Applications of Teichmüller theory to hyperbolic 3-manifolds

Ursula Hamenstädt

March 5, 2007

Chapter 1

Introduction

Let S be a closed oriented surface of genus $g \geq 2$. As is well known, such a surface admits a *hyperbolic metric*, i.e. a Riemannian metric of constant curvature -1. The group of all diffeomorphisms acts via pull-back on the space of all such metrics. The quotient of the space of all hyperbolic metrics on S under the action of the group of all diffeomorphisms of S which are isotopic to the identity is the called the *Teichmüller space* $\mathcal{T}(S)$ of S. It admits a natural structure of a complex manifold which is biholomorphic to a bounded domain in $\mathbb{C}^{3g-3}$.

Let $\pi_1(S)$ be the fundamental group of S. To every point $h \in \mathcal{T}(S)$ corresponds a conjugacy class of an injective homomorphism $\rho : \pi_1(S) \to PSL(2, \mathbb{R})$ with discrete image where as usual, $PSL(2, \mathbb{R})$ is the group of orientation preserving isometries of the *hyperbolic plane* $\mathbf{H}^2$. Two distinct points in $\mathcal{T}(S)$ define distinct conjugacy classes of such representations. As a consequence, the Teichmüller space of S can be identified with a subset of the quotient of the space of representations of $\pi_1(S)$ into $PSL(2, \mathbb{R})$ under the action of $PSL(2, \mathbb{R})$ by conjugation. It turns out that the set of these representations is open and dense in the representation variety.

Now the group $PSL(2, \mathbb{R})$ is naturally a subgroup of the group $PSL(2, \mathbb{C})$ of all orientation preserving isometries of hyperbolic 3-space $\mathbf{H}^3$. Thus every point in $\mathcal{T}(S)$ determines up to conjugacy a discrete torsion free subgroup of $PSL(2, \mathbb{C})$ and hence a 3-dimensional hyperbolic quotient manifold which is diffeomorphic to $S \times \mathbb{R}$. However, not every hyperbolic 3-manifold diffeomorphic to $S \times \mathbb{R}$ is of this form. Indeed, by a classical theorem of Bers (see [**?**]), there is an open subset of the space of all conjugacy classes of representations $\pi_1(S) \to PSL(2, \mathbb{C})$ which can naturally be parametrized by $\mathcal{T}(S) \times \mathcal{T}(S)$ and where the representations obtained from the points in $\mathcal{T}(S)$ via the embedding $PSL(2, \mathbb{R}) \to PSL(2, \mathbb{C})$ correspond to the diagonal in $\mathcal{T}(S) \times \mathcal{T}(S)$.

The image groups of these representations can be characterized by geometric propertis of the corresponding quotients of $\mathbf{H}^3$. Namely, for every subgroup Γ of $PSL(2,\mathbb{C})$, the *limit set* of Γ is the set of all accumulation points in the ideal boundary S^2 of $\mathbf{H}^3$ of an orbit Γx, where $x \in \mathbf{H}^3$ is any fixed point. This limit set Λ does not depend on the choice of x, and its is a closed Γ-invariant subset of S^2. The representation $\rho : \pi_1(S) \to PSL(2,\mathbb{C})$ corresponds to a point $(g,h) \in \mathcal{T}(S) \times \mathcal{T}(S)$ if and only if the limit set of $\rho(\pi_1(S))$ is a topological circle. The *domain of discontinuity* $\Omega = S^2 - \Lambda$ then consists of two topological discs in S^2, and Γ acts on these discs as a group of biholomorphic automorphisms. By the Riemann mapping theorem, each of these discs is biholomorphic to the unit disc in $\mathbb{C}$ and hence the restriction of Γ to these discs preserves a hyperbolic metric. In other words, the action of Γ on these discs defines two conjugacy classes of homomorphisms of $\pi_1(S)$ into $PSL(2,\mathbb{R})$ and hence two points in Teichmüller space. This pair of points is just the pair $(g,h) \in \mathcal{T}(S) \times \mathcal{T}(S)$ mentioned above. The points (g,h) are called the *end invariants* of the hyperbolic 3-manifold $\mathbf{H}^3/\rho(\pi_1(S))$.

As hyperbolic metrics on S can degenerate, representations of $\pi_1(S)$ into the group $PSL(2,\mathbb{C})$ can also degenerate. However, the latter case is more complicated. Namely, there are "degenerate" representations of $\pi_1(S)$ into $PSL(2,\mathbb{C})$ which are discrete and injective and such that moreover the quotient manifold $M = \mathbf{H}^3/\rho(\pi_1(S))$ is homeomorphic to $S \times \mathbb{R}$. A particularly nice class of such examples arise as infinite cyclic coverings of closed hyperbolic 3-manifolds which fibre over the circle. It turns out that one can associate to such a manifold M a pair of end invariants as well, one for each of the two ends of M. Such an invariant can either be a point in Teichmüller space for S or a point in a topological space which describes a particular family of degenerations of marked hyperbolic structures on surfaces.

If in addition the injectivity radius of M is bounded from below by a positive constant, then the invariants associated to the degenerated ends are points in the *Gromov boundary* of the *curve graph* of S. We explain this fact in detail in Section 6 below. The *ending lamination conjecture* which was in a much more general form formulated by Thurston and proved by Minsky [**?**] and Brock, Canary and Minksy states that a hyperbolic 3-manifold M which is homeomorphic to $S \times \mathbb{R}$ is determined up to marked isometry by its end invariants. Under the additional assumption that the injectivity radius of M is bounded from below by a positive constant, the solution of this conjecture was earlier established by Minksy [**?**].

The purpose of these lecture notes is to present a proof of this special case together with most of the background material. We only assume that the reader is familiar with basic topology and geometry. The proof given here is a variation of the original one. It uses some of the tools which were developed after the paper [**?**] was published and is meant to serve as an introduction to the ideas which lead to the proof of the general case.

Chapter 2

Foundations

2.1 Mapping class groups

Throughout, S denotes a *closed* oriented surface of genus $g \geq 2$. As in the introduction, the Teichmüller space $\mathcal{T}(S)$ of S is defined to be the quotient of the space of all hyperbolic metrics on S under the action of the group $\mathrm{Diff}_0^+(S)$ of all diffeomorphisms of S which are isotopic to the identity. Note that $\mathrm{Diff}_0^+(S)$ is a normal subgroup of the group $\mathrm{Diff}^+(S)$ of all orientation preserving diffeomorphisms of S.

Definition 2.1.1. The group $\mathrm{Diff}^+(S)/\mathrm{Diff}_0^+(S)$ of isotopy classes of orientation preserving diffeomorphisms of S is called the *mapping class group* $\mathrm{Mod}(S)$ of S.

By definition, the mapping class group $\mathrm{Mod}(S)$ acts on $\mathcal{T}(S)$. Moreover, two (equivalence classes of) hyperbolic metrics $g, h \in \mathcal{T}(S)$ are contained in the same orbit of the action of $\mathrm{Mod}(S)$ if and only if the hyperbolic metrics g, h on S are isometric.

A diffeomorphism of S maps the fundamental group $\pi_1(p, S)$ of all homotopy classes of loops based at the point $p \in S$ isomorphically onto the fundamental group $\pi_1(\varphi(p), S)$ based at the point $\varphi(p)$. Thus φ induces an isomorphism of the fundamental group $\pi_1(S)$, determined by a fixed choice of a basepoint, which is well defined up to a change of the basepoint, i.e. up to an inner automorphism of $\pi_1(S)$. Moreover, if ψ is isotopic to φ then the conjugacy classes of the isomorphisms of $\pi_1(S)$ induced by φ and ψ coincide. Thus if we denote by $\mathrm{Out}(\pi_1(S))$ the *outer automorphism group* of $\pi_1(S)$ which is the quotient of the group of all automorphisms of $\pi_1(S)$ under the normal subgroup of *inner automorphisms*, then there is a homomorphism $\mathrm{Mod}(S) \to \mathrm{Out}(\pi_1(S))$. Its image is contained in the subgroup $\mathrm{Out}^+(\pi_1(S))$ of index 2 of all outer automorphisms which act trivially on the second cohomology group of $\pi_1(S)$. The following result is due to Dehn (see [?] for more and for references).

Theorem 2.1.2. *The natural homomorphism* $\mathrm{Mod}(S) \to \mathrm{Out}^+(\pi_1(S))$ *is an isomorphism.*

The theorem of Dehn relates in a direct way algebraic information- the purely algebraically defined group $\mathrm{Out}^+(\pi_1(S))$- to more geometric information, namely isometry classes of hyperbolic metrics on S).

A particularly important family of mapping classes is described in the next example.

Example 2.1.3. Dehn twists
Let $\varphi : A = \mathbb{R}/\mathbb{Z} \times [-1, 1] \to S$ be a smooth orientation preserving embedding of a standard annulus into S whose *core curve* $\gamma : t \to \varphi(t, 0)$ is not contractible. Let $\theta : \mathbb{R} \to \mathbb{R}$ be a smooth non-decreasing function which satisfies $\theta(t) = 0$ for $t \le -1$ and $\theta(t) = 2\pi$ for $t \ge 1$. Define a diffeomorphism $D_\gamma : S \to S$ by

$$D_\gamma(s) = \begin{cases} s & \text{if } s \notin \varphi(A); \\ \varphi(x + \theta(y), y) & \text{if } s = \varphi(x, y). \end{cases}$$

It is easy two see that the isotopy class of this diffeomorphism only depends on the free homotopy class of the core curve γ and therefore the above definition determines an element of $\mathrm{Mod}(S)$ which is called a *positive Dehn twist* along γ (Chapter 4 of [?]). The inverse of a positive Dehn twist is called a *negative Dehn twist*. If ζ is any compact arc in the annulus $\varphi(A)$ which connects the two boundary components, then the arc $D_\gamma(\zeta)$ is not homotopic to ζ with fixed endpoints. This implies that a positive Dehn twist about γ is a non-trivial element of $\mathrm{Mod}(S)$.

The following result is due to Dehn and Lickorich (see Theorem 4.6 in [?]). For its formulation, we use the following.

Definition 2.1.4. A closed curve c on S is called *simple* if it does not have self-intersections. A simple closed curve c is *non-separating* (or *separating*) if the surface obtained from S by cutting S open along c is connected (or disconnected).

The theorem of Dehn and Lickorich can now be stated as follows.

Theorem 2.1.5. $\mathrm{Mod}(S)$ *is generated by Dehn twists along* $3g - 1$ *non-separating simple closed curves.*

2.2 Quadratic differentials

Let Σ be a closed *Riemann surface*, i.e. a closed surface equipped with a complex structure. This means that Σ can be covered by charts U_i with coordinate functions $\varphi_i : U_i \to \mathbb{C}$ such that the coordinate changes are biholomorphic. Using these charts, we can define.

Definition 2.2.1. A *quadratic differential* for Σ is given by a set of holomorphic functions f_i in local holomorphic coordinates z_i which satisfy the transition rules

$$f_i(z_i)dz_i^2 = f_j dz_j^2 \text{ where } \frac{dz_j}{dz_i}dz_i = dz_j.$$

Thus a quadratic differential is nothing else but a holomorphic section of the tensor product of the cotangent bundle of Σ with itself.

A quadratic differential q has finitely many zeros. Near a regular point x where q does not vanish we can define a special coordinate z by requiring that in this coordinate, the differential is just the usual holomorphic quadratic differential dz^2. Such a coordinate is obtained by chosing first a local representative fdz^2 and replacing z by the function w defined by $w(z) = \int_x^z \sqrt{f}dz$ where $\sqrt{f}$ is any locally defined branch of the square root function and where x is a fixed point contained in the coordinate chart. Such a coordinate is unique up to translation and multiplication with -1.

Now translations and reflections are euclidean isometries and hence the quadratic differential defines via pull-back of the canonical charts near each regular point an euclidean Riemannian metric which is compatible with the complex structure. This euclidean structure extends to the zeros of q as follows.

For the moment h be an arbitrary euclidean metric on the complement of finitely many points $x_1, \dots, x_s \in S$. A *singular point* x for h of *cone angle* $p\pi$ for an integer $p \geq 3$ is a point $x \in \{x_1, \dots, x_s\}$ which admits a neighborhood isometric to the union of p closed circular euclidean half-discs $D_1, \dots, D_p$ glued with orientation reversing isometries along their boundaries so that the total angle at the point x equals $p\pi$. For example, a singular point of cone angle 3π is obtained by cutting the euclidean plane open along the positive real line and gluing a half-space along the boundary of the split plane. We call a singularity of cone angle $p\pi$ for some $p \geq 3$ a *p-pronged singularity.* A 3-pronged singularity is shown on the right hand side of Figure 2.A.

Definition 2.2.2. A *singular euclidean metric* on S consists of an euclidean metric on the complement of finitely many points $x_1, \dots, x_\ell \in S$ and such that every point x_i is a singularity of cone angle $p_i\pi$ for some integer $p_i \geq 3$.

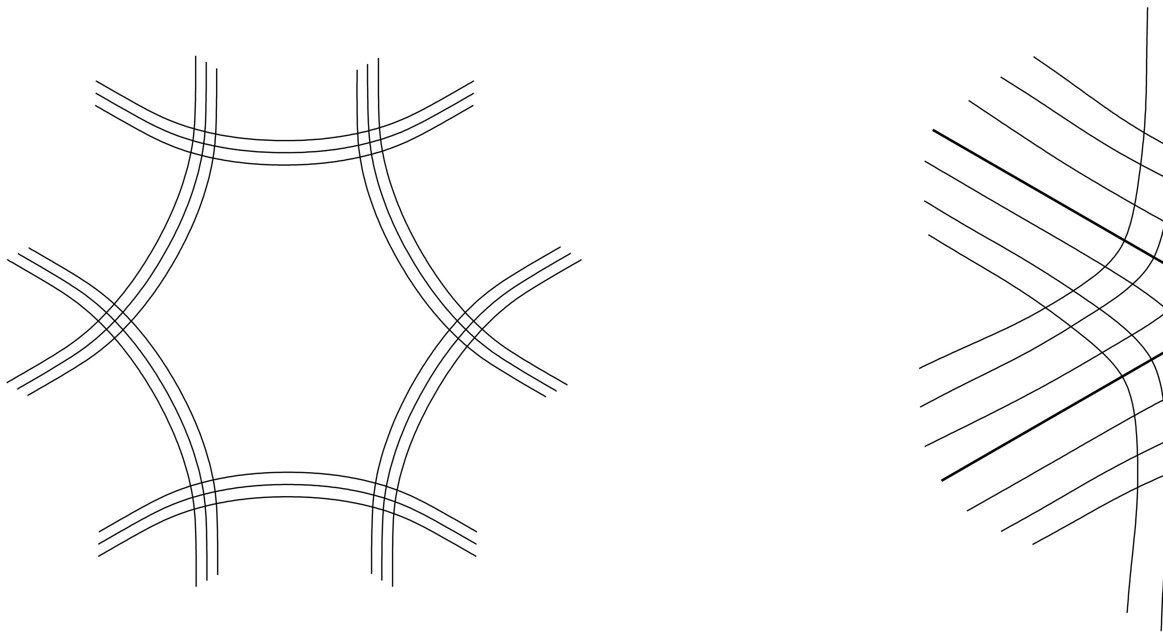

Figure 2.1

By the above consideration, a holomorphic quadratic differential q on the Riemann surface Σ defines near every regular point for q an euclidean metric. Near a zero of q there is a holomorphic chart such that in this chart the differential q can be written in the form $z^p dz^2$ for some $p \geq 1$. An easy calculation shows that the zero is a $p+2$-pronged singularity for the euclidean metric on the complement of the zeros of q defined by the quadratic differential. In other words, every holomorphic quadratic differential on Σ defines a singular euclidean metric. The conformal class of this metric is just the class of the complex structure.

Vice versa, let now h be any singular euclidean metric on S as in Definition 2.2.2. We claim that such a singular euclidean metric determines a marked complex structure and hence a hyperbolic metric on S as follows. If z is a regular point for this metric, then any orientation preserving isometric chart is declared to be holomorphic. Since an orientation preserving (local) isometry of the complex plane is biholomorphic, transition functions for these charts are biholomorphic and hence we obtain a complex structure on the complement of the singular points. At a p-pronged singularity x we can define a chart about x using the map $z \to z^{p/2}$ where x corresponds to the origin in the complex plane, and this chart is locally biholomorphic equivalent to the isometric charts away from x. By uniformization, the complex structure determines a unique hyperbolic metric on S. Note that a constant multiple of a singular euclidean metric is again a singular euclidean metric defining the same complex structure. Thus in the sequel we normalize all singular euclidean metrics to have area one.

If h is the singular euclidean metric defined by a holomorphic quadratic differential q then there is covering of the complement of the singular points by euclidean charts with transition functions of the form $z \to \pm z + c$ for some $c \in \mathbb{C}$. A surface equipped with a singular euclidean metric with this property is called a *half-translation surface*. Vice versa, every half-translation surface is a Riemann surface equipped with a holomorphic quadratic differential. The *area* of a quadratic differential is the area of the underlying singular euclidean metric.

Every quadratic differential q defines two natural transverse singular foliations on S whose set of singular points corresponds precisely to the set of singular points of q. The *horizontal foliation* is the pull-back of the foliation of $\mathbb{C}$ into lines parallel to the real axis under the canonical charts which determine the half-tranlation structure. A vector X tagent to a regular point on S is tangent to the horizontal foliation if and only if $q(X) \geq 0$. The *vertical foliation* is the pull-back of the foliation of $\mathbb{C}$ into lines parallel to the imaginary axis under the canonical charts. A vector tangent to a regular point on S is tangent to the vertical foliation if and only if $q(X) \leq 0$. Each singular point for q with cone angle $p\pi$ is a singular point of index $2 - p$ for the line field tangent to the foliation. In particular, by the Gauss-Bonnet theorem, the sum of the orders of the zeros of q equals $4g - 4$.

Chapter 3

The curve graph

3.1 Combinatorial geometry

3.1.1 Mapping tori

In this motivating section we discuss some elementary relations between the geometry and the topology of some particularly simple 3-manifolds constructed from surfaces in the following way.

Definition 3.1.1. The *mapping torus* M_ψ of a diffeomorphism $\psi \in \mathrm{Diff}^+(S)$ is defined by

$$M_\psi = S \times [0,1]/\sim \ \text{ where } (x,1) \sim (\psi(x),0).$$

It is not difficult to see that up to orientation preserving diffeomorphism, the mapping torus M_ψ of a diffeomorphism $\psi \in \mathrm{Diff}^+(S)$ only depends on the isotopy class of ψ. In particular, we can talk about the mapping torus of a mapping class, and we write M_φ for the mapping torus of a mapping class $\varphi \in \mathrm{Mod}(S)$. The mapping torus M_φ of φ has the following properties.

1. M_φ is a $K(\pi,1)$-space.
2. There is an exact sequence

   $$0 \to \pi_1(S) \to \pi_1(M_\varphi) \to \mathbb{Z} \to 0$$

   i.e. $\pi_1(S)$ is a normal subgroup of $\pi_1(M_\varphi)$ with factor group $\mathbb{Z}$.

3. There is a natural surjection $M_\varphi \to S^1$, i.e. M_φ *fibres over the circle.*

The surjection $M_\varphi \to S^1$ is obtained by mapping the equivalence class of $(x,t) \in S \times [0,1]$ to the equivalence class of t in the circle $S^1 = [0,1]/\sim$ obtained by identifying the two endpoints of the unit interval. By construction, the manifold M_φ admits a regular covering $\hat{M}$ with fundamental group $\pi_1(S)$ and deck transformation group $\mathbb{Z}$. This covering space is naturally diffeomorphic to $S \times \mathbb{R}$.

The above discussion motivates the following **basic question**: For a given mapping class $\varphi \in \mathrm{Mod}(S)$, how are
- the combinatorial properties of φ
- the "geometry" of the mapping torus M_φ (for a smooth Riemannian metric)
- the topology of M_φ
related?

3.1.2 Geometric models

One method to approach this question is to construct an "easy to understand" combinatorial model for the geometry of M_φ and relate this model to the topology of M_φ. A combinatorial model for the geometry of M_φ can for example be obtained from a geometric model for the fundamental group of M_φ. A geometric correspondence between two geometric models is given by a *quasi-isometry* which we define next.

Definition 3.1.2. An *L-quasi-isometric embedding* of a metric space (X,d) into a metric space (Y,d) is a map $F : X \to Y$ such that

$$d(x,y)/L - L \leq d(Fx,Fy) \leq Ld(x,y) + L$$

for all $x,y \in X$. F is an *L-quasi-isometry* if moreover for every $y \in Y$ there is some $x \in X$ with $d(F(x),y) \leq L$. An L- quasi-isometric embedding of a connected subset of the real line into a metric space Y is called an *L-quasi-geodesic* in Y.

A typical example of two quasi-isometric metric spaces can be obtained as follows. Let d,d' be two distance functions on a set X. If there is a number $L > 0$ such that

$$d(x,y)/L \leq d'(x,y) \leq Ld(x,y) \text{ for all } x,y \in X$$

then the identity map $(X,d) \to (X,d')$ is an L-quasi-isometry.

An action of a group Γ on a locally compact topological space X is called *proper* if for every compact subset K of X there are only finitely many elements

$g \in \Gamma$ with $gK \cap K \neq \emptyset$. A *finitely presented* group Γ admits a proper cocompact action on a simply connected locally compact complete metric space X. Namely, such a metric space can be obtained as the universal covering X of a finite simplicial complex with fundamental group Γ, equipped with the piecewise euclidean metric which is defined by the simplicial structure. There is a fundamental relation between the structure of Γ and the geometry of X. To explain this relation we first collect some basic properties of finitely generated groups.

Namely, let Γ be a finitely generated group with a finite *symmetric* generating set $\mathcal{G}$. This means that $\mathcal{G}$ contains with an element h also its inverse h^{-1}. Then every element of Γ can be represented as a product of elements in $\mathcal{G}$. In other words, it can be written as a *word* in the letters $h \in \mathcal{G}$. The *word norm* $|g|$ of $g \in \Gamma$ is the minimal length of such a word representing g; it depends on $\mathcal{G}$. Clearly $|gh| \leq |g||h|$ for all g, h and therefore

$$d(g,h) = |g^{-1}h|$$

defines a Γ-invariant metric on Γ. This means that d is a distance function which satisfies $d(ug, uh) = d(g, h)$ for all $u, g, h \in \Gamma$.

Up to quasi-isometry, this metric does not depend on the choice of the generating set. This is immediate from the following lemma.

Lemma 3.1.3. *The word norms $|\,|$ and $|\,|'$ of any two symmetric generating sets for Γ are equivalent: There is a constant $c > 0$ such that $|g|/c \leq |g|' \leq c|g|$ for all $g \in \Gamma$.*

Proof. Let $\mathcal{G}, \mathcal{G}'$ be two different symmetric generating sets for Γ defining word norms $|\,|, |\,|'$. Then every $h \in \mathcal{G}$ can be written as a word in the generators $\mathcal{G}'$. Since $\mathcal{G}$ is finite, there is a number $m > 0$ such that every $h \in \mathcal{G}$ can be written as such a word of length at most m. Now if $g \in \Gamma$ and $|g| = \ell$ then $g = h_1 \cdots h_\ell$ with $h_i \in \mathcal{G}$. Replacing each h_i by a word of length at most m in the generators $\mathcal{G}'$ yields a representation of g as a word in $\mathcal{G}'$ whose length is at most ℓm. Thus $|g|' \leq m|g|$. By symmetry, the lemma follows. $\square$

The above lemma says that finitely generated groups Γ admit a geometric model which is unique up to quasi-isometry, i.e. up to an L-quasi-isometry for some (unspecified) $L > 1$.

