

Coarse and fine geometry of the Thurston metric

David Dumas, Anna Lenzhen, Kasra Rafi, and Jing Tao




Figure 0: *In-envelopes in Teichmüller space; see Remark 5.5.*

1 Introduction

In this paper we study the geometry of the Thurston metric on the Teichmüller space $\mathcal{T}(S)$ of hyperbolic structures on a surface S . Some of our results on the coarse geometry of this metric apply to arbitrary surfaces S of finite type; however, we focus particular attention on the case where the surface is a once-punctured torus, $S = S_{1,1}$. In that case, our results provide a detailed picture of the infinitesimal, local, and global behavior of the geodesics of the Thurston metric on $\mathcal{T}(S_{1,1})$, as well as an analogue of Royden's theorem (cf. [Roy71]).

Thurston's metric

Recall that Thurston's metric $d_{\text{Th}} : \mathcal{T}(S) \times \mathcal{T}(S) \rightarrow \mathbf{R}$ is defined by

$$d_{\text{Th}}(X, Y) = \sup_{\alpha} \log \left(\frac{\ell_{\alpha}(Y)}{\ell_{\alpha}(X)} \right) \quad (1)$$

where the supremum is over all simple closed curves α in S and $\ell_\alpha(X)$ denotes the hyperbolic length of the curve α in X . This function defines a forward-complete asymmetric Finsler metric, introduced by Thurston in [Thu86c]. In the same paper, Thurston introduced two key tools for understanding this metric which will be essential in what follows: *stretch paths* and *maximally-stretched laminations*.

The maximally stretched lamination $\Lambda(X, Y)$ is a chain-recurrent geodesic lamination which is defined for any pair of distinct points $X, Y \in \mathcal{T}(S)$. Typically $\Lambda(X, Y)$ is just a simple curve, in which case that curve uniquely realizes the supremum defining d_{Th} . In general $\Lambda(X, Y)$ can be a more complicated lamination that is constructed from limits of sequences of curves that asymptotically realize the supremum. The precise definition is given in Section 2.6 (or [Thu86c, Section 8], where the lamination is denoted $\mu(X, Y)$).

Stretch paths are geodesics constructed from certain decompositions of the surface into ideal triangles. More precisely, given a hyperbolic structure $X \in \mathcal{T}(S)$ and a complete geodesic lamination λ one obtains a parameterized stretch path, $\text{stretch}(X, \lambda, \cdot) : \mathbf{R} \rightarrow \mathcal{T}(S)$, with $\text{stretch}(X, \lambda, 0) = X$ and which satisfies

$$d_{\text{Th}}(\text{stretch}(X, \lambda, s), \text{stretch}(X, \lambda, t)) = t - s \quad (2)$$

for all $s, t \in \mathbf{R}$ with $s < t$.

Thurston showed that there also exist geodesics in $\mathcal{T}(S)$ that are concatenations of segments of stretch paths along different geodesic laminations. The abundance of such “chains” of stretch paths is sufficient to show that d_{Th} is a geodesic metric space, and also that it is not uniquely geodesic—some pairs of points are joined by more than one geodesic segment.

Envelopes

The first problem we consider is to quantify the failure of uniqueness for geodesic segments with given start and end points. For this purpose we consider the set $E(X, Y) \subset \mathcal{T}(S)$ that is the union of all geodesics from X to Y . We call this the *envelope* (from X to Y).

Based on Thurston’s construction of geodesics from chains of stretch paths, it is natural to expect that the envelope would admit a description in terms of the maximally-stretched lamination $\Lambda(X, Y)$ and its completions. We focus on the punctured torus case, because here the set of completions is always finite.

In fact, a chain-recurrent lamination on $S_{1,1}$ (such as $\Lambda(X, Y)$, for any $X \neq Y \in \mathcal{T}(S_{1,1})$) is either

- (a) A simple closed curve,
- (b) The union of a simple closed curve and a spiral geodesic, or
- (c) A measured lamination with no closed leaves

These possibilities are depicted in Figure 1. See [BZ04] for more details.

We show that the geodesic from X to Y is unique when $\Lambda(X, Y)$ is of type (b) or (c), and when it has type (a) the envelope has a simple, explicit description. More precisely, we have:




Figure 1: *The three types of chain-recurrent laminations on $S_{1,1}$.*

Theorem 1.1 (Structure of envelopes for the punctured torus).

- (i) For any $X, Y \in \mathcal{T}(S_{1,1})$, the envelope $E(X, Y)$ is a compact set.
- (ii) $E(X, Y)$ varies continuously in the Hausdorff topology as a function of X and Y .
- (iii) If $\Lambda(X, Y)$ is not a simple closed curve, then $E(X, Y)$ is a segment on a stretch path (which is then the unique geodesic from X to Y).
- (iv) If $\Lambda(X, Y) = \alpha$ is a simple closed curve, then $E(X, Y)$ is a geodesic quadrilateral with X and Y as opposite vertices. Each edge of the quadrilateral is a stretch path along a completion of a chain-recurrent geodesic lamination properly containing α .

In the course of proving the theorem above, we write explicit equations for the edges of the quadrilateral-type envelopes in terms of Fenchel-Nielsen coordinates (see (21)–(22)). Also note that in part (iv) of the theorem, a chain-recurrent lamination properly containing α has multiple completions, but they all give the same stretch path (see Corollary 2.3).

This theorem also highlights a distinction between two cases in which the d_{Th} -geodesic from X to Y is unique—the cases (b) and (c) discussed above. In case (b) the geodesic to Y is unique but some initial segment of it can be chained with another stretch path and remain geodesic: The boundary of a quadrilateral-type envelope from X with maximally-stretched lamination α furnishes an example of this. In case (c), however, a geodesic that starts along the stretch path from X to Y is entirely contained in that stretch path (see Proposition 5.2).

Figure 0 can also be seen as an illustration of this theorem: It shows regions in $\mathcal{T}(S_{1,1})$ bounded by pairs of stretch rays from rational points on the circle at infinity to the hexagonal punctured torus. Such *in-envelopes* are limiting cases of the envelopes of type (iv) where X

is replaced by a lamination. These are defined precisely and studied in Section 5. Figure 0 is discussed in more detail in Remark 5.5.

Short curves

Returning to the case of an arbitrary surface S of finite type, in section 3 we establish results on the coarse geometry of Thurston metric geodesic segments. This study is similar in spirit to the one of Teichmüller geodesics in [Raf05], in that we seek to determine whether or not a simple curve α becomes short along a geodesic from X to Y . As in that case, a key quantity to consider is the amount of twisting along α from X to Y , denoted $d_\alpha(X, Y)$ and defined in Section 2.8.

For curves that *interact* with the maximally-stretched lamination $\Lambda(X, Y)$, meaning they belong to the lamination or intersect it essentially, we show that becoming short on a geodesic with endpoints in the thick part of $\mathcal{T}(S)$ is equivalent to the presence of large twisting:

Theorem 1.2. *There exists a constant ϵ_0 such that the following statement holds. Let X, Y lie in the ϵ_0 -thick part of $\mathcal{T}(S)$ and let α be a simple curve on S that interacts with $\Lambda(X, Y)$. Then the minimum length ℓ_α of α along any Thurston metric geodesic from X to Y satisfies*

$$\frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha} \stackrel{*}{\underset{+}{\approx}} d_\alpha(X, Y)$$

with implicit constants that depend only ϵ_0 , and where $\text{Log}(x) = \max(1, \log(x))$.

Here $\stackrel{*}{\underset{+}{\approx}}$ means equality up to an additive and multiplicative constant; see Section 2.1. The theorem above and additional results concerning length functions along geodesic segments are combined in Theorem 3.1.

In Section 4 we specialize once again to the Teichmüller space of the punctured torus in order to say more about the coarse geometry of Thurston geodesics. Here every simple curve interacts with every lamination, so Theorem 1.2 is a complete characterization of short curves in this case. Furthermore, in this case we can determine the order in which the curves become short.

To state the result, we recall that the pair of points $X, Y \in \mathcal{T}(S_{1,1})$ determine a geodesic in the dual tree of the Farey tessellation of $\mathbf{H}^2 \simeq \mathcal{T}(S_{1,1})$. Furthermore, this path distinguishes an ordered sequence of simple curves—the *pivots*—and each pivot has an associated *coefficient*. These notions are discussed further in Section 4.

We show that pivots for X, Y and short curves on a d_{Th} -geodesic from X to Y coarsely coincide in an order-preserving way, once again assuming that X and Y are thick:

Theorem 1.3. *Let $X, Y \in \mathcal{T}(S_{1,1})$ lie in the thick part, and let $\mathcal{G}: I \rightarrow \mathcal{T}(S_{1,1})$ be a geodesic of d_{Th} from X to Y . Let ℓ_α denote the minimum of $\ell_\alpha(\mathcal{G}(t))$ for $t \in I$. We have:*

- (i) *If α is short somewhere in \mathcal{G} , then α is a pivot.*
- (ii) *If α is a pivot with large coefficient, then α becomes short somewhere in \mathcal{G} .*

(iii) If both α and β become short in \mathcal{G} , then they do so in disjoint intervals whose ordering in I agrees with that of α, β in $\text{Pivot}(X, Y)$.

(iv) There is an a priori upper bound on ℓ_α for $\alpha \in \text{Pivot}(X, Y)$.

In this statement, various constants have been suppressed (such as those required to make *short* and *large* precise). We show that all of the constants can be taken to be independent of X and Y , and the full statement with these constants is given as Theorem 4.3 below.

We have already seen that there may be many Thurston geodesics from X to Y , and due to the asymmetry of the metric, reversing parameterization of a geodesic from X to Y does not give a geodesic from Y to X . On the other hand, the notion of a pivot is symmetric in X and Y . Therefore, by comparing the pivots to the short curves of an arbitrary Thurston geodesic, Theorem 4.3 establishes a kind of symmetry and uniqueness for the *combinatorics* of Thurston geodesic segments, despite the failure of symmetry or uniqueness for the geodesics themselves.

Rigidity

A Finsler metric on $\mathcal{T}(S)$ gives each tangent space $T_X\mathcal{T}(S)$ the structure of a normed vector space. Royden showed that for the Teichmüller metric, this normed vector space uniquely determines X up to the action of the mapping class group [Roy71]. That is, the tangent spaces are isometric (by a linear map) if and only if the hyperbolic surfaces are isometric.

We establish the corresponding result for the Thurston's metric on $\mathcal{T}(S_{1,1})$ and its corresponding norm $\|\cdot\|_{\text{Th}}$ (the *Thurston norm*) on the tangent bundle.

Theorem 1.4. *Let $X, Y \in \mathcal{T}(S_{1,1})$. Then there exists an isometry of normed vector spaces*

$$(T_X\mathcal{T}(S_{1,1}), \|\cdot\|_{\text{Th}}) \rightarrow (T_Y\mathcal{T}(S_{1,1}), \|\cdot\|_{\text{Th}})$$

if and only if X and Y are in the same orbit of the extended mapping class group.

The idea of the proof is to recognize lengths and intersection numbers of curves on X from features of the unit sphere in $T_X\mathcal{T}(S)$. Analogous estimates for the shape of the cone of lengthening deformations of a hyperbolic one-holed torus were established in [Gué15]. In fact, Theorem 1.4 was known to Guéritaud and can be derived from those estimates [Gué16]. We present a self-contained argument that does not use Guéritaud's results directly, though [Gué15, Section 5.1] provided inspiration for our approach to the infinitesimal rigidity statement.

A local rigidity theorem can be deduced from the infinitesimal one, much as Royden did in [Roy71].

Theorem 1.5. *Let U be a connected open set in $\mathcal{T}(S_{1,1})$, considered as a metric space with the restriction of d_{Th} . Then any isometric embedding $(U, d_{\text{Th}}) \rightarrow (\mathcal{T}(S_{1,1}), d_{\text{Th}})$ is the restriction to U of an element of the extended mapping class group.*

Intuitively, this says that the quotient of $\mathcal{T}(S_{1,1})$ by the mapping class group is “totally unsymmetric”; each ball fits into the space isometrically in only one place. Of course, applying Theorem 1.5 to $U = \mathcal{T}(S_{1,1})$ we have the immediate corollary

Corollary 1.6. *Every isometry of $(\mathcal{T}(S_{1,1}), d_{\text{Th}})$ is induced by an element of the extended mapping class group, hence the isometry group is isomorphic to $\text{PGL}(2, \mathbf{Z})$.*

Here we have used the usual identification of the mapping class group of $S_{1,1}$ with $\text{GL}(2, \mathbf{Z})$, whose action on $\mathcal{T}(S_{1,1})$ factors through the quotient $\text{PGL}(2, \mathbf{Z})$.

The analogue of Corollary 1.6 for Thurston’s metric on higher-dimensional Teichmüller spaces was established by Walsh in [Wal14] using a characterization of the horofunction compactification of $\mathcal{T}(S)$. Walsh’s argument does not apply to the punctured torus, however, because it relies on Ivanov’s characterization (in [Iva97]) of the automorphism group of the curve complex (a result which does not hold for the punctured torus).

Passing from the infinitesimal (i.e. norm) rigidity to local or global statements requires some preliminary study of the smoothness of the Thurston norm. In Section 6.1 we show that the norm is locally Lipschitz continuous on $T\mathcal{T}(S)$ for any finite type hyperbolic surface S . By a recent result of Matveev-Troyanov [MT17], it follows that any d_{Th} -preserving map is differentiable with norm-preserving derivative. This enables the key step in the proof of Theorem 1.5, where Theorem 1.4 is applied to the derivative of the isometry.

Additional notes and references

In addition to Thurston’s paper [Thu86c], an exposition of Thurston’s metric and a survey of its properties can be found in [PT07]. Prior work on the coarse geometry of the Thurston metric on Teichmüller space and its geodesics can be found in [CR07] [LRT12] [LRT15]. The notion of the maximally-stretched lamination for a pair of hyperbolic surfaces has been generalized to higher-dimensional hyperbolic manifolds [Kas09] [GK17] and to vector fields on \mathbf{H}^2 equivariant for convex cocompact subgroups of $\text{PSL}(2, \mathbf{R})$ [DGK16].

Acknowledgments

The authors thank the American Institute of Mathematics for hosting the workshop “Lipschitz metric on Teichmüller space” and the Mathematical Sciences Research Institute for hosting the semester program “Dynamics on Moduli Spaces of Geometric Structures” where some of the work reported here was completed. The authors gratefully acknowledge grant support from NSF DMS 0952869 and DMS 1709877 (DD), NSF DMS 1611758 (JT), NSERC RGPIN 435885 (KR), and from NSF DMS 1107452, 1107263, 1107367 “RNMS: GEometric structures And Representation varieties” (the GEAR Network). The authors also thank François Guéritaud for helpful conversations related to this work, and specifically for suggesting the statement of Theorem 6.8. Finally, the authors thank the anonymous referees for their careful reading of the paper and for helpful comments and corrections.

2 Background

2.1 Approximate comparisons

We use the notation $a \stackrel{*}{\approx} b$ to mean that quantities a and b are equal up to a uniform multiplicative error, i.e. that there exists a positive constant K such that $K^{-1}a \leq b \leq Ka$. Thus

for example $a \overset{*}{\asymp} 1$ means that a is bounded above and below by positive constants. Similarly, the notation $a \overset{*}{\asymp} b$ means that $a \leq Kb$ for some K .

The analogous relations up to additive error are $a \overset{\pm}{\asymp} b$, meaning that there exists C such that $a - C \leq b \leq a + C$, and $a \overset{\pm}{\preceq} b$ which means $a \leq b + C$ for some C . Hence $a \overset{\pm}{\asymp} 0$ means that a is bounded above and below by constants.

For equality up to both multiplicative and additive error, we write $a \overset{*}{\overset{\pm}{\asymp}} b$. That is, $a \overset{*}{\overset{\pm}{\asymp}} b$ means that there exist constants K, C such that $K^{-1}a - C \leq b \leq Ka + C$.

Unless otherwise specified, the implicit constants depend only on the topological type of the surface S . When the constants depend on the Riemann surface X , we use the notation $\overset{*}{\asymp}_X$ and $\overset{\pm}{\preceq}_X$ instead.

For functions f, g of a real variable x we write $f \sim g$ to mean that $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$.

2.2 Surfaces, curves, and laminations

Throughout this paper S denotes an oriented surface of finite type, i.e. the complement of a finite subset P of the interior of \bar{S} , a compact oriented surface with boundary. Elements of P are the *punctures*.

