2, \mathbb{R})$, and define

$$\mathrm{Dev}|_{(\mathrm{U_n}\cap\gamma)} = \mathrm{T_1} \circ \mathrm{T_2} \circ \ldots \circ \mathrm{T_n} \circ \phi_{\mathrm{n}}|_{(\mathrm{U_n}\cap\gamma)}, \tag{3.1}$$

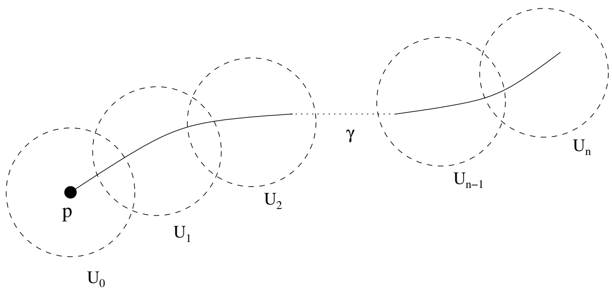

Fig. 2.

noting that Dev is now well-defined on $(U_0 \cup \ldots \cup U_n) \cap \gamma$, and thus on the whole of γ. We set

$$\mathrm{Dev}(\gamma) = \mathrm{Dev}(\gamma(1)) \in \mathbb{H}.$$

At this point, it is straightforward, although not terribly amusing, to show that $\mathrm{Dev}(\gamma)$ depends only on the initial chart $\phi_0 : U_0 \to \mathbb{H}$ and the homotopy class of γ; this is carefully explained in ([5], Prop. B.1.3), for instance. Thus we have obtained a well-defined map

$$\mathrm{Dev} : \tilde{\mathrm{X}} \to \mathbb{H},$$

which is a local isometry with respect to the natural hyperbolic structure on $\tilde{X}$ (as we will see, if X is complete then Dev will be a global isometry). As a consequence, we obtain that any two choices of initial chart produce developing maps which differ by an element of $\mathrm{PSL}(2, \mathbb{R})$.

3.2.2. *Two technical lemmas*

Having introduced the developing map, we continue towards a proof of Theorem 3.8. Following the strategy of [5, 24], the proof is based on the two results we now present.

Lemma 3.9: *Suppose X is equipped with a complete hyperbolic structure. Then its universal cover $\tilde{X}$ is also complete.*

Proof: Let $(\tilde{z}_n)_n \subset \tilde{X}$ be a Cauchy sequence. As the covering map $\pi : \tilde{X} \to X$ does not increase distances, then $(\pi(\tilde{z}_n))_n$ is a Cauchy sequence and thus converges to a point $z \in X$, since X is complete. Let U be an evenly covered open neighbourhood of z. Since $(\tilde{z}_n)_n$ is Cauchy, all but finitely many elements of $(\tilde{z}_n)$ belong to exactly one of the preimages $\tilde{U}$ of U and thus converge to the preimage of z contained in $\tilde{U}$. □

Lemma 3.10: *Let X be a surface equipped with a complete hyperbolic structure. Then, the developing map* $\mathrm{Dev} : \tilde{\mathrm{X}} \to \mathbb{H}$ *is a surjective covering map.*

Proof: Since Dev is a local homeomorphism by construction, it suffices to prove that Dev satisfies the *path-lifting property.* This is, we want to establish that, for all $z_0 \in \mathrm{Dev}(\tilde{\mathrm{X}})$, all $\tilde{z}_0 \in \mathrm{Dev}^{-1}(\mathrm{z}_0)$, and all piecewise differentiable paths $\gamma : [0,1] \to \mathbb{H}$ with $\gamma(0) = z_0$, there exists a path $\tilde{\gamma} : [0,1] \to \tilde{X}$ such that $\tilde{\gamma}(0) = \tilde{z}_0$ and $\mathrm{Dev} \circ \tilde{\gamma} = \gamma$.

Let $z_0 \in \mathrm{Dev}(\tilde{\mathrm{X}})$, $\tilde{z}_0 \in \mathrm{Dev}^{-1}(\mathrm{z}_0)$, and $\gamma : [0,1] \to \mathbb{H}$ a piecewise differentiable path with $\gamma(0) = z_0$. Consider

$$t_0 = \sup\left\{ t \in [0,1] \mid \exists \tilde{\gamma} : [0,t] \to \tilde{X} \text{ with } \tilde{\gamma}(0) = \tilde{z}_0 \text{ and } \mathrm{Dev} \circ \tilde{\gamma} = \gamma|_{[0,\mathrm{t}]} \right\}.$$

We want to show that $t_0 = 1$. First, note that, since Dev is a local isometry, then $t_0 > 0$. Consider, for all $t < t_0$, the lift $\tilde{\gamma} : [0,t] \to \tilde{X}$ of $\gamma : [0,t] \to \mathbb{H}$ and observe that $\tilde{\gamma}$ is unique, again because Dev is a local isometry. Let t_n be an increasing sequence converging to t_0. Then $(\tilde{\gamma}(t_n))_n$ is a Cauchy sequence in $\tilde{X}$; otherwise $\tilde{\gamma}([0,t_0))$ would have infinite length, which is impossible; see ([5], Prop. B.1.3) for details. Therefore $(\tilde{\gamma}(t_n))_n$ converges, by Lemma 3.9, and thus we define $\tilde{\gamma}(t_0)$ to be this limit. Finally, since Dev is an isometry in a neighbourhood of $\tilde{\gamma}(t_0)$, it follows that $t_0 = 1$, as claimed. Therefore, Dev is a covering map.