A *length space* is a metric space with the additional property that the distance between any two points is the infimum of the lengths of arcs connecting the points. Particular examples are manifolds equipped with a distance function defined by a Riemannian metric. By the theorem of Hopf Rinow (see p.40 of [**?**]), a locally compact complete length space is *geodesic* and *proper*. This means that any two points can be connected by a minimal geodesic, and closed subsets of bounded diameter are compact.

We indicated above that if Γ is a finitely presented group then there are proper length spaces X on which Γ acts isometrically, properly and cocompactly. Then for a point $x \in X$ which is not fixed by any nontrivial element of Γ, the map which associates to $g \in \Gamma$ the point $gx \in X$ defines an embedding of Γ into X and hence gives rise to a geometric model for Γ. This model turns out to be quasi-isometric to the model given by a word norm. This is the statement of a result of Švarc-Milnor which we present next (see also p. 140 of [**?**]).

Proposition 3.1.4. *(Švarc-Milnor): Let X be a locally compact length space and let Γ be any group which admits a proper isometric cocompact action on X. Then Γ is finitely generated, and for every $x \in X$ the orbit map $g \in \Gamma \to gx \in X$ is a quasi-isometry between Γ and X.*

Proof. Let d be the distance on X. The group Γ acts on the length space (X, d) properly, isometrically and cocompactly. This implies that (X, d) is complete and hence proper. Moreover, there is compact subset K of X such that $\cup_{g\in\Gamma} gK = X$ (see Section 8.4 of [**?**] for these simple facts).

Let D be the diameter of K and let B be the closed $3D$-neighborhood of K in X. Then B is compact and therefore the subset $\mathcal{G}$ of Γ consisting of all elements $g \in \Gamma$ with $gB \cap B \neq \emptyset$ is finite and moreover symmetric.

We claim that $\mathcal{G}$ generates Γ. Namely, fix a point x in K, let $g \in \Gamma$ and let $\gamma : [0, d(x, gx)] \to X$ be a geodesic connecting x to gx. Let $m \geq 0$ be the largest integer such that $mD \leq d(x, gx)$. For every $\ell \leq m$ there is some $g_\ell \in \Gamma$ such that $\gamma(\ell D) \in g_\ell(K)$, and hence with $d(\gamma(\ell D), g_\ell x) \leq D$. Write also $g_{m+1} = g$. Then $d(g_\ell x, g_{\ell+1}x) \leq 3D$ for all $\ell \leq m$ and consequently $g_\ell^{-1} g_{\ell+1} \in \mathcal{G}$ for all ℓ. Since $g_{m+1} = g$ this shows that $\mathcal{G}$ generates Γ.

Moreover, if $|\,|$ is the word norm on Γ induced by $\mathcal{G}$ then for every $g \in \Gamma$ we have $|g| \leq d(x, gx)/D + 1$. Therefore the map $F : \Gamma \to X$ defined by $F(g) = gx$ satisfies

$$d(Fg, Fh) = d(gx, hx) = d(x, g^{-1}hx) \geq D|g^{-1}h| - D. \tag{3.1}$$

Note that F is equivariant with respect to the action of Γ on itself by left translation and the isometric action of Γ on X.

On the other hand, since $\mathcal{G}$ is finite there is a constant $L > 0$ such that $d(gx, x) \leq L$ for all $g \in \mathcal{G}$ (in fact, we can choose $L = 8D$, however this is of no importance here). If $h = g_1 \dots g_k \in \Gamma$ can be represented as a word in the alphabet $\mathcal{G}$ of length k then

$$\begin{aligned} d(hx, x) &\leq d(hx, g_1 \dots g_{k-1}x) + d(g_1 \dots g_{k-1}x, x) \\ &= d(g_k x, x) + d(g_1 \cdots g_{k-1}x, x) \leq kL \end{aligned}$$

by an inductive application of the triangle inequality and the fact that the action of Γ is isometric. This shows that the map F is coarsely Lipschitz: For all $g, h \in \Gamma$ we have

$$d(F(g), F(h)) = d(x, g^{-1}hx) \leq L|g^{-1}h|. \tag{3.2}$$

(Note that this argument is valid for *every* metric space X which admits an isometric action by Γ.)

As a consequence of (3.1) and (3.2), the orbit map F is a quasi-isometric embedding of Γ equipped with the word norm defined by $\mathcal{G}$ into X. Since every point in X is at distance at most D from a point in $F\Gamma$, the map F is in fact a quasi-isometry. $\square$

By Proposition 3.1.4, if $\tilde{M}$ is the universal covering of a closed manifold M, then $\tilde{M}$ equipped with the distance induced by the lift of any length metric on M is quasi-isometric to the fundamental group of M. In particular, any two such distances are quasi-isometric.

3.1.3 Hyperbolic geodesic metric spaces

The geometry of the fundamental group $\pi_1(S)$ of a closed surface S of genus $g \geq 2$ has some particularly nice properties which we discuss in this subsection.

Definition 3.1.5. For a number $\delta > 0$, a geodesic metric space X is called *δ-hyperbolic* (in the sense of Gromov) if for every geodesic triangle in X with sides a, b, c the side a is contained in the δ-neighborhood of $b \cup c$.

A geodesic metric space is called *hyperbolic* (in the sense of Gromov) if it is δ-hyperbolic for some $\delta > 0$. We refer to the book [**?**] for the basic results on hyperbolic geodesic metric spaces and for references. Of particular importance is the following property.

Proposition 3.1.6. *Let X be a hyperbolic geodesic metric space. Then for every $L > 1$ there is a constant $R(L) > 0$ only depending on the hyperbolicity constant such that every L-quasi-geodesic of bounded diameter is contained in the $R(L)$-tubular neighborhood of a geodesic.*

As a fairly immediate corollary of Proposition 3.1.6 one obtains

Corollary 3.1.7. *Let X, Y be two quasi-isometric geodesic metric spaces. Then X is hyperbolic if and only if Y is hyperbolic.*

For a metric space X and $x \in X$, the *Gromov product* of two points $y, z \in X$ with respect to x is defined by

$$(y|z)_x = \frac{1}{2}(d(y,x) + d(z,x) - d(y,z)). \tag{3.3}$$

Let now X be hyperbolic. We say that a sequence of points $(y_i) \subset X$ *converges at infinity* if $(y_i|y_j)_x \to \infty$ $(i, j \to \infty)$. This does not depend on the choice of the point x. Two such sequences $(y_i), (z_j)$ are called *equivalent* if $(y_i|z_j)_x \to \infty$. It follows from the properties of a hyperbolic geodesic metric space that this notion of equivalence defines an equivalence relation for sequences converging at infinity which does not depend on the choice of the basepoint x. The space of equivalence classes is called the *Gromov boundary* ∂X of X. There is a topology on the union $X \cup \partial X$ which restricts to the given topology on X and such that X is an open and dense subset of $X \cup \partial X$. Moreover, this topology is natural in the sense that an isometry of X extends to a homeomorphism of $X \cup \partial X$. If X is proper then $X \cup \partial X$ is compact. The following useful proposition is fairly immediate from the definitions. We refer to the book [**?**] for its proof.

Proposition 3.1.8. *Let X, Y be hyperbolic geodesic metric spaces. Then every quasi-isometry $F : X \to Y$ extends to a homeomorphism $f : \partial X \to \partial Y$. If F is equivariant with respect to an isometric action of a group Γ on X and Y then the map f is Γ-equivariant as well.*

The following two examples are fundamental for our discussion. For later reference, note that any length metric on a compact space X induces a geodesic metric on the universal covering of X.

Example 3.1.9. The hyperbolic space
The n-dimensional hyperbolic space $\mathbf{H}^n$ is δ-hyperbolic for some $\delta > 0$. Its Gromov boundary $\partial\mathbf{H}^n$ is just the standard sphere S^{n-1}, and $\mathbf{H}^n \cup \partial\mathbf{H}^n$ is homeomorphic to a closed ball in $\mathbb{R}^n$.

A closed surface S of genus $g \geq 2$ admits a Riemannian metric of constant curvature -1 and hence its fundamental group $\pi_1(S)$ acts properly discontinuously, freely, isometrically and cocompactly on the hyperbolic plane $\mathbf{H}^2$. Thus by Proposition 3.1.4, $\pi_1(S)$ is quasi-isometric to $\mathbf{H}^2$. Now if q is any geodesic metric on a closed surface S of genus $g \geq 2$ inducing the usual topology then by Proposition 3.1.4, the universal covering $\tilde{S}$ of S equipped with the lift $\tilde{q}$ of q is quasi-isometric to the group $\pi_1(S)$ and hence to $\mathbf{H}^2$. In particular, every geodesic γ for the metric $\tilde{q}$ is an L-quasi-geodesic with respect to the hyperbolic metric for a number $L > 0$ only depending on q. By Proposition 3.1.6, this quasi-geodesic is contained in a uniformly bounded neighborhood of a hyperbolic geodesic. Moreover, if γ is one sided infinite then γ converges to a point in the Gromov boundary S^1 of $\mathbf{H}^2$.

Example 3.1.10. Mapping tori
Let M_φ be the mapping torus of $\varphi \in \mathrm{Mod}(S)$. Choose a hyperbolic metric σ

on S and a Riemannian metric on M_φ so that $(S,\sigma) \to S \times \{\frac{1}{2}\} \to M$ is an isometric embedding. By Proposition 3.1.4, the fundamental group of M_φ is equivariantly quasi-isometric to the universal covering $\tilde{M}_\varphi$ of M_φ. The space $\tilde{M}_\varphi$ is diffeomorphic to $\mathbb{R}^3$ and admits an isometric action by the fundamental group $\pi_1(S)$ of S, with quotient manifold $\hat{M}$ which is diffeomorphic to $S \times \mathbb{R}$. However, the quasi-isometry class of the metric on $\hat{M}$ depends on φ. Namely, if φ is the identity, then $\hat{M}$ is quasi-isometric to a Riemannian product $(S,\sigma) \times \mathbb{R}$ and hence for every fixed constant $\kappa > 0$ the number of free homotopy classes of curves which admit a representative of length at most κ is *finite*. On the other hand, for every closed curve c on S and for each $i \in \mathbb{Z}$, the minimal length of a closed curve in $\hat{M}$ which is freely homotopic to $\varphi^i(c)$ is *uniformly bounded, independent of* i. In particular, if there is a closed curve c on S such that the curves $\varphi^i c$ are pairwise not freely homotopic, then there are infinitely many free homotopy classes of curves which admit a representative of uniformly bounded length. This implies that $\hat{M}$ is not quasi-isometric to a Riemannian product $(S,\sigma) \times \mathbb{R}$.

This example motivates the following

Question: Can we recover from the "collection of shorts curves" in $\hat{M}$ the mapping class φ and hence the topology of M_φ?

To approach this question, we first equip the collection of all *simple* closed curves on the surface S with an additional structure.

3.2 The complex of curves

In this section we introduce a geometric space constructed from free homotopy classes of simple closed curves on S on which the mapping class group acts as a group of isometries. This space was introduced by Harvey to investigate the topology of the mapping class group. It turned out later that its geometry provides a powerful tool to obtain a better understanding of the Teichmüller space and the structure of hyperbolic 3-manifolds.

Definition 3.2.1. The *complex of curves* of the surface S is the simplicial complex whose *vertex set* $\mathcal{C}(S)$ is the set of all nontrivial free homotopy classes of non-oriented simple closed curves on S. A collection $\alpha_1, \dots, \alpha_k \subset \mathcal{C}(S)$ spans a simplex if and only if the curves $\alpha_1, \dots, \alpha_k$ can be realized disjointly. The *curve graph* $\mathcal{CG}(S)$ is the one-skeleton of the complex of curves.

In the sequel we often do not distinguish between a simple closed curve and its free homotopy class in our notation. The curve graph admits a natural structure of a *metric graph* with vertex set $\mathcal{C}(S)$ and where two simple closed curves are connected by an edge of length one if and only if they can be realized disjointly. We then can define the distance $d(\alpha, \beta)$ between two simple closed curves $\alpha, \beta \in \mathcal{C}(S)$ to be the infimum of the lengths of any path connecting α to β. Note that if this infimum is finite then it is an integer and hence the infimum is assumed. As a consequence, if $\mathcal{C}(S)$ is connected then d is a geodesic metric. Moreover, the mapping class group naturally acts on free homotopy classes of simple closed curves preserving disjointness and hence it acts on $\mathcal{CG}(S)$ as a group of simplicial isometries. In the spirit of the theorem of Švarc-Milnor, we therefore can try to relate the geometry of $\mathcal{C}(S)$ to the geometry of the Teichmüller space and the mapping class group with the goal to obtain a geometric understanding of the $\mathbb{Z}$-covers of mapping tori introduced in Section 3.1.

3.2.1 Distance estimates

As a first step toward an understanding of the geometry of the curve graph we show that the curve graph is connected, and we estimate the distance between any two of its vertices. For this we use the following auxiliary tool.

Definition 3.2.2. The *intersection number* $i(\alpha, \beta)$ between two simple closed curves α, β on S is the minimal number of intersection points between any two curves freely homotopic to α, β.

Let σ be a hyperbolic metric on S. Then every free homotopy class α on S can be represented by a unique closed geodesic for σ, and this geodesic is up to parametrization the unique curve of minimal length in this class. This implies in particular that if α can be represented by a simple closed curve, then the geodesic representative of α is simple. Moreover, the geodesic representatives of two classes α, β intersect in precisely $i(\alpha, \beta)$ points. We refer to Section 2 of [**?**] for a proof of this well known fact.

We use intersection numbers to show.

Lemma 3.2.3. *The curve graph is connected, and* $d(\alpha, \beta) \le i(\alpha, \beta) + 1$ *for all* $\alpha, \beta \in \mathcal{C}(S)$.

Proof. We show by induction on $m \ge 0$ that any two curves $\alpha, \beta \in \mathcal{C}(S)$ with $i(\alpha, \beta) \le m$ can be connected by a path in $\mathcal{CG}(S)$ of length at most $m + 1$.

If $i(\alpha, \beta) = 0$ then α, β can be realized disjointly and hence they are joined by an edge of length one. Thus the case $m = 0$ is immediate from the definitions. Assume that the claim holds true whenever $i(\alpha, \beta) \le m - 1$ for some $m \ge 1$. Let $\alpha, \beta \in \mathcal{C}(S)$ be such that $i(\alpha, \beta) = m$. Up to changing α, β by a homotopy we may assume that α, β intersect transversely in precisely $i(\alpha, \beta) = m$ points. Then the intersection of β with $S - \alpha$ consists of a collection of $i(\alpha, \beta)$ simple pairwise disjoint arcs with endpoints on α which are not homotopic with fixed endpoints to subarcs of α.

We distinguish two cases.

Case 1: There is a component β_0 of $\beta \cap (S - \alpha)$ which begins and ends on the same side of α in S.

Let ζ be the boundary of a small tubular neighborhood of $\alpha \cup \beta_0$. Then ζ is disjoint from α, moreover ζ is a union of three distinct connected components $\zeta_1, \zeta_2, \zeta_3$. Up to changing the numbering, ζ_1 is freely homotopic to α. The components ζ_2, ζ_3 are freely homotopic to the concatenation of a proper subarc of α with the arc β_0. They are both homotopically nontrivial since otherwise the arc β_0 is homotopic with fixed endpoints to a subarc of α. Since ζ_2 is disjoint from α, there is an edge in $\mathcal{CG}(S)$ with endpoints α and ζ_2. Moreover, an intersection point between ζ_2 and β corresponds to an intersection point between α and β which is distinct from the endpoints of β_0. Thus we have $i(\zeta_2, \beta) \leq i(\alpha, \beta) - 2 = m - 2$. By the induction hypothesis, ζ_2 can be connected to β by a path in $\mathcal{CG}(S)$ of length at most $m - 1 < m$ and hence α can be connected to β by a path of length at most $m + 1$.

Case 2: $\beta \cap (S - \alpha)$ does not contain any component which begins and ends on the same side of α.

We distinguish two subcases. In the first subcase, all the components of $\beta \cap (S - \alpha)$ are freely homotopic relative to α. Since β does not have self-intersections, its free homotopy class is prime and hence in this case there is only a single such arc. Then a tubular neighborhood of $\alpha \cup \beta$ is a torus with one boundary component ζ, and ζ is a simple closed separating curve in S which is disjoint from both α and β. Thus α and β can be connected in $\mathcal{CG}(S)$ by a path of length $2 = i(\alpha, \beta) + 1$.

If not all components of $\beta \cap (S - \alpha)$ are freely homotopic relative to α then there is a subarc $\tilde{\beta} = \beta_0 \cup \beta_1$ of β which begins and ends on α and whose intersection with $S - \alpha$ consists of two components β_0, β_1 with endpoints on α which are not freely homotopic relative to α. The boundary of a tubular neighborhood of $\alpha \cup \tilde{\beta}$ contains a component ζ which is homotopic to the composition of β_0, a subarc α_0 of α joining the endpoint of β_0 to the endpoint of β_1, the arc β_1^{-1} and a subarc α_1 of α which is disjoint from α_0 and joins the starting point of β_1 to the starting point of β_0. The simple closed curve ζ is disjoint from α, and the intersection number between ζ and β is at most $i(\alpha, \beta) - 1$. As above, this yields the induction step.

As a consequence, $\mathcal{CG}(S)$ is connected, and $d(\alpha, \beta) \leq i(\alpha, \beta) + 1$ for all $\alpha, \beta \in \mathcal{C}(S)$. $\square$

The next proposition is due to Bowditch [**?**] and relates intersection numbers to the distance in the curve graph in a more precise way.

Proposition 3.2.4. *There is a number $\kappa = \kappa(S) > 0$ such that we have $d(\alpha,\beta) \leq \kappa \log i(\alpha,\beta) + \kappa$ for all $\alpha,\beta \in \mathcal{C}(S)$.*

Proof. By Lemma 3.2.3, we have $d(\alpha,\beta) \leq i(\alpha,\beta) + 1$ for all $\alpha,\beta \in \mathcal{C}(S)$.

Let α,β be arbitrary simple closed curves, realized with the minimal number of intersections. We claim that there is a simple closed curve $\gamma \in \mathcal{C}(S)$ with $\max\{i(\alpha,\gamma), i(\beta,\gamma)\} \leq \sqrt{3i(\alpha,\beta)}$ provided that $i(\alpha,\beta) \geq 9$.

For this let $n = i(\alpha,\beta) \geq 9$ and let b be the integer part of $\sqrt{n}$. Then $\sqrt{n} - 1 \leq b \leq \sqrt{n}$ and hence if $n < b + 3$ then $n - 3 < \sqrt{n}$ and $n^2 - 7n + 9 < 0$ which implies that $n < 9$. Thus by our assumption $n \geq 9$ we have $n \geq b + 3$.

Choose a subarc β' of β containing exactly $b + 2$ points in $\alpha \cap \beta$. Let moreover $\alpha = \alpha_1 \cup \cdots \cup \alpha_\ell$ be a decomposition of α into $\ell \leq b + 1$ subarcs α_i containing each at most $b+1$ points of $\alpha \cap \beta$. Such a decomposition exists since $n = i(\alpha,\beta) \leq (b + 1)^2$. There is a number i such that $|\alpha_i \cap \beta'| \geq 2$. Let x, y be two intersections points of α_i with β' which are closest along α_i. This means that the subarc α_0 of α_i connecting x to y does not contain a point of $\alpha \cap \beta'$ in its interior. Let β_0 be the subarc of β' connecting y to x. Then $\gamma_0 = \alpha_0 \cup \beta_0$ is a simple closed curve. Since no subarc of β with endpoints on α is homotopic with fixed endpoints to a subarc of α, the curve γ_0 is homotopically nontrivial.

By construction, every intersection point of γ_0 and the curve α corresponds to an intersection point of β_0 and α. Since β_0 is a subarc of β', the number of such intersection points does not exceed $b + 2 \leq 2\sqrt{i(\alpha,\beta)}$. Moreover, every essential intersection point of γ_0 and the curve β is an essential intersection of α_0 and β and hence the number of these intersections is not bigger than $b + 1 \leq 2\sqrt{i(\alpha,\beta)}$. Thus γ_0 satisfies our above requirement.

Using this procedure, we can construct inductively for all $\alpha,\beta \in \mathcal{C}(S)$ with $i(\alpha,\beta) \geq 9$ a sequence $(\gamma_i) \subset \mathcal{C}(S)$ of length at most $\log_2(4i(\alpha,\beta))$ connecting α to β and such that for each i of distance between γ_i and γ_{i+1} is at most 10. This implies the proposition. ☐

3.2.2 Additional basic properties

In this subsection we derive some additional easy properties of the complex of curves which will be useful in later sections. First recall that a *pair of pants* is a planar surface diffeomorphic to an open disc with two closed discs removed from its interior.

Lemma 3.2.5. 1. *The complex of curves is locally infinite and of dimension* $3g-4$.

2. *A simplex of maximal dimension in the complex of curves is a* pants decomposition *of* S*: After cutting* S *open along the curves of the simplex we obtain* $2g-2$ *pairs of pants.*

Proof. To show that $\mathcal{CG}(S)$ is locally infinite, choose any non-separating simple closed curve α on S. Then $S-\alpha$ is a surface of genus $g-1\geq 1$ with 2 boundary circles. Such a surface contains infinitely many free homotopy classes of simple closed curves, and each such free homotopy class defines an element $\beta\in\mathcal{C}(S)$ with $d(\alpha,\beta)=1$ which is connected to α by an edge.

Let σ be a simplex in $\mathcal{C}(S)$. Then σ is spanned by $k>0$ vertices $\alpha_1,\dots,\alpha_k$, and these vertices are simple closed curves which can be realized disjointly. We have to show that $k\leq 3g-3$. For this we show by induction on $n=2g-2+m$ that a *bordered oriented* surface $S_{g,m}$ of genus $g\geq 0$ with $m\geq 0$ boundary circles and $2g-2+m\geq 1$ contains at most $\max\{3g-3+2m,0\}$ simple closed curves which are *essential*, i.e. non-contractible and not freely homotopic into the boundary, and which are not mutually freely homotopic. Moreover, a system of such curves of maximal cardinality decomposes $S_{g,m}$ into pairs of pants.

In the case $n=1$ we either have $g=0$ and $m=3$ and $S_{g.m}$ is a pair of pants, or $g=1,m=1$ and $S_{g,m}$ is a bordered torus with connected boundary. Cutting $S_{1,1}$ open along any essential simple closed curve results in a pair of pants, so there is nothing to show. Now assume that the claim is known for some $n\geq 1$ and consider a surface $S=S_{g,m}$ with $2g-2+m=n+1$. Let α be a simple closed curve on S and assume first that α is non-separating. Then the bordered surface $S-\alpha=S'$ obtained from S by cutting S open along α is connected. Clearly S' has $m+2$ boundary components. A calculation of the Euler characteristic shows that the genus of S' equals $g-1$. By induction, S' contains at most $3g-6+2m+2=3g-3+2m-1$ essential simple closed pairwise disjoint mutually not freely homotopic curves and hence S contains at most $3g-3+2m$ such curves. Moreover, a system of such curves of maximal cardinality decomposes S' into pairs of pants and hence the union of such a system with the curve α decomposes S into pairs of pants. If α is separating then $S-\alpha$ has two connected components S_1,S_2 to which we can apply the induction hypothesis as before. This shows that the dimension of the complex of curves equals $3g-4$ as claimed and also implies the second part of the lemma. □

In the next lemma, we look more closely at the embedded graph in S defined by two simple closed curves α,β with $d(\alpha,\beta)\geq 3$.