A *multicurve* is a closed 1-manifold on S defined up to homotopy such that no connected component is homotopic to a point, a puncture, or boundary of S . A connected multicurve will just be called a *curve*. Note that with our definition, there are no curves on the two- or three-punctured sphere, so we will ignore those cases henceforth. The *geometric intersection number* $i(\alpha, \beta)$ between two curves is the minimal number of intersections between representatives of α and β . If we fix a hyperbolic metric on S , then every (multi)curve has a unique geodesic representative, and $i(\alpha, \beta)$ is just the number of intersections between the geodesic representative of α and the geodesic representative of β . For any curve α on S , we denote by D_α the *left* Dehn twist about α .

Fix a complete hyperbolic metric of finite area on S , so that the boundary components (if any) are geodesic. A *geodesic lamination* λ on S is a closed subset which is a disjoint union of simple complete geodesics. These geodesics are called the *leaves* of λ . Two different hyperbolic metrics on S determine canonically isomorphic spaces of geodesic laminations, so the space of geodesic laminations $\mathcal{GL}(S)$ depends only on the topology of S . This is a compact metric space equipped with the metric of Hausdorff distance on closed sets. The closure of the set of multicurves in $\mathcal{GL}(S)$ is the set of *chain-recurrent* laminations.

We will call a geodesic lamination *maximal chain-recurrent* if it is chain-recurrent and not properly contained in another chain-recurrent lamination. A geodesic lamination is *complete* if its complementary regions in S are ideal triangles. Note that all chain-recurrent laminations are necessarily compactly supported. Thus, when S has punctures, a chain-recurrent lamination can never be complete. For a given geodesic lamination λ , we refer to any complete lamination containing λ as a *completion* (of λ).

In the case of the punctured torus $S_{1,1}$, the maximal chain-recurrent laminations are types (b) and (c) in Figure 1. Case (b), i.e. a curve and a spiraling geodesic, will be especially important in the sequel, and so we introduce the following notation for these laminations: Given a curve

α , let $\alpha_0^+ = \alpha \cup \delta$ where the geodesic δ spirals toward α in each direction, turning to the *left* as it does so. Similarly we define α_0^- to be the union of α and a spiraling leaf that turns right. (Adding a leaf that turns opposite ways on its two ends yields a non-chain-recurrent lamination.)

The motivation for this sign convention for α_0^\pm is that it is compatible with a common way to describe simple curves on $S_{1,1}$ in terms of *slope* while regarding α as vertical. More precisely, consider an oriented curve $\vec{\eta}$ with $i(\eta, \alpha) = 1$, and let $\vec{\alpha}$ denote the orientation of α so that the homology classes $[\vec{\eta}], [\vec{\alpha}]$ give a positive ordered basis of $H_1(S_{1,1})$ with respect to the orientation of $S_{1,1}$. If a simple curve $\gamma \neq \alpha$ has homology class $q[\vec{\eta}] + p[\vec{\alpha}]$ for some orientation, then $p/q \in \mathbf{Q}$ is the slope of γ (relative to that basis). We consider α itself to have slope $1/0 = \infty \in \mathbf{QP}^1$ and this exhibits a bijection between \mathbf{QP}^1 and the set of simple curves on $S_{1,1}$.

Now, a sequence of simple curves distinct from α whose slopes go to $+\infty$ have Hausdorff limit α_0^+ , while a sequence with slopes going to $-\infty$ has Hausdorff limit α_0^- . Thus α_0^+ (resp. α_0^-) is approximated by curves of large positive (resp. negative) slope.

All of the maximal chain-recurrent laminations on $S_{1,1}$ have a single complementary region, which is a punctured bigon. Such a lamination therefore has exactly three completions, corresponding to the three ways to add leaves that cut the bigon into ideal triangles shown in Figure 2. (For more detail on classifying laminations on the punctured torus, we refer the reader to [BZ04].)

A convenient way to distinguish among the completions of a maximal chain-recurrent lamination λ on the punctured torus is to use the hyperelliptic involution. This is an involutive orientation-preserving isometry ι that preserves every simple closed geodesic, and thus every chain-recurrent lamination. The action of ι on the complementary bigon of a maximal chain-recurrent lamination exchanges the two spikes, and therefore the only completion which is ι -invariant is the one with leaves going to both spikes, i.e. type (i) in Figure 2. We call this the *canonical completion* of λ .

We denote the canonical completion of α_0^+ by α^+ , and that of α_0^- by α^- . Thus $\alpha^\pm = \alpha_0^\pm \cup w \cup w'$ where w and w' are leaves emanating from the puncture and spiraling into α . For example, α^+ is shown in Figure 3.

The *stump* of a geodesic lamination (in the terminology of [Thé07]) is its maximal compactly-supported sublamination that admits a transverse measure of full support.

2.3 Teichmüller space

Let $\mathcal{T}(S)$ be the Teichmüller space of complete finite-area hyperbolic structures on S . We will only consider $\mathcal{T}(S)$ in cases where S has no boundary. The space $\mathcal{T}(S)$ is homeomorphic to $\mathbf{R}^{6g-6+2n}$, if S has genus g and n punctures. Given $X \in \mathcal{T}(S)$ and a curve α on S , we denote by $\ell_\alpha(X)$ the length of the geodesic representative of α on X . For brevity we refer to $\ell_\alpha(X)$ as the *length of α on X* .

For any $\epsilon > 0$, we will denote by $\mathcal{T}_\epsilon(S)$ the set of points in $\mathcal{T}(S)$ on which every curve has length at least ϵ ; this is the ϵ -*thick part* of Teichmüller space.




Figure 2: The three ways to complete a maximal chain-recurrent lamination on $S_{1,1}$ by adding two leaves in its complementary bigon.

A positive real number ϵ is called a (two-dimensional) *Margulis number* if two distinct curves on a hyperbolic surface of length less than ϵ are necessarily disjoint. Fix a Margulis number $\epsilon_M < 1$ such that that for any curve α of length less than ϵ_M , the shortest curve β that intersects α has $i(\alpha, \beta) \leq 2$. It follows from the collar lemma that any sufficiently small ϵ_M has this property.

2.4 Shearing of ideal triangles

Let \mathbf{H} denote the upper half plane model of the hyperbolic plane, with ideal boundary $\partial\mathbf{H} = \mathbf{R} \cup \{\infty\}$. In this section, we will define the shearing of two ideal triangles in \mathbf{H} which share an ideal vertex. This is a specific case of the more general shearing defined in [Bon96, Section 2].

Two distinct points $x, y \in \partial\mathbf{H}$ determine a geodesic $[x, y]$ and three distinct points $x, y, z \in \partial\mathbf{H}$ determine an ideal triangle $\Delta(x, y, z)$. Recall that an ideal triangle in \mathbf{H} has a unique inscribed circle which is tangent to all three sides of the triangle. Each tangency point is called the *midpoint* of the side.

Let $\gamma = [\gamma^+, \gamma^-]$ be a geodesic in \mathbf{H} . Suppose two ideal triangles Δ and Δ' lie on different sides of γ . We allow the possibility that γ is an edge of Δ or Δ' (or both). Suppose Δ is asymptotic to γ^+ and the Δ' is asymptotic to γ^- . Let m be the midpoint along the side of Δ closest to γ . The pair γ^+ and m determine a horocycle that intersects γ at a point p . Let m' and p' be defined similarly using Δ' and γ^- . We say p' is to the left of p (relative to Δ and Δ') if the path along the horocycle from m to p and along γ from p and p' turns left; p' is to the right of p otherwise. Note that p' is to the left of p if and only if p is to the left of p' . The shearing $s_\gamma(\Delta, \Delta')$ along γ relative to the two triangles is the signed distance between p and




Figure 3: Leaves of α^+ (the canonical completion of α_0^+) shown in the torus cut open along α .

p' , where the sign is positive if p' is to the left of p and negative otherwise. Note that this sign convention gives $s_\gamma(\Delta, \Delta') = s_\gamma(\Delta', \Delta)$.

2.5 Shearing coordinates in Teichmüller space

Given any complete geodesic lamination λ , there is an embedding $s_\lambda: \mathcal{T}(S) \rightarrow \mathbf{R}^N$ by the *shearing coordinates* relative to λ , where $N = \dim \mathcal{T}(S)$. The image of this embedding is an open convex cone. Details of the construction of this embedding can be found in [Bon96] and [Thu86c, Section 9].

Using the shearing of ideal triangles discussed above, we will define the shearing coordinates in the case where λ is the canonical completion of a maximal chain-recurrent lamination on $S_{1,1}$ with finitely many leaves. That is, we consider $\lambda = \alpha^+$ or $\lambda = \alpha^-$ for a simple curve α , and describe the map $s_\lambda: \mathcal{T}(S_{1,1}) \rightarrow \mathbf{R}^2$.

We begin with an auxiliary map $s_\lambda^0: \mathcal{T}(S_{1,1}) \rightarrow \mathbf{R}^4$ which records a shearing parameter for each leaf of λ , and then we identify the 2-dimensional subspace of \mathbf{R}^4 that contains the image in this specific situation.

Let l be a leaf of λ and fix a lift \tilde{l} of l to $\tilde{X} = \mathbf{H}$. If l is a non-compact leaf, then l bounds two ideal triangles in X , which admit lifts Δ and Δ' with common side \tilde{l} . If $l = \alpha$ is the compact leaf, then we choose Δ and Δ' to be lifts of the two ideal triangles complementary to λ that lie on different sides of \tilde{l} and which are each asymptotic to one of the ideal points of \tilde{l} . Now define $s_l(X) = s_\gamma(\Delta, \Delta')$, and let the $s_\lambda^0: \mathcal{T}(S_{1,1}) \rightarrow \mathbf{R}^4$ be the map defined by

$$s_\lambda^0(X) = (s_\delta(X), s_\alpha(X), s_w(X), s_{w'}(X)).$$

We claim that in fact, $s_w(X) = s_{w'}(X) = 0$ and that $s_\delta(X) = \mp \ell_\alpha(X)$ for $\lambda = \alpha^\pm$. It will then follow that s_λ^0 takes values in a 2-dimensional linear subspace of \mathbf{R}^4 , allowing us to equivalently consider the embedding $s_\lambda: \mathcal{T}(S_{1,1}) \rightarrow \mathbf{R}^2$ defined by

$$s_\lambda(X) = (\ell_\alpha(X), s_\alpha(X)).$$

To establish the claim, cut the surface X open along α to obtain a pair of pants which is further decomposed by w, w', δ into a pair of ideal triangles. The boundary lengths of this

hyperbolic pair of pants are ℓ_α , ℓ_α , and 0. Gluing a pair of ideal triangles along their edges but with their edge midpoints shifted by signed distances a, b, c gives a pair of pants with boundary lengths $|a + b|, |b + c|, |a + c|$, and with the signs of $a + b, b + c, a + c$ determining the direction in which the seams spiral toward those boundary components (this is discussed in more detail in [Thu86a, Section 3.9]). Specifically, a positive sum corresponds to the seam turning to the right while approaching the corresponding boundary geodesic, and a negative sum corresponds to the seam turning to the left. Applying this to our situation, and recalling that for $\lambda = \alpha^+$ all spiraling leaves turn left when approaching the boundary of the pair of pants, and we obtain

$$s_w(X) + s_\delta(X) = s_{w'}(X) + s_\delta(X) = -\ell_\alpha$$

and

$$s_w(X) + s_{w'}(X) = 0.$$

This gives $s_w(X) = s_{w'}(X) = 0$ and $s_\delta(X) = -\ell_\alpha(X)$. For the case $\lambda = \alpha^-$ the equations are the same except that $-\ell_\alpha$ is replaced by ℓ_α , and the solution becomes $s_w(X) = s_{w'}(X) = 0$ and $s_\delta(X) = \ell_\alpha$.

Finally, we consider the effect of the various choices made in the construction of $s_\lambda(X)$. The coordinate $\ell_\alpha(X)$ is of course canonically associated to X , and independent of any choices. For $s_\alpha(X)$, however, we had to choose a pair of triangles Δ, Δ' on either side of the lift $\tilde{\alpha}$. In this case, different choices differ by finitely many moves in which one of the triangles is replaced by a neighbor on the other side of a lift of w, w' , or δ . Each such move changes the value of $s_\alpha(X)$ by adding or subtracting one of the values $s_w(X), s_{w'}(X)$, or $s_\delta(X)$; this is the additivity of the shearing cocycle established in [Bon96, Section 2]. By the computation above each of these moves actually adds 0 or $\pm\ell_\alpha(X)$. Hence $s_\alpha(X)$ is uniquely determined up to addition of an integer multiple of $\ell_\alpha(X)$.

2.6 The Thurston metric

For a pair of points $X, Y \in \mathcal{T}(S)$, in the introduction we defined the quantity

$$d_{\text{Th}}(X, Y) = \sup_{\alpha} \log \frac{\ell_\alpha(Y)}{\ell_\alpha(X)}$$

where the supremum is taken over all simple curves. Another measure of the difference of hyperbolic structures, in some ways dual to this length ratio, is

$$L(X, Y) = \inf_f \log L_f$$

where L_f is the Lipschitz constant, and where the infimum is taken over Lipschitz maps $f : X \rightarrow Y$ in the preferred homotopy class. Thurston showed:

Theorem 2.1. *For all $X, Y \in \mathcal{T}(S)$ we have $d_{\text{Th}}(X, Y) = L(X, Y)$, and this function is an asymmetric metric, i.e. it is positive unless $X = Y$ and it obeys the triangle inequality.*

Denote by $\bar{d}_{\text{Th}}(X, Y) = \max\{d_{\text{Th}}(X, Y), d_{\text{Th}}(Y, X)\}$. The topology of $\mathcal{T}(S)$ is compatible with \bar{d}_{Th} , so by $X_i \rightarrow X$ we will mean $\bar{d}_{\text{Th}}(X_i, X) \rightarrow 0$. By the Hausdorff distance on closed sets in $\mathcal{T}(S)$ we will mean with respect to the metric \bar{d}_{Th} .

Thurston showed that the infimum Lipschitz constant is realized by a homeomorphism from X to Y . Any map which realizes the infimum is called *optimal*.

Further, Thurston constructs a chain-recurrent lamination $\Lambda(X, Y)$ such that there exists a $e^{d_{\text{Th}}(X, Y)}$ -Lipschitz map in the preferred homotopy class from a neighborhood of $\Lambda(X, Y)$ in X to an neighborhood of the same lamination in Y , multiplying arc length along $\Lambda(X, Y)$ by a factor of $e^{d_{\text{Th}}(X, Y)}$, and so that $\Lambda(X, Y)$ is the largest chain-recurrent lamination with this property. We call $\Lambda(X, Y)$ the *maximally-stretched lamination* (from X to Y). The same lamination is also characterized in terms of optimal maps: $\Lambda(X, Y)$ is the largest chain-recurrent lamination such that every optimal map from X to Y multiplies arc length on $\Lambda(X, Y)$ by a factor of $e^{d_{\text{Th}}(X, Y)}$.

The length ratio for simple curves extends continuously to $\mathcal{PML}(S)$, which is compact. Therefore, the length-ratio supremum is always realized by some measured lamination. Any measured lamination that realizes the supremum has support contained in the stump of $\Lambda(X, Y)$.

Suppose a parameterized path $\mathcal{G}: [0, d] \rightarrow \mathcal{T}(S)$ is a geodesic from X to Y (parameterized by unit speed). Then the following holds: for any $s, t \in [0, d]$ with $s < t$ and for any arc ω contained in the geometric realization of $\Lambda(X, Y)$ on X , the arc length of ω is stretched by a factor of e^{t-s} under an optimal map from $\mathcal{G}(s)$ to $\mathcal{G}(t)$. We will sometimes denote $\Lambda(X, Y)$ by $\lambda_{\mathcal{G}}$.

2.7 Stretch paths

Certain geodesics of Thurston's metric can be described using shearing coordinates. Let λ be a complete geodesic lamination and $X \in \mathcal{T}(S)$. For any $t \in \mathbf{R}$ let $\text{stretch}(X, \lambda, t)$ be the unique point in $\mathcal{T}(S)$ such that

$$s_{\lambda}(\text{stretch}(X, \lambda, t)) = e^t s_{\lambda}(X).$$

Letting t vary, we have that $\text{stretch}(X, \lambda, t)$ is a parameterized path in $\mathcal{T}(S)$ that maps to an open ray from the origin in \mathbf{R}^N under the shearing coordinates. This is the *stretch path along λ from X* .