The fact that we can lift paths from $\mathbb{H}$ to $\tilde{X}$ quickly implies that Dev is surjective. Indeed, let $z \in \mathbb{H}$, and choose $z_0 \in \mathrm{Dev}(\tilde{\mathrm{X}})$ and a path $\gamma : [0,1] \to \mathbb{H}$ with $\gamma(0) = z_0$. Denote by $\tilde{\gamma} : [0,1] \to \tilde{X}$ the lift of γ. Then $\mathrm{Dev}(\tilde{\gamma}(1)) = \mathrm{z}$, as desired. □

Since Dev is a surjective covering map and $\mathbb{H}$ is simply-connected, we deduce that $\tilde{X}$ is homeomorphic to $\mathbb{H}$. This, together with the fact that Dev is a local isometry, implies:

Corollary 3.11: *Let X be a surface equipped with a complete hyperbolic structure. Then, then universal cover $\tilde{X}$ of X is isometric to* $\mathbb{H}$. □

3.2.3. *Holonomy*

Let X be a surface equipped with a hyperbolic structure, and choose a basepoint p on X. If we consider closed paths based at p in the construction

(3.1) of the developing map above, we obtain a map

$$\mathrm{Hol} : \pi_1(\mathrm{X}, \mathrm{p}) \to \mathrm{PSL}(2, \mathbb{R}),$$

defined by $\mathrm{Hol}([\gamma]) = \mathrm{T}_1 \circ \mathrm{T}_2 \circ \ldots \circ \mathrm{T}_\mathrm{n} \in \mathrm{PSL}(2, \mathbb{R})$. By definition, the map Hol is a homomorphism, and is commonly referred to as the *holonomy representation* of $\pi(X, p)$. Again, Hol only depends on the choice of initial chart $\phi_0 : U_0 \to \mathbb{H}$, and any two choices produce conjugate homomorphisms. Moreover, we have:

Lemma 3.12: $\mathrm{Hol} : \pi_1(\mathrm{X}, \mathrm{p}) \to \mathrm{PSL}(2, \mathbb{R})$ *is injective.*

Proof: Let $[\gamma] \in \pi_1(X, p)$ and suppose that $\mathrm{Hol}([\gamma]) = \mathrm{Id} \in \mathrm{PSL}(2, \mathbb{R})$. Then, the developing image of γ is a loop based at p. Since $\mathbb{H}$ is simply-connected, we can find a homotopy between this loop and the trivial loop. Finally, since Dev is a covering map, we may lift such homotopy to a homotopy between γ and the trivial loop, as desired. □

Therefore, the holonomy representation gives an identification of $\pi_1(X, p)$ with a Fuchsian group $\Gamma = \mathrm{Hol}(\pi_1(\mathrm{X}, \mathrm{p}))$. Since $\pi_1(X, p)$ acts on $\tilde{X}$ freely and properly discontinously, the same holds for the action of Γ on $\mathbb{H}$, by Proposition 2.6. Therefore, Γ is a torsion-free Fuchsian group, again by Proposition 2.6.

Proof of Theorem 3.8. Let X be a surface equipped with a complete hyperbolic structure. By Corollary 3.11, its universal cover $\tilde{X}$ is isometric to $\mathbb{H}$ via the developing map. Moreover, the holonomy map gives an identification $\pi_1(X, p)$ with a torsion-free Fuchsian group Γ, and thus the result follows. □

4. Teichmüller space

In this section we introduce the Teichmüller space $\mathcal{T}(S)$ of an orientable surface S of genus $g \geq 1$, the space of distinct geometric structures on S. In order to keep the exposition as simple as possible, we restrict our attention to closed surfaces only. In this case, Remark 3.7 gives that S carries a hyperbolic (resp. Euclidean) structure if and only if $g \geq 2$ (resp. $g = 1$).

In addition, we will focus solely on topological aspects of Teichmüller space, with the ultimate goal of proving, in Section 5, that the natural action of the *mapping class group* on Teichmüller space is properly discontinuous. In particular, we will not make reference to the various different metrics on Teichmüller space. We refer the reader to [1, 5, 13, 14, 16, 17], and the references therein, for a detailed exposition of Teichmüller spaces.

4.1. *Two definitions*

We now give two equivalent definitions of the Teichmüller space of a surface S of genus $g \geq 2$, one as the set of distinct hyperbolic structures on S and the other as the set of conjugacy classes of discrete faithful representations of $\pi_1(S)$ into $\mathrm{PSL}(2, \mathbb{R})$.

Definition 4.1: (Teichmüller space of a hyperbolic surface, I) Let S be a closed topological surface of genus $g \geq 2$. The Teichmüller space $\mathcal{T}(S)$ of S is

$$\mathcal{T}(S) = \{(X, f)\}/\sim,$$

where

- X is S equipped with a hyperbolic structure,
- $f : S \to X$ is a homeomorphism, called the *marking*, and
- $(X, f) \sim (Y, g)$ if and only if there is an isometry $\iota : X \to Y$ such that $\iota \circ f$ is homotopic to g.

In order to reduce notation, we will denote points $[(X, f)] \in \mathcal{T}(S)$ simply by X whenever we do not need to make explicit reference to the marking. We now present an equivalent definition of Teichmüller space which, in particular, will allow us to define a natural topology on $\mathcal{T}(S)$.

Definition 4.2: (Teichmüller space of a hyperbolic surface, II) Let S be a closed surface of genus $g \geq 2$. The Teichmüller space of S is

$$\mathcal{T}(S) = \mathcal{DF}(\pi_1(S), \mathrm{PSL}(2, \mathbb{R}))/\,\mathrm{PSL}(2, \mathbb{R}),$$

the set of discrete, faithful representations of $\pi_1(S)$ into $\mathrm{PSL}(2, \mathbb{R})$, up to conjugation.