Lemma 3.2.6. *If* $\alpha,\beta\in\mathcal{C}(S)$ *and if* $d(\alpha,\beta)\geq 3$ *then* α,β jointly fill up S, *i.e.* $S-(\alpha\cup\beta)$ *is a union of topological discs.*

Proof. Assume that $\alpha, \beta \in \mathcal{C}(S)$ are simple closed curves which do not jointly fill up S and which satisfy $d(\alpha,\beta) \geq 2$. Choose a hyperbolic metric on S and assume that α, β are geodesics with respect to this metric. Since α, β can not be realized disjointly, $\alpha \cup \beta$ is connected. The complementary components of $\alpha \cup \beta$ in S are oriented hyperbolic surfaces with convex piecewise geodesic boundary. Here convex means that the oriented exterior angle at any corner of the boundary is non-negative.

Since α, β do not jointly fill up S, there is a component A of $S-(\alpha\cup\beta)$ which is not simply connected. Its boundary is a disjoint union of topological circles. If there is a boundary circle γ of A which is contractible in S then γ bounds an embedded disc $D \subset S$. The disc D is naturally oriented and equipped with a hyperbolic metric with piecewise geodesic boundary. Since $S-\gamma$ is disconnected, either $A \subset D$ or $A \cap D = \emptyset$. Now if $A \cap D = \emptyset$ then the orientations of γ as the oriented boundary of D and as a subset of the oriented boundary of A are distinct. Therefore the boundary γ of D is piecewise geodesic, and at each of its corners the exterior angle is *negative*. Since the metric on the disc D is negatively curved, this violates the Gauss Bonnet formula. Thus $A \subset D$ and hence since A is not simply connected by assumption, there is a point $x \in D-A$. However, this means that $\alpha \cup \beta \subset S$ is disconnected which is impossible.

As a consequence, a boundary component of A is a simple closed curve γ which is not contractible in S. Up to isotopy, this curve is disjoint from $\alpha \cup \beta$ and hence we have $d(\gamma,\alpha) = 1 = d(\gamma,\beta)$ and therefore $d(\alpha,\beta) \leq 2$. This shows the lemma. ☐

Finally we investigate the action of the mapping class group $\mathrm{Mod}(S)$ on the curve graph.

Lemma 3.2.7. $\mathrm{Mod}(S)$ *acts on* $\mathcal{CG}(S)$ *as a group of simplicial isometries. There are only* $[g/2]+1$ *distinct orbits of vertices.*

Proof. The mapping class group $\mathrm{Mod}(S)$ naturally acts on the free homotopy classes of simple closed curves respecting intersection numbers. Thus $\mathrm{Mod}(S)$ acts on $\mathcal{CG}(S)$ as a group of simplicial isometries preserving the subset of non-separating curves. Now a simple closed curve $\alpha \in \mathcal{C}(S)$ is non-separating if and only if $S-\alpha$ is an oriented surface of genus $g-1$ with two boundary components. Thus if β is another non-separating simple closed curve then there is a diffeomorphism of $S-\alpha$ onto $S-\beta$ which can be extended to a diffeomorphism of S mapping α to β. Consequently the action of $\mathrm{Mod}(S)$ on non-separating curves is transitive.

On the other hand, if $\alpha \in \mathcal{C}(S)$ is separating, then it decomposes S into a surface S_1 of genus $g' \in \{1,\dots,[g/2]\}$ with one boundary circle and a surface S_2 of genus $g-g'$ with one boundary circle. As before, for any other separating curve β inducing a decomposition of S into surfaces with the same genus there is an orientation preserving diffeomorphism of $S - \alpha$ onto $S - \beta$ which induces a diffeomorphism of S mapping α to β. $\square$

3.3 The Hausdorff topology on compact subsets of S

In Section 3.2 we introduced the curve graph $\mathcal{CG}(S)$, we showed that it is a connected geodesic metric space, and we estimated the distance between any two of its vertices. In particular, by Lemma 3.2.3, curves whose distance in $\mathcal{CG}(S)$ is large have large intersection numbers. However, so far we do not know whether the diameter of $\mathcal{CG}(S)$ is infinite.

Our next goal is to show that this is indeed the case. The basic idea is to choose a hyperbolic metric σ on S and look at sequences of simple closed geodesics on (S, σ) whose lengths with respect to the hyperbolic metric are unbounded. We then can try to take a limit of such a sequence and obtain information on distances in $\mathcal{CG}(S)$ by analyzing this limiting object.

It turns out that one can make sense of this simple idea using the following definition.

Definition 3.3.1. The *Hausdorff distance* $d_H(A, B)$ of two closed subsets A, B of a proper metric space X is defined by

$$d_H(A, B) = \inf\{\epsilon > 0 \mid A \subset U_\epsilon(B), B \subset U_\epsilon(A)\}$$

where $U_\epsilon(A)$ is the ϵ-neighborhood of A.

Note that the Hausdorff distance of two unbounded closed sets in X may be infinite. However, the distance between two compact subsets of X is always finite, and the Hausdorff distance is indeed a distance for the family of all compact subsets of X. Namely, symmetry and the triangle inequality are obvious. Moreover, if $d_H(A, B) = 0$ then A is contained in $\cap_{\epsilon>0} U_\epsilon(B)$. Since B is closed, we have $B = \cap_{\epsilon>0} U_\epsilon(B)$ and hence $A \subset B$. By symmetry, we conclude that $B \subset A$ and hence $A = B$.

If X is compact then the topology on the space $\mathcal{K}(X)$ of compact subsets of X defined by the Hausdorff distance does not depend on the metric defining the topology and is called the *Hausdorff topology.* We have.

Lemma 3.3.2. *Let X be a proper metric space. If a sequence $\{A_i\} \subset \mathcal{K}(X)$ converges in the Hausdorff topology to a compact set $A \in \mathcal{K}(X)$ then the following conditions are satisfied.*

1. *If $x \in X$ is such that there exists a subsequence $\{A_{i_j}\}$ of the sequence $\{A_i\}$ and there are points $x_j \in A_{i_j}$ with $x_j \to x$ then $x \in A$.*

2. *For every $x \in A$ there exist $x_i \in A_i$ for all i such that $x_i \to x$.*

If X is compact then these two conditions are also sufficent for convergence of $\{A_i\} \subset \mathcal{K}(X)$ to $A \in \mathcal{K}(X)$.

Proof. Let X be a proper metric space and assume that the sequence $\{A_i\} \subset \mathcal{K}(X)$ of compact subset of X converges with respect to the Hausdorff distance to a compact set A. We have to show that the two properties stated in the lemma are satisfied. Namely, by definition, for every $\epsilon > 0$ there is a number $i(\epsilon) > 0$ such that $d_H(A_i, A) < \epsilon$ for all $i \geq i(\epsilon)$ and hence $A_i \subset U_\epsilon(A)$ for every $i \geq i(\epsilon)$. Let $x_{i_j} \in A_{i_j}$ for a subsequence $\{A_{i_j}\}$ of $\{A_i\}$ and assume that $x_{i_j} \to x$ $(j \to \infty)$. Since $A = \cap_{\epsilon>0} U_\epsilon(A)$ the point x is necessarily contained in A.

Similarly, let $x \in A$ and assume without loss of generality that $i(\frac{1}{m}) < i(\frac{1}{m+1})$ for all $m > 0$. Since $A \subset U_\epsilon(A_i)$ for every $i \geq i(\epsilon)$ and all $\epsilon > 0$ there is for every $i \in [i(\frac{1}{m}), i(\frac{1}{m+1}))$ some $x_i \in A_i$ with $d(x, x_i) < 1/m$. Then the sequence $\{x_i\}$ converges to x by construction. In other words, the requirements (1) and (2) above are satisfied.

Now assume that X is compact and that $\{A_i\}$ is a sequence of compact subsets of X which satisfies the first property stated in the lemma for some compact subset A of X. We claim that for every $\epsilon > 0$ there is some $i(\epsilon) > 0$ with $A_i \subset U_\epsilon(A)$ for all $i \geq i(\epsilon)$. Namely, otherwise there is some $\epsilon > 0$ and there is a subsequence i_j and for every j there is some $x_{i_j} \in A_{i_j}$ with $d(x_{i_j}, A) \geq \epsilon$. Up to passing to a subsequence we may assume that the sequence $\{x_{i_j}\}$ converges to some point $x \in X - A$. However, this violates the first property above. The second property implies that $A \subset U_\epsilon(A_i)$ for all sufficiently large i. $\square$

The two conditions stated above make sense for any sequence of closed subsets of a proper metric space X (or even an arbitrary second countable locally compact topological space) and can be used to define a topology on the space of closed subsets of X. In Section 4.3 we will use a similar construction which allows us to take limits of hyperbolic 3-manifolds.

In the case that the metric space X is compact we can say more (compare also Section 4.3).

Lemma 3.3.3. *The space $\mathcal{K}(X)$ of compact subsets of a compact metric space X equipped with the Hausdorff distance is compact.*

Proof. Since the topology on $\mathcal{K}(X)$ is metrizable it suffices to show that $\mathcal{K}(X)$ is sequentially compact.

Thus let $\{A_i\}$ be any sequence of compact subsets of X. Let $\mathcal{U} = \{U_k\}$ be a countable basis for the topology of X. Such a basis exists since X is compact and metrizable. Let $k > 0$; if $U_k \cap A_i \neq \emptyset$ for infinitely many i then we can find a subsequence $\{A_{i_j}\}$ of the sequence $\{A_i\}$ such that $A_{i_j} \cap U_k \neq \emptyset$ for all but finitely many j. Using diagonalization we can construct a subsequence of the sequence $\{A_i\}$ which satisfies (1) and (2) in Lemma 3.3.2 and hence converges. $\square$

As an immediate consequence of Lemma 3.3.3 we obtain.

Corollary 3.3.4. *Let X, Y be compact metric spaces and let $\varphi : X \to Y$ be continuous. Then φ induces a continuous map from the space $\mathcal{K}(X)$ of compact subsets of X into the space $\mathcal{K}(Y)$ of compact subsets of Y.*

Proof. Since X is compact, a continuous map $\varphi : X \to Y$ is closed and hence it induces a map $\varphi_* : \mathcal{K}(X) \to \mathcal{K}(Y)$. By continuity of φ, the image under φ_* of a sequence of compact subsets of X which satisfies the properties 1) and 2) in Lemma 3.3.2 satisfies these properties as well. From this continuity of φ_* is immediate. $\square$

Example 3.3.5. Filling sequences
Let X be a compact metric space and let $Q = \{x_1, x_2, \dots\} \subset X$ be any countable dense set. Then for each i the set $A_i = \{x_1, \dots, x_i\} \subset X$ is compact. The sets A_i converge as $i \to \infty$ in the Hausdorff topology to X. If X is proper and non-compact and if $Q = \{x_1, x_2, \dots\} \subset X$ is a dense subset, then the sets A_i defined above satisfy the two conditions in Lemma 3.3.2 with $A = X$, however they do not converge with respect to the topology defined by the Hausdorff distance.

3.4 Geodesic laminations

3.4.1 The space of geodesic laminations

The goal of this section is to describe all Hausdorff limits of sequences of simple closed geodesics on the closed surface S equipped with a fixed hyperbolic metric. For this call an infinite geodesic on S *simple* if it does not have self-intersections. We need the following definition.

Definition 3.4.1. A *geodesic lamination* on S is a *closed* subset of S foliated into simple geodesics. A *sublamination* of a geodesic lamination λ is a geodesic lamination which is a closed subset of λ.

Example 3.4.2. Geodesic multicurves
A simple closed geodesic or, more generally, a simple *geodesic multicurve*, i.e. a disjoint union of simple closed geodesics, is a geodesic lamination.

A simple geodesic which is contained in a geodesic lamination λ is called a *leaf* of λ.

Geodesic laminations were introduced by Thurston and proved to be powerful tools for the investigation of Teichmüller spaces and Kleinian groups. The book of Casson and Bleiler [**?**] and the notes of Canary, Epstein and Green [**?**] contain well written accounts on geodesic laminations and some of their early applications. In this section we only present those properties of geodesic laminations which are needed in the later sections.

Let λ be a geodesic lamination on S. Then every $x \in \lambda$ is contained in precisely one leaf of λ. Thus we can associate to a point $x \in \lambda$ the tangent line at x of this leaf, viewed as a point in the *projectivized tangent space PTS of S*. This is the smooth fibre bundle over S whose fibre at x consists of the one-dimensional subspaces of the tangent space of S at x. The canonical projection $\pi : PTS \to S$ is a smooth surjection. By compactness of PTS, this projection is also closed.

Lemma 3.4.3. *Let λ be a geodesic lamination on S. Then the map $\lambda \to PTS$ which associates to $x \in \lambda$ the tangent line at x of the leaf of λ through x is continuous.*

Proof. (Compare Lemma 3.1 of [**?**].) Let λ be a geodesic lamination and let $\varphi : \lambda \to PTS$ be the map which associates to $x \in \lambda$ the tangent line at x of the leaf of λ through x. If φ is not continuous then there is a sequence $\{x_i\} \subset \lambda$ converging to $x \in \lambda$ such that the tangent lines $\varphi(x_i)$ of λ at x_i do not converge to the tangent line $\varphi(x)$ at x. Since PTS is compact and the projection π is continuous, after passing to a subsequence we may assume that the sequence $\{\varphi(x_i)\}$ converges to a tangent line $v \neq \varphi(x)$ at the basepoint x. Then the geodesic in S which is tangent to v intersects the leaf of λ through x transversely at x. By smooth dependence of a geodesic on its initial velocity, this implies that for sufficiently large i the leaf of λ through x_i intersects the leaf of λ through x transversely which is impossible. $\square$

Lemma 3.4.3 can be used for another way to look at a geodesic lamination λ. Namely, the *geodesic foliation* of the projectivized tangent bundle of S is the foliation whose leaves are the tangent lines of the geodesics in S. Call a subset A of PTS *saturated* for the geodesic foliation if for every point $x \in A$ the entire leaf of the geodesic foliation through x is contained in A.

Lemma 3.4.4. *There is a one-to-one correspondence between geodesic laminations on S and closed subset A of PTS which are saturated for the geodesic foliation and such that the restriction of the canonical projection π to A is injective.*

Proof. Let λ be a geodesic lamination and for $x \in \lambda$ let $\varphi(x) \in PTS$ be the tangent line of the leaf of λ passing through x. By Lemma 3.4.3, the map φ is continuous, moreover it is injective. Thus $\varphi(\lambda)$ is a compact subset of PTS which clearly has the properties stated in the lemma.

Vice versa, by definition, a closed subset A of PTS which is saturated for the geodesic foliation and such that the restriction of π to A is injective consists of tangent lines of simple geodesics. Moreover, the geodesic foliation of PTS restricts to a foliation of A which projects to a foliation of $\pi(A)$. Thus $\pi(A)$ is a closed subset of S which is foliated into simple geodesics and hence it is a geodesic lamination. $\square$

Next we have a closer look at the space of all geodesic laminations.

Lemma 3.4.5. *1. The space $\mathcal{L}(S)$ of geodesic laminations on S equipped with the Hausdorff topology is compact.*

2. The closure in S of any non-empty disjoint union of simple geodesic is a geodesic lamination.

Proof. (See Lemma 3.2 and Theorem 3.4 of [**?**].) Since by Lemma 3.3.3 the Hausdorff topology on compact subsets of S is compact and metrizable, for the first part of the lemma it is enough to show the following. Let $\{\lambda_i\} \subset \mathcal{L}(S)$ be a sequence of geodesic laminations which converges in the Hausdorff topology to a compact subset λ of S. Then λ is a geodesic lamination.

For this let $\Lambda_i \subset PTS$ be the space of tangent lines of the leaves of λ_i. By Lemma 3.4.3, Λ_i is a compact subset of the compact metrizable space PTS. After passing to a subsequence we may assume that Λ_i converges as $i \to \infty$ to a compact subset Λ of PTS. Since Λ_i is saturated for the geodesic foliation, by properties 1) and 2) in Lemma 3.3.2 the same is true for Λ. Namely, by the second property, if $v \in \Lambda$ then there is a sequence of points $v_i \in \Lambda_i$ with $v_i \to v$. Since tangent lines of geodesics depend smoothly on their initial velocities, the first property shows that any compact subarc of the leaf of the geodesic foliation through v is a Hausdorff limit of compact subarcs of the leaves through the points v_i.

By Corollary 3.3.4, the image of Λ under the canonical projection $\pi : PTS \to S$ coincides with the limit λ of the sequence $\{\lambda_i\}$. Thus by Lemma 3.4.4, to show that λ is a geodesic lamination we are left with showing that the restriction of π to Λ is injective. However, if this is not the case then there are two distinct points $v, w \in \Lambda$ which are mapped to the same point in S. Then the geodesics in S which are tangent to v, w intersect transversely at $\pi(v) = \pi(w)$. Since transverse intersection is an open condition, geodesics defined by nearby tangent lines also intersect transversely. Since by the second property in Lemma 3.3.2 both points v, w are limits of points $v_i, w_i \in \Lambda_i$, this violates the assumption that the restriction of π to Λ_i is injective for each i. This shows that $\pi(\Lambda)$ is a geodesic lamination as claimed and completes the proof of the first part of the lemma. The second part follows from the same argument. $\square$

There is yet another description of geodesic laminations. For its formulation, recall that the hyperbolic plane $\mathbf{H}^2$ admits a natural compactification by attaching the circle S^1 which is just the Gromov boundary of $\mathbf{H}^2$. Let $\Delta \subset S^1 \times S^1$ be the diagonal and let $\iota : S^1 \times S^1 \to S^1 \times S^1$ be the involution exchanging the two factors. An unoriented geodesic in the hyperbolic plane $\mathbf{H}^2$ is just a point in the Möbius band $(S^1 \times S^1 - \Delta)/\iota$ of unordered pairs of distinct points in S^1.

We say that two unordered pairs of distinct points $x \neq y, x' \neq y' \in S^1$ are *not separated* under the action of $\pi_1(S)$ as a group of isometries of $\mathbf{H}^2$ if for every $g \in \pi_1(S)$ the points gx', gy' are contained in the closure of one of the two components of $S^1 - \{x, y\}$. This is equivalent to stating that the projection to S of the geodesic connecting x to y does not intersect the projection of the geodesic connecting x' to y' transversely. If $x = x', y = y'$ then this in turn is equivalent to stating that the projection to S of the geodesic connecting x to y is simple. We have.

Lemma 3.4.6. *A geodesic lamination is a closed $\pi_1(S)$-invariant subset of the Möbius band $(S^1 \times S^1 - \Delta)/\iota$ consisting of unordered pairs of points which are not separated under the action of $\pi_1(S)$.*

Proof. By Lemma 3.4.4, the space of tangent lines of a geodesic lamination lifts to a closed $\pi_1(S)$-invariant subset A of the projectivized tangent bundle $PT\mathbf{H}^2$ of the hyperbolic plane $\mathbf{H}^2$ which is saturated for the geodesic foliation. The projection $\Pi : PT\mathbf{H}^2 \to (S^1 \times S^1 - \Delta)/\iota$ which associates to a point $v \in PT\mathbf{H}^2$ the unordered pair of endpoints in $S^1 = \partial\mathbf{H}^2$ of the geodesic which is tangent to v is continuous, closed and equivariant with respect to the action of $\pi_1(S)$. Thus the projection $\Pi(A) \subset (S^1 \times S^1 - \Delta)/\iota$ of the set A is closed and $\pi_1(S)$-invariant. Moreover, points in $\Pi(A)$ are geodesics in $\mathbf{H}^2$ whose projections to S are simple and pairwise disjoint. We observed above that this is equivalent to stating that any two such pairs of points are not separated under the action of $\pi_1(S)$.

Similarly, the preimage under the projection Π of a closed $\pi_1(S)$-invariant set of unordered pairs of distinct points in S^1 is a closed $\pi_1(S)$-invariant subset of $PT\mathbf{H}^2$ which is saturated for the geodesic foliation and consists of lifts of simple geodesics in S. $\square$

3.4.2 Examples

To construct interesting geodesic laminations on any closed hyperbolic surface S we first have a closer look at the structure of an individual geodesic lamination. We begin with the following observation (see Lemma 3.3 of [**?**]).

Lemma 3.4.7. *A geodesic lamination on S is nowhere dense in S.*

Proof. Let λ be a geodesic lamination on S. Then every point $x \in \lambda$ is contained in a unique leaf of λ. By Lemma 3.4.3, the tangent line at x of this leaf depends continuously on x. Since the Euler characteristic of S is negative, there is no continuous line field on S. Hence any geodesic lamination is a proper subset of S.

Similarly, assume that there is a geodesic lamination λ on S which contains an open subset U of S. By continuity, every leaf of λ passing through U is entirely contained in the set V of all points in λ which admit a neighborhood in S contained in λ. Thus V contains with a point $x \in \lambda$ the entire leaf of λ through x. As a consequence, the closure A of the subset V of λ is a geodesic lamination, and A is also the closure in S of an open subset of S. Then A is a hyperbolic surface of finite volume whose boundary consists of simple geodesics, moreover A is foliated by simple geodesics. However, the Euler characteristic of such a surface is negative and hence it does not admit a continuous line field. This is a contradiction to the assumption that λ contains an open subset of S and shows that λ is indeed nowhere dense in S. ☐

It follows from Lemma 3.4.7 that the complement of a geodesic lamination λ in the hyperbolic surface S is a nontrivial hyperbolic surface of finite area whose completion S' has geodesic boundary. By the Gauß Bonnet theorem, the area of each connected component of the surface S' is a multiple of π. (In fact, the Lebesgue measure of a geodesic lamination on S vanishes [**?**] and hence the area of S' coincides with the area of S.) Since the area of S equals $4\pi(g-1)$, the number of connected components of $S - \lambda$ is bounded from above by $4g-4$, and the area of each connected component does not exceed $4\pi(g-1)$.

If a connected component S_0 of the surface S' is simply connected, then a developing map of S_0 into the hyperbolic plane is injective, and the area bound implies that S_0 is an ideal polygon with at most $4g-2$ vertices, i.e. the convex hull of at most $4g-2$ points in the boundary S^1 of $\mathbf{H}^2$. We refer to [**?**] for a more detailed discussion of the structure of geodesic laminations.

Definition 3.4.8. A geodesic lamination is called *maximal* if its complementary regions are all ideal triangles. The geodesic lamination λ *fills up* S if $S - \lambda$ is a union of ideal polygons. The geodesic lamination λ is called *minimal* if every half-leaf of λ is dense in λ. A *minimal component* of a geodesic lamination λ

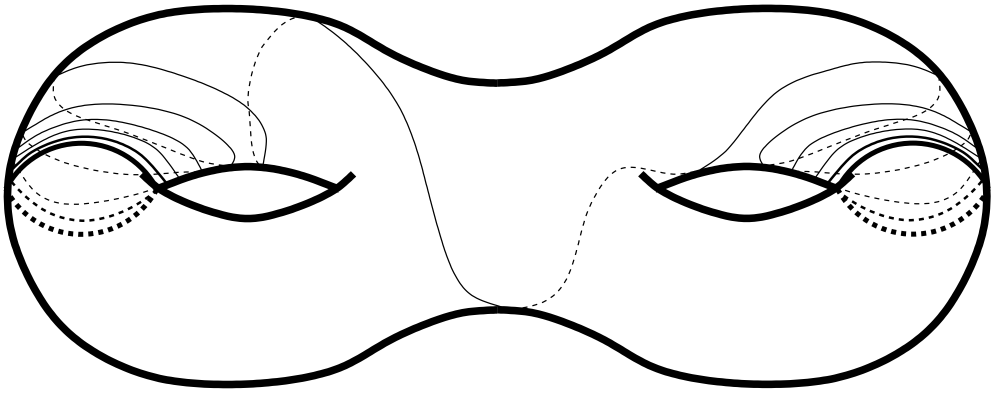

Figure 3.3

is a minimal sublamination of λ. A half-leaf ℓ of a geodesic lamination *spirals* about a minimal component μ if $\ell \not\subset \mu$ and if the closure of ℓ equals $\ell \cup \mu$.