Thurston showed that the path $t \mapsto \text{stretch}(X, \lambda, t)$ is a geodesic in $\mathcal{T}(S)$ in the sense of (2). Note that we always consider the stretch path to be oriented in the direction of increasing t , which is natural since the asymmetry of the metric implies that the same path parameterized in the opposite direction may not be geodesic.

Also, if $\lambda_0 \subset \lambda$ is the largest chain-recurrent sublamination, then λ_0 is the maximally-stretched lamination for any pair of points $\text{stretch}(X, \lambda, s)$ and $\text{stretch}(X, \lambda, t)$ with $s < t$.

Removing the point X from a stretch path from X leaves two (open) *stretch rays*; of these, the one corresponding to $t > 0$ is a stretch ray *starting at X* and that with $t < 0$ is the one *ending at X* .

Thurston used stretch paths to show that $\mathcal{T}(S)$ equipped with the Thurston metric is a geodesic metric space. We summarize his results below. See the statement and proof of [Thu86c, Theorem 8.5] for more details.

Theorem 2.2 ([Thu86c]). *For any $X, Y \in \mathcal{T}(S)$, let $\Lambda(X, Y)$ be the maximally-stretched lamination from X to Y . Let λ be any completion of $\Lambda(X, Y)$. Then there exists a geodesic \mathcal{G} from X to Y consisting of a finite concatenation of stretch path segments*

$$\mathcal{G} = \mathcal{G}_1 \cdots \mathcal{G}_n,$$

where \mathcal{G}_1 is a segment of $\text{stretch}(X, \lambda, t)$, and all other \mathcal{G}_i 's stretch along some complete lamination containing $\Lambda(X, Y)$. Furthermore, such a geodesic can be chosen so that if X_i is the initial point of \mathcal{G}_i , then for all $i > 1$ we have $\Lambda(X_i, Y) \supsetneq \Lambda(X_{i-1}, Y)$. In particular, we can always take $n \leq 2|\chi(S)|$.

In general, geodesics of the Thurston metric from X to Y are not unique. But when $\Lambda(X, Y)$ is maximal chain-recurrent, then there is a unique geodesic. This statement follows from Theorem 2.2 but it is not explicitly stated in [Thu86c]. For completeness, we provide a proof:

Corollary 2.3. *Given $X, Y \in \mathcal{T}(S)$, suppose $\Lambda(X, Y)$ is maximal chain-recurrent. Let λ be a completion of $\Lambda(X, Y)$. Then $\text{stretch}(X, \lambda, t)$ is the unique geodesic from X to Y . In particular, for the punctured torus $S_{1,1}$, the three completions of $\Lambda(X, Y)$ give rise to the same stretch path in $\mathcal{T}(S_{1,1})$.*

Proof. We first show that the stretch path for λ connects X to Y , i.e. $\text{stretch}(X, \lambda, t) = Y$ for some t . By Theorem 2.2, there is a geodesic path \mathcal{G} from X to Y consisting of a concatenation of segments along stretch paths $\mathcal{G}_1, \dots, \mathcal{G}_n$, where \mathcal{G}_1 is a segment of $\{\text{stretch}(X, \lambda, t) : t \geq 0\}$. Let X_i be the initial point of \mathcal{G}_i . If $n \geq 2$, then $\Lambda(X, Y) = \Lambda(X_1, Y) \subsetneq \Lambda(X_2, Y)$ by Theorem 2.2. But this is impossible since $\Lambda(X, Y)$ is maximal chain-recurrent, so $n = 1$ and Y lies on \mathcal{G} .

Now suppose \mathcal{G} is any geodesic from X to Y . Let Z be a point on \mathcal{G} . We have $\Lambda(X, Y) \subset \Lambda(X, Z)$. Since $\Lambda(X, Y)$ is maximal chain-recurrent, $\Lambda(X, Z) = \Lambda(X, Y)$. By the previous discussion, we can connect X to Z by a segment of $\text{stretch}(X, \lambda, t)$. Since this true for all Z in \mathcal{G} , the geodesic \mathcal{G} must be a segment of $\text{stretch}(X, \lambda, t)$. \square

2.8 Twisting

There are several notions of twisting which we will define below. While these notions are defined for different classes of objects, in cases where several of the definitions apply, they are equal up to an additive constant.

Let A be an annulus. Fix an orientation of the core curve α of A . For any simple arc ω in A with endpoints on different components of ∂A , we orient ω so that the algebraic intersection number $\omega \cdot \alpha$ is equal to one. Given an ordered pair of simple arcs ω and ω' , the choice of the orientation above allows us to assign a sign to each intersection point in the interior of A between ω and ω' . The sum $\omega \cdot \omega'$ of these signed intersections is called the *algebraic intersection number* between ω and ω' . Note that $\omega \cdot \omega'$ is independent of the choice of the orientation of α . Also note that we do not consider intersections between ω and ω' in the boundary of A . With our choice, we always have $\omega \cdot D_\alpha(\omega) = 1$, where as above D_α denotes the left Dehn twist about α .

Now let S be a surface and α is a simple closed curve on S . Let $\widehat{S} \rightarrow S$ be the covering space associated to $\pi_1(\alpha) < \pi_1(S)$. Then \widehat{S} has a natural Gromov compactification that is homeomorphic to a closed annulus. By construction, the core curve $\widehat{\alpha}$ of this annulus maps homeomorphically to α under this covering map.

Let λ and λ' be two geodesic laminations (possibly curves) on S , both intersecting α transversely. We define their (signed) twisting relative to α as $\text{twist}_\alpha(\lambda, \lambda') = \min \widehat{\omega} \cdot \widehat{\omega}'$, where $\widehat{\omega}$ is a lift of a leaf of λ and $\widehat{\omega}'$ is a lift of a leaf of λ' , with both lifts intersecting $\widehat{\alpha}$, and the minimum is taken over all such leaves and their lifts. Note that for any two such lifts ω and ω' (still intersecting $\widehat{\alpha}$) the quantity $\widehat{\omega} \cdot \widehat{\omega}'$ exceeds $\text{twist}_\alpha(\lambda, \lambda')$ by at most 2.

Next we define the twisting of two hyperbolic metrics X and Y on S relative to α . Let \widehat{X}, \widehat{Y} denote the lifts of these hyperbolic structures to \widehat{S} . Using the hyperbolic structure \widehat{X} , choose a geodesic $\widehat{\omega}$ that is orthogonal to the geodesic in the homotopy class of $\widehat{\alpha}$. Let $\widehat{\omega}'$ be a geodesic constructed similarly from \widehat{Y} . We set $\text{twist}_\alpha(X, Y) = \min \widehat{\omega} \cdot \widehat{\omega}'$, where the minimum is taken over all possible choices for $\widehat{\omega}$ and $\widehat{\omega}'$. Similar to the previous case, this minimum differs from the intersection number $\widehat{\omega} \cdot \widehat{\omega}'$ for a particular pair of choices by at most 2.

Finally, we define $\text{twist}_\alpha(X, \lambda)$, the twisting of a lamination λ about a curve α on X . This is defined if λ contains a leaf that intersects α transversely. Let $\widehat{\omega}$ be a geodesic of \widehat{X} orthogonal to the geodesic homotopic to $\widehat{\alpha}$. Let ω' be any leaf of λ intersecting α , and let $\widehat{\omega}'$ be a lift of this leaf to \widehat{X} which intersects $\widehat{\alpha}$. Then $\text{twist}_\alpha(X, \lambda) = \min \widehat{\omega} \cdot \widehat{\omega}'$, with the minimum taken over all choices of ω' , $\widehat{\omega}'$, and $\widehat{\omega}$.

Each type of twisting defined above is *signed*. In some cases the absolute value of the twisting is the relevant quantity; we use the notation $d_\alpha(\cdot, \cdot) = |\text{twist}_\alpha(\cdot, \cdot)|$ for the corresponding *unsigned twisting* in each case.

The following way to compute the unsigned twisting $d_\alpha(X, \lambda) = |\text{twist}_\alpha(X, \lambda)|$ will be useful in the sequel. Consider the universal cover $\widetilde{X} \cong \mathbf{H}$. Let $\widetilde{\alpha}$ be a lift of α and let $\widetilde{\omega}'$ be a lift of a leaf of λ intersecting $\widetilde{\alpha}$. Let L be the length of the orthogonal projection of $\widetilde{\omega}'$ to $\widetilde{\alpha}$ and let ℓ be the length of the geodesic representative of α on X . Let $\widetilde{\omega}$ be an orthogonal geodesic of $\widetilde{\alpha}$. There is a loxodromic isometry T of \mathbf{H} associated to α that preserves $\widetilde{\alpha}$, and applying powers of this isometry to $\widetilde{\omega}$ gives a family of orthogonal geodesics to $\widetilde{\alpha}$ which meet it at points spaced by distance ℓ . Then $d_\alpha(X, \lambda)$ is the number of these translates that intersect $\widetilde{\omega}'$, as each such translate gives one intersection in the quotient $\widehat{X} = \mathbf{H}/\langle T \rangle$ considered above. Therefore, this number is between $(\lfloor L/\ell \rfloor - 1)$ and $\lfloor L/\ell \rfloor$, and $d_\alpha(X, \lambda) \stackrel{\pm}{\approx} L/\ell$ with additive error at most 2 (see also [Min96, Section 3] for more details).

3 Twisting parameter along a Thurston geodesic

In this section, S is any oriented surface of finite type and $\mathcal{T}(S)$ is the associated Teichmüller space.

Recall that $\mathcal{T}_\epsilon(S)$ denotes the ϵ -thick part of $\mathcal{T}(S)$. Consider two points $X, Y \in \mathcal{T}_\epsilon(S)$. Recall that we say a curve α *interacts* with a geodesic lamination λ if α is a leaf of λ or if α intersects λ essentially. Suppose α is a curve that interacts with $\Lambda(X, Y)$. Let $\mathcal{G}: [0, T] \rightarrow \mathcal{T}$ be any geodesic from X to Y , and let $\ell_\alpha = \min_t \ell_\alpha(t)$. We are interested in curves which become

short somewhere along \mathcal{G} . We call an interval of time $[a, b] \subset [0, T]$ the *active interval* for α along \mathcal{G} if $[a, b]$ is the maximal such interval with $\ell_\alpha(a) = \ell_\alpha(b) = \epsilon$. Note that any curve which is sufficiently short somewhere on \mathcal{G} has a nontrivial active interval.

The main goal of this section is to prove the following theorem, which in particular establishes Theorem 1.2. As in the introduction we use the notation $\text{Log}(x) = \min(1, \log(x))$. Denote $X_t = \mathcal{G}(t)$.

Theorem 3.1. *There exists a constant ϵ_0 such that the following statement holds. Let $X, Y \in \mathcal{T}_{\epsilon_0}(S)$ and α be a curve that interacts with $\Lambda(X, Y)$. Let \mathcal{G} be any geodesic from X to Y and $\ell_\alpha = \min_t \ell_\alpha(t)$. Then*

$$d_\alpha(X, Y) \stackrel{*}{\asymp} \frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha}.$$

If $\ell_\alpha < \epsilon_0$, then $d_\alpha(X, Y) \stackrel{\pm}{\asymp} d_\alpha(X_a, X_b)$, where $[a, b]$ is the active interval for α . Further, for all sufficiently small ℓ_α , the twisting $d_\alpha(X_t, \lambda)$ is uniformly bounded for all $t \leq a$ and $\ell_\alpha(t) \stackrel{*}{\asymp} e^{t-b} \ell_\alpha(b)$ for all $t \geq b$. All errors in this statement depend only on ϵ_0 .

Note that if α is a leaf of $\Lambda(X, Y)$, then it does not have an active interval because its length grows exponentially along \mathcal{G} , and the theorem above says that in this $d_\alpha(X, Y)$ is uniformly bounded. If α crosses a leaf of $\Lambda(X, Y)$, then $d_\alpha(X, Y)$ is large if and only if α gets short along any geodesic from X to Y . Moreover, the minimum length of α is the same for any geodesic from X to Y , up to a multiplicative constant. Further, the theorem says that, essentially, all of the twisting about α occurs in the active interval $[a, b]$ of α .

Before proceeding to the proof of the theorem, we need to introduce a notion of horizontal and vertical components for a curve that crosses a leaf of $\Lambda(X, Y)$ and analyze how their lengths change in the active interval. This analysis will require some lemmas from hyperbolic geometry.

Lemma 3.2. *Let ω and ω' be two disjoint geodesics in \mathbf{H} with no endpoint in common. Let $p \in \omega$ and $p' \in \omega'$ be the endpoints of the common perpendicular between ω and ω' . Let $x \in \omega$ be arbitrary and let $x' \in \omega'$ be the point on the same side of $[p, p']$ as x such that $d_{\mathbf{H}}(x, p) = d_{\mathbf{H}}(x', p')$. Then*

$$\sinh \frac{d_{\mathbf{H}}(p, p')}{2} \cosh d_{\mathbf{H}}(x, p) = \sinh \frac{d_{\mathbf{H}}(x, x')}{2}. \quad (3)$$

For any $y \in \omega'$, we have

$$\sinh d_{\mathbf{H}}(p, p') \cosh d_{\mathbf{H}}(x, p) \leq \sinh d_{\mathbf{H}}(x, y) \quad (4)$$

and

$$d_{\mathbf{H}}(x', y) \leq d_{\mathbf{H}}(x, y). \quad (5)$$

Proof. We refer to Figure 4 for the proof. Equation (3) is well known, as the four points x, x', p', p form a Saccheri quadrilateral. The point $y' \in \omega'$ closest to x has $\angle xy'p' = \pi/2$, so x, y', p', p forms a Lambert quadrilateral and the following identity holds

$$\sinh d_{\mathbf{H}}(p, p') \cosh d_{\mathbf{H}}(x, p) = \sinh d_{\mathbf{H}}(x, y')$$




Figure 4: Saccheri and Lambert quadrilaterals.

Equation (4) follows since $d_{\mathbf{H}}(x, y') \leq d_{\mathbf{H}}(x, y)$. For (5), set $A = \angle xx'y$ and $B = \angle x'xy$ and consider the triangle $\triangle xx'y$. Depending on which side of x' the point y is, A is obtuse or acute. In any case, $A \geq B$. It is a standard fact that the side opposite the bigger angle in a triangle is longer. Hence $d_{\mathbf{H}}(x, y) \geq d_{\mathbf{H}}(x', y)$. \square

In this section we will often use the following elementary estimates for hyperbolic trigonometric functions. The proofs are omitted.

Lemma 3.3.

- (i) If $0 \leq x \leq 1$ or $0 \leq \sinh(x) \leq 1$, we have $\sinh(x) \leq 2x$.
- (ii) For all $x \geq 0$, $\frac{1}{2}e^x \leq \cosh(x) \leq e^x$ and $x \leq \sinh(x) \leq \frac{1}{2}e^x$.
- (iii) For all $x \geq 1$, we have $\sinh(x) \geq \frac{1}{4}e^x$.
- (iv) For all $x \geq 1$, we have

$$\log(2) \leq \operatorname{arcsinh}(x) - \log(x) \leq \log(3)$$

and

$$0 \leq \operatorname{arccosh}(x) - \log(x) \leq \log(2).$$

Now consider $X \in \mathcal{T}(S)$ and a geodesic lamination λ on X . If α crosses a leaf ω of λ , define $V_X(\omega, \alpha)$ to be a shortest arc with endpoints on ω that, together with an arc $H_X(\omega, \alpha)$ of ω , form a curve homotopic to α . Thus $V_X(\omega, \alpha)$ and $H_X(\omega, \alpha)$ meet orthogonally and α passes through the midpoints of both of these arcs (see Figure 5). If α is a leaf of λ , then we set $H_X(\omega, \alpha) = \alpha$ and let $V_X(\omega, \alpha)$ be the empty set.

Define h_X and v_X to be the lengths of $H_X(\omega, \alpha)$ and $V_X(\omega, \alpha)$ respectively. By considering the right triangles formed by these curves and α (which have hypotenuse along α), it is immediate that

$$\max(h_X, v_X) \leq \ell_\alpha(X) \leq h_X + v_X. \quad (6)$$

The quantities h_X and v_X can be computed in the universal cover $\tilde{X} \cong \mathbf{H}$ as follows. Let $\tilde{\omega}$ and $\tilde{\alpha}$ be intersecting lifts of ω and α to \mathbf{H} . Let ϕ be the hyperbolic isometry with axis $\tilde{\alpha}$ and translation length $\ell_\alpha(X)$. Set $\tilde{\omega}' = \phi(\tilde{\omega})$ and let ψ be the hyperbolic isometry taking $\tilde{\omega}$ to $\tilde{\omega}'$ with axis perpendicular to the two geodesics. Since ϕ and ψ both take ω to ω' , their composition $\psi^{-1}\phi$ is a hyperbolic isometry with axis $\tilde{\omega}$. The quantity v_X is the translation length of ψ and h_X is the translation length of $\psi^{-1}\phi$. For the latter, this means that $h_X = d_{\mathbf{H}}(\psi(q), \phi(q))$ for any $q \in \tilde{\omega}$.