The set $\mathcal{DF}(\pi_1(S), \mathrm{PSL}(2, \mathbb{R}))/\,\mathrm{PSL}(2, \mathbb{R})$ is called the $\mathrm{PSL}(2, \mathbb{R})$-*character variety* of $\pi_1(S)$. The equivalence of Definitions 4.1 and 4.2 is essentially contained in the statement of Theorem 3.8. Indeed, a point $[(X, f)] \in \mathcal{T}(S)$ determines a conjugacy class of faithful representations of $\pi_1(X) \cong \pi_1(S)$ into $\mathrm{PSL}(2, \mathbb{R})$ via the holonomy map. Conversely, given $\rho \in \mathcal{DF}(\pi_1(S), \mathrm{PSL}(2, \mathbb{R}))$, then $X = \mathbb{H}/\rho(\pi_1(S))$ comes equipped with a natural hyperbolic structure, by Example 3.5. Now, ρ induces a homotopy equivalence $h : S \to X$ which is then homotopic to a homeomorphism $f : S \to X$, the desired marking. Finally, any two conjugate representations produce isometric surfaces.

Remark 4.3: (Topology on $\mathcal{T}(S)$) Observe that, in light of Definition 4.2, $\mathcal{T}(S)$ carries a natural topology as a quotient of $\mathrm{PSL}(2,\mathbb{R})^{2g}$, since

$$\mathcal{DF}(\pi_1(S), \mathrm{PSL}(2,\mathbb{R})) \subset \mathrm{Hom}(\pi_1(S), \mathrm{PSL}(2,\mathbb{R})) = \mathrm{PSL}(2,\mathbb{R})^{2g}.$$

Example 4.4: (Teichmüller space of the torus) If S has genus 1, we define $\mathcal{T}(S)$ as the set of distinct Euclidean structures of *unit area* on S, by performing the obvious changes in Definition 4.1. By the same reasoning as above, we may identify $\mathcal{T}(S)$ with the set of *marked* torsion-free Euclidean lattices, modulo Euclidean isometries and scalings; the term "marked" means that every lattice has a specified ordered pair of generators. Up to isometry and scaling, we can arrange for one of the generators of the lattice to be 1, and the other one to lie above the x-axis. In this way have identified $\mathcal{T}(S)$ with the upper-half plane $\mathbb{H}$. For our purposes, this is just an identification as topological spaces; that said, Teichmüller spaces carry a natural metric, the so-called *Teichmüller metric*, for which the Teichmüller space of the torus, equipped with this metric, is isometric to the hyperbolic plane $(\mathbb{H}, d_{\mathbb{H}})$.

4.2. *Fenchel-Nielsen coordinates.*

For a general surface S, the definition of Teichmüller space does not give a very clear insight on the structure of $\mathcal{T}(S)$. This will change once we introduce the so-called *Fenchel-Nielsen coordinates* for Teichmüller space. In terms of these coordinates, a point $X \in \mathcal{T}(S)$ will correspond to $6g-6$ real numbers; half of these correspond to the *lengths*, measured in X, of the curves in a fixed pants decomposition, and the other half correspond to the *twist* with which different pants have been glued to obtain the structure X. Before we define these coordinates, we need to introduce a few notions.

4.2.1. *Length functions.*

Again, S denotes a closed surface of genus $g \geq 2$. Let γ be a homotopically non-trivial simple closed curve on S, and thus a non-trivial element of $\pi_1(S)$.

We claim that, given $X \in \mathcal{T}(S)$, there exists a unique simple closed geodesic in X that is homotopic to γ. To see this, we first regard $\pi_1(X) \cong \pi_1(S)$ as a subgroup of $\mathrm{PSL}(2,\mathbb{R})$, using the holonomy map. Under this identification, γ corresponds to a hyperbolic isometry $\bar{\gamma}$; otherwise it would be elliptic, which is impossible since X is a surface; or parabolic, which is also impossible since parabolic isometries have zero translation distance and X is compact. Now, γ is homotopic to the simple closed geodesic contained

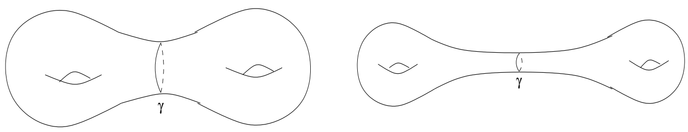

Fig. 3. Two different points in the Teichmüller space of a surface of genus 2

in projection of the axis of $\bar{\gamma}$. The uniqueness of the simple closed geodesic is obtained along similar lines; see Prop. 1.3 of [13].

Let $\mathcal{C}(S)$ be the set of homotopy classes of simple closed curves on S. For simplicity, we will refer to the elements of $\mathcal{C}(S)$ simply as *curves*, and we will often blur the distinction between a curve and any of its representatives.

Given $\gamma \in \mathcal{C}(S)$, the *length function* of γ is the function

$$l_{\cdot}(\gamma) : \mathcal{T}(S) \to \mathbb{R}_+$$

given by

$$l_{[(X,f)]}(\gamma) = \mathrm{length}_X(f(\gamma)),$$

where $\mathrm{length}_X(f(\gamma))$ denotes the length of the unique geodesic representative of $f(\gamma)$ in X. For simplicity, we will denote $l_{[(X,f)]}(\gamma)$ simply by $l_X(\gamma)$.

Example 4.5: If $X, Y \in \mathcal{T}(S)$ are such that $\{l_X(\gamma)\}_{\gamma \in \mathcal{C}(S)} \neq \{l_Y(\gamma)\}_{\gamma \in \mathcal{C}(S)}$, then $X \neq Y$. In particular, the two surfaces in Figure 3 represent different different points in the Teichmüller space of the closed surface of genus 2.

As we will see, length functions are central to the definition of Fenchel-Nielsen coordinates; in addition, they are used to define the so-called *Thurston's compactification* of $\mathcal{T}(S)$, see the article [23] in this volume.