Example 3.4.9. Simple geodesics
Every simple closed geodesic is a minimal geodesic lamination. If μ is a sublamination of a geodesic lamination λ then a minimal component of μ is a minimal component of λ. In particular, every simple closed geodesic contained in λ is a minimal component of λ.

Example 3.4.10. Maximal geodesic laminations whose minimal components are simple closed curves
Let $c = c_1 \cup \cdots \cup c_k$ be a simple geodesic multicurve. We claim that there is a maximal geodesic lamination μ such that c consists precisely of the minimal components of μ. Moreover, every half-leaf ℓ of μ either is contained in one of the components of c or it *spirals* about one of the components c_i, i.e. the closure of ℓ equals $\ell \cup c_i$.

Namely, let $S_1, \ldots, S_\ell$ be the connected components of $S - c$. The completion S_i' of each of these components S_i is a compact hyperbolic surface with geodesic boundary. For each i, choose a triangulation T_i of S_i' by geodesic arcs with endpoints on the boundary of S_i'. Then each boundary component of S_i' contains at least one vertex of this triangulation. The triangulations of S_i' obtained from T_i by applying a sequence of Dehn multi-twists about each boundary curve in the direction determined by the orientation of the component converge in the Hausdorff topology to a "triangulation" of S_i into ideal triangles. The sides of each such triangle are infinite simple geodesics which spiral about the components of c contained in the boundary of S_i. Thus the union of the sides of all these triangles with the multicurve c is a geodesic lamination on S with the required properties.

So far we have not seen any example of a minimal geodesic lamination which is distinct from a simple closed geodesic. To construct a minimal geodesic lamination which fills up S we use singular euclidean metrics on S.

Example 3.4.11. A minimal geodesic lamination which fills up S
In this example we use singular euclidean metrics on S to construct explicitly a minimal geodesic lamination which fills up S.

Consider for a moment the standard flat torus $T = \mathbb{C}/\mathbb{Z}^2$, i.e. the quotient of the complex plane under the action of the standard square lattice $\mathbb{Z}^2$. Any foliation of $\mathbb{C}$ by parallel lines projects to a geodesic foliation of T. Let $\mathcal{F}$ be the projection to T of a foliation of $\mathbb{C}$ by lines with irrational slope. Then every leaf of $\mathcal{F}$ is dense in T. Moreover, the foliation $\mathcal{F}$ is geodesic with respect to the standard euclidean metric on T.

Let c be a small compact geodesic arc in T which is horizontal with respect to the standard euclidean coordinates and whose endpoints lie on two distinct leaves of $\mathcal{F}$. Cut T open along the interior of c. We obtain a torus T_1 with connected piecewise geodesic boundary ∂T_1 which consists of two copies c_1, c_2 of c joined at the endpoints. Let T_2 be a second copy of T_1 whose boundary consists of the copies c_1', c_2' of c and glue T_1 to T_2 crosswise isometrically along the boundaries in such a way that a point on the arc c_1 is identified with the point on the arc c_2' to which it corresponds. The result of this glueing is an oriented surface S of genus 2. The surface S admits a natural branched covering onto the torus T with two branch points whose images in T are the endpoints of c. The euclidean metric on the disjoint union of the tori T_1, T_2 defines a singular euclidean metric q on S with two singularities of cone angle 4π at the two branch points. Locally near a singular point, the metric is obtained by slitting the euclidean plane open along the positive real half-line and gluing a second copy of the slit plane along the boundary.

The foliation $\mathcal{F}$ on T lifts to a singular geodesic foliation $\mathcal{F}_0$ for the metric q on S. Near any regular point of q, this foliation is a foliation in the usual sense. At a singular point x, the foliation $\mathcal{F}_0$ is singular: There are four half-leaves of $\mathcal{F}_0$ issuing from x. A leaf ℓ of the foliation $\mathcal{F}$ on T which does not pass through an endpoint of c lifts to a leaf ℓ_0 of the foliation $\mathcal{F}_0$ on S. This leaf ℓ_0 is a biinfinite geodesic for the singular euclidean metric q without self-intersections. Each time the leaf ℓ crosses through the arc c, the leaf ℓ_0 passes from the copy $T_i \subset S$ of the slit torus to the copy T_{3-i}. Since ℓ is dense in T, the leaf ℓ_0 is dense in S.

The geodesic ℓ_0 lifts to a geodesic $\tilde{\ell}$ in the universal covering $\tilde{S}$ of S equipped with the lift $\tilde{q}$ of the metric q. Since $\tilde{q}$ is the lift of the geodesic metric q on the closed surface S, the universal covering $(\tilde{S}, \tilde{q})$ is $\pi_1(S)$-equivariantly quasi-isometric to the group $\pi_1(S)$. On the other hand, every choice of a marked hyperbolic metric σ on S determines a $\pi_1(S)$-equivariant quasi-isometry between

$\pi_1(S)$ and the hyperbolic plane $\mathbf{H}^2$ (see Example 3.1.9 in Section 3.1). Hence $\tilde{\ell}$ is a biinfinite quasi-geodesic in $\mathbf{H}^2$. Such a quasi-geodesic is contained in a uniformly bounded neighborhood of a hyperbolic geodesic $\tilde{\gamma}$ and hence it converges in forward and backward direction to the two distinct endpoints $x \neq y \in S^1 = \partial\mathbf{H}^2$. Since the projection ℓ_0 of ℓ to S does not have self-intersections, the points x, y are not separated under the action of $\pi_1(S)$. In particular, the hyperbolic geodesic $\tilde{\gamma}$ connecting the same endpoints in S^1 projects to a simple geodesic in the hyperbolic surface (S, σ). By Lemma 3.4.5, its closure in S is a geodesic lamination λ on (S, σ).

We claim that the complementary components of this geodesic lamination are simply connected. For this assume otherwise. Since the completion of $S - \lambda$ is a hyperbolic surface with geodesic boundary, in this case there is a simple closed geodesic α in (S, σ) which does not intersect λ. A lift of α to $\mathbf{H}^2$ is a geodesic with endpoints $a \neq b \in S^1$. Since α does not intersect λ, the endpoints in S^1 of *any* lift $\tilde{\ell}$ to $\mathbf{H}^2$ of the q-geodesic ℓ_0 are contained in the closure of the same component of $S^1 - \{a, b\}$.

Now let c be a simple closed geodesic for the singular euclidean metric q which is freely homotopic to α. Any curve of minimal length for q in the free homotopy class of α is such a geodesic. Such a geodesic has only finitely many intersections with the singular set of q, and its restriction to the regular set is an euclidean geodesic. Since ℓ_0 does not pass through a singular point of q, any two intersection points between c and ℓ_0 are transverse and essential. This means that a subarc ℓ' of ℓ_0 with endpoints in c is not homotopic with fixed endpoints to an arc which either is contained in c or has fewer transverse intersection points (compare the discussion in Section 2 of [**?**]). However, since ℓ_0 is dense in S, the geodesic c intersects ℓ_0 transversely and hence there are lifts of ℓ_0 to $(\tilde{S}, \tilde{q})$ (identified by a quasi-isometry with $\mathbf{H}^2$) whose endpoints are contained in different components of $S^1 - \{a, b\}$. This is a contradiction which implies that the lamination λ fills up S.

Using the procedure in the above example inductively, we can construct a minimal geodesic lamination which fills up for a closed oriented surface of every genus $g \geq 2$. Namely, to obtain a lamination on a surface S_0 of genus 3, cut the torus T along a second horizontal arc c' which is disjoint from c and such that the four endpoints of c, c' are contained in distinct leaves of $\mathcal{F}$. Then we can glue a copy of T slit open along c' to the surface S equipped with the piecewise euclidean metric defined above and obtain a geodesic lamination on S_0 which is minimal and fills up S_0 by the argument presented above. We leave the details of this construction to the reader (see also [**?**, **?**, **?**]). To summarize this discussion, we have.

Lemma 3.4.12. *Every closed oriented hyperbolic surface S admits a minimal geodesic lamination which fills up S.*

Example 3.4.13. Properties of minimal geodesic laminations which fill up S
If λ is a geodesic lamination which fills up S then λ is a sublamination of a maximal geodesic lamination. Namely, if there is a complementary component of λ which is an ideal polygon P with $p \geq 4$ sides then we can connect two non-adjacent ideal vertices of P by a simple geodesic ℓ which cuts the polygon into two subpolygons with fewer sides each. Note that each of the two subrays of ℓ is asymptotic to the two sides of P which converge to the same ideal vertex. Inductively we can in this way subdivide all complementary polygons into triangles. Moreover, every geodesic lamination μ which contains λ as a sublamination can be obtained in this way. In particular, if ℓ is any half-leaf of μ, then the closure of ℓ in S contains the closure of a half-leaf of λ which is asymptotic to ℓ. As a consequence, if λ is minimal and fills up S and if λ is a sublamination of a geodesic lamination μ then the closure of every half-leaf of μ contains λ. In other words, every sublamination of μ contains λ as a sublamination.

3.4.3 The curve graph is of infinite diameter

The following lemma relates simple closed curves on S to geodesic laminations. It appeared first in [**?**].

Lemma 3.4.14. *A minimal geodesic lamination can be approximated in the Hausdorff topology by simple closed geodesics.*

Proof. The *unit tangent bundle* T^1S of the hyperbolic surface S admits a natural Riemannian metric, the so-called *Sasaki metric*, such that the projection $T^1S \to S$ is a Riemannian submersion. By hyperbolic trigonometry, for every $\epsilon > 0$ there is a number $\delta = \delta(\epsilon) > 0$ with the following property. Let $T > 1$ and let $\gamma : [0, T] \to S$ be a geodesic parametrized by arc length such that the distance in T^1S between $\gamma'(0)$ and $\gamma'(T)$ is at most δ. Connect $\gamma(T)$ to $\gamma(0)$ by a geodesic arc of length at most δ. Then the resulting closed curve $\tilde{\gamma}$ is freely homotopic to a closed geodesic on S whose Hausdorff distance to γ (as subsets of S) is at most ϵ.

Let λ be a minimal geodesic lamination which is not a simple closed geodesic and let ℓ be a leaf of λ. Assume that ℓ is parametrized by arc length. Choose a transverse geodesic arc c through $\ell(0)$ with endpoints in the complementary regions of λ. For $\epsilon > 0$ let $\delta > 0$ be as above. Since every half-leaf of λ is dense in λ, the point $\ell(0)$ is not isolated in $c \cap \lambda$. Moreover, there is a number $T > 0$ such that the Hausdorff distance between λ and $\ell[0, T]$ (or $\ell[-T, 0]$) is at most $\epsilon/2$. By continuous dependence of the tangent line of a leaf of λ on its basepoint established in Lemma 3.3.4 and by smooth dependence of a geodesic on its initial tangent line, there is a neighborhood c_0 of $\ell(0)$ in c such that the

Hausdorff distance between λ and every subarc of length T of every half-leaf of λ with one endpoint in c_0 is at most ϵ.

Using once more Lemma 3.3.4, there is a subarc I of c_0 containing $\ell(0)$ in its boundary with the following properties.

1. $\ell(0)$ is an accumulation point of $\lambda \cap I$.
2. The first return time to I of any half-leaf of λ beginning at a point in I is not smaller than T.
3. If $x, y \in \lambda \cap I$ and if v, w are unit tangents of the leaves of λ passing through x, y which point to the same side of c then the distance in $T^1(S)$ between v and w is at most δ.

Let $k > 0$ be the smallest number such that $\ell(k) \in I$. By the choice of I, we have $k > T$. If the orientations of the intersections of ℓ and c at $\ell(0)$ and $\ell(k)$ coincide then the distance in $T^1 S$ between $\ell'(0)$ and $\ell'(k)$ is at most δ. Connect the point $\ell(k)$ to $\ell(0)$ by an embedded subarc of I of length at most δ which is disjoint from $\ell[0, k]$. By the choice of δ and T, the Hausdorff distance between the geodesic representative of the resulting simple closed curve and the geodesic lamination λ is not bigger than 2ϵ.

Otherwise choose a number $m > k$ such that $\ell(m)$ is the first intersection point of $\ell(k, \infty)$ with the subarc of I bounded by $\ell(0)$ and $\ell(k)$. Such an intersection point exists since $\ell(0)$ is an accumulation point of $I \cap \lambda$. If the orientations of the intersection of ℓ and c at $\ell(0), \ell(m)$ coincide then we can as before connect $\ell(m)$ with $\ell(0)$ by an embedded subarc of I to obtain a simple closed curve whose geodesic representative is of Hausdorff distance at most 2ϵ to λ. Otherwise the orientations of the intersection of ℓ with c coincide at the points $\ell(k), \ell(m)$ and once again, we find a simple closed geodesic with the required properties. $\square$

We use Lemma 3.4.14 to show.

Proposition 3.4.15. *The diameter of $\mathcal{CG}(S)$ is infinite.*

Proof. We follow Luo (see [**?**]) and argue by contradiction: Assume that $\mathrm{diam}(\mathcal{CG}(S)) = D < \infty$.

Let μ be a minimal geodesic lamination which fills S. Such a geodesic lamination was constructed in Example 3.4.11 in this section. Let $\alpha_0 \in \mathcal{C}(S)$ be any fixed simple closed curve. By Lemma 3.4.14, there is a sequence $(\alpha_i)_{i \geq 0} \subset \mathcal{C}(S)$ which converges to μ in the Hausdorff topology as $i \to \infty$. Since $d(\alpha_0, \alpha_i) \leq$

D, after passing to a subsequence we may assume that $d(\alpha_0, \alpha_i) = N \leq D$ for all i.

Since $\mathcal{CG}(S)$ is a geodesic metric graph, we can find simple closed curves $\beta_i \in \mathcal{C}(S)$ $(i > 0)$ such that $d(\alpha_0, \beta_i) = N - 1, d(\beta_i, \alpha_i) = 1$ for all i. Since the space of geodesic laminations equipped with the Hausdorff topology is compact, after passing to a subsequence we may assume that $\beta_i \to \mu'$ in the Hausdorff topology where μ' is a geodesic lamination. We claim that necessarily μ is a sublamination of μ'. To see this note that $d(\beta_i, \alpha_i) = 1$ means that $\beta_i \cup \alpha_i$ is a geodesic lamination. Thus after passing to another subsequence we may assume that the sequence $(\beta_i \cup \alpha_i)$ converges in the Hausdorff topology to a geodesic lamination μ''. However, $\beta_i \to \mu', \alpha_i \to \mu$ and therefore $\mu'' \supset \mu \cup \mu'$. But μ is minimal and fills S by assumption and hence by Example 3.4.13, every geodesic lamination ξ with the property that $\mu \cup \xi$ is a geodesic lamination contains μ as a sublamination. In other words, we have $\mu' \supset \mu$.

Repeat the above argument with the sequence $(\beta_i) \subset \mathcal{C}(S)$. After $N - 1$ steps we conclude that $\alpha_0 \cup \mu$ is a lamination and hence α_0 contains μ as a sublamination which is impossible. $\square$

We conclude this section with a closer look at the action of the mapping class group on $\mathcal{CG}(S)$. Note that this action is by simplicial isometries.

Definition 3.4.16. A mapping class $\varphi \in \mathrm{Mod}(S)$ is called *pseudo-Anosov* if the action on $\mathcal{CG}(S)$ of the subgroup $< \varphi >$ of $\mathrm{Mod}(S)$ generated by φ has unbounded orbits. A mapping class φ is *periodic* if $| < \varphi > | < \infty$. A mapping class which is neither periodic nor pseudo-Anosov is called *reducible*.

Pseudo-Anosov mapping classes can be constructed explicitly. For example, if $\alpha, \beta \in \mathcal{C}(S)$ are simple closed curves with $d(\alpha, \beta) \geq 3$ then the composition of a positive Dehn twist about α with a negative Dehn twist about β is a pseudo-Anosov mapping class. The paper [**?**] contains a detailed account of this construction of Thurston. We do not need pseudo-Anosov mapping classes in the sequel, and we refer to [**?**, **?**] for details on Thurston's classification of mapping classes.

Example 3.4.17. The lengths of simple closed curves on covers of mapping tori
If $\varphi \in \mathrm{Mod}(S)$ is periodic then the mapping torus M_φ of φ has a *finite* covering which is diffeomorphic to $S \times S^1$. If $k > 1$ is the smallest number such that $\varphi^k = 1$ then this covering can be chosen to have k sheets. The fundamental group of the covering is just the normal subgroup of the fundamental group of M_φ which is the preimage of the subgroup $k\mathbb{Z}$ of $\mathbb{Z}$ under the natural projection $\pi_1(M_\varphi) \to \mathbb{Z} = \pi_1(M_\varphi)/\pi_1(S)$ discussed in Section 3.1.

More precisely, by the solution of the *Nielsen realization problem* [**?**], the mapping class φ can be represented by a biholomorphic automorphism φ_0 of a complex structure on S. Then φ_0^k is the identity map and $S \to S/ < \varphi_0 >$ is a branched holomorphic covering. In particular, the manifold M_φ admits a smooth foliation by circles as follows. In the realization of M_φ as a quotient $S \times [0,1]/ \sim$ where $(x,1) \sim (\varphi_0(x), 0)$, this foliation is the projection of the second factor foliation of $S \times [0,1]$. Each circle of the foliation corresponds to an orbit for the action of φ_0 on S. The singular leaves of the foliation correspond to the fixed points of φ_0. A closed 3-manifold which admits a foliation by circles is called a *Seifert fibered space*.

If we denote as before by $\hat{M}$ the $\mathbb{Z}$-cover of M_φ with fundamental group $\pi_1(S)$, then we conclude as in Example 3.1.10 that for *every* smooth metric on $\hat{M}$ and every $L > 0$ the number of elements in $\pi_1(\hat{M}) = \pi_1(S)$ which can be realized by a curve of length at most L in $\hat{M}$ is *finite*.

If $\varphi \in \mathrm{Mod}(S)$ is pseudo-Anosov then for every smooth metric on the mapping torus M_φ of φ there is some $L > 0$ such that the number of elements of $\pi_1(\hat{M}) = \pi_1(S)$ which can be realized by a curve of length at most L is *infinite*. Such a collection of curves can be obtained as the orbit of a fixed curve under the group $< \varphi >< \mathrm{Mod}(S)$ which is infinite by definition of a pseudo-Anosov mapping class.

3.5 Annuli and singular euclidean metrics

In the previous sections we used combinatorial tools to study the curve graph and its geometry. In this section we begin to investigate the relation between the geometry of the curve graph and the geometry of Teichmüller space $\mathcal{T}(S)$ of S.

3.5.1 Moduli of annuli

On the one hand, the Teichmüller space $\mathcal{T}(S)$ is the space of all complex structures on S up to marking preserving biholomorphisms. On the other hand, it can also be identified with the space of all isotopy classes of hyperbolic metrics on S. This correspondence is obtained by associating to a marked complex structure on S the unique marked hyperbolic metric whose *conformal class* defines the complex structure. Simple closed curves on S can be used to relate these structures.

Definition 3.5.1. 1. The *extremal length* of a class $[\gamma]$ of curves in a Riemann surface S is defined to be

$$\ell_\sigma([\gamma]) = \sup_{\sigma'} \frac{\inf_{\gamma' \in [\gamma]} (\ell(\gamma'))^2}{\mathrm{Area}(\sigma')}$$

where the supremum is taken over all metrics σ' whose *conformal class* defines the complex structure and where $\ell(\gamma')$ is the length of γ' with respect to σ'.

2. The *modulus* of a closed annulus A equipped with a complex structure is the extremal length of the class of arcs joining the two boundary components.

Thus the modulus of an annulus is invariant under biholomorphisms. The modulus of the standard round annulus $\{z \mid r \leq |z| \leq 1\} \subset \mathbb{C}$ equals $-\frac{1}{2\pi} \log r$.

Define the *width* w of a closed annulus A in a closed surface equipped with a metric σ to be the distance between the two boundary components of A. If the metric σ is contained in the conformal class of a complex structure on S and if the σ-area of S equals one then the modulus of $A \subset (S, \sigma)$ is bounded from below by w^2. On the other hand, if A is a complex annulus of modulus $m(A) > 0$ then the extremal length of the class of simple closed curves homotopic to either boundary component equals $1/m(A)$ (compare [**?**]).

A *core curve* of an annulus A is a simple closed curve contained in A which is freely homotopic to a boundary component. Call an annulus A in a surface S *essential* if it has a core curve which is not contractible in S. The modulus of an essential annulus in a hyperbolic surface can be related to the hyperbolic metric as follows.

Lemma 3.5.2. *1. For every $r > 0$ there is a number $m(r) > 0$ with the following property. Let c be a simple closed geodesic in a closed hyperbolic surface S. If the length of c is not bigger than r then c is the core curve of an embedded annulus in S of modulus at least $m(r)$.*

2. Let c be a simple closed geodesic in a closed hyperbolic surface of length $r > 0$. Then the modulus of any embedded annulus in S with core curve freely homotopic to c is at most π/r.

Proof. The fist part of the lemma follows from the collar lemma of hyperbolic geometry (see [**?**] for a precise statement and for references). Namely, if c is a simple closed geodesic of length $\ell(c) = r$ in a closed hyperbolic surface S then the tubular neighborhood N of c of radius $\mathrm{arcsinh}(1/\sinh(r/2))$ is an embedded annulus in S. The area of the hyperbolic metric equals $4\pi(g-1)$. Hence N is an embedded annulus in S with core curve c and modulus bounded from below by $(\mathrm{arcsinh}(1/\sinh(r/2)))^2/4\pi(g-1)$.

To show the second part, let c be a simple closed geodesic in a closed hyperbolic surface S and let $A \subset S$ be an embedded annulus with core curve freely homotopic to c. Then A is biholomorphic to a standard round annulus in $\mathbb{C}$. In particular, the universal covering of A is biholomorphic to the unit disc and hence this annulus can be equipped with a complete hyperbolic metric. There is a unique simple closed geodesic γ for this metric which is a core curve for the annulus. If $\ell > 0$ is the length of γ, then the modulus of A equals π/ℓ (see [?]).

On the other hand, it follows from the Schwarz lemma that a holomorphic map between two hyperbolic surfaces does not increase the Poincaré metric. In particular, the holomorphic embedding of the annulus A into S maps the geodesic γ for the complete hyperbolic metric on A to a curve in S whose hyperbolic length is not bigger than the length ℓ of γ. Since a geodesic for the hyperbolic metric on S minimizes the length in its free homotopy class, this implies that the length of the geodesic c is bounded from above by $\pi/\text{modulus}(A)$. This shows the second part of the lemma. □

Lemma 3.5.2 says that we can detect a short simple closed geodesic for a hyperbolic metric on S by looking for an embedded essential annulus in S with large modulus.