In the following, let $X_t = \mathcal{G}(t)$ be a geodesic segment and let $\lambda = \lambda_{\mathcal{G}}$. Let α be a curve that interacts with λ . We will refer to $V_{X_t}(\omega, \alpha)$ and $H_{X_t}(\omega, \alpha)$ as the *vertical* and *horizontal* components of α at X_t . We are interested in the lengths $h_t = h_{X_t}$ and $v_t = v_{X_t}$ of the horizontal and vertical components of α as functions of t . We will show that v_t decreases super-exponentially, while h_t grows exponentially. These statements are trivial if α is a leaf of λ , so we will always assume that α crosses a leaf ω of λ .

Lemma 3.4. *Suppose α crosses a leaf ω of λ . For any $t \geq s$,*

$$h_t \geq e^{t-s} (h_s - v_s).$$

Proof. In \mathbf{H} , choose a lift $\tilde{\alpha}$ of the geodesic representative of α on X_s and a lift $\tilde{\omega}$ of ω that crosses $\tilde{\alpha}$. Let $\tilde{\omega}' = \phi_s(\tilde{\omega})$ where ϕ_s is the hyperbolic isometry with axis $\tilde{\alpha}$ and translation length $\ell_s(\alpha)$. Let ψ_s be the hyperbolic isometry taking $\tilde{\omega}$ to $\tilde{\omega}'$ with axis perpendicular to the two geodesics. Let $p \in \tilde{\omega}$ be the point lying on the axis of ψ_s . By definition,

$$v_s = d_{\mathbf{H}}(p, \psi_s(p)) \quad \text{and} \quad h_s = d_{\mathbf{H}}(\psi_s(p), \phi_s(p)).$$

The configuration of points and geodesics in \mathbf{H} constructed above is depicted in Figure 5; it may be helpful to refer to this figure in the calculations that follow. Note that for brevity the subscript s is omitted from the labels involving ψ, ϕ in the figure.

Let $f : X_s \rightarrow X_t$ be an optimal map and let $\tilde{f} : \mathbf{H} \rightarrow \mathbf{H}$ be a lift of f . Since f is an e^{t-s} -Lipschitz map such that distances along leaves of λ are stretched by a factor of exactly e^{t-s} , the images $\tilde{f}(\tilde{\omega})$ and $\tilde{f}(\tilde{\omega}')$ are geodesics and

$$d_{\mathbf{H}}(\tilde{f}\psi_s(p), \tilde{f}\phi_s(p)) = e^{t-s}h_s \quad \text{and} \quad d_{\mathbf{H}}(\tilde{f}(p), \tilde{f}\psi_s(p)) \leq e^{t-s}v_s.$$

Let ψ_t be the hyperbolic isometry taking $\tilde{f}(\tilde{\omega})$ to $\tilde{f}(\tilde{\omega}')$ with axis their common perpendicular. Let ϕ_t be the hyperbolic isometry corresponding to $f\alpha$ taking $\tilde{f}(\tilde{\omega})$ to $\tilde{f}(\tilde{\omega}')$. Note that $\phi_t\tilde{f} = \tilde{f}\phi_s$, since \tilde{f} is a lift of f . But ψ_s and ψ_t do not necessarily correspond to a conjugacy class of $\pi_1(S)$, so \tilde{f} need not conjugate ψ_s to ψ_t .

By definition,

$$h_t = d_{\mathbf{H}}(\psi_t\tilde{f}(p), \phi_t\tilde{f}(p)) = d_{\mathbf{H}}(\psi_t\tilde{f}(p), \tilde{f}\phi_s(p)).$$

By Lemma 3.2(5),

$$d_{\mathbf{H}}(\tilde{f}\psi_s(p), \psi_t\tilde{f}(p)) \leq d_{\mathbf{H}}(\tilde{f}\psi_s(p), \tilde{f}(p)).$$




Figure 5: Estimating h_t .

Using the triangle inequality and the above equations, we obtain the conclusion.

$$\begin{aligned}
 h_t &\geq d_{\mathbf{H}}\left(\phi_t \tilde{f}(p), \tilde{f}\psi_s(p)\right) - d_{\mathbf{H}}\left(\tilde{f}\psi_s(p), \psi_t \tilde{f}(p)\right) \\
 &\geq d_{\mathbf{H}}\left(\phi_t \tilde{f}(p), \tilde{f}\psi_s(p)\right) - d_{\mathbf{H}}\left(\tilde{f}\psi_s(p), \tilde{f}(p)\right) \\
 &= d_{\mathbf{H}}\left(\tilde{f}\phi_s(p), \tilde{f}\psi_s(p)\right) - d_{\mathbf{H}}\left(\tilde{f}\psi_s(p), \tilde{f}(p)\right) \\
 &\geq e^{t-s}h_s - e^{t-s}v_s
 \end{aligned}$$

□

Lemma 3.5. *Suppose α crosses a leaf ω of λ . There exists $\epsilon_v > 0$ such that if $v_a \leq \epsilon_v$, then for all $t \geq a$, we have:*

$$v_t \leq e^{-Ae^{t-a}}, \quad \text{where } A > 0 \text{ and } A \stackrel{\pm}{\geq} \log \frac{1}{v_a}$$

and where the additive error is at most $\log 4 + 1$.

Proof. We refer to Figure 6. As before, choose a lift $\tilde{\alpha}$ to \mathbf{H} of the geodesic representative of α on X_a and a lift $\tilde{\omega}$ of ω that crosses $\tilde{\alpha}$. Let $\tilde{\omega}' = \phi(\tilde{\omega})$ where ϕ is the hyperbolic isometry with axis $\tilde{\alpha}$ and translation length $\ell(\alpha)$. Let $p \in \tilde{\omega}$ and $p' \in \tilde{\omega}'$ be the endpoints of the common perpendicular between $\tilde{\omega}$ and $\tilde{\omega}'$, so $v_a = d_{\mathbf{H}}(p, p')$.

We assume $v_a < \frac{1}{2}$. Let $[x, y] \subset \tilde{\omega}$ and $[x', y'] \subset \tilde{\omega}'$ be segments of the same length with midpoints p and p' , such that $[x, x']$ and $[y, y']$ have length 1 and are disjoint from $[p, p']$. By (3) from Lemma 3.2,

$$d_{\mathbf{H}}(x, p) = \operatorname{arccosh} \frac{\sinh 1/2}{\sinh v_a/2}$$

We can apply Lemma 3.3(i) and (iv), which give

$$\left| d_{\mathbf{H}}(x, y) - 2 \log \frac{1}{v_a} \right| \leq 2 \log 4. \quad (7)$$




Figure 6: Bounding v_t from above.

In particular, v_a is small if and only if $d_{\mathbf{H}}(x, y)$ is large. Let ϵ_v be small enough so that $d_{\mathbf{H}}(x, y) \geq 4$.

Let $f : X_a \rightarrow X_t$ be an optimal map and $\tilde{f} : \mathbf{H} \rightarrow \mathbf{H}$ a lift of f . Let $r \in \tilde{f}(\tilde{\omega})$ and $r' \in \tilde{f}(\tilde{\omega}')$ be the endpoints of the common perpendicular between $\tilde{f}(\tilde{\omega})$ and $\tilde{f}(\tilde{\omega}')$, so $v_t = d_{\mathbf{H}}(r, r')$. Without a loss of generality, assume that r is farther away from $\tilde{f}(x)$ than $\tilde{f}(y)$. This means

$$d_{\mathbf{H}}(\tilde{f}(x), r) \geq \frac{1}{2} d_{\mathbf{H}}(\tilde{f}(x), \tilde{f}(y)). \quad (8)$$

We also have

$$d_{\mathbf{H}}(\tilde{f}(x), \tilde{f}(y)) = e^{t-a} d_{\mathbf{H}}(x, y) \quad \text{and} \quad d_{\mathbf{H}}(\tilde{f}(x), \tilde{f}(x')) \leq e^{t-a}. \quad (9)$$

By (4) from Lemma 3.2,

$$\sinh d_{\mathbf{H}}(r, r') \cosh d_{\mathbf{H}}(\tilde{f}(x), r) \leq \sinh d_{\mathbf{H}}(\tilde{f}(x), \tilde{f}(x')).$$

Incorporating (8) and (9) to the above inequality yields

$$\sinh d_{\mathbf{H}}(r, r') \leq \frac{\sinh e^{t-a}}{\cosh(\frac{1}{2} e^{t-a} d_{\mathbf{H}}(x, y))}$$

Now use Lemma 3.3(ii) to obtain

$$d_{\mathbf{H}}(r, r') \leq e^{-e^{t-a} \left(\frac{d_{\mathbf{H}}(x, y)}{2} - 1 \right)}.$$

Setting $A = \frac{d_{\mathbf{H}}(x, y)}{2} - 1$ and applying (7) we have that $A > 0$ and $|A - \log \frac{1}{v_a}| \leq \log 4 + 1$. This finishes the proof. \square

Lemma 3.6. *Suppose α crosses a leaf ω of λ . Let ϵ_v be the constant from Lemma 3.5. If $[a, b]$ is an interval of times with $\ell_\alpha(a) = \ell_\alpha(b) = \epsilon < \epsilon_v$, then $\ell_\alpha(t) \stackrel{*}{\asymp} \epsilon$ for all $t \in [a, b]$ with the multiplicative error at most $12e$.*

Proof. Let $t \in [a, b]$.

Suppose first that $v_t > \frac{1}{2}h_t$. Here one can replace $\frac{1}{2}$ by any other number in $(0, 1)$. Then $\ell_\alpha(t) \leq 3v_t$. By Lemma 3.5,

$$v_t \leq e^{-A} \stackrel{*}{\asymp} v_\alpha$$

where the multiplicative error is at most $4e$, and since v_α is bounded above by ϵ , we have $\ell_\alpha(t) \stackrel{*}{\asymp} \epsilon$, with error at most $12e$.

Now suppose $v_t \leq \frac{1}{2}h_t$. Then by Lemma 3.4

$$h_b \geq e^{b-t}(h_t - v_t) \geq \frac{1}{2}e^{b-t}h_t \geq \frac{1}{2}h_t$$

Hence

$$\ell_\alpha(t) \leq h_t + v_t \leq \frac{3}{2}h_t \leq 3h_b \leq 3\epsilon.$$

This finishes the proof. □

For our purposes, an important consequence of Lemma 3.6 is that if the curve is short enough at the endpoints of an interval, then its length will be below ϵ_M throughout that interval. Specifically, fix $\epsilon_0 > 0$ so that

$$\epsilon_0 < \min\left(\frac{\epsilon_M}{12e}, \epsilon_v\right),$$

where ϵ_M is the Margulis number chosen in Section 2.3 and ϵ_v is the constant from Lemma 3.5. Then as an immediate corollary of Lemma 3.6 we have:

Corollary 3.7. *If $[a, b]$ is an interval such that $\ell_\alpha(a) = \ell_\alpha(b) = \epsilon_0$, then $\ell_\alpha(t) < \epsilon_M$ for all $t \in [a, b]$. □*

Next we study the relationship between the relative twisting $d_\alpha(X, \lambda)$ and the length of $V_X(\omega, \alpha)$ and $\ell_\alpha(X)$.

Lemma 3.8. *Suppose α crosses a leaf ω of λ . Fix $X = X_t$ and let $\ell = \ell_\alpha(X)$ and v be the length of $V_X(\omega, \alpha)$. Then the following statements hold.*

(i) *If $\ell \leq \epsilon_M$, then*

$$\ell d_\alpha(X, \lambda) \stackrel{\pm}{\asymp} 2 \log \frac{\ell}{v}.$$

(ii) *If $d_\alpha(X, \lambda) \geq \frac{4\epsilon_M}{\ell} + 2$, then*

$$v \leq 2e^{-\frac{\ell}{4}d_\alpha(X, \lambda)}.$$

Proof. The reader may find it helpful to look at Figure 5 for this proof.

Let B be the angle between $\tilde{\alpha}$ and $\tilde{\omega}$. Let L be the length of the projection of $\tilde{\omega}$ to $\tilde{\alpha}$. Recall that $d_\alpha(X, \lambda) \stackrel{\pm}{\asymp} \frac{L}{\ell}$ with additive error at most 2. Since $\ell \leq \epsilon_M$, this implies

$$\ell d_\alpha(X, \lambda) \stackrel{\pm}{\asymp} L \quad (10)$$

with additive error at most $2\epsilon_m$. By hyperbolic geometry, L satisfies

$$1 = \cosh \frac{L}{2} \cdot \sin B. \quad (11)$$

To find $\sin B$, denote by ϕ the hyperbolic isometry with axis $\tilde{\alpha}$ and with translation length ℓ . Let $\tilde{\omega}' = \phi(\tilde{\omega})$. Denote by x the intersection of $\tilde{\alpha}$ and $\tilde{\omega}$ and set $x' = \phi(x)$. Let $p \in \tilde{\omega}$ and $p' \in \tilde{\omega}'$ be the points on the common perpendicular between $\tilde{\omega}$ and $\tilde{\omega}'$. That is, $p' = \psi(p)$ where ψ is the translation along an axis perpendicular to $\tilde{\omega}$ such that $\psi(\tilde{\omega}) = \tilde{\omega}'$. By construction, $d_{\mathbf{H}}(x, x') = \ell$ and $d_{\mathbf{H}}(p, p') = v$. Then the intersection point of $[p, p']$ and $[x, x']$ is the midpoint of both. Thus $\sin B$ can be found from

$$\sin B \sinh \frac{\ell}{2} = \sinh \frac{v}{2}. \quad (12)$$

Combining (11) and (12) we obtain

$$L = 2 \operatorname{arccosh} \frac{\sinh \ell/2}{\sinh v/2}. \quad (13)$$

When $\ell \leq \epsilon_M < 1$, we can apply Lemma 3.3(i) and (iv) to simplify (13), obtaining

$$L \stackrel{\pm}{\asymp} 2 \log \frac{\ell}{v}$$

which in combination with (10) gives

$$\ell d_\alpha(X, \lambda) \stackrel{\pm}{\asymp} 2 \log \frac{\ell}{v},$$

with the additive error in the latter estimate at most $2 \log 4 + 2\epsilon_M$.

Now we consider the upper bound on v under the assumption $d_\alpha(X, \lambda) \geq \frac{4\epsilon_M}{\ell} + 2$. By (10) we have $L \geq \ell d_\alpha(X, \lambda) - 2\epsilon_M$ and incorporating this with (13) gives

$$\frac{\sinh \ell/2}{\sinh v/2} \geq \cosh \left(\frac{\ell}{2} d_\alpha(X, \lambda) - \epsilon_M \right).$$

Therefore,

$$\frac{v}{2} \leq \sinh \frac{v}{2} \leq \frac{\sinh \ell/2}{\cosh \left(\frac{\ell}{2} d_\alpha(X, \lambda) - \epsilon_M \right)} \leq \frac{e^{\frac{\ell}{2}}}{e^{\frac{\ell}{2} d_\alpha(X, \lambda) - \epsilon_M}} = e^{\frac{\ell}{2} + \epsilon_M - \frac{\ell}{2} d_\alpha(X, \lambda)}$$

where the third inequality above uses the fact that $\frac{\ell}{2} d_\alpha(X, \lambda) - \epsilon_M > 0$ to apply Lemma 3.3(ii). Furthermore, our assumed lower bound on $d_\alpha(X, \lambda)$ gives

$$\frac{\ell}{2} + \epsilon_M - \frac{\ell}{2} d_\alpha(X, \lambda) \leq -\frac{\ell}{4} d_\alpha(X, \lambda)$$

and substituting this into the previous bound on $\frac{v}{2}$ we find

$$v \leq 2e^{-\frac{\ell}{4}d_\alpha(X,\lambda)}$$

which completes the proof. \square

The following lemma implies that the length of the vertical component does not decrease too quickly along a geodesic ray if the curve starts out being approximately vertical and remains short throughout the ray.