Let $X \in \mathcal{T}(S)$ and $\gamma \in \mathcal{C}(S)$. Let $\rho : \pi_1(X) \to \mathrm{PSL}(2,\mathbb{R})$ be the holonomy representation of $\pi_1(X)$, noting that $\rho(\gamma)$ is a hyperbolic element of $\mathrm{PSL}(2,\mathbb{R})$. Recall from (2.1) that the trace and translation distance of $\rho(\gamma)$ are related by $\mathrm{tr}^2(\rho(\gamma)) = 4\cosh^2(\frac{l(\rho(\gamma))}{2})$. Also, note that $l(\rho(\gamma)) = l_X(\gamma)$, and so we deduce that length functions are continuous:

Lemma 4.6: *(Length functions are continuous) For every $\gamma \in \mathcal{C}(S)$, the function $l_{\cdot}(\gamma) : \mathcal{T}(S) \to \mathbb{R}_+$ is continuous.* $\square$

4.2.2. *Multicurves and pants decompositions.*

Given two curves $\gamma, \gamma' \subset S$, the *intersection number* of γ and γ', denoted by $i(\gamma, \gamma')$, is the minimal cardinality of $\gamma \cap \gamma'$ among all representatives of γ and γ'. If $i(\gamma, \gamma') = 0$, we say that γ and γ' are *disjoint*. We say that two curves γ, γ' *fill* the surface if $S \setminus (\gamma \cup \gamma')$ is a union of topological disks; equivalently, if any non-trivial curve on S intersects at least one of γ or γ'.

A *multicurve* on S is a collection of pairwise distinct, pairwise disjoint curves; such a collection is necessarily finite, and consists of at most $3g-3$ curves. A multicurve μ that is maximal with respect to inclusion is called a *pants decomposition* of S; note that $S \setminus \mu$ has exactly $2g-2$ components, and that the closure of each of them is homeomorphic to a sphere with three boundary components, or *pair of pants.*

4.2.3. *The Teichmüller space of a pair of pants.*

We start by stating a well-known result in hyperbolic geometry, namely that a right-angled hyperbolic hexagon is determined by the lengths of any three non-consecutive sides. By a *marked* hyperbolic hexagon H we mean a hexagon in $\mathbb{H}$, together with a distinguished vertex, and a labelling $s_1, \ldots, s_6$ of the sides of H, in such way that the sides occur in that order when travelling counterclockwise along H from the distinguised vertex. Denote by l_i the hyperbolic length of the side s_i. We have:

Lemma 4.7: *Let $a, b, c > 0$. There exists a marked right-angled hyperbolic hexagon $H \subset \mathbb{H}$ such that $l_1 = a$, $l_3 = b$, and $l_5 = c$. Moreover, any two such marked hexagons are isometric via an element of* $\mathrm{PSL}(2, \mathbb{R})$ *sending one distinguished vertex to the other.*

The proof of Lemma 4.7 is an exercise in hyperbolic geometry; see ([13], Prop. 10.4) for details. Armed with Lemma 4.7, we are now in a position to understand the Teichmüller space $\mathcal{T}(\mathcal{P})$ of a pair of pants $\mathcal{P}$; the definition of $\mathcal{T}(P)$ is analogous to Definition 4.1, now considering hyperbolic structures with geodesic boundary, and requiring the isometry and the homotopies to fix the boundary pointwise. Denoting the three boundary components of $\mathcal{P}$ by $\gamma_1, \gamma_2, \gamma_3$, we have:

Lemma 4.8: *The map*

$$F : \mathcal{T}(\mathcal{P}) \to \mathbb{R}^3_+,$$

given by $F(X) = (l_X(\gamma_1), l_X(\gamma_2), l_X(\gamma_3))$, is a homeomorphism.

Proof: (Sketch) *(i) F is onto:* Let $(a,b,c) \in \mathbb{R}^3_+$. By Lemma 4.7, up to the action of $\mathrm{PSL}(2,\mathbb{R})$ there exists a unique marked right-angled hyperbolic hexagon H such that s_1, s_3 and s_5 have length $a/2, b/2$ and $c/2$, respectively. Now, glue two copies of H along s_2, s_4, s_6, obtaining a hyperbolic structure with geodesic boundary on $\mathcal{P}$, such that the lengths of the three boundary components are equal to a, b, c, respectively.

(ii) F is injective: Consider $X \in \mathcal{T}(\mathcal{P})$ and let $F(X) = (a,b,c) \in \mathbb{R}^3_+$. For each $i \neq j$ there exists a unique geodesic arc A_{ij} from γ_i to γ_j, perpendicular to both γ_i and γ_j. Then $X \setminus (A_{12} \cup A_{23} \cup A_{13})$ has two connected components, and the closure of each is a right-angled hyperbolic hexagon. Since a, b, c are fixed, we know the lengths of three non-consecutive sides of each hexagon. Then, by Lemma 4.7, each hexagon is determined up to isometry and, therefore, so is the hyperbolic structure on X.

(iii) F is continuous. Finally, to see that F is continuous one first needs to modify Definition 4.2 to accommodate for surfaces with boundary. Once this is done, the continuity of F follows immediately from the definition; see ([13], Prop. 10.4) for details. $\square$

4.2.4. *The coordinates*

Let S be a closed orientable surface of genus $g \geq 2$. We want to define a homeomorphism

$$F : \mathcal{T}(S) \to \mathbb{R}^{3g-3}_+ \times \mathbb{R}^{3g-3}.$$

Fix a pants decomposition $\gamma_1, \ldots, \gamma_{3g-3}$ on S, and fix an orientation for each of the curves. Let $X \in \mathcal{T}(S)$. The first $3g-3$ coordinates of $F(X)$, known as the *length parameters* of X, are simply the lengths $l_X(\gamma_1), \ldots, l_X(\gamma_{3g-3})$ in X of the curves γ_i. The other $3g-3$ coordinates $\tau_1(X), \ldots, \tau_{3g-3}(X)$, known as the *twist parameters* of X, are slightly more complicated. There are many (equivalent) ways of defining them; see, for instance, [5, 6, 16, 17, 26]. One way to do it is as follows:

Each of the curves γ_i is contained in a unique component S_i of $X \setminus (\bigcup_{j \neq i} \gamma_j)$ whose closure is homeomorphic to either a torus with one boundary component, or a sphere with four boundary components. Choose a curve β_i that is contained in S_i and intersects γ_i minimally; see Figure 4. In addition, in each pair of pants of $S_i \setminus \gamma_i$ we consider the unique geodesic arc that is entirely contained in that pair of pants, has endpoints on γ_i and is perpendicular to γ_i. Denote by $(A^i_j)_j$ the collection of arcs obtained in this way, observing that $(A^i_j)_j$ has exactly one element if S_i is home-

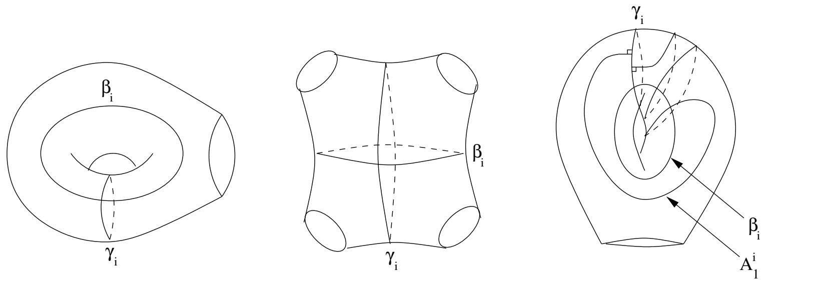

Fig. 4. Left and centre: A curve $\beta_i \subset S_i$ intersecting γ_i minimally, depending on the two possibilities for S_i. Right: The arc A_i when S_i is homeomorphic to a torus with one boundary component.

omorphic to a torus with one boundary component, and that it has two otherwise; see Figure 4 for an example of the former case.

Choose a basepoint $p \in \gamma_i \cap A^i_1$, and observe that $\pi_1(S_i, p)$ is generated by elements that have representatives which are entirely contained in $\gamma_i \cup (\cup_j A^i_j)$. Therefore, we may homotope β_i onto a curve β'_i contained in $\gamma_i \cup (\cup_j A^i_j)$; moreover, by tightening β'_i if necessary, we may assume that β'_i does not backtrack along γ_i. Then define $\tau_i(X)$ as the signed length of the segment of β'_i that runs along γ_i and contains p; the sign is positive if β'_i runs along γ_i in the sense given by the fixed orientation on γ_i, and negative otherwise.

Once we have defined Fenchel-Nielsen coordinates, we may state our promised theorem:

Theorem 4.9: *The map*

$$F : \mathcal{T}(S) \to \mathbb{R}^{3g-3}_+ \times \mathbb{R}^{3g-3},$$

given by

$$F(X) = (l_X(\gamma_1), \ldots, l_X(\gamma_{3g-3}), \tau_1(X), \ldots, \tau_{3g-3}(X)),$$

is a homeomorphism.

Proof: The map F is continous since it is defined in terms of length functions, which are continuous by Lemma 4.6. Also, F is bijective because it admits an inverse, which may intuitively be described as follows: given a tuple

$$(l_1, \ldots, l_{3g-3}, \tau_1, \ldots, \tau_{3g-3}) \in \mathbb{R}^{3g-3}_+ \times \mathbb{R}^{3g-3},$$

one first constructs $2g-2$ hyperbolic pairs of pants whose bounday components have length prescribed by the l_i, and then one glues the pairs of pants along the boundaries according to the twist parameters τ_i; see ([13], Thm. 10.6) for details. □

5. Mapping class groups

In this section we introduce the mapping class group of a surface and discuss some of its elements. We then describe how the mapping class group acts on Teichmüller space, and prove that this action is properly discontinuous. We refer the reader to [13, 14, 18] for a thorough discussion on mapping class groups.

5.1. *Definition and examples*

Let S be an orientable surface of genus $g \geq 1$. Again, for simplicity, we restrict our attention to the case where S is closed.

Definition 5.1: The *mapping class group* $\mathrm{Mod}(S)$ of S is the group of homotopy classes of orientation-preserving homeomorphisms of S; in other words,

$$\mathrm{Mod}(S) = \mathrm{Homeo}^+(S)/\,\mathrm{Homeo}_0(S),$$

where $\mathrm{Homeo}_0(S)$ denotes the connected component of $\mathrm{Homeo}^+(S)$ containing the identity. Elements of $\mathrm{Mod}(S)$ are called *mapping classes*. We will also need to consider the *extended mapping class group* $\mathrm{Mod}^{\pm}(S)$, that is, the group of all homeomorphisms of S up to homotopy.

Example 5.2: (Mapping class group of the torus.) If S is a torus, then $\mathrm{Mod}^{\pm}(S) \cong \mathrm{GL}(2,\mathbb{Z})$. Indeed, given a homeomorphism $g : S \to S$, let $g_* \in \mathrm{GL}(2,\mathbb{Z})$ be the induced automorphism of $\pi_1(S) \cong \mathbb{Z}^2$. Now, homotopic homeomorphisms induce conjugate automorphisms, and thus we have a homomorphism $G : \mathrm{Mod}^{\pm}(S) \to \mathrm{GL}(2,\mathbb{Z})$ given by $G([g]) = [g_*]$. The homomorphism G is clearly surjective; also, if $[g_*]$ is the identity, then g is homotopic to the identity and so G is also injective. Using the same reasoning, plus the fact that orientation-preserving homeomorphism must preserve algebraic intersection number, we obtain that $\mathrm{Mod}(S) \cong \mathrm{SL}(2,\mathbb{Z})$.

The example above is a particular instance of a general result, known as the *Dehn-Nielsen-Baer Theorem*. This result, which we state next, asserts that, if S is closed, the outer automorphism group $\mathrm{Out}(\pi_1(S))$ of $\pi_1(S)$

is isomorphic to the extended mapping class group $\mathrm{Mod}^{\pm}(S)$. The proof follows an argument similar to the one for the torus, but a substantial amount of extra work is required; see ([13], Thm. 8.1).