Hyperbolic metrics on the surface S arise via uniformization of a complex structure on S, while singular euclidean metrics on S as defined in Chapter 2 can be constructed purely combinatorially. We discuss in Section 3.6 how to construct such metrics from two simple closed curves on S whose distance in the curve graph is at least three. Unfortunately, the relation between a singular euclidean metric q on S of area one and the hyperbolic metric in the same conformal class is not so easy to understand. In particular, a simple closed geodesic for the metric q of small length may not be the core curve of an annulus of large modulus, and the largest width with respect to the metric q of an embedded annulus with very short core curve may be arbitrarily small.

The following proposition is essential for what follows. It is due to Masur and Minsky [?] and Bowditch [?] and says that a singular euclidean metric of area one admits at least one annulus of big modulus and big width. The proof we present here is due to Bowditch. We will need a version which is also valid for a singular euclidean metrics with a finite number of singularities of cone angle π.

In the following proposition, we allow the surface S to be of genus 0 or one.

Proposition 3.5.3. *Let S be a closed surface of genus $g \geq 0$ and let $k \geq 0$ be such that $3g - 3 + k \geq 2$. Then there is a number $w > 0$ only depending on g and k such that for every area one singular euclidean metric on a surface S of genus g with k singularities of cone angle π there is an embedded annulus in S of width $\geq w$.*

Proof. The main idea is to use an *isoperimetric inequality* for a singular euclidean metric q of area one on the surface S with a finite set P of singular points of cone angle π which can be stated as follows. Let $D \subset S - P$ be an embedded disc or a once punctured embedded dics with piecewise smooth boundary γ. Let $\ell(\gamma)$ be the length of γ with respect to the metric q. Then the area of D does not exceed $\frac{1}{2}(\ell(\gamma))^2$. We begin with proving the proposition under the assumption that this isoperimetric inequality holds true.

A *spine* of S is an embedded graph G in S such that $S - G$ is a union of discs or once punctured discs. Let q be a singular euclidean metric on S of area one. By the isoperimetric inequality, the q-length of a spine of S is bounded from below by a universal constant $\eta_0 > 0$.

We saw before that the number of free homotopy classes of pairwise disjoint simple arcs in S with boundary on a simple closed curve c is bounded from above by a constant $k - 1 > 0$ only depending on the genus and the number of punctures of S. Namely, cutting $S - c$ open along such an arc results in a surface S' whose Euler characteristic is strictly bigger than the Euler characteristic of $S - c$. Let $\eta_1 = \eta_0/k$ where as before, $\eta_0 > 0$ is a lower bound for the length of a spine of S. Note that η_1 only depends on the genus and the number of punctures of S.

Let $\eta_2 = \eta_1/(100 + 2k)$ and let c be a simple closed geodesic on the surface S equipped with the metric q whose length is within η_2 of the infimum of all possible lengths of *essential* simple closed curve, i.e. curves which are neither contractible nor homotopic into a puncture.

For $\epsilon > 0$ let $N_\epsilon(c)$ be the tubular neighborhood of radius ϵ about c. Let

$$\epsilon_0 = \sup\{\epsilon \mid N_\epsilon(c) = \text{annulus}\}.$$

If $\epsilon_0 \geq \eta_2$ then the annulus $N_{\eta_0}(c)$ satisfies the requirement in the proposition. Otherwise note that there are two points on the boundary of $N_{\epsilon_0}(c)$ which coincide in S and therefore there is a geodesic arc in N_{η_2} of length $2\epsilon_0 < 2\eta_2$ with both endpoints on c. Choose a maximal set $\beta_1, \dots, \beta_m \subset N_{\eta_2}$ of such simple geodesic arcs with endpoints on c which are pairwise disjoint, of length at most $2\eta_2$ which are mutually not freely homotopic relative to c and such that the image of the fundamental group of $\pi_1(N - P)$ in $\pi_1(N - S)$ coincides with the image of the fundamental group of $\sigma = c \cup \beta_1 \cup \dots \cup \beta_m$. Note that by the above consideration, we may assume that $m \leq k$.

If all components of $(S - P) - N$ are discs or once punctured discs, then $\pi_1(S - P)$ coincides with the image of the fundamental group of σ and σ is a spine. Then the length of σ is at least η_0. Since the length of each of the arcs β_i is at most $2\eta_2$, this implies by the choice of η_2 that the length $\ell(c)$ of c is not smaller than $100\eta_2$.

Now if c is non-separating then at least one of the arcs β_i connects opposite sides of c. Namely, if all the arcs β_i connect the same side of c then the intersection number of every closed curve on S with c is even. However, since c is non-separating, this is impossible.

Now assume that β_i is one of these geodesic arcs which connects opposite sides of c. Connect the endpoints of this arc with a subarc of c of length at most $\eta_1/2$. The resulting simple closed curve c_0 has one essential intersection with c and hence it is neither contractible not freely homotopic to into a puncture. The length of c_0 does not exceed $\ell(c)/2+2\eta_2$ which contradicts the assumption that the length of c is within η_2 of the minimal length of an essential simple closed curve in $S-P$.

If c is separating then by our assumption that $3g-3+k\geq 2$, each of the sides of c is connected by one of the arcs β_i and one side is connected by at least two such arcs. The there are two such arcs β_1,β_2 with endpoints a,b which cut out a subarc $\hat{c}$ of c of length at most $\ell(c)/3$. We may assume that no endpoint of $\beta_1\cup\beta_2$ is contained in $\hat{c}$. Let $\hat{a},\hat{b}$ be the second endpoints of β_1,β_2 and let $\tilde{c}$ be a subarc of c of length at most $\ell(c)/c$ which connects $\hat{b}$ to $\hat{a}$. Then the closed curve c_0 which is the concatentation of $\beta_1,\hat{c}_0,\beta_2,\tilde{c}$ has length at most $5\ell(c)/6+4\eta_0$, and it is essential. The curve c_0 may have an essential self-intersection, but in this case we may remove the self-intersection by exchanging crosswise the glueings of the component arcs and obtain in this way a simple closed curve with the required properties. However, the length of this curve is smaller than $\ell(c)-2\eta_2$ which is a contradiction. As a consequence, the graph σ is not a spine for $S-P$ and hence there exists a nontrivial essential simple closed curve d in $S-N_{\eta_2}(c)$.

Choose a smooth distance decreasing function

$$f:S\to[0,\infty) \text{ with } f(c)=0, f(d)=\eta_2.$$

Let $m=3g-2$ and for some $\ell\in\{1,\dots,3g-2\}$ let a be regular value of f contained in the interval $(\frac{(2\ell-1)\eta_2}{2m},\frac{2\ell\eta_2}{2m})$. Then $f^{-1}[0,a]\subset S$ is an embedded subsurface with smooth boundary which is different from a disc.

Let S_0 be the connected component of this surface which contains the curve c in its interior. Then the boundary of S_0 is a union of topological circles contained in $f^{-1}(a)$. If each of these circles is contractible in S or homotopic into a puncture then each of these circles bounds an embedded disc in S or a once punctured disc. But this just means that either S_0 or $S-S_0$ is a union of discs or once punctured discs which is impossible since $c\subset S_0, d\subset S-S_0$. Thus there is at least one component γ_ℓ of the boundary of S_0 which is not homotopic to zero in S, and this curve is a loop contained in $f^{-1}(\frac{(2\ell-1)\eta_2}{2m},\frac{2\ell\eta_2}{2m})$.

Since f is distance decreasing, for all $\ell\neq u$ the distance between the preimages $f^{-1}(\frac{(2\ell-1)\eta_2}{2m},\frac{2\ell\eta_2}{2m})$ and $f^{-1}(\frac{(2u-1)\eta_2}{2m},\frac{2u\eta_2}{2m})$ is at least $\eta_2/2m$. In particular, the distance in S between any two distinct of the curves γ_u, γ_ℓ is at least $\eta_2/2m$. Since there are at most $3g-3+k$ free homotopy classes of mutually disjoint simple closed essential curves on S, there are numbers $1 \leq u < \ell \leq 3g-2+k$ such that the curves γ_u, γ_ℓ are freely homotopic. Then these curves bound an annulus whose width is bounded from below by $\eta_2/2m$. This shows the proposition under the assumption that the isoperimetric inequality holds true.

To show the isoperimetric inequality, note that an embedded disc or a once punctured embedded disc in $S - P$ with piecewise smooth boundary lifts to an embedded disc in the universal covering $\tilde{S}$ of S with piecewise smooth boundary. Thus by standard approximation, it is enough to show the isoperimetric inequality for compact discs D with *smooth* boundary in the plane $\tilde{S}$ equipped with a singular euclidean metric $\tilde{q}$ with isolated singularities of cone angles $p\pi$ for some integers $p \geq 3$ and with at most one singularity of cone angle π contained in D.

Thus let γ be a smooth Jordan curve in $\tilde{S}$ bounding a disc D. Via a small deformation we may assume that γ does not pass through a singular point for the metric $\tilde{q}$. By our assumption on the cone angles of the singular points, for every number $\epsilon > 0$ there is a deformation of the metric $\tilde{q}$ to a smooth Riemannian metric g on the plane so that the following conditions are satisfied.

1. The metric g is of non-positive curvature away from small neighborhoods of the punctures.
2. Near γ, the metric g coincides with the singular euclidean metric $\tilde{q}$.
3. The difference between the g-area of the disc D and the area of D with respect to the singular euclidean metric $\tilde{q}$ is at most ϵ and with at most one cone point of cone angle π in the interior.

As a consequence, it is sufficient to prove the isometric inequality for discs with smooth boundary in simply connected 2-dimensional manifolds $\tilde{S}$ of non-positive curvature.

Let $D \subset \tilde{S}$ be such an embedded disc with smooth boundary γ. If D is not a convex subset of $\tilde{S}$ then there is a geodesic arc ζ in $\tilde{S}$ with both endpoints on γ whose interior is disjoint from D. The union of ζ with one of the two subarcs of γ with the same endpoints bound a disc D' whose interior is disjoint from D. Denote this subarc of γ by γ_1.

Since ζ is a geodesic, the length of γ_1 is not smaller than the length of ζ. Then $D \cup D'$ is a disc in $\tilde{S}$ with boundary $(\gamma - \gamma_1) \cup \zeta$. In particular, the length of the boundary of $D \cup D'$ is not bigger than the length of the boundary of D, and the area of $D \cup D'$ is bigger than the area of D. Thus for the purpose of the isoperimetric inequality, we may assume without loss of generality that the disc

D is convex. Via a small deformation we may in fact assume that D is strictly convex, i.e. every geodesic segment in $\tilde{S}$ which connects two distinct points in the boundary of D intersects γ transversely.

Let d be the distance in $\tilde{S}$ defined by the Riemannian metric. Let x_0 be the singular point in D with cone angle π and let x_1 be a point on γ whose distance to x_0 is minimal. Assume that the boundary γ of D is parametrized by arc length on $[0, \ell]$ with $\gamma(0) = x_1$ where $\ell = \ell(\gamma)$ is the length of γ. Since the distance R between x_0 and x_1 is minimal, comparison together with the fact that the curvature of $\tilde{S} - x_0$ is non-positve implies that $\ell \geq \pi R$ and hence the distance between x_0 and any point on γ is at most ℓ.

For each $t \in [0, \ell]$ let $\varphi(\cdot, t) : [0, 1] \to \tilde{S}$ be the geodesic parametrized proportional to arc length which connects x_0 to $\gamma(t)$. Since D is convex by assumption, the map $\varphi : [0, 1] \times [0, \ell] \to \tilde{S}$ restricts to a diffeomorphism of $(0, 1] \times (0, \ell)$ onto $D - \{x_0\}$. Let area* be the area of D with respect to the euclidean cone metric determined by the requirement that the arcs $s \to \varphi(s, t)$ are geodesic segments of length $d(x_0, \gamma(t))$ parametrized proportional to arc length and that the length of the tangent of the curve γ coincides with its length in the metric $\tilde{q}$. Since γ is parametrized by arc length, infinitesimally near a segment $s \to \varphi(s, t)$, this metric is given in the coordinates (s, t) by the quadratic form

$$d(x_0, \gamma(t))^2 ds^2 + a(t)s^2 dt^2 + b(t)sdsdt$$

with $a(t) \leq 1$.

Since the maximum of the distance between x_0 and a point on γ does not exceed ℓ, we obtain the estimate

$$\text{area}^* = \int_0^1 \int_0^\ell |ds \wedge dt| dsdt \leq \int_0^1 \int_0^\ell d(x_0, \gamma(t)) sdsdt \leq \frac{(\ell(\gamma))^2}{2}.$$

However, by comparison, the area of the disc D with respect to the nonpositively curved metric on $\tilde{S}$ does not exceed area*. This completes the proof of the isoperimetric inequality and hence completes the proof of the proposition. (With more effort, it can be shown that a simply connected surface of non-positive curvature satisfies the isoperimetric inequality area$(D) \leq \frac{1}{4\pi}(\text{length}(\partial D))^2$ for all embedded discs D which is sharp for round discs in the euclidean plane.) □

3.5.2 Short curves

If q is any singular euclidean metric on S of area one and if $A \subset (S, q)$ is an embedded annulus of width $w > 0$ then the extremal length of the families of curves homotopic to the boundary of A is at most $1/w^2$. Thus by Proposition

3.5.3, there is a number $\chi_0 > 0$ only depending on the genus of S with the following property. Every singular euclidean metric of area one admits a simple closed curve of length at most χ_0 which is a core curve of an annulus of width at least w. By Bers's result (see [?]), via enlarging the number χ_0 we may also assume that for every hyperbolic metric on S there is a pants decomposition of S consisting of simple closed curves of length at most χ_0.

For an area one singular euclidean metric q on S define a closed curve c on S to be *q-short* if c is freely homotopic to a curve of q-length at most χ_0. Similarly, we call a closed curve on a hyperbolic surface *short* if it is freely homotopic to a curve of length a most χ_0.

Recall that a singular euclidean metric q on the marked surface S defines a marked complex structure and hence a marked hyperbolic metric.

Corollary 3.5.4. *Let q be a singular euclidean metric on S of area one.*

1. *For every $\ell > 0$ the diameter in $\mathcal{CG}(S)$ of the set of simple closed curves of q-length at most ℓ is bounded from above by a universal constant not depending on q.*
2. *The distance in $\mathcal{CG}(S)$ between a q-short curve and a short curve for the hyperbolic metric defined by q is bounded from above by a constant $k > 0$ only depending on the genus of S.*

Proof. Let $A \subset S$ be an embedded anulus in S whose width with respect to the singular euclidean metric q is at least w where $w > 0$ is as in Proposition 3.5.3. Let $\delta \subset A$ be a core curve of A of length at most χ_0.

If c is any simple closed curve then for every essential intersection of c with δ (i.e. an intersection which can not be removed by a homotopy) there is a subarc of c connecting the two boundary curves of the annulus A. The length of this subarc is at least w, and subarcs corresponding to different essential intersection points are disjoint. Thus the length ℓ of c is not smaller than $wi(c, \delta)$ and hence $i(c, \delta) \leq \ell/w$. Lemma 3.2.3 then implies that the distance in $\mathcal{CG}(S)$ between δ and c is at most $\ell/w + 1$. This shows the first part of the corollary.

To show the second part, recall from the discussion in the beginning of this section that the modulus of the annulus A is bounded from below by w^2. By Lemma 3.5.2, the (hyperbolic) length of the hyperbolic geodesic c on S which is freely homotopic to δ is uniformly bounded from above. Moreover, c is the core curve of a hyperbolic annulus whose width is uniformly bounded from below. Using the discussion in the previous paragraph, a short curve on S with respect to the hyperbolic metric has uniformly bounded intersection number with c. From this the second part of the corollary follows. ∎

Define a map $\Psi : \mathcal{T}(S) \to \mathcal{C}(S)$ by associating to a marked hyperbolic metric σ on S a simple closed curve $\Psi(\sigma)$ which is short for σ. The following corollary is a restatement of the second part of Corollary 3.5.4.

Corollary 3.5.5. *Let q be a singular euclidean metric on S of area one and let $c \in \mathcal{C}(S)$ be a q-short curve on S. Then $d(c, \Psi(\sigma)) \le \chi_1$ where σ is the hyperbolic metric defined by q and $\chi_1 > 0$ only depends on the genus of S.*

The map Ψ is *coarsely equivariant* with respect to the action of the mapping class group $\mathrm{Mod}(S)$ on Teichmüller space $\mathcal{T}(S)$ and on the curve graph which means the following.

Corollary 3.5.6. *There is a constant $\chi_2 > 0$ such that*

$$d(\Psi(g\sigma), g\Psi(\sigma)) \le \chi_2 \text{ for all } g \in \mathrm{Mod}(S), \text{ all } \sigma \in \mathcal{T}(S).$$

Proof. If $\sigma \in \mathcal{T}(S)$ and $g \in \mathrm{Mod}(S)$ then $g\Psi(\sigma)$ and $\Psi(g\sigma)$ are both simple closed curves whose $g\sigma$-length does not exceed χ_0. □

3.6 Hyperbolicity of the curve graph

In this section we use singular euclidean metrics on S to show that the curve graph $\mathcal{GC}(S)$ of S is a hyperbolic geodesic metric space. The stratey is to construct for any two simple closed curves $\alpha, \beta \in \mathcal{C}(S)$ with $d(\alpha, \beta) \ge 3$ a "path" $\eta(\alpha, \beta) : [0, 1] \to \mathcal{CG}(S)$ connecting α to β and such that this family of paths satisfies the thin triangle condition as in the definition of a hyperbolic geodesic metric space. In a second step, we then show that these paths are necessarily contained in uniformly bounded neighborhoods of geodesics in $\mathcal{C}(S)$. For the constuction of the path $\eta(\alpha, \beta)$ we first construct from α, β a family of singular euclidean metrics of area one on S and define the points on $\eta(\alpha, \beta)$ to be short curves for these metrics.

3.6.1 Weighted geodesic multicurves and singular euclidean metrics

In Example 3.4.11 we used singular euclidean metrics to construct minimal geodesic laminations on S which fill up S. On the other hand, we can construct from a pair $\alpha, \beta \in \mathcal{C}(S)$ with $d(\alpha, \beta) \ge 3$ a singular euclidean metric on S. This subsection is devoted to this construction which we carry out more generally for *weighted multicurves*. Such a weighted multicurve is a multicurve $\alpha = \alpha_1 \cup \cdots \cup \alpha_k$, i.e a disjoint union of simple closed mutually not freely homotopic curves, together with the assignment of a positive weight $a_1, \dots, a_k \in (0,\infty)$ to each component of α.

If α, β are weighted multicurves then the *intersection number* $i(\alpha,\beta)$ is defined by the requirement that it extends the usual intersection number bilinearly. Thus if $\alpha = \cup_i a_i\alpha_i$ and $\beta = \cup_j b_j\beta_j$ with $\alpha_i, \beta_j \in \mathcal{C}(S)$ then

$$i(\alpha,\beta) = \sum_{i,j} a_i b_j i(\alpha_i,\beta_j).$$

Similarly, if γ, δ are *any* closed curves on S then we can define the intersection number $i(\gamma,\delta)$ to be the minimal number of intersection points between any closed curves freely homotopic to γ, δ and extend this definition of an intersection number bilinearly to (not necessarily disjoint) finite unions of weighted closed curves.

Now let $\alpha = \cup_i a_i\alpha_i, \beta = \cup_j b_j\beta_j$ be weighted multicurves which *jointly fill up* S. This means that the components of $S - (\bigcup_i \alpha_i \cup \bigcup_j \beta_j)$ are simply connected. For example, by Lemma 3.2.6, the condition is always satisfied if there are components $\alpha_i, \beta_j \in \mathcal{C}(S)$ with $d(\alpha_i,\beta_j) \geq 3$. Assume that $i(\alpha,\beta) = 1$. Choose representatives of the free homotopy classes of α, β with the minimal number of intersection points and such that each such intersection point is transverse. We observed in Section 3.2 that geodesic representatives with respect to a hyperbolic metric have this property. Then α, β decompose S into polygons. Each such polygon has an even number of sides, and these sides are subarcs of α, β in alternating order.

Place each intersection point of α_i, β_j in the interior of an euclidean rectangle with sides of length a_i, b_j such that the sides of length b_j are parallel to α_i and the sides of length a_i are parallel to β_j. These rectangles can be aligned in such a way that any two of them meet at most along a side or a vertex and that they cover S. More precisely, if P is a complementary polygon of $S - (\alpha \cup \beta)$ with $2m \geq 4$ sides then each of the $2m$ vertices of P is contained in a rectangle, and these rectangles are arranged in such a way that they all meet at a common vertex contained in the interior of P and that they cover P.

Since the area of each rectangle corresponding to an intersection point between α_i and β_j equals $a_i b_j$, the area of the resulting singular euclidean metric on S equals one. The singular points of this metric are p-pronged singularities for some $p \geq 3$, and they are in one-to-one correspondence with the complementary polygons of $\alpha \cup \beta$ which have more than four sides. Each of the curves α_i, β_j is the core curve of an embedded closed euclidean annulus A_i, B_j of width a_i, b_j. For $i \neq j$, the annuli A_i, A_j only intersect along the boundary, and $\cup_i A_i = S$.

Each euclidean rectangle in the above construction has two sides parallel to α and two sides are parallel to β. If we call the sides parallel to α *horizontal* and the sides parallel to β *vertical*, then the gluing of the rectangles respects the horizontal and the vertical sides. As a consequence, the natural foliation of a rectangle into horizontal (or vertical) arcs extends across the boundary of the rectangles to a foliation of the complement of the singular points which we call the *horizontal foliation* (or the *vertical foliation*). At a singular point of cone angle $p\pi$ for some $p \geq 3$ the foliations have standard singularities, with p half-leaves issuing from the singular point.

We use these facts to estimate the length of a closed geodesic in S equipped with the above singular euclidean metric.

Lemma 3.6.1. *Let q be the singular euclidean metric of area one constructed from the weighted multicurves α, β with $i(\alpha, \beta) = 1$. Then the length of a closed geodesic δ for q is contained in the interval*

$$[\max\{i(\alpha, \delta), i(\beta, \delta)\}, i(\alpha, \delta) + i(\beta, \delta)].$$

Proof. Let δ be any closed geodesic for the metric q. We show first that the length of δ is not smaller than $i(\delta, \alpha)$. Since the case $i(\delta, \alpha) = 0$ is trivial we may assume that $i(\delta, \alpha) > 0$.

By construction, every component α_j of α is the core curve of an embedded euclidean annulus A_j for the metric q of width a_j. Two distinct such annuli only intersect along their boundaries. Then to each essential intersection point of δ with a component α_j of α corresponds a subarc of δ crossing the annulus A_j. The length of this arc is not smaller than the width a_j of A_j. Moreover, two such subarcs corresponding to two different essential intersection points between δ and α_j correspond to parameter intervals for some parametrization of δ with disjoint interior. But this just means that each essential intersection of δ with α_j contributes at least a_j to the length of δ. Therefore the q-length of δ is not smaller than $i(\alpha, \delta)$.

By assumption, δ is a geodesic for the singular euclidean metric q. Thus away from the singular points, the tangent δ' of δ is defined and can be decomposed as $\delta' = \delta'_h + \delta'_v$ where δ'_h is the component of δ' in the horizontal direction and δ'_v is the component of δ' in the vertical direction. Since these directions are orthogonal, the length of δ equals

$$\int \sqrt{|\delta'_h|^2 + |\delta'_v|^2}$$

where $|\,|$ is just the euclidean norm. On the other hand, since δ is a geodesic, if δ is not a multiple of the core curve of the annulus A_j then each connected component of the intersection of δ with A_j is an euclidean geodesic connecting the two boundary components of the euclidean cylinder. Hence the angle of this component with the vertical direction is constant, and the length of the vertical component of its tangent equals a_j. Thus by the above argument, we have

$$i(\alpha, \delta) = \int |\delta'_v|$$

from which the lemma follows. □

As an immediate corollary, we obtain.