Lemma 3.9. *Suppose α crosses a leaf ω of λ . There exists $A > 0$ with $A \stackrel{\pm}{\asymp} \log \frac{1}{\epsilon_0}$ such that the following holds. If $\ell_\alpha(a) = \epsilon_0$ and $v_a \geq \frac{\epsilon_0}{4}$, and if $\ell_\alpha(t) < \epsilon_M$ for all $t \geq a$, then we have*

$$v_t \stackrel{*}{\asymp} e^{-Ae^{t-a}}.$$

Proof. Let β be a shortest curve at time a that intersects α . Recall that ϵ_M was chosen so that $\ell_\alpha(a) < \epsilon_M$ implies $i(\alpha, \beta) \in \{1, 2\}$. We will give the proof in the case $i(\alpha, \beta) = 1$, with the other case being essentially the same. Since α is short for all $t > a$, the part of β in a collar neighborhood of α has length that can be estimated in terms of the length of α and the relative twisting of X_t and β (see [CRS08, Lemma 7.3]), giving a lower bound for the length of β itself:

$$\ell_\beta(t) \stackrel{\pm}{\asymp} d_\alpha(X_t, \beta)\ell_\alpha(t) + 2 \log \frac{1}{\ell_\alpha(t)}.$$

On the other hand, since $\ell_\alpha(a) = \epsilon_0$ and $v_a \geq \frac{\epsilon_0}{4}$, applying Lemma 3.8 to X_a tells us that $d_\alpha(X_a, \lambda)$ is bounded. Hence $|d_\alpha(X_t, \lambda) - d_\alpha(X_t, \beta)| \stackrel{\pm}{\asymp} 1$ which means that we can write

$$\ell_\beta(t) \stackrel{\pm}{\asymp} d_\alpha(X_t, \lambda)\ell_\alpha(t) + 2 \log \frac{1}{\ell_\alpha(t)}.$$

The length of β cannot grow faster than the length of λ , therefore

$$d_\alpha(X_t, \lambda)\ell_\alpha(t) + 2 \log \frac{1}{\ell_\alpha(t)} \stackrel{\pm}{\asymp} e^{t-a}\ell_\beta(a).$$

Applying Lemma 3.8 again, now to X_t , we have

$$2 \log \frac{\ell_\alpha(t)}{v_t} + 2 \log \frac{1}{\ell_\alpha(t)} \stackrel{\pm}{\asymp} d_\alpha(X_t, \lambda)\ell_\alpha(t) + 2 \log \frac{1}{\ell_\alpha(t)} \stackrel{\pm}{\asymp} e^{t-a}\ell_\beta(a)$$

which implies

$$v_t \stackrel{*}{\asymp} e^{-\frac{1}{2}e^{t-a}\ell_\beta(a)}.$$

The claim now follows from the fact that $\ell_\beta(a) \stackrel{\pm}{\asymp} 2 \log \frac{1}{\epsilon_0}$. \square

Theorem 3.10. *Suppose α crosses a leaf ω of λ . Let $[a, b]$ be an interval such that $\ell_\alpha(a) = \ell_\alpha(b) = \epsilon_0$. Then*

$$d_\alpha(X_a, X_b) \stackrel{*}{\asymp} e^{b-a}.$$

The length of α is minimum in the interval $[a, b]$ at a time $t_\alpha \in [a, b]$ satisfying

$$t_\alpha - a \stackrel{\pm}{\asymp} \text{Log}(b - t_\alpha), \quad (14)$$

and the minimum length is $\ell_\alpha(t_\alpha) \stackrel{*}{\asymp} e^{-(b-t_\alpha)}$.

Furthermore, if $(b - a)$ is sufficiently large, then $\log(b - t_\alpha) > 1$ and so (14) also holds with Log replaced by \log .

In some of the preceding lemmas we indicated the dependence of multiplicative and additive errors on ϵ_0 . However, since ϵ_0 is a fixed constant, we will ignore such dependence in most cases from now on.

Proof. We split the proof into two cases, depending on whether the interval $[a, b]$ is “short” or “long”. More precisely we consider the cases $(b - a) \leq Q$ and $(b - a) > Q$ for some positive real Q , the *threshold*. The implicit constants in the approximate comparisons we derive in each case will depend on Q , and at various points in the long-interval case it will be necessary to assume Q is sufficiently large (i.e. greater than some universal constant). At the end we can fix any Q large enough to satisfy all of those assumptions.

First we consider the short-interval case, $(b - a) \leq Q$. Here, all of the claims of approximate equality in the theorem will hold because all of the quantities in question are bounded. Since $t_\alpha \in [a, b]$, both $(t_\alpha - a)$ and $\text{Log}(b - t_\alpha)$ are nonnegative and bounded above, i.e. $t_\alpha - a \stackrel{\pm}{\asymp} \text{Log}(b - t_\alpha) \stackrel{\pm}{\asymp} 0$.

The surface X_{t_α} admits maps from X_a and to X_b with bounded Lipschitz constant (at most e^Q). Since α has length ϵ_0 on both X_a and X_b , this shows that $\ell_\alpha(t)$ is bounded above and below by positive constants depending on Q for all $t \in [a, b]$, i.e. that $\ell_\alpha(t_\alpha) \stackrel{*}{\asymp} 1$. Since $1 \geq e^{-(b-t_\alpha)} \geq e^{-(b-a)} \geq e^{-Q}$, we also have $e^{-(b-t_\alpha)} \stackrel{*}{\asymp} 1$, and thus $\ell_\alpha(t_\alpha) \stackrel{*}{\asymp} e^{-(b-t_\alpha)}$.

To obtain the bound on $d_\alpha(X_a, X_b)$ in the short-interval case, we recall from [Min96] that the rate at which $d_\alpha(X_a, \cdot)$ can change is bounded with the bound depending on the length of α . As noted above we have upper and lower bounds for the length of α along the geodesic between X_a to X_b , hence $d_\alpha(X_a, X_b) \stackrel{\pm}{\asymp} 0$. We are assuming an upper bound on $(b - a)$, so this implies $d_\alpha(X_a, X_b) \stackrel{\pm}{\asymp} e^{b-a}$.

Now we turn to the long-interval case, $(b - a) > Q$. First we require $Q > \log(2)$, so that $e^{b-a} > 2$. It follows that $h_a - v_a \leq \epsilon_0/2$; to see this, assume for contradiction that $h_a - v_a > \epsilon_0/2$. Then Lemma 3.4 gives

$$h_b \geq \frac{\epsilon_0}{2} e^{b-a} > \epsilon_0,$$

while (6) gives

$$h_b \leq \ell_\alpha(b) = \epsilon_0,$$

a contradiction.

Now, since $h_a - v_a \leq \epsilon_0/2$ and $h_a + v_a \geq \ell_\alpha(a) = \epsilon_0$, we find $\frac{\epsilon_0}{4} \leq v_a \leq \epsilon_0$, i.e. at time $t = a$ the curve is nearly perpendicular to λ , and $\ell_\alpha(a)/v_a \stackrel{*}{\asymp} 1$. Applying Lemma 3.8 we obtain

$$\epsilon_0 d_\alpha(X_a, \lambda) = \ell_\alpha(a) d_\alpha(X_a, \lambda) \stackrel{\pm}{\asymp} 0.$$

Dividing by ϵ_0 we obtain $d_\alpha(X_a, \lambda) \stackrel{\pm}{\asymp} 0$.

By Corollary 3.7 we have $\ell_\alpha(t) < \epsilon_M$ for all $t \in [a, b]$. Using this, the bounds of Lemma 3.5 and Lemma 3.9 show that there are $A, B > 0$ such that

$$e^{-Be^{t-a}} \stackrel{*}{\asymp} v_t \leq e^{-Ae^{t-a}}, \text{ for all } t \in [a, b]. \quad (15)$$

(And in fact those lemmas show $A, B \stackrel{\pm}{\asymp} \log \frac{1}{\epsilon_0}$.) Taking the logarithm of (15) gives

$$\log \frac{1}{v_t} \stackrel{*}{\mp} e^{t-a}, \text{ for all } t \in [a, b], \quad (16)$$

where the additive error comes from the multiplicative error in (15), and the multiplicative error from the constants A, B .

We claim that for Q sufficiently large there exists $s \in [a, b]$ such that $h_s = 2v_s$. Indeed, if $h_s < 2v_s$ for all $s \in [a, b]$, then, since $h_b + v_b \geq \epsilon_0$, we have $\frac{1}{3}\epsilon_0 \leq v_b \leq \epsilon_0$. Using (16) with $t = b$ this gives an upper bound on e^{b-a} , which is a contradiction if Q is large enough. On the other hand, if $h_s > 2v_s$ for all $s \in [a, b]$, then Lemma 3.4 implies that h_b is large if $(b-a)$ is sufficiently large. Specifically, by taking Q larger than a universal constant we would have $h_b > \epsilon_0$, contradicting that $\epsilon_0 = \ell_\alpha(b) \geq h_b$. Thus by requiring Q to be large enough so that both of these arguments apply, we have $h_s = 2v_s$ for some $s \in [a, b]$. For the rest of the proof, let s denote any such point in the interval.

Since v_s and h_s are comparable, it follows from (6) that $v_s \stackrel{*}{\asymp} \ell_\alpha(s)$. Since e^{b-s} is the Lipschitz constant from X_s to X_b , we have $\ell_\alpha(s)e^{b-s} \geq \ell_\alpha(b) = \epsilon_0$. In particular $\epsilon_0 \stackrel{*}{\asymp} v_s e^{b-s}$. On the other hand, Lemma 3.4 gives

$$h_b \geq \frac{1}{2}h_s e^{b-s} = v_s e^{b-s}.$$

Thus $v_s e^{b-s} \leq h_b \leq \ell_\alpha(b) = \epsilon_0$. All together, we obtain

$$v_s \stackrel{*}{\asymp} e^{-(b-s)}. \quad (17)$$

Now using (16) with $t = s$ and (17) together we find

$$e^{s-a} \stackrel{*}{\mp} (b-s). \quad (18)$$

From this, it follows that

$$\text{Log}(b-s) \stackrel{\pm}{\asymp} (s-a). \quad (19)$$

Indeed, if $\log(b-s) \geq 1$ then $\log(b-s) = \text{Log}(b-s)$ and (19) is the result of taking the logarithm of (18). Otherwise $\log(b-s) < 1$, in which case $\text{Log}(b-s) = 1$ and (18) gives a uniform upper bound on $(s-a)$, so (19) holds simply because both sides are nonnegative and bounded.

Finally, since v_t is essentially decreasing double-exponentially, h_t is increasing exponentially and $\ell_\alpha(t) \geq \max\{v_t, h_t\}$, it follows that $t_\alpha \stackrel{\pm}{\asymp} s$. This gives us the order of the minimal length of α , which is approximated by $\ell_\alpha(s) \stackrel{*}{\asymp} v_s$. Also, using $t_\alpha \stackrel{\pm}{\asymp} s$, we find that equation (19) also holds if we replace s by t_α , which gives (14).

To complete the long-interval case we estimate $d_\alpha(X_a, X_b)$. By Lemma 3.8 and (16),

$$\epsilon_0 d_\alpha(X_b, \lambda) \stackrel{\pm}{\asymp} 2 \log \frac{\epsilon_0}{v_b} \stackrel{*}{\asymp} e^{b-a},$$

and we can absorb the additive error in the multiplicative error since the expression on the right is bounded away from 0. Since v_t decreases double-exponentially, v_b is very small compared to ϵ_0 for $(b - a)$ large, so $\epsilon_0 d_\alpha(X_b, \lambda)$ is bounded away from 0. Dividing by ϵ_0 (and absorbing this into the multiplicative error as well) we find $d_\alpha(X_b, \lambda) \stackrel{*}{\asymp} e^{b-a}$. Since $d_\alpha(X_a, \lambda) \stackrel{\pm}{\asymp} 0$, this is the desired estimate.

Fixing a value for the threshold Q large enough to satisfy all of the conditions derived in the long-interval analysis above, the estimates in both parts of the proof become uniform (i.e. no longer depend on an additional parameter).

It only remains to prove the final claim from the statement of the theorem. For this, we show $(b - t_\alpha)$ can be made larger than a given constant just by assuming that $(b - a)$ is sufficiently large. Suppose for contradiction that $(b - t_\alpha)$, and hence also $(b - s)$, can be bounded with $(b - a)$ arbitrarily large. Then $e^{s-a} \stackrel{*}{\asymp} e^{b-a}$ is large while $(b - s)$ is bounded, contradicting (18). \square

Note that Theorem 3.10 highlights an interesting contrast between the behavior of Thurston metric geodesics and that of Teichmüller geodesics: Along a Teichmüller geodesic, a curve α achieves its minimum length near the midpoint of the interval in which α is short (see [Raf14, Section 3]), and this minimum is on the order of $d_\alpha(X, Y)^{-1}$. However, for a Thurston metric geodesic, the minimum length occurs much closer to the start of the interval (assuming the interval is sufficiently long) since $(t_\alpha - a)$ is only on the order of $\log(b - t_\alpha)$. In addition, the minimum length on the Thurston geodesic is larger than in the Teichmüller case, though only by a logarithmic factor.

To exhibit this difference, Figure 7 shows a Teichmüller geodesic segment and a stretch path segment (for lamination β^+) joining the same pair of points in the upper half plane model of $\mathcal{T}(S_{1,1})$. Here β is a simple closed curve. In this model, the imaginary part of a point $z \in \mathbf{H}$ is approximately $\pi/\ell_\alpha(z)$, where α is a curve which has approximately the same length at both endpoints but which becomes short somewhere along each path. Thus the expected (and observed) behavior of the Thurston geodesic is that its maximum height is lower than that of the Teichmüller geodesic, but that this maximum height occurs closer to the starting point for the Thurston geodesic. Further properties of Thurston geodesics in the punctured torus case are explored in the next section.

Continuing toward the proof of Theorem 3.1, we show:

Lemma 3.11. *Suppose α crosses a leaf ω of λ . There exists a constant $C > 0$ such that if $\ell_\alpha(s) \geq \epsilon_0$ and $d_\alpha(X_s, \lambda) \geq C$, then $\ell_\alpha(t) \stackrel{*}{\asymp} e^{t-s} \ell_\alpha(s)$ for all $t \geq s$.*

Proof. By Lemma 3.8, if $d_\alpha(X_s, \lambda) > \frac{4\epsilon_M}{\ell_\alpha(s)} + 2$, then

$$v_s \leq 2e^{-\frac{\ell_\alpha(s)}{4} d_\alpha(X_s, \lambda)}.$$




Figure 7: A Teichmüller geodesic (blue) and a stretch path (red) in the Teichmüller space $\mathcal{T}(S_{1,1}) \simeq \mathbf{H}$ of the punctured torus. Both geodesic segments start at $X = -16.302 + i$ and end at $Y = i$, and each has its midpoint marked.