Theorem 5.3: *(Dehn-Nielsen-Baer) Let S be a closed surface of genus $g \geq 1$. Then* $\mathrm{Mod}^{\pm}(S) \cong \mathrm{Out}(\pi_1(S))$. □

5.1.1. *Examples of mapping classes*

We now give some examples of mapping classes:

Example 5.4: (Finite order.) If $\psi : S \to S$ is a finite order homeomorphism, then its homotopy class $[\psi]$ is a finite order mapping class. Conversely, it is not difficult to see that every finite order mapping class is represented by a finite order homeomorphism; essentially, since Teichmüller space is contractible, every mapping class of finite order must have a fixed point, see ([13], Thm. 7.1) for details. More generally, a celebrated result of Kerkchoff [22] states that every finite subgroup of $\mathrm{Mod}(S)$ is realized by a finite group of surface homeomorphisms.

Example 5.5: (Dehn twist.) Consider the annulus $A = ([0, n] \times [0, 1])/\sim$, where $(0, y) \sim (n, y)$. Let $T : A \to A$ be the affine homeomorphism of A that takes the vector $(0, 1)$ to the vector $(n, 1)$, so

$$T = \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix}.$$

Let $\gamma \in \mathcal{C}(S)$ and let N_γ be a regular neighbourhood of γ. Choose an orientation-preserving homeomorphism $h : A \to N_\gamma$. The *right Dehn twist* t_γ about γ is defined as

$$t_\gamma(x) = \begin{cases} hTh^{-1}(x), & x \in N_\gamma \\ x, & x \notin N_\gamma \end{cases}$$

Observe that t_γ is only well-defined as a mapping class.

One of the many reasons why Dehn twists constitute an important type of mapping class is because they generate the mapping class group. In fact, one has more:

Theorem 5.6: *(Dehn-Lickorish.)* $\mathrm{Mod}(S)$ *is generated by finitely many Dehn twists.*

We refer the reader to ([13], Ch. 4) for a proof of Theorem 5.6, and for explicit examples of Dehn twists that generate $\mathrm{Mod}(S)$.

We now introduce another important type of mapping classes, namely *pseudo-Anosov* mapping classes, by means of an example due to Thurston [25].

Example 5.7: (Pseudo-Anosov.) Let α and β be two curves that fill S, and choose representatives of α and β that realize $i(\alpha, \beta)$. Since α and β fill S, every connected component of $S \setminus (\alpha \cup \beta)$ is a topological disk. We thus obtain a cell decomposition of S whose vertices are precisely the $i(\alpha, \beta)$ points of intersection between α and β, and whose 1-cells are contained in $\alpha \cup \beta$.

Now consider the dual cell complex $\mathcal{D}$ of this cell decomposition. The complex $\mathcal{D}$ is also a cell decomposition of S, whose 2-cells correspond precisely to the $i(\alpha, \beta)$ points of intersection of α and β. By deeming each 2-cell of $\mathcal{D}$ to be a Euclidean square, we obtain a *singular Euclidean structure* on S: away from the vertices of $\mathcal{D}$ the metric is locally Euclidean, and at the vertices of $\mathcal{D}$ there are *cone singularities*, each with cone angle $k\pi$ for some $k \geq 2$: this is explained in more detail in Leininger's article [23].

We choose geodesic representatives of α and β in this singular Euclidean structure, so that α and β bisect each square through the midpoints of parallel edges, and intersect each other at the centre of each square. See Figure 5 for a example on a closed surface of genus 2.

Now, the Dehn twists t_α and t_β act as affine transformations of this singular Euclidean structure, namely by the matrices

$$t_\alpha = \begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad t_\beta = \begin{pmatrix} 1 & 0 \\ -n & 1 \end{pmatrix},$$

where $n = i(\alpha, \beta)$. Therefore

$$t_\alpha t_\beta^{-1} = \begin{pmatrix} 1+n^2 & n \\ n & 1 \end{pmatrix}.$$

The matrix of $t_\alpha t_\beta^{-1}$ has two real eigenvalues $\lambda, \lambda^{-1} \neq 1$. The corresponding eigenvectors are orthogonal, and thus determine a pair of orthogonal *singular foliations* of the surface, with singularities at the vertices of $\mathcal{D}$; again, see Leininger's article [23] for a thorough explanation of this. Now, $t_\alpha t_\beta^{-1}$ preserves these foliations, and expands along one by a factor of λ and contracts along the other by a factor of λ^{-1}.

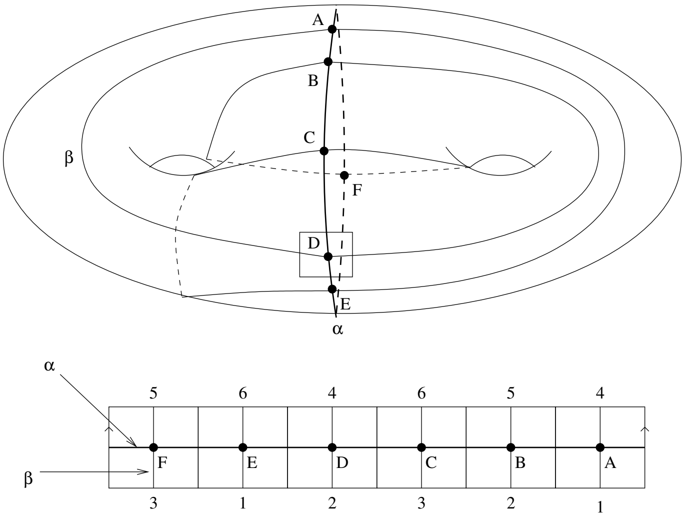

Fig. 5. Top: Two curves α and β that fill a genus 2 surface S; here $i(\alpha, \beta) = 6$. Bottom: The dual cell decomposition of S determined by α and β. We obtain S by gluing the top and bottom of the rectangle, and then the vertical sides according to the labelling given.