Corollary 3.6.2. *There is a unversal constant $\kappa_0 > 0$ with the following property. Let α, β be weighted multicurves with $i(\alpha, \beta) = 1$ which jointly fill up S. Then there is a simple closed curve δ which is short for the singular euclidean metric defined by α, β and which satisfies*

$$i(\delta, \rho) \leq \kappa_0 \big(i(\alpha, \rho) + i(\beta, \rho)\big) \textit{ for all } \rho \in \mathcal{C}(S).$$

Proof. Let q be the singular euclidean metric of area one defined by α, β and let δ be a geodesic for q which is the core curve of an annulus of q-width at least w where $w > 0$ is as in Proposition 3.5.3. By the discussion in Section 3.5, the curve δ is q-short. Moreover, we observed in the proof of Corollary 3.5.4 that the q-length of any closed curve ρ on S is not smaller than $wi(\rho, \delta)$. On the other hand, by Lemma 3.6.1, if ρ is a geodesic for q then the length of ρ is bounded from above by $i(\alpha, \rho) + i(\beta, \rho)$. The corollary follows with $\kappa_0 = 1/w$. □

3.6.2 A path family in $\mathcal{CG}(S)$

We use Corollary 3.6.2 to construct for a universal constant $\kappa_1 > 0$ and for any two simple closed curves $\alpha, \beta \in \mathcal{C}(S)$ with $d(\alpha, \beta) \geq 3$ a (non-continuous) map $\tilde{\eta}(\alpha, \beta) : \mathbb{R} \to \mathcal{C}(S)$ with the following properties.

1. $\tilde{\eta}(\alpha, \beta)(t) = \alpha$ for all sufficiently small $e^t \in \mathbb{R}$ and $\tilde{\eta}(\alpha, \beta)(t) = \beta$ for all sufficiently large $e^t \in \mathbb{R}$.
2. $d(\tilde{\eta}(\alpha, \beta)(s), \tilde{\eta}(\alpha, \beta)(t)) \leq \kappa_1$ if $|t - s| \leq 1$.
3. $i(\alpha, \tilde{\eta}(\alpha, \beta)(t)) i(\tilde{\eta}(\alpha, \beta)(t), \beta) \leq \kappa_1$ for all t.

Namely, for $t \in \mathbb{R}$ consider the euclidean metric q_t of area one determined by the pair $(e^{-t}\alpha, e^t\beta / i(\alpha, \beta))$. Define

$$\tilde{\eta}(\alpha, \beta)(t) = \text{ a } q_t - \text{short simple closed curve.}$$

By Lemma 3.6.1, for sufficiently small e^t the curve α is q_t-short, and for e^t very large, the curve β is q_t-short. Thus we may assume that $\tilde{\eta}(\alpha,\beta)(t) = \alpha$ for small e^t and $\tilde{\eta}(\alpha,\beta)(t) = \beta$ for large e^t. Therefore the first property above is satisfied.

By Lemma 3.6.1, for every simple closed curve c on S the length of a geodesic representative of c with respect to the metric q_t is not bigger than $2e^{|t-s|}$ times the length of a geodesic representative of c for the metric q_s. In particular, for $|t-s| \leq 1$ the q_t-length of $\tilde{\eta}(\alpha,\beta)(s)$ is bounded from above by $8\chi_0$. On the other hand, by the first part of Corollary 3.5.4, the diameter in the curve graph $\mathcal{CG}(S)$ of the set of curves whose q_t-length is bounded from above by $8\chi_0$ is bounded from above by a universal constant, independent of t. This shows that the second property above is satisfied as well.

The third property follows from the fact that by Lemma 3.6.1 and the choice of $\tilde{\eta}(\alpha,\beta)(t)$, for every $t \in \mathbb{R}$ the quantities

$$e^{-t} i(\alpha, \tilde{\eta}(\alpha,\beta)(t)) \text{ and } e^t i(\tilde{\eta}(\alpha,\beta)(t), \beta)$$

are bounded from above by χ_0.

The following lemma is due to Bowditch [**?**] and shows that the "paths" $\tilde{\eta}(\alpha,\beta)$ depend coarsely continuously on α,β.

Lemma 3.6.3. *There is a number $\kappa_2 > 0$ with the following property. Let $\alpha,\beta_1,\beta_2 \in \mathcal{C}(S)$ and assume that $d(\beta_1,\beta_2) = 1$ and that $d(\alpha,\beta_i) \geq 3$. Then the Hausdorff distance between $\tilde{\eta}(\alpha,\beta_1)(\mathbb{R})$ and $\tilde{\eta}(\alpha,\beta_2)(\mathbb{R})$ is at most κ_2.*

Proof. Let $t \in \mathbb{R}$ and consider the weighted multi-curve $e^{-t}\alpha$. Via multiplying $\beta_1, beta_2$ with a positive constant we may assume that $i(e^{-t}\alpha,\beta_i) = 1$. Then $\gamma = \frac{1}{2}(\beta_1 * \beta_2)$ is a weighted geodesic mulit-curve with $i(e^{-t}\alpha,\gamma) = 1$, and $e^{-1}\alpha,\gamma$ define a singular euclidean metric q_t of area one on S.

Let δ be a q_t-short curve on S. Then $\max\{i(e^{-t}\alpha,\delta), i(\gamma,\delta)\} \leq \chi_0$ and hence $\max\{i(\beta_1,\delta), i(\beta_2,\delta))\} \leq 2\chi_0$. As a consequence, the length of δ for both singular euclidean metrics defined by $e^{-t}\alpha,\beta_1$ and $e^{-t}\alpha,\beta_2$ does not exceed $4\chi_0$. Since the diameter in $\mathcal{C}(S)$ of the set of all curves whose length with respect to a fixed singular euclidean metric of area one does not exceed $4\chi_0$ is uniformly bounded and since $t \in \mathbb{R}$ was arbitrary, this shows the lemma. ☐

Denote by $U_r(B)$ the r-neighborhood of a subset B of the curve graph of S. The next corollary is due to Bowditch [**?**] and shows that the family of maps $\tilde{\eta}(\alpha,\beta)$ constructed above satisfies the thin triangle condition.

Lemma 3.6.4. *There is a number $\kappa_3 > 0$ and for $\alpha,\beta,\gamma \in \mathcal{C}(S)$ whose pairwise distances are at least 3 there is $\delta \in \mathcal{C}(S)$ such that*

$$\delta \in U_{\kappa_3}(\tilde{\eta}(\alpha,\beta)(\mathbb{R})) \cap U_{\kappa_3}(\tilde{\eta}(\beta,\gamma)(\mathbb{R})) \cap U_{\kappa_3}(\tilde{\eta}(\gamma,\alpha)(\mathbb{R})).$$

Proof. Let $a, b, c > 0$ be such that

$$abi(\alpha, \beta) = bci(\beta, \gamma) = cai(\gamma, \alpha) = 1.$$

Let $\delta \in \mathcal{C}(S)$ be such that

$$i(\delta, \rho) \leq 2\kappa_0 \max\{ai(\alpha, \rho), bi(\beta, \rho)\} \text{ for all } \rho \in \mathcal{C}(S)$$

where $\kappa_0 > 0$ is as in Corollary 3.6.2. Then

$$i(\delta, \gamma) \leq 2\kappa_0 \max\{ai(\alpha, \gamma), bi(\beta, \gamma)\} = 2\kappa_0/c$$

and similarly $i(\delta, \alpha) \leq 2\kappa_0/a$ and therefore

$$\max\{ci(\delta, \gamma), ai(\delta, \alpha)\} \leq 2\kappa_0.$$

Thus by Lemma 3.6.1, the length of δ with respect to the singular euclidean metric defined by $a\alpha, c\gamma$ is at most $4\kappa_0$. By Corollary 3.5.4, the diameter of the set of all curves whose length with respect to a fixed singular euclidean metric of area one on S is at most $\max\{\chi_0, 4\kappa_0\}$ is bounded from above by a universal constant $\kappa_2 > 0$. This shows that $\delta \in U_\kappa(\tilde{\eta}(\alpha, \beta)$. The same argument is also valid for the metric defined by $b\beta, c\gamma$ and the metric defined by $a\alpha, b\beta$ whence the lemma. □

As a consequence, we obtain.

Proposition 3.6.5. *There is a number $\kappa > 0$ and for any two curves $\alpha, \beta \in \mathcal{C}(S)$ there is a path $\eta(\alpha, \beta) : [0, 1] \to \mathcal{CG}(S)$ with the following properties.*

1. *If $d(\alpha, \beta) = 1$ then $\operatorname{diam}(\eta(\alpha, \beta)[0, 1]) \leq \kappa$.*
2. *If $\alpha, \beta \in \mathcal{C}(S)$ and if $s < t \in [0, 1]$ then for the Hausdorff-distance d_H for subsets of $\mathcal{CG}(S)$ we have*

$$d_H(\eta(\alpha, \beta)[s, t], \eta(\eta(\alpha, \beta)(s), \eta(\alpha, \beta)(t))[0, 1]) \leq \kappa.$$

3. *For α, β, γ, $\eta(\alpha, \beta)[0, 1] \subset U_\kappa(\eta(\beta, \gamma)[0, 1] \cup \eta(\gamma, \alpha)[0, 1])$ where $U_\kappa(A)$ denotes the κ-neighborhood of a set A.*

Proof. If $\alpha, \beta \in \mathcal{C}(S)$ jointly fill up S then define $\eta(\alpha, \beta)$ to be a reparametrization of $\tilde{\eta}(\alpha, \beta)[-k, k]$ where $k > 0$ is sufficiently large that $\tilde{\eta}(\alpha, \beta)(t) = \alpha$ for all $t \leq -k$ and $\tilde{\eta}(\alpha, \beta)(t) = \beta$ for $t \geq k$.

Curves α, β with $d(\alpha, \beta) = 2$ intersect and fill a subsurface S_0 of S. This means that S_0 is a connected bordered subsurface of S containing $\alpha \cup \beta$ in its interior, and $S_0 - \alpha \cup \beta$ is a union of topological discs and annuli homotopic to a boundary component of S_0. Thus for every $t > 0$ we can define the $(e^{-t}\alpha, e^t\beta)$-length of a curve δ in S to be $(e^t i(\alpha,\delta) + e^{-t} i(\beta,\delta))/i(\alpha,\beta)$. The diameter in $\mathcal{C}(S)$ of the set of all short curves is uniformly bounded as before. With this definition, the length of δ vanishes if and only if δ is freely homotopic to a curve contained in $S - S_0$.

Now let $\alpha_1, \alpha_2, \beta$ be any curves with $d(\alpha_i, \beta) \geq 2$. Choose $a_1 > 0, a_2 > 0$ such that $a_1 i(\alpha_1, \beta) = a_2 i(\alpha_2, \beta)$. As before, we can do the above construction with the weighted multicurves $a_1\alpha_1 \cup a_2\alpha_2, \beta$. We conclude that for every t a curve which is short for the pseudo-metric q_t defined by $(e^{-t}(a_1\alpha_1 \cup a_2\alpha_2), e^t\beta)$ is short for the metrics defined by $e^{-t}a_i\alpha_i, e^t\beta$ for $i = 1, 2$. Thus if we define the curve $\eta(c,d)$ by associating to t a short curve for q_t then the curve system $\eta(c,d)$ varies coarsely continuously with c, d: If $d(c,c') \leq 1$ then the Hausdorff distance between $\eta(c,d)$ and $\eta(c',d)$ is uniformly bounded. $\square$

3.6.3 Hyperbolicity

Recall from Definition 3.1.5 the definition of a hyperbolic geodesic metric space, and recall from Definition 3.1.2 the definition of a quasi-geodesic. Our goal is to show that the maps $\eta(\alpha,\beta)$ defined in the preceding subsection are quasi-geodesics up to parametrization. The following extension of the notion of quasi-geodesic is adapted to this idea.

Definition 3.6.6. Let X be a geodesic metric space. For a number $L \geq 1$ and a connected subset J of the real line, a map $\gamma : J \to X$ is called an *unparametrized L-quasi-geodesic* if there is a connected subset $I \subset \mathbb{R}$ and there is a homeomorphism $\varphi : I \to J$ such that $c \circ \varphi : I \to X$ is an L-quasi-geodesic.

Hyperbolicity of the curve graph now follows from Proposition 3.6.5 and the following criterion.

Proposition 3.6.7. *Let (X,d) be a geodesic metric space. Assume that there is a number $D > 0$ and for every pair of points $x, y \in X$ there is a map $\eta(x,y) : [0,1] \to X$ connecting $\eta(x,y)(0) = x$ to $\eta(x,y)(1) = y$ so that the following three conditions are satisfied.*

1. *If $d(x,y) \leq 1$ then the diameter of $\eta(x,y)[0,1]$ is at most D.*
2. *For $x, y \in X$ and $0 \leq s \leq t \leq 1$, the Hausdorff distance between $\eta(x,y)[s,t]$ and $\eta(\eta(x,y)(s), \eta(x,y)(t))[0,1]$ is at most D.*
3. *For any $x, y, z \in X$ the set $\eta(x,y)[0,1]$ is contained in the D-neighborhood of $\eta(x,z)[0,1] \cup \eta(z,y)[0,1]$.*

Then (X,d) is δ-hyperbolic for a number $\delta > 0$ only depending on D, and there is a number $L > 1$ such that each of the curves $\eta(x,y)$ is an unparametrized L-quasi-geodesics.

Proof. Let (X,d) be a geodesic metric space. Assume that there is a number $D > 0$ and there is a family of paths $\eta(x,y) : [0,1] \to X$, one for every pair of points $x, y \in X$, which satisfy the hypotheses in the statement of the proposition. To show hyperbolicity for X it is then enough to show the existence of a constant $\kappa > 0$ such that for all $x, y \in X$ and every geodesic $\nu : [0,\ell] \to X$ connecting $\nu(0) = x$ to $\nu(\ell) = y$, the Hausdorff-distance between $\nu[0,\ell]$ and $\eta(x,y)[0,1]$ is at most κ. Namely, if this is the case then for every geodesic triangle with sides a, b, c the side a is contained in the $3\kappa + D$-neighborhood of $b \cup c$.

To show the existence of such a constant $\kappa > 0$, let $x, y \in X$ and let for the moment $c : [0, 2^k] \to X$ be *any* path of length $\ell(c) \le 2^k$ parametrized proportional to arc length and connecting $c(0) = x$ to $c(2^k) = y$. Write $\eta_1 = \eta(c(0), c(2^{k-1}))$ and write $\eta_2 = \eta(c(2^{k-1}), c(2^k))$. By the third property stated in the proposition, the D-neighborhood of $\eta_1 \cup \eta_2$ contains $\eta(c(0), c(2^k))$. Repeat this construction with the points $c(2^{k-2}), c(3 \cdot 2^{k-2})$ and the arcs η_1, η_2.

Inductively we conclude that the path $\eta(c(0), c(2^k))$ is contained in the $(\log_2 \ell(c))D$-neighborhood of a path $\tilde{c} : [0, 2^k] \to \mathcal{C}(S)$ whose restriction to each interval $[m-1, m]$ $(m \le 2^k)$ equals up to parametrization the arc $\eta(c(m-1), c(m))$. Since $d(c(m-1), c(m)) \le 1$, by the first requirement in the proposition, the diameter of each of the sets $\eta(c(m-1), c(m))[0,1]$ is bounded from above by D and therefore the arc $\eta(c(0), c(2^k))$ is contained in the $(\log_2 \ell(c))D + D$-neighborhood of $c[0, 2^k]$.

Now let $c : [0,k] \to X$ be a geodesic connecting $c(0) = x$ to $c(k) = y$ which is parametrized by arc length. Let $t > 0$ be such that $\eta(x,y)(t)$ has maximal distance to $c[0,k]$, say that this distance equals χ. Choose some $s > 0$ such that $d(c(s), \eta(x,y)(t)) = \chi$ and let $t_1 < t < t_2$ be such that $d(\eta(x,y)(t), \eta(x,y)(t_u)) = 2\chi$ $(u = 1, 2)$. In the case that there is no $t_1 \in [0,t)$ (or $t_2 \in (t,1]$) with $d(\eta(x,y)(t), \eta(x,y)(t_1)) \ge 2\chi$ (or $d(\eta(x,y)(t), \eta(x,y)(t_2)) \ge 2\chi$) we choose $t_1 = 0$ (or $t_2 = 1$).

By our choice of χ, there are numbers $s_u \in [0,k]$ such that

$$d(c(s_u), \eta(x,y)(t_u)) \le \chi \, (u = 1, 2).$$

Then the distance between $c(s_1)$ and $c(s_2)$ is at most 6χ. Compose the subarc $c[s_1, s_2]$ of c with a geodesic connecting $\eta(x,y)(t_1)$ to $c(s_1)$ and a geodesic connecting $c(s_2)$ to $\eta(x,y)(t_2)$. We obtain a curve ν of length at most 8χ. By our above observation, the $(\log_2(8\chi))D + D$-neighborhood of this curve contains the arc $\eta(\eta(x,y)(t_1), \eta(x,y)(t_2))$.

However, the Hausdorff distance between $\eta(x,y)[t_1,t_2]$ and $\eta(\eta(x,y)(t_1),\eta(x,t)(t_2))$ is at most D and therefore the $(\log_2(8\chi))D+2D$-neighborhood of the arc ν contains $\eta(x,y)[t_1,t_2]$. But the distance between $\eta(t)$ and our curve ν equals χ by construction and hence we have $\chi \le (\log_2(8\chi))D + 2D$. In other words, χ is bounded from above by a universal constant $\kappa > 0$, and $\eta(x,y)$ is contained in the κ-neighborhood of the geodesic c.

A similar argument also shows that the 3κ-neighborhood of $\eta(x,y)$ contains $c[0,k]$. Namely, by the above consideration, for every $t \le 1$ the set $A(t) = \{s \in [0,k] \mid d(c(s),\eta(x,y)(t)) \le \kappa\}$ is a non-empty closed subset of $[0,k]$. The diameter of the sets $A(t)$ is bounded from above by 2κ.

Assume to the contrary that $c[0,k]$ is not contained in the 3κ-neighborhood of $\eta(x,y)$. Then there is a subinterval $[a_1,a_2] \subset [0,k]$ of length $a_2 - a_1 \ge 4\kappa$ such that $d(c(s),\eta(x,y)[0,1]) > \kappa$ for every $s \in (a_1,a_2)$. Since $d(c(s),c(t)) = |s-t|$ for all s,t we conclude that for every $t \in [0,1]$ the set $A(t)$ either is entirely contained in $[0,a_1]$ or it is entirely contained in $[a_2,k]$.

Define $C_1 = \{t \in [0,1] \mid A(t) \subset [0,a_1]\}$ and $C_2 = \{t \in [0,1] \mid A(t) \subset [a_2,1]\}$. Then the sets C_1, C_2 are disjoint and their union equals $[0,1]$; moreover, we have $0 \in C_1$ and $1 \in C_2$. On the other hand, the sets C_i are closed. Namely, let $(t_i) \subset C_1$ be a sequence converging to some $t \in [0,1]$. Let $s_i \in A(t_i)$ and assume after passing to a subsequence that $s_i \to s \in [0,a_1]$. Now $\kappa \ge d(c(s_i),\eta(x,y)(t_i)) \to d(c(s),\eta(x,y)(t))$ and therefore $s \in A(t)$ and hence $t \in C_1$. However, $[0,1]$ is connected and hence we arrive at a contradiction. In other words, the geodesic c is contained in the 3κ-neighborhood of $\eta(x,y)$. This completes the proof of the proposition. □

Recall from Section 3.5 the definition of the map $\Psi : \mathcal{T}(S) \to \mathcal{C}(S)$. As an immediate corollary of Proposition 3.6.5 and Proposition 3.6.7 we obtain.

Theorem 3.6.8 ([?, ?]). *The curve graph is hyperbolic. Moreover, there is a number $L > 1$ such that for any weighted multicurves α,β with $i(\alpha,\beta) = 1$ which jointly fill up S, the following holds. For $t \in \mathbb{R}$ let $\sigma_t \in \mathcal{T}(S)$ be the marked complex structure defined by the singular euclidean metric determined by $(e^{-t}\alpha, e^t\beta)$. Then the assignment which associates to $t \in \mathbb{R}$ the curve $\Psi(\sigma_t) \in \mathcal{C}(S)$ is an unparametrized L-quasi-geodesic.*

3.7 Currents and intersection

3.7.1 Geodesic currents

In the previous sections we constructed from simple closed curves $\alpha,\beta \in \mathcal{C}(S)$ with $d(\alpha,\beta) \ge 3$ and numbers $a,b > 0$ with $abi(\alpha,\beta) = 1$ a singular euclidean metric q on S with the property that the q-length of any closed q-geodesic γ is approimately equal to $ai(\alpha,\gamma)+bi(\beta,\gamma)$. In this section we generalize the intersection number between closed curves which will enable us to get a better understanding of all singular euclidean metrics on S of area one. The results we present here are due to Bonahon [**?**].

Fix a hyperbolic structure on S and let α,β be any closed geodesics on S. Since geodesics intersect in the minimal possible number of points, the intersection number $i(\alpha,\beta)$ between α and β can be calculated as follows.

Let $\tilde\alpha$ be a lift of α to the hyperbolic plane. Then $\tilde\alpha$ is a biinfinite geodesic which is invariant under a deck transformation $g\in\pi_1(S)$ whose conjugacy class defines α. The deck transformation g acts on $\tilde\alpha$ as a translation. Choose a compact connected fundamental domain $J\subset\tilde\alpha$ for this action whose endpoints do not project to an intersection point of α and β in S. Then the distinct intersection points of α and β are in one-to one correspondence to the intersection points between J and the collection of all lifts of β to $\mathbf{H}^2$. As a consequence, if we define a $\pi_1(S)$-invariant locally finite measure $\Psi(\beta)$ on the space of unoriented geodesics $\mathcal{E}=(S^1\times S^1-\Delta)/\iota$ in $\mathbf{H}^2$ by placing a Dirac mass at every point corresponding to a lift of β to $\mathbf{H}^2$ then $i(\alpha,\beta)$ is just the $\Psi(\beta)$-volume of the set of all unoriented geodesics which pass through J. More generally, we can calculate weighted intersection numbers in the same way by simply multiplying Dirac measures by weights.

Weighted Dirac masses on unordered pairs of points corresponding to lifts of a closed geodesic on S can be generalized as follows.

Definition 3.7.1. A *geodesic current* for S is a $\pi_1(S)$-invariant locally finite Borel measure on the space $\mathcal{E}=(S^1\times S^1-\Delta)/\iota$ of unoriented geodesics on $\mathbf{H}^2$.

The space $\mathrm{Curr}(S)$ of geodesic currents can be equipped with the weak*-topology. With this topology, a sequence $\{\xi_i\}$ of geodesic currents converges to $\xi\in\mathrm{Curr}(S)$ if and only if for every continuous function f on $\mathcal{E}$ with compact support we have $\int f d\xi_i\to\int f d\xi$. Note that by the requirement that a current is invariant under the action of $\pi_1(S)$ and by the fact that there is a compact subset $K\subset\mathcal{E}$ such that $\cup_{g\in\pi_1(S)}gK=\mathcal{E}$, the space of geodesic currents μ with $\mu(K)\leq$ const is compact. Note also that the space of geodesic currents is a closed convex subset of a topological vector space. In particular, a current can be multiplied with a positive number and the result is again a current. Moreover, the sum of two currents is a current.