Since $\ell_\alpha(s) \geq \epsilon_0$ in this case, there is a universal constant C so that this estimate applies when $d_\alpha(X_s, \lambda) > C$. Furthermore, we can choose C so that the inequality above gives

$$v_s \leq \frac{1}{3} \ell_\alpha(s)$$

and so $h_s \geq \frac{2}{3} \ell_\alpha(s)$ and $2v_s \leq h_s$. Incorporating Lemma 3.4, we have that for all $t > s$,

$$\ell_\alpha(t) \geq h_t \geq \frac{1}{2} e^{t-s} h_s \geq \frac{1}{3} e^{t-s} \ell_\alpha(s).$$

On the other hand, $\ell_\alpha(t) \leq e^{t-s} \ell_\alpha(s)$. This finishes the proof. \square

Lemma 3.12. *Suppose α interacts with λ . If $\ell_\alpha(t) \geq \epsilon_0$ for all $t \in [a, b]$, then $d_\alpha(X_a, X_b) \not\leq 0$.*

Proof. We first show that for any $s \leq t$, if $\ell_\alpha(t) \not\leq e^{t-s} \ell_\alpha(s)$, then $d_\alpha(X_s, X_t) \not\leq 0$. Let β be the shortest curve at X_s that intersects α . At time t , the length of β satisfies

$$\ell_\beta(t) \geq \ell_\alpha(t) d_\alpha(\beta, X_t) - D \ell_\alpha(t),$$

where $D \geq 0$ is universal. Also, $d_\alpha(\beta, X_s)$ is bounded by the choice of β . Hence we can write

$$\frac{\ell_\beta(t)}{\ell_\alpha(t)} \not\leq d_\alpha(X_s, X_t).$$

Therefore, since $\ell_\beta(t) \leq e^{t-s} \ell_\beta(s)$ and $\ell_\alpha(t) \not\leq e^{t-s} \ell_\alpha(s)$, we have

$$d_\alpha(X_s, X_t) \not\leq \frac{\ell_\beta(s)}{\ell_\alpha(s)}.$$

Let ϵ_B be the Bers constant. If $\ell_\alpha(s) \geq \epsilon_B$, then $\ell_\beta(t) \leq \epsilon_B$ and so $\frac{\ell_\beta(s)}{\ell_\alpha(s)} < 1$. If $\ell_\alpha(s) \leq \epsilon_B$, then $\ell_\beta(s)$ is up to a bounded multiplicative error the width of the collar about α , So in this case, since $\ell_\alpha(s) \geq \epsilon_0$, we have

$$\frac{\ell_\beta(s)}{\ell_\alpha(s)} \stackrel{*}{\asymp} \frac{1}{\epsilon_0} \log \frac{1}{\epsilon_0} \stackrel{*}{\asymp} 1.$$

If α is a leaf of λ , then $\ell_\alpha(b) = e^{b-a}\ell_\alpha(a)$, so the conclusion follows from the paragraph above. Now suppose that α crosses a leaf of λ . Let C be the constant of Lemma 3.11. If $d_\alpha(X_t, \lambda) < C$ for all $t \in [a, b]$, then we are done. Otherwise, there is an earliest time $t \in [a, b]$ such that $d_\alpha(X_t, \lambda) \geq C$. It is immediate that $d_\alpha(X, X_t) \stackrel{\pm}{\asymp} 0$. By Lemma 3.11, $\ell_\alpha(b) \stackrel{*}{\asymp} e^{b-t}\ell_\alpha(t)$, so $d_\alpha(X_t, X_b) \stackrel{\pm}{\asymp} 0$ by the above paragraph. The result follows. \square

We will now prove the theorem stated at the beginning of this section.

Proof of Theorem 3.1. If $\ell_\alpha \geq \epsilon_0$, then by Lemma 3.12,

$$d_\alpha(X, Y) \stackrel{\pm}{\asymp} 0 \stackrel{\pm}{\asymp} \frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha}.$$

Now suppose $\ell_\alpha < \epsilon_0$ and let $[a, b]$ be the active interval for α . From Theorem 3.10, the minimal length ℓ_α occurs at $t_\alpha \in [a, b]$ satisfying $t_\alpha - a \stackrel{\pm}{\asymp} \text{Log}(b - t_\alpha)$, and $\ell_\alpha \stackrel{*}{\asymp} e^{-(b-t_\alpha)}$. We then have

$$\begin{aligned} d_\alpha(X_a, X_b) &\stackrel{*}{\asymp} e^{b-a} = e^{b-t_\alpha} e^{t_\alpha-a} \\ &\stackrel{*}{\asymp} e^{b-t_\alpha} e^{\text{Log}(b-t_\alpha)} \end{aligned}$$

If $(b-a)$ is large enough so that Theorem 3.10 gives $\text{Log}(b-t_\alpha) = \log(b-t_\alpha)$, then this shows $d_\alpha(X_a, X_b) \stackrel{*}{\asymp} e^{b-t_\alpha} (b-t_\alpha) \stackrel{*}{\asymp} \frac{1}{\ell_\alpha} \log(\frac{1}{\ell_\alpha})$, and since $\ell_\alpha \leq \epsilon_0$, we have $\frac{1}{\ell_\alpha} \log \frac{1}{\ell_\alpha} \stackrel{\pm}{\asymp} \frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha}$ with equality for ℓ_α small enough. By Lemma 3.12, $d_\alpha(X, X_a)$ and $d_\alpha(X_b, Y)$ are both uniformly bounded. Thus $d_\alpha(X, Y) \stackrel{\pm}{\asymp} d_\alpha(X_a, X_b)$ and the estimate on $d_\alpha(X, Y)$ from Theorem 3.10 follows in this case.

Otherwise, $(b-a)$ is bounded above by a universal constant, in which case we will show $d_\alpha(X, Y) \stackrel{\pm}{\asymp} \frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha}$ by showing that both sides are uniformly bounded. First, the upper bound on $(b-a)$ gives a positive lower bound on ℓ_α (which is already bounded above by ϵ_0) and so $\frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha} \stackrel{\pm}{\asymp} 0$. On the other hand, using the bound on $(b-a)$, Theorem 3.10 gives $d_\alpha(X_a, X_b) \stackrel{\pm}{\asymp} 0$, and as before $d_\alpha(X, Y) \stackrel{\pm}{\asymp} d_\alpha(X_a, X_b)$. We conclude $d_\alpha(X, Y) \stackrel{\pm}{\asymp} 0$, as required.

For the last statement of Theorem 3.1, let C be the constant of Lemma 3.11. By assumption $\ell_\alpha(t) > \epsilon_0$ for all $t \leq a$. If there exists $t \leq a$ such that $d_\alpha(X_t, \lambda) \geq C$, then $\ell_\alpha(t_\alpha) \stackrel{*}{\asymp} e^{t_\alpha-t}\ell_\alpha(t)$, where t_α is the time of the minimal length of ℓ_α . This is impossible for all sufficiently small ℓ_α . Finally, since $d_\alpha(X_a, X_b) \stackrel{*}{\asymp} \frac{1}{\ell_\alpha} \text{Log} \frac{1}{\ell_\alpha}$, for all sufficiently small ℓ_α , we can guarantee that $d_\alpha(X_b, \lambda) \geq C$. The final conclusion follows by Lemma 3.11. \square

Recall that two curves that intersect cannot both have lengths less than ϵ_M at the same time. Therefore, if α and β intersect and $\ell_\alpha < \epsilon_0$ and $\ell_\beta < \epsilon_0$, then their active intervals must be disjoint. This defines an ordering of α and β along \mathcal{G} . In the next section, we will focus on the torus $S_{1,1}$ and show that the order of α and β along \mathcal{G} will always agree with their order in the projection of $\mathcal{G}(t)$ to the Farey graph.




Figure 8: *The Farey graph.*

4 Coarse description of geodesics in $\mathcal{T}(S_{1,1})$

4.1 Farey graph

See [Min99] for background on the Farey graph.

Let $S_{1,1}$ be the once-punctured torus and represent its universal cover by the hyperbolic plane \mathbf{H} . Identify the ideal boundary $\partial\mathbf{H}$ with $\mathbf{R} \cup \{\infty\}$. The point ∞ is considered an extended rational number with reduced form $1/0$. As in Section 2.2, fix a positive ordered basis for $H_1(S_{1,1})$ and use this to associate a slope $p/q \in \mathbf{QP}^1 = \mathbf{Q} \cup \{\infty\}$ to every simple curve. In this section we pass freely between a rational number and the associated simple curve.

Given two curves $\alpha = p/q$ and $\beta = r/s$ in reduced fractions, their geometric intersection number is $|ps - rq|$. Form a graph with vertex set \mathbf{QP}^1 as follows: Connect p/q and r/s by an edge if $|ps - rq| = 1$. The resulting graph \mathcal{F} is called the *Farey graph*, which is also the curve graph of $S_{1,1}$. This graph embeds naturally in $\mathbf{H} \cup \partial\mathbf{H}$, with its edges realized as hyperbolic geodesics (see Figure 8). These geodesics cut \mathbf{H} into ideal triangles; this is the *Farey tessellation*. In this tessellation, each edge bounds exactly two ideal triangles with zero relative shearing. Thus each edge of \mathcal{F} is equipped with a well-defined midpoint.

Let α denote the curve with slope $1/0$. The action of D_α on curves distinct from α corresponds to the mapping of slopes $m \mapsto m+1$. Let $\beta_0 \in \mathcal{F}$ be any curve with $i(\alpha, \beta_0) = 1$. The associated Dehn twist family about α is $\beta_n = D_\alpha^n(\beta_0)$. Then $\{\beta_n\}_{n \in \mathbf{Z}}$ is exactly the set of vertices of \mathcal{F} that are connected to α by an edge, or equivalently, the set of curves with slope in \mathbf{Z} .

4.2 Markings and pivots

A *marking* on $S_{1,1}$ is an unordered pair of curves $\{\alpha, \beta\}$ such that $i(\alpha, \beta) = 1$. Given a marking $\{\alpha, \beta\}$, there are four markings that are obtained from μ by an *elementary move*,

namely:

$$\{\alpha, D_\alpha(\beta)\}, \quad \{\alpha, D_\alpha^{-1}(\beta)\}, \quad \{\beta, D_\beta(\alpha)\}, \quad \{\beta, D_\beta^{-1}(\alpha)\}.$$

Note that the set of markings on $S_{1,1}$ can be identified with the set of edges of \mathcal{F} , and two edges differ by an elementary move if and only if they bound a common triangle in the Farey tessellation of \mathbf{H} . Denote by \mathcal{MG} the graph with markings as vertices and an edge connecting two markings that differ by an elementary move. Then \mathcal{MG} has the following property.

Lemma 4.1. *For any $\mu, \mu' \in \mathcal{MG}$, there exists a unique geodesic connecting them.*

Proof. Each edge of \mathcal{F} separates \mathbf{H} into two disjoint half-spaces. Let $E(\mu, \mu')$ be the set of edges in \mathcal{F} that separate the interior of μ from the interior of μ' . Set $\overline{E}(\mu, \mu') = E(\mu, \mu') \cup \{\mu, \mu'\}$. Every $\nu \in E(\mu, \mu')$ disconnects μ from μ' , and thus must appear in every geodesic from μ and μ' . Conversely, any $\nu \in \mathcal{MG}$ lying on a geodesic from μ to μ' must lie in $\overline{E}(\mu, \mu')$. For each $\nu \in \overline{E}(\mu, \mu')$, let H_ν be the half-space in \mathbf{H} containing the interior of μ' . There is a linear order on $\overline{E}(\mu, \mu') = \{\mu_1 < \mu_2 < \dots < \mu_n\}$ induced by the relation $\mu_i < \mu_{i+1}$ if and only if $H_{\mu_i} \supset H_{\mu_{i+1}}$. The sequence $\mu = \mu_1, \mu_2, \dots, \mu_n = \mu'$ is the unique geodesic path in \mathcal{MG} from μ to μ' . \square

Given two markings μ and μ' and a curve α , let n_α be the number of edges in $\overline{E}(\mu, \mu')$ containing α . We say α is a *pivot* for μ and μ' if $n_\alpha \geq 2$, and n_α is the *coefficient* of the pivot. Let $\text{Pivot}(\mu, \mu')$ be the set of pivots for μ and μ' . This set is naturally linearly ordered as follows. Given $\alpha \in \text{Pivot}(\mu, \mu')$, let e_α be the last edge in $\overline{E}(\mu, \mu')$ containing α . Then for $\alpha, \beta \in \text{Pivot}(\mu, \mu')$, we set $\alpha < \beta$ if e_α appears before e_β in g .

Recall that in Section 2.8 we defined the unsigned twisting (along α) for a pair of curves β, β' ; this is denoted $d_\alpha(\beta, \beta')$. Generalizing this, we define unsigned twisting for the pair of markings μ, μ' by

$$d_\alpha(\mu, \mu') = \min_{\beta \subset \mu, \beta' \subset \mu'} d_\alpha(\beta, \beta'),$$

where β is a curve in μ and β' is a curve in μ' . Similarly we define $d_\alpha(\beta, \mu') = \min_{\beta' \subset \mu'} d_\alpha(\beta, \beta')$. In terms of these definitions, we have:

Lemma 4.2 ([Min99]). *For any $\mu, \mu' \in \mathcal{MG}$ and curve α , we have $n_\alpha \stackrel{\pm}{\geq} d_\alpha(\mu, \mu')$. For $\alpha, \beta \in \text{Pivot}(\mu, \mu')$, if $\alpha < \beta$, then $d_\alpha(\beta, \mu') \stackrel{\pm}{\geq} 1$ and $d_\beta(\mu, \alpha) \stackrel{\pm}{\geq} 1$. Conversely, if n_α is sufficiently large and $d_\alpha(\beta, \mu') \stackrel{\pm}{\geq} 1$, then $\alpha < \beta$.*

Identify $\mathcal{T}(S_{1,1})$ with \mathbf{H} in the usual way. Under this identification, if e is an edge of \mathcal{F} with endpoints α and β , then the set points along e correspond to the set of surfaces on which α and β are the shortest curves and they intersect perpendicularly. The midpoint of e correspond to when the two curves have the same length. This length is a uniform constant independent of the edge e .

For any $X \in \mathcal{T}(S_{1,1})$, there exists an ideal triangle Δ in the Farey tessellation of \mathbf{H} containing X . The three vertices of Δ correspond to the three shortest curves on X . We will define a *short marking* on X as follows. If X has at least 2 systoles, then let A be the set of systoles on X . If X has a unique systole, then let A be set consisting of the systole plus the second shortest curves on X . In either case, A is a subset of the vertices of Δ , so A has cardinality

at most 3 and every pair of curves in A correspond to an edge in Δ . A short marking on X is any pair of curves in A . Note that in our definition, there is either a unique marking or three short markings on X . This implies that, given $X, Y \in \mathcal{T}(S_{1,1})$, there are well-defined short markings μ_X and μ_Y on X and Y such that $d_{\text{Mg}}(\mu_X, \mu_Y)$ is minimal among all short markings on X and Y . By Lemma 4.1, the geodesic from μ_X to μ_Y is unique. Note that any edge of $E(\mu_X, \mu_Y)$ separates μ_X from μ_Y , and hence it separates X from Y . We will denote by $\text{Pivot}(X, Y) = \text{Pivot}(\mu_X, \mu_Y)$ and refer to $\text{Pivot}(X, Y)$ as the set of pivots for X and Y .

Given $X, Y \in \mathcal{T}(S_{1,1})$, we have that $d_\alpha(X, Y) \stackrel{\pm}{\asymp} d_\alpha(\mu_X, \mu_Y)$.

Let ϵ_0 be the constant of the previous section. The following statements establish Theorem 1.3 of the introduction.

Theorem 4.3. *Suppose $X, Y \in \mathcal{T}_{\epsilon_0}(S_{1,1})$ and let $\mathcal{G}(t)$ be any geodesic from X to Y , parameterized by an interval $I \subset \mathbf{R}$. Let $\ell_\alpha = \inf_t \ell_\alpha(t)$. There are positive constants ϵ_1 , C_1 , and C_2 such that*

- (i) *If $\ell_\alpha \leq \epsilon_1$, then $\alpha \in \text{Pivot}(X, Y)$ and $d_\alpha(X, Y) \geq C_1$.*
- (ii) *If $d_\alpha(X, Y) \geq C_2$, then $\ell_\alpha \leq \epsilon_1$ and $\alpha \in \text{Pivot}(X, Y)$.*
- (iii) *Suppose α and β are distinct curves such that there exist $s, t \in I$ with $\ell_\alpha(s) \leq \epsilon_1$ and $\ell_\beta(t) \leq \epsilon_1$. Then $\alpha < \beta$ in $\text{Pivot}(X, Y)$ if and only if $s < t$.*
- (iv) *For any $\alpha \in \text{Pivot}(X, Y)$, $\ell_\alpha \stackrel{*}{\asymp} 1$.*

Proof. The proof will show that any sufficiently small ϵ_1 works. We first require $\epsilon_1 < \epsilon_0$ where ϵ_0 is the constant selected in the previous section.

Let $\lambda = \Lambda(X, Y)$. On the torus, every curve α interacts with λ . If λ contains α , then $\ell_\alpha(t) = e^t \ell_\alpha(X)$. But this implies $\ell_\alpha = \ell_X(\alpha) \geq \epsilon_0$ and $d_\alpha(X, Y) \stackrel{\pm}{\asymp} 0$ by Lemma 3.12. Thus, we may assume that α crosses a leaf of λ . By Theorem 3.1, $d_\alpha(X, Y) \stackrel{\pm}{\asymp} \frac{1}{\ell_\alpha} \log \frac{1}{\ell_\alpha}$. Since $d_\alpha(X, Y) \stackrel{\pm}{\asymp} d_\alpha(\mu_X, \mu_Y) \stackrel{\pm}{\asymp} n_\alpha$ (the latter by Lemma 4.2), we can select ϵ_1 small enough and $C_1 > 0$ so that $\ell_\alpha \leq \epsilon_1$ implies that $d_\alpha(X, Y) \geq C_1$ and that $n_\alpha \geq 2$, i.e. α is a pivot. This gives (i). Using the same approximate equalities, if $d_\alpha(X, Y)$ is large we find that ℓ_α is small, and we can select C_2 satisfying (ii).