As mentioned before, the mapping class $t_\alpha t_\beta^{-1}$ is an example of a *pseudo-Anosov* mapping class. In general, a pseudo-Anosov mapping class comes equipped with a pair of orthogonal *singular foliations* of the surface (see Leininger's article [23] for a detailed exposition of foliations on surfaces), and expands along one by a fixed factor λ and contracts along the other by a factor of λ^{-1}. It is not difficult to see that, as a consequence, a pseudo-Anosov mapping class does not fix any non-trivial simple closed curves on S.

Nielsen-Thurston's classification of mapping classes. A mapping class ϕ may fix a non-trivial multicurve on the surface (e.g. a Dehn twist) or it may not (e.g. a pseudo-Anosov mapping class). In the former case, ϕ is said to be *reducible*; in the latter case, ϕ is called *irreducible.* Observe that a finite order mapping class may be reducible or irreducible; see ([13], Ch. 13.2.2) for specific examples. The celebrated *Nielsen-Thurston classification* of elements of $\mathrm{Mod}(S)$ asserts that every irreducible mapping class of infinite order is pseudo-Anosov. Namely:

Theorem 5.8: *(Nielsen-Thurston classification) Let $\phi \in \mathrm{Mod}(S)$. Then ϕ is either periodic, reducible or pseudo-Anosov.*

We point the reader to [13, 14] and the references therein for a detailed exposition on the Nielsen-Thurston classification of mapping classes.

Using the Nielsen-Thurston classification, one may give a complete description of the structure of a general mapping class, as we now briefly explain; see ([13], Ch. 13) for more details. A *reduction system* for $\phi \in \mathrm{Mod}(S)$ is a multicurve $\mu \subset S$ such that $\phi(\mu) = \mu$. Now, ϕ fixes a canonically defined multicurve μ_ϕ on S, namely the intersection of all *maximal* (with respect to inclusion) reduction systems for ϕ; following [7], the multicurve μ_ϕ is called the *canonical reduction system* of ϕ. The mapping class ϕ may permute the elements of μ_ϕ, as well as the connected components of $S \setminus \mu_\phi$. However, there exists $n \in \mathbb{N}$ such that ϕ^n does not permute the components of μ_ϕ or $S \setminus \mu_\phi$; observe that n is uniformly bounded above in terms of the genus of S. Then, ϕ^n acts as a power of a Dehn twist about each component of μ_ϕ , and the restriction of ϕ^n to each connected component S' of $S \setminus \mu_\phi$ is either the identity or a pseudo-Anosov mapping class of $\mathrm{Mod}(S')$.

5.2. *The action of* $\mathbf{Mod}(S)$ *on* $\mathcal{T}(S)$

The mapping class group acts naturally on Teichmüller space, namely if $\psi \in \mathrm{Mod}(S)$ and $[(X, f)] \in \mathcal{T}(S)$ then $\psi([(X, f)]) = [(X, f \circ g^{-1})]$, where g denotes any representative of ψ. In terms the character variety $\mathcal{D}F(\pi_1(S), \mathrm{PSL}(2, \mathbb{R}))$, the action of $\mathrm{Mod}(S)$ on $\mathcal{T}(S)$ is given by Theorem 5.3. As mentioned above, the main goal of this section is to prove the following result:

Theorem 5.9: $\mathrm{Mod}(S)$ *acts on* $\mathcal{T}(S)$ *properly discontinuously.*

Proof: Suppose, for contradiction, that there exist a compact set $K \subset \mathcal{T}(S)$ and a sequence $(\psi_n)_n$ of *distinct* elements of $\mathrm{Mod}(S)$ such that

$$\psi_n(K) \cap K \neq \emptyset, \text{ for all } n \in \mathbb{N}.$$

Thus, there exists a sequence $(X_n)_n$ of elements of K such that $\psi_n(X_n) \in K$ for all $n \in \mathbb{N}$. Let α and β be two curves that fill S. On the one hand, Lemma 4.6 implies that there exists $R = R(K) > 0$ such that $l_\alpha(X) + l_\beta(X) \leq R$ for all $X \in K$. On the other hand, we will show that, up to relabelling α and β,

$$\lim_{n\to\infty} l_{\psi_n^{-1}(\alpha)}(X_n) = \infty.$$

Having showed this, we will obtain the desired contradiction since, by definition, $l_{\psi_n^{-1}(\alpha)}(X_n) = l_\alpha(\psi_n(X_n))$.

Claim 1. *At least one of* $(\psi_n^{-1}(\alpha))_n$ *and* $(\psi_n^{-1}(\beta))_n$ *has a non-constant subsequence.*

Proof of Claim 1. Indeed, suppose this were not the case. Then, up to taking a subsequence, there are simple closed curves α' and β' on S such that $\psi_n^{-1}(\alpha) = \alpha'$ and $\psi_n^{-1}(\beta) = \beta'$ for all n. Therefore, up to the action of $\mathrm{Mod}(S)$, we may assume that $\psi_n^{-1}(\alpha) = \alpha$ and $\psi_n^{-1}(\beta) = \beta$ for all n. Now, α and β fill S, and so each component of $S \setminus (\alpha \cup \beta)$ is a topological disk. Since ψ_n fixes both α and β, its action on S is determined by the induced permutation on the set of disks of $S \setminus (\alpha \cup \beta)$. As S is compact, there are only finitely many such disks, and we have a contradiction to the ψ_n being pairwise distinct. Thus our claim follows.

Hence, up to relabelling and extracting a subsequence, we may assume that $(\psi_n^{-1}(\alpha))_n$ is a sequence of distinct curves on S; in order to simplify notation, we will write $\alpha_n = \psi_n^{-1}(\alpha)$. Next, we claim:

Claim 2. *There exists a pants decomposition* P *of* S *such that* $i(\alpha_n, \gamma) \to \infty$ *for some* $\gamma \in P$.