As in Section 5, let $PT\mathbf{H}^2$ be the projectivized tangent space of $\mathbf{H}^2$. There is a natural continuous $\pi_1(S)$-equivariant projection $\Pi:PT\mathbf{H}^2\to(S^1\times S^1-\Delta)/\iota=\mathcal{E}$ which associates to a tangent line in $\mathbf{H}^2$ the unordered pair of endpoints of the geodesic tangent to this line. The preimage of any point in $\mathcal{E}$ is just a leaf of the geodesic foliation. Thus each point $v\in PT\mathbf{H}^2$ has a neighborhood which is naturally homeomorphic to the product of a neighborhood of $\Pi(v)$ in $(S^1 \times S^1 - \Delta)/\iota$ with an interval in the real line.

The metric on $\mathbf{H}^2$ induces a $\pi_1(S)$-invariant length element on each leaf of the geodesic foliation. This length element can be viewed as a family dt of locally finite Borel measures on the leaves. The product of a geodesic current μ with this length element defines a $\pi_1(S)$-invariant locally finite product measure on $PT\mathbf{H}^2$. Since $\pi_1(S)$ acts properly and cocompactly on $PT\mathbf{H}^2$, this measure induces a finite measure $\rho(\mu)$ on the projectivized tangent bundle PTS of the closed surface S. We call the total mass of $\rho(\mu)$ the *length* of μ. With this definition, the length of the current defined by the Dirac masses of the lifts of a closed geodesic α in S is just the length of α in S.

Lemma 3.7.2. *The function which associates to $\mu \in \mathrm{Curr}(S)$ its length is continuous.*

Proof. For the proof of the lemma, simply observe that the assignment which associates to a current μ its product with the length element on the leaves of the geodesic foliation is continuous with respect to the weak*-topology. □

Since PTS is compact, the space of Borel measures on PTS of total mass at most one is compact with respect to the weak$*$ topology. Thus we obtain.

Corollary 3.7.3. *The space of geodesic currents of length at most one is compact.*

As a consequence, if $\mu \in \mathrm{Curr}(S)$ has an *atom*, i.e. a point with positive measure, then the length of the leaf of the geodesic foliation in PTS defined by this atom is finite. In other words, each atomic measure in the space $\mathrm{Curr}(S)$ is supported on the lifts of closed geodesics in S. Note also that every geodesic current ξ can be written in the form $\xi = a\xi_0$ where $a > 0$ and where ξ_0 is a current of length one.

Recall that the space of weighted closed geodesics embeds into the space of geodesic currents. We have.

Lemma 3.7.4. *Any geodesic current can be approximated in the weak*-topology by sums of weighted closed geodesics.*

Proof. The lemma is not needed in the sequel, and we only give a sketch of the proof.

The unit tangent bundle T^1S of S is a two-fold covering of PTS. The covering map associates to a unit tangent vector the tangent line it defines. The *geodesic flow* Φ^t acts on T^1S by associating to a unit tangent vector v

and a time $t \in \mathbb{R}$ the unit tangent $\Phi^t v = \gamma_v'(t)$ at t of the geodesic γ_v in S with initial velocity $\gamma_v'(0) = v$. A geodesic current lifts to a Φ^t-invariant Borel measure on T^1S. Thus it is enough to show that every Φ^t-invariant Borel probability measure on T^1S can be approximated in the weak*-topology by sums of φ^t-invariant measures supported on closed orbits.

Since every Φ^t-invariant Borel probability measure on T^1S can be represented as a generalized convex combination of ergodic invariant probability measures (i.e. measures μ with the property that every Φ^t-invariant Borel subset of T^1S either has full measure or zero measure), it is enough to show that every ergodic invariant probability measure can be approximated by measures supported on closed orbits.

Now if μ is an ergodic invariant probability measure on T^1S, then by the Birkhoff ergodic theorem, there is a point $v \in T^1S$ such that μ is a weak*-limit of measures

$$\frac{1}{T}\int_0^T \delta_{\Phi^t v}$$

where δ_w is the Dirac measure at w. Moreover, by the Poincaré recurrence theorem, the orbit of v is necessarily recurrent, i.e. it passes through every neighborhood U of v for arbitrarily large times. Now as in the proof of Lemma 3.4.14, if U is sufficiently small and if T is sufficiently large with $\Phi^T v \in U$, then the unit tangent of the geodesic representative of the free homotopy class defined as a concatenation of the geodesic arc $\gamma_v[0,T]$ with a short arc connecting $\gamma_v(T)$ to $\gamma_v(0)$ is contained in a small neighborhood of the arc $\{\Phi^t v \mid t \in [0,T]\}$. An application of this observation to longer and longer subsegments of the orbit through v returning to smaller and smaller neighborhoods of v gives an approximation as required. This yields the lemma. $\square$

Write again $\mathcal{E} = (S^1 \times S^1 - \Delta)/\iota$. Note that $\pi_1(S)$ acts on $\mathcal{E} \times \mathcal{E}$ as a group of homeomorphisms. Let moreover $\mathcal{D}$ be the complement of the diagonal in the bundle $PT\mathbf{H}^2 \oplus PT\mathbf{H}^2$ over $\mathbf{H}^2$. Thus a point in $\mathcal{D}$ is a point $x \in \mathbf{H}^2$ together with a pair of transverse lines in the tangent space of $\mathbf{H}^2$ at x. We have.

Lemma 3.7.5. *There is a $\pi_1(S)$-equivariant embedding $\mathcal{D} \to \mathcal{E} \times \mathcal{E}$.*

Proof. There is a natural $\pi_1(S)$-equivariant continuous map $PT\mathbf{H}^2 \oplus PT\mathbf{H}^2 \to \mathcal{E} \times \mathcal{E}$ which maps a pair of tangent lines to the corresponding pair of unoriented geodesics. Since two distinct geodesics in $\mathbf{H}^2$ intersect in at most one point, the restriction of this map to the set $\mathcal{D}$ is injective. $\square$

A pair of currents $(\mu, \nu) \in \mathrm{Curr}(S) \times \mathrm{Curr}(S)$ defines a locally finite $\pi_1(S)$-invariant product measure on $\mathcal{E} \times \mathcal{E}$. This measure restricts to a $\pi_1(S)$-invariant measure on $\mathcal{D}$ via the embedding $\mathcal{D} \to \mathcal{E} \times \mathcal{E}$. By invariance, the measure projects to a measure on the complement $\mathcal{D}_0$ of the diagonal in the bundle $PTS \oplus PTS$ over the closed surface S.

Definition 3.7.6. The intersection number $i(\mu,\nu)$ between two currents $\mu,\nu \in \mathrm{Curr}(S)$ is the total mass of the measure induced from $\mu \times \nu$ on $\mathcal{D}_0 \subset PTS \oplus PTS$.

Note that by construction, the intersection form is symmetric and convex bilinear. This means that $i(\alpha,\beta) = i(\beta,\alpha)$, that $i(a\alpha,\beta) = ai(\alpha,\beta)$ for all α,β, all $a > 0$ and $i(\alpha,\beta_1+\beta_2) = i(\alpha,\beta_1) + i(\alpha,\beta_2)$.

We observe.

Lemma 3.7.7. *The intersection number between any two currents is finite.*

Proof. Let K be a compact fundamental domain for the action of $\pi_1(S)$ on $\mathbf{H}^2$. For any geodesic current μ, the μ-mass of the set $G(K)$ of all geodesics passing through K is finite. By construction, the intersection number $i(\mu,\nu)$ does not exceed $\mu(G(K))\nu(G(K))$. □

Note that the intersection number between the geodesic current α defined by a closed geodesic in S and a geodesic current β equals the β-mass of the set of all geodesics which intersect transversely a fundamental domain in the axis of an element $g \in \pi_1(S)$ defining the free homotopy class of α. In particular, if β is the current defined by a closed geodesic then $i(\alpha,\beta) = i(\beta,\alpha)$ equals the usual geometric intersection number between α and β as defined earlier.

The following proposition is due to Bonahon [**?**] and is a very important tool for the understanding of Teichmüller space and Kleinian groups.

Proposition 3.7.8. *The intersection form is a continuous with respect to the weak*-topology.*

Proof. The intersection number $i(\alpha,\beta)$ of two geodesic currents $\alpha,\beta \in \mathrm{Curr}(S)$ is an evaluation of the measure $\alpha \times \beta$ on the complement of the diagonal $\mathcal{D}_0$ in the bundle $PTS \oplus PTS$ obtained via the embedding in Lemma 3.7.5. The measure $\alpha \times \beta$ on $\mathcal{D}_0$ depends continuously on α and β in the weak*-topology. Since $\mathcal{D} \subset PT\mathbf{H}^2 \times PT\mathbf{H}^2$ is open and $\pi_1(S)$-invariant, this implies that the intersection form is lower semi-continuous: If $\{(\alpha_i,\beta_i)\} \subset \mathrm{Curr}(S) \times \mathrm{Curr}(S)$ is any sequence converging to (α,β), then we have

$$i(\alpha,\beta) \leq \liminf_{i\to\infty} i(\alpha_i,\beta_i).$$

Thus we only have to show that

$$i\alpha,\beta) \geq \limsup_{i\to\infty} i(\alpha_i,\beta_i)$$

holds true as well.

Now if a sequence of points $(x_i, y_i) \in TPS \oplus TPS$ converges to the diagonal and if $(\tilde{x}_i, \tilde{y}_i) \in PT\mathbf{H}^2 \times PT\mathbf{H}^2$ are lifts of (x_i, y_i) which are contained in a fixed compact fundamental domain for the action of $\pi_1(S)$ then the images of (x_i, y_i) under the embedding $\mathcal{D} \to \mathcal{E} \times \mathcal{E}$ converge to the diagonal in $\mathcal{E} \times \mathcal{E}$. This immediately implies the following. Let Δ be the diagonal in $\mathcal{E} \times \mathcal{E}$. If $(\alpha_i, \beta_i) \to (\alpha, \beta)$ in $\mathrm{Curr}(S) \times \mathrm{Curr}(S)$ and if $\alpha \times \beta(\Delta) = 0$ then $i(\alpha_i, \beta_i) \to i(\alpha, \beta)$.

However, a product measure $\alpha \times \beta$ on $\mathcal{E} \times \mathcal{E}$ only gives positive mass to the diagonal if it has atoms. Moreover, if $\alpha = \alpha_0 + \alpha_1, \beta = \beta_0 + \beta_1$ where α_1, β_1 are the atomic parts of α, β then $\alpha \times \beta(\Delta) = \alpha_0 \times \beta_0(\Delta)$. Since atomic currents correspond to sums of weighted closed geodesics, it is enough to show the following. Let α, β be two currents defined by weighted closed geodesics on S of length one. If $\alpha_i \to \alpha, \beta_i \to \beta$ then $i(\alpha_i, \beta_i) \to i(\alpha, \beta)$. Since the length function for geodesic currents is continuous, we may assume that the length of each of the currents α_i, β_i equals one.

Let $\epsilon_0 > 0$ be smaller than half the injectivity radius of S. For $\epsilon < \epsilon_0$ and a closed set A of S let $U_\epsilon(A)$ be the ϵ-neighborhood of A in the surface S. Choose points $x_0 \in \alpha, y_0 \in \beta$ which are not intersection points between α and β or self-intersection points of α or β. Let $\epsilon_2 < \epsilon_1 < \epsilon_0$ be sufficiently small that the following holds true. Let c_0, d_0 be compact geodesic arcs of length $2\epsilon_2$ with midpoints x_0, y_0 which are orthogonal to α, β and which intersect $\alpha \cup \beta$ only at x_0 (or y_0). If $m \geq 1$ and if γ, η are geodesic arcs with both endpoints on c_0, d_0 which are entirely contained in $U_{\epsilon_1}(\alpha), U_{\epsilon_2}(\beta)$ and which have precisely $m - 1$ intersection points with c_0, d_0 in their interior then γ and η intersect in either $mi(\alpha, \beta)$ or in $mi(\alpha, \beta) + 1$ points (the constant 1 in this estimate stems from the fact that we may have $\alpha = \beta$, and in this case nearby lifts of γ, η in the universal covering $\mathbf{H}^2$ may cross with a very small intersection angle).

Let $V(\alpha), V(\beta)$ be a small two-dimensional submanifold of PTS which is transverse to the geodesic foliation, contains the tangent lines v_0, w_0 of α, β at x_0, y_0 in its interior and consists of tangent lines v, w with footpoints on c_0, d_0 and such that every v, w is the tangent line of a geodesic arc $\gamma(v), \gamma(w)$ with the above properties. We view $\gamma(v), \gamma(w)$ as a compact subsegment of a leaf of the geodesic foliation on PTS. Denote by $W(\alpha), W(\beta) \subset PTS$ the union of these subsegments. Since $(\alpha_i, \beta_i) \to (\alpha, \beta)$, for every $\epsilon > 0$ there is some $i(\epsilon) > 0$ such that for every $i > i(\epsilon)$ the α_i-mass of $W(\alpha)$ and the β_i-mass of $W(\beta)$ is at least $1 - \epsilon$. Then the intersection $i(\alpha_i, \beta_i)$ is bounded from above by $(1 - \epsilon)^2 i(\alpha, \beta) + \epsilon^2 \kappa$ where $\kappa > 0$ is a universal constant. Since $\kappa > 0$ was arbitrary, the claim follows. This completes the proof of the proposition. $\square$

The *support* μ of a geodesic current is the smallest $\pi_1(S)$-invariant closed subset B of $\mathcal{E}$ such that $\mu(\mathcal{E} - B) = 0$. The space of geodesic currents with vanishing self-intersection number can now be characterized as follows.

Proposition 3.7.9. *For a geodesic current μ, the following are equivalent.*

1. $i(\mu, \mu) = 0$.

2. The support of μ *is a geodesic lamination.*

Proof. Let μ be a geodesic current whose support is a geodesic lamination λ. Then no leaf of λ intersects any leaf of λ transversely. Thus if we view λ as a subset of $\mathcal{E}$, then $\lambda \times \lambda$ is contained in the complement of the image of the set $\mathcal{D}$ under the embedding $\mathcal{D} \to \mathcal{E} \times \mathcal{E}$. By the definition of the intersection number, this implies that $i(\mu, \mu) = 0$

On the other hand, if the support $\text{supp}(\mu)$ of μ is *not* a geodesic lamination then there are density points x, x for μ so that the pair (x, y) is contained in the image of $\mathcal{D}$ under the embedding described in Lemma 3.7.5. But this just means that the $\mu \times \mu$-mass of the image of $\mathcal{D}$ is positive. By definition, we conclude that $i(\mu, \mu) > 0$. □

A translation invariant transverse measure for a geodesic lamination λ is an assignment of a finite positive Borel measure to each compact arc which intersects λ transversely, with endpoints in the complementary components of λ, and which is invariant under transverse translation of the arc. Proposition 3.7.9 shows that a geodesic lamination equipped with a transverse translation invariant measure is nothing else but a geodesic current of vanishing self-intersection.

Definition 3.7.10. A *measured geodesic lamination* on S is a geodesic lamination together with a *translation invariant transverse measure.*

By continuity of the intersection form, the subset of $\text{Curr}(S)$ of all measured geodesic laminations is closed.

The restriction of the weak*-topology to the space of geodesic currents can also be described as follows. A sequence $\{\lambda_i\}$ of measured geodesic laminations converges to a measured geodesic lamiantion λ if for each compact arc $\rho \subset S$ which is transverse to λ and whose endpoints are contained in complementary components for λ there is some $i_0 > 0$ such that ρ is transverse to λ_i for every $i \geq i_0$, with endpoints in complementary components of λ_i, and such that the measures disposed on ρ by λ_i converge as $i \to \infty$ weakly to the measure disposed on ρ by λ.

Example 11: Weighted simple multicurves
A weighted simple multicurve is a measured geodesic lamination such that the transverse measure associates to each transverse arc the weighted Dirac measures of the intersection points of the arc with c.

Definition 3.7.11. Two measured geodesic laminations λ, μ *jointly fill up* S if $i(\mu, \zeta) + i(\lambda, \zeta) > 0$ for every $\zeta \in \mathcal{ML}$.

Lemma 3.7.12. *1. The set of pairs $(\lambda, \mu) \in \mathcal{ML} \times \mathcal{ML}$ which jointly fill up S is an open subset of $\mathcal{ML} \times \mathcal{ML}$.*

2. If μ, λ are weighted simple closed geodesics then μ, λ jointly fill up if and only if the distance of their supports in $\mathcal{C}(S)$ is at least 3.

Proof. To show the first part of the lemma, note that the condition is invariant under scaling, so it is enough to show the following. Let $\mathcal{A} \subset \mathcal{ML}$ be the compact space of all measured geodesic laminations of length one on S. Then the set of all $\lambda, \mu \in \mathcal{A} \times \mathcal{A}$ so that the function $\beta \in \mathcal{A} \to i(\lambda, \beta) + i(\mu, \beta) \in [0, \infty)$ assumes only positive values is open. However, this is immediate from compactness and the continuity of the intersection form.

To show the second part of the lemma, note first that if $d(\alpha, \beta) \leq 2$ then there is a simple closed curve γ which is disjoint from both α, β and hence α, β do not jointly fill up S. On the other hand, if $d(\alpha, \beta) \geq 3$ the α, β define a singular euclidean metric q of area one. Let h be the corresponding hyperbolic metric. Then the identity map $(S, q) \to (S, h)$ is bilipschitz. In particular, there is a constant $\delta > 0$ such that the q-length of every closed curve c is not smaller than $\delta \ell_h(c)$. In other words, by Lemma 3.6.1, we have $\delta \ell_h(c) \leq (i(\alpha, c) + i(\beta, c))/i(\alpha, \beta)$. By continuity and the fact that closed curves are dense in the space of all currents we conclude that for every measured geodesic lamination μ the value $i(\alpha, \mu) + i(\beta, \mu)$ is not smaller than $\delta / i(\alpha, \beta)$ times the h-length of μ. In particular, this function is positive. ☐

3.8 Cobounded Teichmüller geodesics

3.8.1 Quadratic differentials

In this section we relate the geometry of the curve graph to the geometry of Teichmüller space equipped with the *Teichmüller metric*. We begin with a short discussion of those properties of *quadratic differentials* needed to describe the relevant properties of the Teichmüller metric.

Recall from Section 7 that any two weighted multicurves α, β which jointly fill up S and satisfy $i(\alpha, \beta) = 1$ define a singular euclidean metric $q = q(\alpha, \beta)$ of area one on S. There is a tiling of S by euclidean rectangles for q. This tiling induces a pair of orthogonal geodesic foliations on the complement of the singular points called the *horizontal* and the *vertical* foliation. Using these foliations we can distinguish a family of isometric charts on the complement of the singular points of q into the complex plane which are determined by the requirement that they preserve the singular euclidean metric and map horizontal line segments to line segments in the plane parallel to the real axis. Coordinate changes for these charts are translations in $\mathbb{C}$ or the composition of a translation with the reflection $z \to -z$. In particular, these charts define the complex structure underlying q.

The holomorphic quadratic differential dz^2 on $\mathbb{C}$ (which can be viewed as the standard trivialization of the tensor product of the cotangent bundle of the complex plane with itself) is invariant under such coordinate changes and hence it defines a holomorphic quadratic differential on the complement in S of the singular points. Near a singular point of cone angle $p\pi$, equipped with the holomorphic coordinate given by the function $\zeta \to \zeta^{p/2}$, the quadratic differential dz^2 is represented in the form $\zeta \to \frac{p+2}{2}\zeta^p d\zeta^2$ and hence admits a natural holomorphic extension, with a zero of order $p-2$ at the singular point. In other words, the singular euclidean metric determines a complex structure $h \in \mathcal{T}(S)$ together with a *holomorphic quadratic differential* for h. For each $t \in \mathbb{R}$ the pair of weighted multicurves $e^{-t}\alpha, e^t\beta$ defines a complex structure and a holomorphic quadratic differential which are obtained from the structure for (α, β) by simply postcomposing the above family of isometric charts with the matrix

$$A(t) = \begin{pmatrix} e^t & \\ & e^{-t} \end{pmatrix}$$

On the other hand, every holomorphic quadratic differential ζ on a Riemann surface S determines a family of distinguished holomorphic charts on the complement of the zeros of S whose transition functions are translations or the composition of translations with the reflection. Namely, locally near any regular point for ζ there is a holomophic coordinate z such that in these coordinates, the quadratic differential is represented in the form dz^2. The coordinate is unique up to translation and reflection. The area of the quadratic differential is the area of the corresponding singular euclidean metric. Postcomposition of these distinguished charts with the family of matrices $A(t)$ $(t \in \mathbb{R})$ defines a flow Φ^t on the bundle of quadratic differentials over Teichmüller space which is called the *Teichmüller flow.*

A *Teichmüller geodesic* on S parametrized by arc length is the projection to S of a flow line of the Teichmüller flow. In particular, every pair (α, β) of weighted geodesic multicurves (α, β) with $i(\alpha, \beta) = 1$ which jointly fill up S defines a holomorphic quadratic differential $q(\alpha, \beta)$ of area one.

The following result is due to Teichmüller and to Jenkins-Strebel.

Theorem 3.8.1. *1. Every pair of points in $\mathcal{T}(S)$ can be connected by a unique Teichmüller geodesic.*

2. *Quadratic differentials defined by weighted simple geodesic multicurves α, β which jointly fill up S and which satisfy $i(\alpha, \beta) = 1$ are dense in the bundle of all area one quadratic differentials over Teichmüller space.*

More precisely, a pair (λ, μ) of measured geodesic laminations which jointly fill up S in the sense of Definition 3.7.11 defines an area one singular euclidean metric on S and hence a complex structure as well as a *Teichmüller geodesic*

$$t \to \text{ the complex structure defined by } (e^t\lambda, e^{-t}\mu)$$

in Teichmüller space $\mathcal{T}(S)$. This area one singular euclidean metric is the limit of a sequence of area one singular euclidean metrics defined by weighted geodesic multicurves which approximate λ, μ. Moreover, by continuity and Lemma 3.6.1, the length of a simple closed geodesic γ for this metric is contained in the interval

$$[\max\{i(\gamma, \lambda), i(\gamma, \mu)\}, i(\gamma, \lambda) + i(\gamma, \mu)].$$

The measured geodesic lamination λ is called *vertical* and the measured geodesic lamination μ is called *horizontal*.

Recall from Section 6 the definition of the map $\Psi : \mathcal{T}(S) \to \mathcal{C}(S)$ which associates to $x \in \mathcal{T}(S)$ a simple closed x-geodesic $\Psi(x)$ whose length is at most χ_0 (where $\chi_0 > 0$ is a Bers constant for S). We have.

Lemma 3.8.2. *There exists a number $L > 0$ such that*

$$d(\Psi x, \Psi y) \leq Ld(x, y) + L \forall x, y \in \mathcal{T}(S)$$

where d is the Teichmüller distance on $\mathcal{T}(S)$.