We now fix our constants ϵ_1 , C_1 and C_2 so (i) and (ii) are satisfied. By fixing these constants, we can now ignore the dependence of any additive or multiplicative errors on them.

For (iii), suppose $\ell_\alpha \leq \epsilon_1$ and $\ell_\beta \leq \epsilon_1$. By (i), they are both pivots. Let $[a, b]$ be the active interval for α . Recall that this is the longest interval such that $\ell_\alpha(a) = \ell_\beta(b) = \epsilon_0$. Recall that by Corollary 3.7 we have $\ell_\alpha(t) < \epsilon_M$ for all $t \in [a, b]$. Similarly, let $[c, d]$ be the active interval for β . On the torus, two curves always intersect, so α and β cannot be simultaneously shorter than ϵ_M , so $[a, b]$ and $[c, d]$ must be disjoint. By Lemma 4.2, $\alpha < \beta$ if and only if $d_\alpha(\beta, \mu_Y) \stackrel{\pm}{\asymp} 1$. By Theorem 3.1 and Lemma 3.12, $b < c$ if and only if $d_\alpha(X_c, Y) \stackrel{\pm}{\asymp} 1$. Since β is ϵ_0 -short on X_c , we have $d_\alpha(\beta, \mu_Y) \stackrel{\pm}{\asymp} d_\alpha(X_c, Y)$. This finishes (iii).

Before we prove (iv), we introduce some notation. For each curve α , let $H_\alpha \subset \mathcal{T}(S_{1,1})$ be the set of hyperbolic structures where $\ell_\alpha(X) \leq \epsilon_1$. Since $\epsilon_1 < \epsilon_M$, the sets H_α and H_β are

disjoint if $\alpha \neq \beta$. Let e be an edge of \mathcal{F} and denote its endpoints by α and β . The segment of e outside of H_α and H_β is a closed interval containing the midpoint of e . Along this interval, the length of α and β is uniformly bounded (by a constant that depends only on ϵ_1).

To prove (iv), let $\alpha \in \text{Pivot}(X, Y)$ and assume $\ell_\alpha > \epsilon_1$. Let $e \in E(\mu_X, \mu_Y)$ be an edge containing α . Let β be the other curve of e . The edge e separates X and Y , so any geodesic $\mathcal{G}(t)$ from X to Y must cross e at some point X_t . If $\ell_\beta(t) > \epsilon_1$, then neither α or β is ϵ_1 -short on X_t , so X_t lies in the segment of e outside of H_α and H_β . Hence $\ell_\alpha(t) \stackrel{*}{\asymp} 1$ by the discussion in the previous paragraph. On the other hand, if $\ell_\beta(t) \leq \epsilon_1$, then β is a pivot by (i). Either $\alpha < \beta$ or $\beta < \alpha$ in $\text{Pivot}(X, Y)$. If $\alpha < \beta$, then $d_\beta(X, \alpha) \stackrel{\pm}{\asymp} 1$ by Lemma 4.2. Let $[a, b]$ be the active interval for β . By Theorem 3.1 we have $d_\beta(X, X_a) \stackrel{\pm}{\asymp} 1$, and d_β satisfies the triangle inequality up to additive error (by [MM00, Equation 2.5]), so we conclude $d_\beta(\alpha, X_a) \stackrel{\pm}{\asymp} 1$. This, together with $\ell_\beta(a) = \epsilon_0$, yields $\ell_\alpha(a) \stackrel{*}{\asymp} 1$. If $\beta < \alpha$, then the same argument using X_b and Y in place of X_a and X also yields $\ell_\alpha(b) \stackrel{*}{\asymp} 1$. This concludes the proof. \square

5 Envelopes in $\mathcal{T}(S_{1,1})$

5.1 Fenchel-Nielsen coordinates along stretch paths in $\mathcal{T}(S_{1,1})$

We now focus on the once-punctured torus $S_{1,1}$, and on the completions α^\pm of the maximal chain-recurrent laminations containing a simple closed curve α discussed in Section 2.2.

Consider the curve α as a pants decomposition of $S_{1,1}$ and define $\tau_\alpha(X)$ to be the Fenchel-Nielsen twist coordinate of X relative to α . Note that $\tau_\alpha(X)$ is well defined up to a multiple of $\ell_\alpha(X)$, and after making a choice at some point, $\tau_\alpha(X)$ is well defined. The Fenchel-Nielsen theorem states that the pair of functions $(\log \ell_\alpha(\cdot), \tau_\alpha(\cdot))$ define a diffeomorphism of $\mathcal{T}(S_{1,1}) \rightarrow \mathbf{R}^2$.

Each α^\pm defines a foliation \mathcal{F}_α^\pm on $\mathcal{T}(S_{1,1})$ whose leaves are the α^\pm -stretch paths. In the α^\pm shearing coordinate system, the image of $\mathcal{T}(S_{1,1})$ in \mathbf{R}^2 is a convex cone, and the foliation \mathcal{F}_α^\pm are by open rays from the origin.

In this section we denote a point on the α^\pm stretch path through X by $X_t^\pm = \text{stretch}(X, \alpha^\pm, t)$. The function $\log \ell_\alpha(X_t^\pm) = \log \ell_\alpha(X) + t$ is smooth in t . Our first goal is to establish the following theorem.

Theorem 5.1. *For any simple closed curve α on $S_{1,1}$ and any point $X = X_0 \in \mathcal{T}(S_{1,1})$, the functions $\tau_\alpha(X_t^\pm)$ are smooth in t . Further,*

$$\tau_\alpha(X_t^+) > \tau_\alpha(X_t^-) \quad \text{and} \quad \frac{d}{dt} \tau_\alpha(X_t^+) \Big|_{t=0} > \frac{d}{dt} \tau_\alpha(X_t^-) \Big|_{t=0}.$$

That is, the pair of foliations \mathcal{F}_α^+ and \mathcal{F}_α^- are smooth and transverse.

We proceed to prove smoothness of $\tau_\alpha(X_t^+)$. Recall that the α^+ shearing embedding is $s_{\alpha^+}(X) = (\ell_\alpha(X), s_\alpha(X))$ where $s_\alpha(X)$ was defined in Section 2.5, and that like τ_α , the function s_α is defined only up to adding an integer multiple of $\ell_\alpha(X)$. To further lighten our notation, we will often write $\ell_\alpha(t)$ instead of $\ell_\alpha(X_t^+)$, and $s_\alpha(t)$ for $s_\alpha(X_t^+)$.




Figure 9: Computing the Fenchel-Nielsen twist along the α^+ stretch path.

We also denote $\tau_\alpha(0)$ and $\ell_\alpha(0)$ by τ_0 and ℓ_0 respectively. Note that the values of τ_α and ℓ_α do not depend on the choice of α^+ or α^- but the values of the shearing coordinates do.

We know, from the description of stretch paths in Section 2.7, that

$$s_\alpha(t) = s_\alpha(0)e^t \quad \text{and} \quad \ell_\alpha(t) = \ell_0 e^t.$$

We can now compute $\tau_\alpha(X_t^+)$ as follows, referring to Figure 9. Fix a lift $\tilde{\alpha}$ of α to be the imaginary line (shown in blue in Figure 9) in the upper half-plane \mathbf{H} . Now develop the picture on both sides of $\tilde{\alpha}$. Since we are considering α^+ , all the triangles on the left of $\tilde{\alpha}$ are asymptotic to ∞ and all the triangles on the right of $\tilde{\alpha}$ are asymptotic to 0. Below we will choose some normalization, but first note that the hyperelliptic involution exchanges the two complementary triangles T and T' of α^+ while preserving α as a set. Let $\iota: \mathbf{H} \rightarrow \mathbf{H}$ be a lift of this involution chosen to preserve $\tilde{\alpha}$, which therefore has the form $\iota(z) = -e^c/z$ for some $c = c_t \in \mathbf{R}$. Notice that ι exchanges the two sides of $\tilde{\alpha}$ and that it fixes a unique point $ie^{c/2}$ in \mathbf{H} .

To fix the shearing coordinate $s_\alpha(t)$, we make the choice of triangles in \mathbf{H} required by the construction of Section 2.4. Choose two triangles Δ_l and Δ_r in \mathbf{H} separated by $\tilde{\alpha}$ so that one is a lift of T and the other is a lift of T' and $\iota(\Delta_r) = \Delta_l$. Let \tilde{w} be the edge of Δ_l that is a lift of w , namely, $\tilde{w} = [x, \infty]$ for some $x < 0$. Let $\phi_\alpha(z) = e^{\ell_\alpha} z$ be the isometry associated to $\tilde{\alpha}$ oriented toward ∞ . The image $\phi_\alpha(\tilde{w}) = [e^{\ell_\alpha} x, \infty]$ is another lift of w . Let $\tilde{\delta}$ be the lift of δ that is asymptotic to \tilde{w} and $\phi_\alpha(\tilde{w})$. By applying a further dilation to the picture if necessary, we can assume that $\tilde{\delta} = [x-1, \infty]$. Now, the geodesic $\tilde{w}' = [x, x-1]$ is a lift of w' .

With our normalization, the midpoint of $[x, \infty]$ associated to Δ_l is the point $(x, 1)$. Recall

that $s_\delta = -\ell_\alpha$ in this case, which is the shearing between triangles $\Delta_l = [x, x-1, \infty]$ and $[x-1, e^{\ell_\alpha}x, \infty]$. This means that their midpoints on $\tilde{\delta}$ have y -coordinates with ratio e^{ℓ_α} , i.e.

$$\frac{|e^{\ell_\alpha}x - (x-1)|}{|(x-1) - x|} = e^{\ell_\alpha}$$

from which it follows that

$$x = -\coth(\ell_\alpha/2).$$

Let h_l be the endpoint on $\tilde{\alpha}$ of the horocycle based at infinity containing the midpoint of \tilde{w} considered as an edge of Δ_l . Let $h_r = \iota(h_l)$. By construction, $h_l = 1$ and $h_r = e^c$. We can normalize so that $s_\alpha = c$.

To visualize the Fenchel-Nielsen twist parameter $\tau_+(t)$ at X_t about α , consider the shortest geodesic arc β with both endpoints on α intersecting perpendicularly (so β only intersects α twice). By symmetry, this arc intersects δ at a point q that is equidistant to the midpoints of δ associated to T and T' . We choose a lift $\tilde{\beta}$ that passes through $\tilde{q} = (x-1, e^{\ell_\alpha/2})$. Let p_l be the endpoint on $\tilde{\alpha}$ of the lift of β that passes through $\tilde{\delta}$. Since $\tilde{\beta}$ is perpendicular to $\tilde{\alpha}$, we have q and p_l lie on a Euclidean circle centered at the origin. Using the Pythagorean theorem, we obtain:

$$p_l = \sqrt{(x-1)^2 + (e^{\ell_\alpha/2})^2} = e^{\ell_\alpha/2} \coth \frac{\ell_\alpha}{2}.$$

Let $ip_r = \iota(ip_l) = -e^c/ip_l = ie^c/p_l$. Up to an integral multiple of ℓ_α , the twisting $\tau_\alpha(t)$ is the signed distance between ip_l and ip_r ; that is

$$\begin{aligned} \tau_\alpha &= \log \frac{p_r}{p_l} \pmod{\ell_\alpha} \\ &= \log \frac{e^c}{e^{\ell_\alpha} \coth^2(\frac{\ell_\alpha}{2})} \pmod{\ell_\alpha} \\ &= c - 2 \log \coth \frac{\ell_\alpha}{2} \pmod{\ell_\alpha} \\ &= s_\alpha - 2 \log \coth \frac{\ell_\alpha}{2} \pmod{\ell_\alpha} \end{aligned}$$

In particular, at $t = 0$, we obtain

$$\tau_0 = s_\alpha(0) - 2 \log \coth \frac{\ell_0}{2} \pmod{\ell_0}. \quad (20)$$

As we mentioned previously, $s_\alpha(t) = s_\alpha(0)e^t$ and $\ell_\alpha(t) = \ell_0e^t$. Hence,

$$\tau_\alpha(t) = e^t s_\alpha(0) - 2 \log \coth \frac{e^t \ell_0}{2} \pmod{\ell_\alpha}.$$

Solving for τ_0 using (20), we obtain

$$\tau_\alpha(X_t^+) = e^t \tau_0 + 2e^t \log \coth \frac{\ell_0}{2} - 2 \log \coth \frac{e^t \ell_0}{2} \pmod{\ell_\alpha}. \quad (21)$$

Now let X_t^- be the stretch path starting from X associated to α^- . The computation in this case is similar; in fact $2\tau = s_\alpha^+ + s_\alpha^- \pmod{\ell_\alpha}$. Thus

$$\tau_\alpha(X_t^-) = e^t \tau_0 - 2e^t \log \coth \frac{\ell_0}{2} + 2 \log \coth \frac{e^t \ell_0}{2} \pmod{\ell_\alpha}. \quad (22)$$

This shows that $\tau(X_t^+)$ and $\tau(X_t^-)$ are both smooth functions of t . Note that $\tau(X_t^+) - \tau(X_t^-)$ is well-defined. By a simple computation, we see that $\tau(X_t^+) - \tau(X_t^-) > 0$, and

$$\left. \frac{d}{dt} (\tau_\alpha(X_t^+) - \tau_\alpha(X_t^-)) \right|_{t=0} = 4 \log \coth \frac{\ell_0}{2} + 2\ell_0 \tanh \frac{\ell_0}{2} \operatorname{csch}^2 \frac{\ell_0}{2} > 0.$$

This finishes the proof of Theorem 5.1.

5.2 Structure of envelopes in general

For any surface S of finite type and a chain-recurrent lamination λ on S and $X \in \mathcal{T}(S)$, define

$$\operatorname{Out}(X, \lambda) = \{Z \in \mathcal{T}(S) : \lambda = \Lambda(X, Z)\} \quad \text{and} \quad \operatorname{In}(X, \lambda) = \{Z \in \mathcal{T}(S) : \lambda = \Lambda(Z, X)\}.$$

We call these the *out-envelope* and *in-envelope* of X (respectively) in the direction λ .

Proposition 5.2. *The out-envelopes and in-envelopes have the following properties.*

- (i) *If λ is maximal chain-recurrent, then for any completion $\widehat{\lambda}$ of λ , the set $\operatorname{Out}(X, \lambda)$ is the stretch ray starting at X associated to $\widehat{\lambda}$, and the set $\operatorname{In}(X, \lambda)$ is the stretch ray associated to $\widehat{\lambda}$ ending at X .*
- (ii) *The closure of $\operatorname{Out}(X, \lambda)$ consists of points Y with $\lambda \subset \Lambda(X, Y)$. Similarly, the closure of $\operatorname{In}(X, \lambda)$ is the set of points Y with $\lambda \subset \Lambda(X, Y)$.*
- (iii) *If λ is a simple closed curve, then $\operatorname{Out}(X, \lambda)$ and $\operatorname{In}(X, \lambda)$ are open sets.*

Proof. First assume λ is maximal chain-recurrent and let $\widehat{\lambda}$ be a completion of it. By Corollary 2.3, if $\Lambda(X, Y) = \lambda$, then there exists $t > 0$ such that $Y = \operatorname{stretch}(X, \widehat{\lambda}, t)$ and this is the only geodesic from X to Y . That is, any point in $\operatorname{Out}(X, \lambda)$ can be reached from X by following the stretch ray along $\widehat{\lambda}$ starting at X . Similarly if $Y \in \operatorname{In}(X, \lambda)$, then the stretch ray along $\widehat{\lambda}$ starting at Y contains X , or equivalently, the stretch ray along $\widehat{\lambda}$ ending at X contains Y . This is (i).

For the other statements, we use [Thu86c, Theorem 8.4], which shows that if Y_i converges Y , then any limit point of $\Lambda(X, Y_i)$ in the Hausdorff topology is contained in $\Lambda(X, Y)$. Applying this to a point Y in the closure of $\operatorname{Out}(X, \lambda)$, and a sequence $Y_i \in \operatorname{Out}(X, \lambda)$ converging to Y , we obtain $\lambda \subset \Lambda(X, Y)$. For the other direction of (ii), let Y be any point such that $\lambda \subset \Lambda(X, Y)$. To show Y is in the closure of $\operatorname{Out}(X, \lambda)$, we find a point $Z \in \operatorname{Out}(X, \lambda)$ such that $d_{\operatorname{Th}}(Y, Z) = \epsilon$, for any ϵ . Let λ' be any maximal chain-recurrent lamination such that $\lambda = \lambda' \cap \Lambda(X, Y)$, and let $Z = \operatorname{stretch}(Y, \lambda', \epsilon)$. We have $d_{\operatorname{Th}}(Y, Z) = \epsilon$. Since $\lambda = \lambda' \cap \Lambda(X, Y)$, we must have $\Lambda(X, Z) = \lambda$. This shows (ii) for $\operatorname{Out}(X, \lambda)$. The analogous statement for $\operatorname{In}(X, \lambda)$ is proven similarly.