Proof of Claim 2. Choose any pants decomposition Q and suppose that, for all $\gamma \in Q$, $i(\alpha_n, \gamma)$ is uniformly bounded. Therefore the number of arcs of α_n in the complement of Q is bounded independently of n. As a consequence, the curves α_n differ only up to Dehn twisting about some component of Q; more formally, again up to extracting a subsequence, there exists a simple closed curve α' on S such that $\alpha_n = T_n(\alpha')$, for some $T_n \in \mathbb{T}_Q$, the subgroup of $\mathrm{Mod}(S)$ generated by the Dehn twists on the components of Q. As the α_n are pairwise distinct, then $T_n \neq T_m$ for $n \neq m$; moreover, up to extracting a subsequence, there exists a curve $\gamma' \in Q$ such that every T_n is supported on a submulticurve of Q containing γ'. Let γ be a curve in $S - (Q - \gamma')$ such that either $i(\gamma, \gamma') = 1$ if $S - (Q - \gamma')$ contains a one-holed torus, or $i(\gamma, \gamma') = 2$ if $S - (Q - \gamma')$ contains a four-holed sphere; compare with Figure 4, where β_i and γ_i play the role of γ and γ' respectively. Setting $P = (Q - \gamma') \cup \gamma$, we obtain the desired result.

Continuing with the proof of the main result, we may choose a curve γ on S such that $i(\alpha_n, \gamma) \to \infty$, by Claim 2. Now, there exists $\epsilon = \epsilon(K) > 0$ such that, for all $X \in K$, the ϵ-neighborhood of γ in X is an embedded annulus in X; this may be seen explicitly by considering the construction of a hyperbolic pair of pants from hyperbolic hexagons and using that K is compact, and is also an easy consequence of the *Collar Lemma* of Keen

[21] and Halpern [15]. Since $X_n \in K$ for all n, we have

$$l_{\psi_n^{-1}(\alpha)}(X_n) \geq \epsilon \cdot i(\psi_n^{-1}(\alpha), \gamma) \to \infty,$$

which gives the desired contradiction. This finishes the proof of Theorem 5.9. □

References

1. W. Abikoff, The real analytic theory of Teichmüller space. Lecture Notes in Mathematics, 820. *Springer*, Berlin, 1980.
2. J.W. Anderson. *Hyperbolic geometry.* Springer Undergraduate Mathematics Series. Springer-Verlag London, Ltd., London, 1999.
3. W. Ballmann. *Lectures on spaces of nonpositive curvature. With an appendix by Misha Brin.* DMV Seminar, 25. Birkhuser Verlag, Basel, 1995.
4. A. F. Beardon, The geometry of discrete groups. Graduate Texts in Mathematics, 91. *Springer-Verlag*, New York, 1983.
5. R. Benedetti & C. Petronio. *Lectures on hyperbolic geometry.* Universitext. Springer-Verlag, Berlin, 1992. xiv+330 pp.
6. M. Bestvina, K. Bromberg, K. Fujiwara, J. Souto, Shearing coordinates and convexity of length functions. Preprint, 2009.
7. J. Birman, A. Lubotzky and J. McCarthy, *Abelian and solvable subgroups of the mapping class groups*, Duke Math. J. 50 (1983).
8. F. Bonahon, Low-dimensional geometry. *From Euclidean surfaces to hyperbolic knots.* Student Mathematical Library, 49. IAS/Park City Mathematical Subseries. American Mathematical Society, Providence, RI; Institute for Advanced Study (IAS), Princeton, NJ, 2009. xvi+384 pp.
9. M. R. Bridson & A. Haefliger. *Metric spaces of non-positive curvature.* Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 319. Springer-Verlag, Berlin,
10. R. D. Canary, D. B. A. Epstein & P. L. Green. *Notes on notes of Thurston With a new foreword by Canary.* London Math. Soc. Lecture Note Ser., 328, Fundamentals of hyperbolic geometry: selected expositions, 1115, Cambridge Univ. Press, Cambridge, 2006,
11. A. J. Casson & S. A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston.* London Mathematical Society Student Texts, 9. Cambridge University Press, Cambridge, 1988. iv+105 pp
12. M. P. do Carmo. *Riemannian geometry.* Mathematics: Theory & Applications. Birkhäuser Boston, Inc., Boston, MA, 1992
13. B. Farb & D. Margalit. *A primer on mapping class groups.* Princeton University Press, to appear.
14. A. Fathi, F. Laudenbach & V. Poenaru, editors. *Travaux de Thurston sur les surfaces*, second edition, Astérisque 66, Soc. Math. France, Paris (1991)
15. N. Halpern, A proof of the collar lemma. *Bull. London Math. Soc.* **13** (1981), no. 2, 141144.
16. J. H. Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1. Teichmüller theory.* Matrix Editions, Ithaca, NY, 2006.

17. Y. Imayoshi & M. Taniguchi. *An introduction to Teichmüller spaces.* Springer-Verlag, Tokyo, 1992. xiv+279 pp.
18. N. V. Ivanov, Nikolai. *Mapping class groups.* Handbook of geometric topology, 523633, North-Holland, Amsterdam, 2002.
19. G. A. Jones & D. Singerman. *Complex functions. An algebraic and geometric viewpoint.* Cambridge University Press, Cambridge, 1987
20. S. Katok. *Fuchsian groups.* Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1992.
21. L. Keen, Collars on Riemann surfaces. *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pp. 263268. Ann. of Math. Studies, No. 79, Princeton Univ. Press, Princeton, N.J., 1974.
22. S. P. Kerckhoff, The Nielsen realization problem. *Ann. of Math. (2)* **117** (1983), no. 2, 235265.
23. C. J. Leininger, Degeneration of hyperbolic structures on surfaces. Also in this volume.
24. C. M. Series *Hyperbolic Geometry.* Online notes, available from www.maths.warwick.ac.uk/ masbb
25. W. P. Thurston, On the geometry and dynamics of diffeomorphisms of surfaces, *Bull. Amer. Math. Soc*, **19**(2), 1988.
26. W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1.* Princeton Mathematical Series, 35. Princeton University Press, Princeton, NJ, 1997.