Proof. Let $x \neq y$ be two points in Teichmüller space whose distance $d(x, y) = d$ is at most one. Then there is a Teichmüller geodesic connecting x to y of length at most one. This geodesic is given by a family q_t of quadratic differentials $0 \leq t \leq d$. The length of a simple closed curve c for the singular euclidean metric defined q_d does not exceed $2e^d$ times its length for the singular euclidean metric defined by q_0. We observed Corollary 3.5.4 that the distance in $\mathcal{C}(S)$ between $\Psi(x), \Psi(y)$ and a curve which is short for q_0, q_d is uniformly bounded. Moreover, the diameter in $\mathcal{C}(S)$ of the set of all curves of unformly bounded length for q_d is uniformly bounded as well. Thus there is a constant $L \geq 1$ not depending on x, y such that $d(\Psi(x), \Psi(y)) \leq L$. Since the Teichmüller metric is geodesic, the lemma follows. ☐

Lemma 3.8.3. *Let λ be a measured geodesic laminations whose support is minimal and fills up S and let $\gamma(t)$ be a Teichmüller geodesic with vertical measured geodesic lamination λ. Then up to passing to a subsequence, the sequence $\Psi(\gamma(i))$ converges in $\mathcal{PML}$ to a measured geodesic lamination supported in λ.*

Proof. If c is any closed curve then $i(c,\lambda) > 0$ and hence $i(c, e^t\lambda) \to \infty$ $(t \to \infty)$. This implies that $i(\gamma(t), \lambda) \to 0$ $(t \to \infty)$. □

We saw in Section 7 that the images under the map $\Psi : \mathcal{T}(S) \to \mathcal{C}(S)$ of Teichmüller geodesics are unparametrized L-quasi-geodesic for a number $L > 1$ only depending on the genus of the surface. This fact was used by Klarreich [**?**] to describe the Gromov boundary of the curve graph. We include one of her observations here which are needed later on. For this we define.

Definition 3.8.4. A sequence of simple closed geodesics $(c_i) \subset \mathcal{C}(S)$ *converges in the coarse Hausdorff topology* to a *minimal geodesic lamination* λ *which fills up* S if every accumulation point of (c_i) in $\mathcal{PML}$ is supported in λ.

Klarreich [**?**] showed.

Theorem 3.8.5. *Let $\gamma : [0,\infty) \to \mathcal{C}(S)$ be any quasi-geodesic in the curve graph. Then $\gamma(t)$ converges as $t \to \infty$ in the coarse Hausdorff topology to a minimal geodesic lamination which fills up S.*

Proof. Let $\gamma : [0,\infty) \to \mathcal{C}(S)$ be any quasi-geodesic in the curve graph and write $\alpha = \gamma(0)$. By hyperbolicity, for any two sequences i_j, ℓ_k going to infinity, the Gromov products $(\gamma(i_j)|\gamma(\ell_k))_\alpha$ tend to infinity as $j,k \to \infty$. We view each of the curves $\gamma(t)$ as a projective measured geodesic lamination. By compactness of the space $\mathcal{PML}$ of projective measured geodesic laminations, there is a subsequence $\{\gamma(i_j)\}$ which converges in $\mathcal{PML}$ to a projective measured geodesic lamination λ.

We first claim that if $\mu \in \mathcal{PML}$ is a projective measured geodesic lamination which is the limit of another subsequence $\{\gamma(\ell_k)\}$ then λ, μ do not jointly fill up S. Namely, otherwise λ, μ define a Teichmüller geodesic. Moreover, by Lemma 3.7.12, for large enough j, k the simple closed curves $\gamma(i_j), \gamma(\ell_k)$ jointly fill up S. They Teichmüller geodesics $\eta(j,k)$ defined by $\gamma(i_j), \gamma(\ell_k)$ converge uniformly on compact sets to the Teichmüller geodesic η defined by λ, μ. In particular, the Teichmüller geodesics pass through a fixed compact subset of Teichmüller space.

However, the images of the geodesics $\eta(j,k)$ under the map Ψ are uniform unparametrized quasi-geodesics in the curve graph connecting $\gamma(i_j)$ to $\gamma(\ell_k)$. Since the Teichmüller geodesics $\eta(j,k)$ pass through a fixed compact subset of $\mathcal{T}(S)$, their images under Ψ pass through a fixed bounded subset of $\mathcal{CG}(S)$. By hyperbolicity, this implies that the Gromov products $(\gamma(i_j)|\gamma(\ell_k))_\alpha$ are bounded from above by a universal constant. However, since γ is an infinite quasi-geodesic in $\mathcal{C}(S)$, these Gromov products are necessarily unbounded.

A similar arguments shows that the support of λ fills up S. Namely, assume otherwise. Let $j > 0$ be large. Then for large enough $k > j$ the distance in

$\mathcal{C}(S)$ between $\gamma(i_j)$ and $\gamma(\ell_k)$ is at least 3 and hence the simple closed curves $\gamma(i_j)$ and $\gamma(i_k)$ jointly fill up S. By the above consideration, after passing to a subsequence we may assume that these geodesics converges to a Teichmüller geodesic defined by $\gamma(i_j)$ and the measured geodesic lamination λ. Since λ does not fill up S, there is a simple closed curve β on S which does not intersect λ. As $t \to \infty$, the length of β along the geodesic defined by λ and $\gamma(i_j)$ tends to zero. Therefore by continuity, for sufficiently large k the length of β is short somewhere along the Teichmüller geodesic connecting $\gamma(i_j)$ and $\gamma(i_k)$. This means that the image under Ψ of this Teichmüller geodesic passes through a uniformly bounded neighborhood of β. However, since j was arbitrary, this violates the fact that the Gromov products $(\gamma(i_j)|\gamma(\ell_k))$ tend to infinity as $j, k \to \infty$. This shows the proposition. ☐

The *ϵ-thick part* $\mathcal{T}(S)_\epsilon$ of Teichmüller space $\mathcal{T}(S)$ is the subset of all metrics whose *systole* (length of shortest closed geodesic) is at least ϵ.

Definition 3.8.6. For $\epsilon > 0$, a *quasi-convex curve* in $\mathcal{T}(S)_\epsilon$ is a closed subset of $\mathcal{T}(S)$ whose *Hausdorff distance* to the image of a geodesic arc $\zeta : J \to \mathcal{T}(S)_\epsilon$ is at most $1/\epsilon$.

The follwoing technical result is taken from [**?**].

Theorem 3.8.7. *For every $\nu > 1$ there is a constant $\epsilon = \epsilon(\nu) > 0$ with the following properties. Let $J \subset \mathbb{R}$ be a closed connected set of diameter at least $1/\epsilon$ and let $\gamma : J \to \mathcal{T}(S)$ be a ν-quasi-geodesic. If $\Psi \circ \gamma$ is a ν-quasi-geodesic in $\mathcal{CG}(S)$ then $\gamma(J)$ is a quasi-convex curve in $\mathcal{T}(S)_\epsilon$.*

For the proof of this theorem, we need the following simple preparation.

Lemma 3.8.8. *For every $\nu > 1$ there is a number $\epsilon_0 = \epsilon_0(\nu) > 0$ with the following property. Let $\gamma : [0, n] \to \mathcal{T}(S)$ be a ν-quasi-geodesic whose projection $\Phi\gamma$ to $\mathcal{CG}(S)$ is a ν-quasi-geodesic. If $n \geq 1/\epsilon_0$ then $\gamma[0, n] \subset \mathcal{T}(S)_{\epsilon_0}$.*

Proof. For $h \in \mathcal{T}(S)$ and $\alpha \in \mathcal{C}(S)$ let again $\ell_h(\alpha)$ be the length of the unique closed geodesic for h in the free homotopy class defined by α. By a result of Wolpert [**?**], we have

$$d(h, h') \geq |\log \ell_h(\alpha) - \log \ell_{h'}(\alpha)| \forall \alpha \in \mathcal{C}(S), \forall h, h' \in \mathcal{T}(S).$$

Thus if $h \in \mathcal{T}(S) - \mathcal{T}(S)_\epsilon$ then there is a simple closed curve α with $\ell_h(\alpha) \leq \epsilon$, and $\ell_{h'}(\alpha) \leq e^{d(h,h')}\epsilon$ for all $h' \in \mathcal{T}(S)$. This shows the lemma. ☐

Proof of Theorem 3.8.7:

For $\nu > 1$ define a *ν-Lipschitz curve* in $\mathcal{T}(S)$ to be a ν-Lipschitz map $\gamma : J \to \mathcal{T}(S)$ with respect to the standard metric on $\mathbb{R}$ and the Teichmüller metric on $\mathcal{T}(S)$. Since $\mathcal{T}(S)$ is a smooth manifold and the Teichmüller metric is geodesic, every ν-quasi-geodesic $\gamma : J \to \mathcal{T}(S)$ can be replaced by a piecewise geodesic $\zeta : J \to \mathcal{T}(S)$ which is a 2ν-Lipschitz curve and which satisfies $d(\gamma(t), \zeta(t)) \leq 2\nu$ for all $t \in J$. Thus it is enough to show the statement of the lemma for ν-Lipschitz curves $\gamma : J \to \mathcal{T}(S)$ which are ν-quasi-geodesics and such that $\Psi \circ \gamma$ is a ν-quasi-geodesic in $\mathcal{C}(S)$. In the sequel we also assume that the diameter $|J|$ of the set J is bigger than $1/\epsilon_0$ where $\epsilon_0 = \epsilon_0(\nu)$ is as in Lemma 3.8.8; then $\gamma(J) \subset \mathcal{T}(S)_{\epsilon_0}$.

The curve graph $\mathcal{CG}(S)$ is hyperbolic and $\Psi \circ \gamma$ is a ν-quasi-geodesic by assumption. By Theorem 3.8.5, if J is one-sided infinite, say if $[0, \infty) \subset J$, then the points $\Psi(\gamma(t))$ converge as $t \to \infty$ in the coarse Hausdorff topology to a minimal geodesic laminations on S which fill up S.

A simple closed curve $\alpha \in \mathcal{C}(S)$ defines a projective measured lamination which we denote by $[\alpha]$. Similarly, for a measured geodesic lamination $\lambda \in \mathcal{ML}$ we denote by $[\lambda]$ the projective class of λ. Following Mosher [**?**], we say that the projective measured lamination $[\alpha]$ defined by a simple closed curve $\alpha \in \mathcal{C}(S)$ is *realized* at some $t \in J$ if the length of α with respect to the metric $\gamma(t) \in \mathcal{T}(S)$ is at most χ_0. By the considerations in Section 6, the number of projective measured laminations which are realized at a given point $t \in J$ is uniformly bounded, and $[\Psi(\gamma(t))]$ is realized at $\gamma(t)$. Similarly, we say that the projectivization $[\lambda]$ of a measured geodesic lamination λ is realized at an infinite "endpoint" of J if the support of λ equals the minimal geodesic lamination determined by the quasi-geodesic $\Psi\gamma(J)$. The set of projective measured laminations which are realized at an infinite endpoint of J is a nonempty closed subset of $\mathcal{PML}$ (see [**?**, **?**]). We call a projective measured lamination which is realized at a (finite or infinite) endpoint of J an *endpoint lamination*.

Now $\Psi\gamma$ is a ν-quasi-geodesic in $\mathcal{C}(S)$ by assumption and the diameter in $\mathcal{C}(S)$ of the set of all curves of length at most χ_0 with respect to some fixed hyperbolic metric $h \in \mathcal{T}(S)$ is bounded from above by a universal constant. Since any two curves $\alpha, \beta \in \mathcal{C}(S)$ with $d(\alpha, \beta) \geq 3$ *jointly fill up* S, i.e. are such that every simple closed essential curve $\zeta \in \mathcal{C}(S)$ intersects either α or β transversely, by possibly increasing the lower bound for the diameter of J we may assume that any two projective measured laminations $[\alpha], [\beta]$ which are realized at the two distinct endpoints of J jointly fill up S.

There is a 1-1-correspondence between measured geodesic laminations and equivalence classes of *measured foliations* on S. Via this identification, any pair of distinct points $[\lambda] \neq [\mu] \in \mathcal{PML}$ which jointly fill up the surface S define a unique Teichmüller geodesic line. Thus for every ν-quasi-geodesic $\zeta : J \to \mathcal{T}(S)$ with $|J| \geq 1/\epsilon_0$ such that $\Psi\zeta$ is a ν-quasi-geodesic in $\mathcal{C}(S)$, any pair of projective measured laminations $[\lambda], [\mu]$ realized at the two (possibly infinite) endpoints of ζ defines a unique Teichmüller geodesic $\eta([\lambda], [\mu])$.

Choose a number $R > 2\chi_0$ and a smooth function $\sigma : [0,\infty) \to [0,1]$ with $\sigma[0,\chi_0] \equiv 1$ and $\sigma[R,\infty) \equiv 0$. For each $h \in \mathcal{T}(S)$, the number of simple closed geodesics α for h with $\ell_h(\alpha) \leq R$ is bounded from above by a universal constant not depending on h, and the diameter of the subset of $\mathcal{C}(S)$ containing these curves is uniformly bounded as well. Thus we obtain for every $h \in \mathcal{T}(S)$ a finite Borel measure μ_h on $\mathcal{C}(S)$ by defining $\mu_h = \sum_\beta \sigma(\ell_h(\beta))\delta_\beta$ where δ_β denotes the Dirac mass at β. The total mass of μ_h is bounded from above and below by a universal positive constant, and the diameter of the support of μ_h in $\mathcal{C}(S)$ is uniformly bounded as well. Moreover, the measures μ_h depend continuously on $h \in \mathcal{T}(S)$ in the weak*-topology. This means that for every bounded function $f : \mathcal{C}(S) \to \mathbb{R}$ the function $h \to \int f d\mu_h$ is continuous.

We define now a new "distance" function ρ on $\mathcal{T}(S)$ by

$$\rho(h,h') = \int_{\mathcal{C}(S)\times\mathcal{C}(S)} d(\cdot,\cdot)d\mu_h \times d\mu_h'/\mu_h(\mathcal{C}(S))\mu_{h'}(\mathcal{C}(S)).$$

Clearly the function ρ is positive and continuous on $\mathcal{T}(S)\times\mathcal{T}(S)$ and invariant under the action of $\mathcal{M}(S)$. Moreover, it is immediate that there is a universal constant $a > 0$ such that $\rho(h,h')/a - a \leq d(\Psi(h),\Psi(h')) \leq a\rho(h,h') + a$. As a consequence, for every $\nu > 1$ there is a constant $p = p(\nu) > 1$ with the following property. If $\gamma : J \to \mathcal{T}(S)$ is such that $\Psi\gamma$ is a ν-quasi-geodesic, then γ is a p-quasi-geodesic with respect to the "distance" function ρ. By this we mean that

$$\rho(\gamma(s),\gamma(t))/p - p \leq |s-t| \leq p\rho(\gamma(s),\gamma(t)) + p$$

for all $s,t \in J$. Moreover, for every $p > 1$ there is a constant $\nu = \nu(p) > 1$ such that if $\gamma : J \to \mathcal{T}(S)$ is a Lipschitz curve which is a p-quasi-geodesic with respect to ρ, then $\Psi \circ \gamma$ is a ν-quasi-geodesic in $\mathcal{C}(S)$.

Let $h \in \mathcal{T}(S)$ and let $\mu \in \mathcal{ML}$ be a measured geodesic lamination. The product of the transverse measure for μ together with the length element of h defines a measure on the support of μ whose total mass is called the *h-length* of μ; we denote it by $\ell_h(\mu)$. Following Mosher [**?**], for $p > 1$ define Γ_p to be the set of all triples $(\gamma : J \to \mathcal{T}(S), \lambda_+, \lambda_-)$ with the following properties.

1. $0 \in J$ and the diameter of J is at least $1/\epsilon_0$ where $\epsilon_0 = \epsilon_0(\nu(p))$ is as in Lemma 3.8.8.

2. $\gamma : J \to \mathcal{T}(S)$ is a p-Lipschitz curve which is a p-quasi-geodesic with respect to the "distance" ρ.

3. $\lambda_+, \lambda_- \in \mathcal{ML}$ are laminations of $\gamma(0)$-length 1, and the projective measured lamination $[\lambda_+]$ is realized at the right end, the projective measured lamination $[\lambda_-]$ is realized at the left end of γ.

We equip Γ_p with the product topology, using the weak*-topology on $\mathcal{ML}$ for the second and the third component of our triple and the compact-open topology for the arc γ in $\mathcal{T}(S)$. Note that this topology is metrizable.

We follow Mosher (Proposition 3.17 of [?]) and show that the action of $\mathcal{M}(S)$ on Γ_p is cocompact. Namely, recall from Lemma 3.8.8 that there is a constant $\epsilon_0 > 0$ such that for every $(\gamma, \lambda_+, \lambda_-) \in \Gamma_p$ the image of γ is contained in $\mathcal{T}(S)_{\epsilon_0}$. Since $\mathcal{M}(S)$ acts cocompactly on $\mathcal{T}(S)_{\epsilon_0}$ it is therefore enough to show that the subset of Γ_p consisting of triples with the additional property that $\gamma(0)$ is contained in a fixed compact subset A of $\mathcal{T}(S)_{\epsilon_0}$ is compact. Since our topology is metrizable, this follows if every sequence of points $(\gamma, \lambda_+, \lambda_-)$ with $\gamma(0) \in A$ has a convergent subsequence.

However, by the Arzela-Ascoli theorem, the set of p-Lipschitz maps into $\mathcal{T}(S)_{\epsilon_0}$ issuing from a point in A is compact. Moreover, the function ρ on $\mathcal{T}(S) \times \mathcal{T}(S)$ is continuous and invariant under the action of $\mathcal{M}(S)$ and hence if γ_i converges locally uniformly to γ and if γ_i is a p-quasi-geodesic with respect to ρ for all i then the same is true for γ. Since the function on $\mathcal{T}(S) \times \mathcal{ML}$ which assigns to a metric $h \in \mathcal{T}(S)$ and a measured lamination $\mu \in \mathcal{ML}$ the h-length of μ is continuous and since for every fixed $h \in \mathcal{T}(S)$ the set of measured laminations of h-length 1 is compact and naturally homeomorphic to $\mathcal{PML}$, the action of $\mathcal{M}(S)$ on Γ_p is indeed cocompact provided that the following holds: If $(\gamma_i : J_i \to \mathcal{T}(S)_{\epsilon_0})$ is a sequence of p-Lipschitz curves which converge locally uniformly to $\gamma : J \to \mathcal{T}(S)_{\epsilon_0}$, if the projective measured lamination $[\lambda_i]$ is realized at the right endpoint of J_i and if $[\lambda_i] \to [\lambda]$ in $\mathcal{PML}$ $(i \to \infty)$ then $[\lambda]$ is realized at the right endpoint of J.

To see that this is indeed the case, assume first that $J \cap [0, \infty) = [0, b]$ for some $b \in (0, \infty)$. Then for sufficiently large i we have $J_i \cap [0, \infty) = [0, b_i]$ with $b_i \in (0, \infty)$ and $b_i \to b$. Thus $\gamma_i(b_i) \to \gamma(b)$ $(i \to \infty)$ and therefore for sufficiently large i there is only a *finite* number of curves $\alpha \in \mathcal{C}(S)$ whose length with respect to one of the metrics $\gamma_j(b_j), \gamma(b)$ $(j \geq i)$ is at most χ. By passing to a subsequence we may assume that there is a simple closed curve $\alpha \in \mathcal{C}(S)$ with $[\lambda_j] = [\alpha]$ for all large j. The $\gamma_j(b_j)$-length of α is at most χ for all sufficiently large j and hence the same is true for the $\gamma(b)$-length of α by continuity of the length function. As a consequence, the limit $[\lambda] = [\alpha]$ of the sequence $([\lambda_i])$ is realized at the endpoint $\gamma(b)$ of γ.

In the case that $[0, \infty) \subset J$ we argue as before. Assume first that $b_i < \infty$ for all i and that $b_i \to \infty$. Recall that each of the curves $\Psi\gamma_i$ is a uniform quasi-geodesic in $\mathcal{C}(S)$ and that the map Ψ is coarsely Lipschitz. Let $\alpha_i \in \mathcal{C}(S)$ be the simple closed curve such that $[\alpha_i] = [\lambda_i]$. Then for each i, the curve α_i is contained in a ball about $\Psi(\gamma_i(b_i))$ of radius $R > 0$ independent of i and hence as $i \to \infty$, the curves α_i converge to μ in the coarse Hausdorff topology. Since μ is a minimal geodesic lamination which fills up S, the complement of μ in every lamination ζ containing μ as a sublamination consists of a finite number of isolated leaves and therefore every transverse measure supported in ζ is in fact supported in μ. Thus after passing to a subsequence, the projective measured laminations $[\lambda_i]$ converge as $i \to \infty$ to a projective measured lamination supported in μ. But $[\lambda_i] \to [\lambda]$ in $\mathcal{PML}$ by assumption and hence the lamination $[\lambda]$ is realized at the endpoint of γ. Similarly, if $b_i = \infty$ for infinitely many i then the endpoints $\beta_i \in \partial\mathcal{C}(S)$ of the quasi-geodesics $\Psi\gamma_i$ converge in $\partial\mathcal{C}(S)$ to the endpoint of $\Psi\gamma$. As before, this implies that the limit λ of the projective measured lamination λ_i is realized at the endpoint of γ. This shows our above claim and implies that the action of $\mathrm{Mod}(S)$ on Γ_p is indeed cocompact.

Now we follow Section 3.10 of [**?**]. Namely, each point $(\gamma, \lambda_+, \lambda_-) \in \Gamma_p$ determines the geodesic $\eta([\lambda_+], [\lambda_-])$ in $\mathcal{T}(S)$. This geodesic defines a family q_t of quadratic differentials whose horizontal foliation corresponds to the lamination $e^{-t}\lambda_+$ and whose vertical foliation corresponds to $e^t\lambda_-$ (note that the area of these differentials may not be one, however this is of importance for our argument, compare [**?**]).

For $(\gamma, \lambda_+, \lambda_-)$ let $\sigma(\gamma, \lambda_+, \lambda_-)$ be the point on the geodesic $\eta([\lambda_+], [\lambda_-])$ which corresponds to the quadratic differential defined by the measured geodesic laminations λ_+, λ_-. The map taking $(\gamma, \lambda_+, \lambda_-)$ to $(\gamma(0), \sigma(\gamma, \lambda_+, \lambda_-)) \in \mathcal{T}_{g,m} \times \mathcal{T}(S)$ is continuous and equivariant with respect to the natural action of $\mathrm{Mod}(S)$ on Γ_p and on $\mathcal{T}(S) \times \mathcal{T}(S)$. Since the action of $\mathrm{Mod}(S)$ on Γ_p is cocompact, the same is true for the action of $\mathrm{Mod}(S)$ on the image of our map (see [**?**]). Thus the distance between $\gamma(0)$ and $\sigma(\gamma, \lambda_+, \lambda_-)$ is bounded from above by a universal constant $p > 0$.

Let again $(\gamma, \lambda_+, \lambda_-) \in \Gamma_p$. For each $s \in J$ define

$$a_-(s) = \frac{1}{\ell_{\gamma(s)}(\lambda_-)}, \quad a_+(s) = \frac{1}{\ell_{\gamma(s)}(\lambda_+)}$$

where as before, $\ell_{\gamma(s)}(\lambda_\pm)$ is the $\gamma(s)$-length of $\lambda_\pm$. These are continuous functions of $s \in J$. Define for $s \in \mathbb{R}$ the shift $\gamma'(t) = \gamma(t+s)$; then the ordered triple $(\gamma'(0), a_+(s)\lambda_+, a_-(s)\lambda_-)$ lies in the $\mathrm{Mod}(S)$-cocompact set Γ_p and hence the distance between $\gamma(s)$ and a suitably chosen point on the geodesic $\eta([\lambda_+], [\lambda_-])$ is at most p. As a consequence, the arc γ is contained in the p-neighborhood of the geodesic $\eta([\lambda_+], [\lambda_-])$. Since the curve γ is a p-quasi-geodesic, this implies that the Hausdorff distance between $\gamma(J)$ and a subarc of $\eta([\lambda_+], [\lambda_-])$ is uniformly bounded and shows the theorem. $\square$