To obtain (iii), let λ be a simple closed curve, $Y \in \operatorname{Out}(X, \lambda)$, and Y_i is any sequence converging to Y , then any limit point of $\Lambda(X, Y_i)$ is contained in λ . Since λ is a simple closed curve, $\Lambda(X, Y_i) = \lambda$ for all sufficiently large i . This shows $\operatorname{Out}(X, \lambda)$ is open. The same proof also applies to $\operatorname{In}(X, \lambda)$. \square

Let $X, Y \in \mathcal{T}(S)$, and denote $\lambda = \Lambda(X, Y)$. We define the envelope of geodesics from X to Y to be the set

$$\text{Env}(X, Y) = \left\{ Z : Z \in [X, Y] \text{ for some geodesic } [X, Y] \right\}.$$

Proposition 5.3. *For any $X, Y \in \mathcal{T}(S)$, $\text{Env}(X, Y) = \overline{\text{Out}(X, \lambda)} \cap \overline{\text{In}(Y, \lambda)}$.*

Proof. For any $Z \in \text{Env}(X, Y)$, since Z lies on a geodesic from X to Y , λ must be contained in $\Lambda(X, Z)$ and in $\Lambda(Z, Y)$. This shows $\text{Env}(X, Y) \subset \overline{\text{Out}(X, \lambda)} \cap \overline{\text{In}(Y, \lambda)}$. On the other hand, if $Z \in \overline{\text{Out}(X, \lambda)} \cap \overline{\text{In}(Y, \lambda)}$, then $\lambda \subset \Lambda(X, Z)$ and $\lambda \subset \Lambda(Z, Y)$. That is, if μ is the stump of λ , then $d_{\text{Th}}(X, Z) = \log \frac{\ell_\mu(Z)}{\ell_\mu(X)}$ and $d_{\text{Th}}(Z, Y) = \log \frac{\ell_\mu(Y)}{\ell_\mu(Z)}$, so $d_{\text{Th}}(X, Y) = d_{\text{Th}}(X, Z) + d_{\text{Th}}(Z, Y)$. Thus, the concatenation of any geodesic from X to Z and from Z to Y is a geodesic from X to Y . \square

5.3 Structure of envelopes in $\mathcal{T}(S_{1,1})$

In this section, we specialize our study of envelopes to the case of $S = S_{1,1}$, and prove Theorem 1.1 of the introduction. The proof is divided into several propositions.

Proposition 5.4. *Let α be a simple closed curve on $S_{1,1}$. For any $X \in \mathcal{T}(S_{1,1})$, the set $\text{Out}(X, \alpha)$ is an open region bounded by the stretch rays along α^\pm starting at X . Similarly, $\text{In}(X, \alpha)$ is an open region bounded by the stretch rays along α^\pm ending at X .*

Proof. Set $X_t^\pm = \text{stretch}(X, \alpha^\pm, t)$. By Theorem 2.2, for any surface S and any two points $X, Y \in \mathcal{T}(S)$, Thurston constructed a geodesic from X to Y that is a concatenation of stretch paths, where the number of stretch paths needed in the concatenation is bounded by $2|\chi(S)|$, i.e. the number of triangles in an ideal triangulation of S . In our setting where $S = S_{1,1}$, for $Y \in \text{Out}(X, \alpha)$, this would be either a single stretch path or a union of two stretch paths $[X, Z]$ and $[Z, Y]$ where both $\Lambda(X, Z)$ and $\Lambda(Z, Y)$ contain α . By Corollary 2.3, each one of these is a stretch path along either α^+ or α^- . The initial path can be chosen to stretch along α^+ or α^- arbitrarily. Assuming we first stretch along α^- , then there are t_1 and t_2 such that $Z = \text{stretch}(X, \alpha^-, t_1)$, $Y = \text{stretch}(Z, \alpha^+, t_2)$, and $d_{\text{Th}}(X, Y) = t_1 + t_2$. Set $Z_t^- = \text{stretch}(Z, \alpha^-, t)$ and $Z_t^+ = \text{stretch}(Z, \alpha^+, t)$. By the calculations of the previous section, $\tau_\alpha(Z_t^-) < \tau_\alpha(Z_t^+)$. Since $Z_t^- = X_{t+t_1}^-$ and $Z_t^+ = Y$, we have

$$\ell_\alpha(X_{t_1+t_2}^-) = \ell_\alpha(Y) \quad \text{and} \quad \tau_\alpha(X_{t_1+t_2}^-) < \tau_\alpha(Y)$$

Similarly, if we stretch along α^+ first, then there are s_1 and s_2 such that $W = \text{stretch}(X, \alpha^+, s_1)$, $Y = \text{stretch}(W, \alpha^-, s_2)$, and $s_1 + s_2 = t_1 + t_2$. Then $X_{t_1+t_2}^+ = X_{s_1+s_2}^+$ and by the same argument as above

$$\ell_\alpha(X_{t_1+t_2}^+) = \ell_\alpha(Y) \quad \text{and} \quad \tau_\alpha(Y) < \tau_\alpha(X_{t_1+t_2}^+)$$

That is, Y is inside of the sector bounded by the stretch rays X_t^+ and X_t^- for $t > 0$. By replacing geodesics from X to Y by geodesics from Y to X , we obtain the statement for $\text{In}(X, \alpha)$. \square

Remark 5.5 (Visualization of envelopes). Figure 0 (on the title page) illustrates Proposition 5.4 by showing the sets $\text{In}(X, \alpha)$ in the Poincaré disk model of $\mathcal{T}(S_{1,1})$ for X the hexagonal

punctured torus and for several simple curves α , including the three systoles. In the figure, the disk model is normalized so that the origin corresponds to the hexagonal punctured torus. This figure was produced as follows: The Fenchel-Nielsen coordinate computations of (21)–(22) make it straightforward to compute stretch paths passing through a given point in the relative $\mathrm{SL}(2, \mathbf{R})$ character variety of $\pi_1(S_{1,1})$. The software package CP1 [Dum13] allows the computation of the uniformization map from the disk to the relative character variety; this map was numerically inverted using Newton’s method to transport the computed stretch paths to the disk.

By the results of [Thé07], the stretch lines appearing as boundaries of in-envelopes for $\mathcal{T}(S_{1,1})$ are exactly those which limit on rational points on the circle at infinity as $t \rightarrow -\infty$. Thus Figure 0 can be alternatively described as showing regions bounded by the pairs of stretch rays joining several rational points at infinity to the hexagonal punctured torus.

Corollary 5.6. *Given $X, Y \in \mathcal{T}(S_{1,1})$, if $\Lambda(X, Y)$ is a simple closed curve, then $\mathrm{Env}(X, Y)$ is a compact quadrilateral.*

Proof. The statement follows from Proposition 5.4 and the fact that $\mathrm{Env}(X, Y) = \overline{\mathrm{Out}(X, \alpha)} \cap \overline{\mathrm{In}(Y, \alpha)}$. \square

Proposition 5.7. *In $\mathcal{T}(S_{1,1})$, the set $\mathrm{Env}(X, Y)$ varies continuously as a function of X and Y with respect to the topology induced by the Hausdorff distance on closed sets.*

Proof. First suppose $\Lambda(X, Y)$ is a simple closed curve α . By [Thu86c, Theorem 8.4], if $X_i \rightarrow X$ and $Y_i \rightarrow Y$, then $\Lambda(X, Y)$ contains any limit point of $\Lambda(X_i, Y_i)$; thus $\Lambda(X_i, Y_i) = \alpha$ for all sufficiently large i . That is, for sufficiently large i , $\mathrm{Env}(X_i, Y_i)$ is a compact quadrilateral bounded by segments in the foliations \mathcal{F}_α^\pm . Let Z be the *left corner* of $\mathrm{Env}(X, Y)$, i.e. the intersection point of the leaf of \mathcal{F}_α^+ through X and the leaf of \mathcal{F}_α^- through Y . For any neighborhood U of Z , by smoothness and transversality of \mathcal{F}_α^\pm , there is a neighborhood U_X of X and a neighborhood U_Y of Y , such that for all sufficiently large i , $X_i \in U_X$, $Y_i \in U_Y$, and the leaf of \mathcal{F}_α^+ through X_i and the leaf of \mathcal{F}_α^- through Y_i will intersect in U . That is, for all sufficiently large i , the left corner of $\mathrm{Env}(X_i, Y_i)$ lies close to the left corner of $\mathrm{Env}(X, Y)$. A similar argument holds for the right corners. This shows $\mathrm{Env}(X_i, Y_i)$ converges to $\mathrm{Env}(X, Y)$.

Now suppose $\Lambda(X, Y) = \lambda$ is a maximal chain-recurrent lamination and $X_i \rightarrow X$ and $Y_i \rightarrow Y$. Let $\widehat{\lambda}$ be the canonical completion of λ , and let \mathcal{G} be the stretch path along $\widehat{\lambda}$ passing through X and Y . Also let \mathcal{G}_i and \mathcal{G}'_i be the stretch paths along $\widehat{\lambda}$ through X_i and Y_i respectively. Since stretch paths along $\widehat{\lambda}$ foliate $\mathcal{T}(S_{1,1})$, \mathcal{G}_i and \mathcal{G}'_i either coincide or are disjoint. In the backward direction, all stretch paths along $\widehat{\lambda}$ converge to λ (the stump of $\widehat{\lambda}$) in \mathcal{PML} [Pap91]. If they coincide, then $\Lambda(X_i, Y_i) = \lambda$ and $\mathrm{Env}(X_i, Y_i)$ is a segment of \mathcal{G}_i . If they are disjoint, then they divide $\mathcal{T}(S_{1,1})$ into three disjoint regions. Let M_i be the closure of the region bounded by $\mathcal{G}_i \cup \mathcal{G}'_i$; see Figure 10. In the case that $\mathcal{G}_i = \mathcal{G}'_i$, set $M_i = \mathcal{G}_i$. For any geodesic L from X_i to Y_i , since $X_i, Y_i \in M_i$, if L leaves M_i , then it must cross either \mathcal{G}_i or \mathcal{G}'_i at least twice. But two points on a stretch path cannot be connected by any other geodesic in the same direction, so L must be contained entirely in M_i . Therefore, $\mathrm{Env}(X_i, Y_i) \subset M_i$ (see Figure 10). Since \mathcal{G}_i and \mathcal{G}'_i converge to \mathcal{G} , M_i also converges to \mathcal{G} . Therefore $\mathrm{Env}(X_i, Y_i)$ converges to a subset of \mathcal{G} . For any $Z_i \in \mathrm{Env}(X_i, Y_i)$, $d_{\mathrm{Th}}(X_i, Z_i) + d_{\mathrm{Th}}(Z_i, Y_i) = d_{\mathrm{Th}}(X_i, Y_i)$, so by continuity of d_{Th} ,




Figure 10: $\text{Env}(X_i, Y_i)$ is sandwiched between \mathcal{G}_i and \mathcal{G}'_i .

Z_i must converge to a point $Z \in \mathcal{G}$ with $d_{\text{Th}}(X, Z) + d_{\text{Th}}(Z, Y) = d_{\text{Th}}(X, Y)$. In other words, Z lies on the geodesic from X to Y . This shows $\text{Env}(X_i, Y_i)$ converges to $\text{Env}(X, Y)$. \square

We can now assemble the proof of Theorem 1.1: Part (ii) is Proposition 5.7, part (iii) is Proposition 5.2(i), and part (iv) is Corollary 5.6. Part (i) is immediate by Corollary 5.6 for simple closed curves and by part (iii) for the remaining case.

6 Thurston norm and rigidity

In this section we introduce and study Thurston's norm, which is the infinitesimal version of the metric d_{Th} , and prove Theorems 1.4 and 1.5.

6.1 The norm

Thurston showed in [Thu86c] that the metric d_{Th} is Finsler, i.e. it is induced by a norm $\|\cdot\|_{\text{Th}}$ on the tangent bundle. This norm is naturally expressed as the infinitesimal analogue of the length ratio defining d_{Th} :

$$\|v\|_{\text{Th}} = \sup_{\alpha} \frac{d_X \ell_{\alpha}(v)}{\ell_{\alpha}(X)} = \sup_{\alpha} d_X(\log \ell_{\alpha})(v), \quad v \in T_X \mathcal{J}(S) \quad (23)$$

The following regularity of the norm will be needed in our study of isometries of Thurston's metric:

Theorem 6.1. *Let S be a surface of finite hyperbolic type. Then the Thurston norm function $T\mathcal{J}(S) \rightarrow \mathbf{R}$ is locally Lipschitz.*

The Thurston norm is defined as a supremum of a collection of 1-forms; we will deduce its regularity from that of the forms. In preparation for stating a result to that effect, we must introduce some terminology.

Let M be a smooth manifold, let $\pi : V \rightarrow M$ be a vector bundle over M , and let \mathcal{E} be a collection of sections of V . We say that \mathcal{E} is *locally uniformly bounded* if for each $x \in M$ there exists a neighborhood U of x and a compact set $K \subset V$ such that for each $y \in U$ and $e \in \mathcal{E}$ we have $e(y) \in K$. We say that \mathcal{E} is *locally uniformly Lipschitz* if for each $x \in M$ there exists a neighborhood U of x , a local trivialization $\varphi : \pi^{-1}(U) \xrightarrow{\sim} U \times \mathbf{R}^n$, and a constant M so that for each $e \in \mathcal{E}$, if we use the local trivialization φ to regard the section e as a map $U_i \rightarrow \mathbf{R}^n$, then this function is M -Lipschitz. Here we fix any background norm on \mathbf{R}^n in order to define Lipschitz functions to that space; because all such norms are bi-Lipschitz equivalent, the definition of locally uniformly Lipschitz does not depend on that choice.

Lemma 6.2. *Let M be a smooth manifold and \mathcal{E} a collection of 1-forms on M . Suppose that \mathcal{E} , considered as a collection of sections of T^*M , is locally uniformly bounded and locally uniformly Lipschitz. Then the function $E : TM \rightarrow \mathbf{R}$ defined by*

$$E(v) := \sup_{e \in \mathcal{E}} e(v)$$

is locally Lipschitz (assuming it is finite at one point).

Note that “locally Lipschitz” is a well-defined property of a function on a smooth manifold or a section of a vector bundle; it is equivalent to saying that the collection consisting of only that section (or function) is locally uniformly Lipschitz.

Proof. Any linear function $\mathbf{R}^n \rightarrow \mathbf{R}$ is Lipschitz, however the Lipschitz constant is proportional to its norm as an element of $(\mathbf{R}^n)^*$. Thus, for example, a family of linear functions is uniformly Lipschitz only when the corresponding subset of $(\mathbf{R}^n)^*$ is bounded.

For the same reason, if we take a family of 1-forms on M (sections of T^*M) and consider them as fiberwise-linear functions $TM \rightarrow \mathbf{R}$, then in order for these *functions* on TM to be locally uniformly Lipschitz, we must require the *sections* of T^*M to be both locally uniformly Lipschitz and locally uniformly bounded. Here, the compact set K in the definition of locally uniformly bounded ensures that the pointwise norms of the sections in T^*M are bounded in a neighborhood of any point.

Thus the hypotheses on \mathcal{E} are arranged exactly so that the family of functions $TM \rightarrow \mathbf{R}$ of which E is the supremum is locally uniformly Lipschitz.

The supremum of a family of locally uniformly Lipschitz functions is locally Lipschitz or identically infinity. Since the function E is such a supremum, we find that it is locally Lipschitz once it is finite at one point. \square

Proof of Theorem 6.1. By (23), the Thurston norm is a supremum of the type considered in Lemma 6.2. Therefore, it suffices to show that the set

$$d\log \mathcal{C} := \{d\log \ell_\alpha : \alpha \text{ a simple curve} \}$$

of 1-forms on $\mathcal{T}(S)$ is locally uniformly bounded and locally uniformly Lipschitz.

To see this, first recall that length functions extend continuously from curves to the space $\mathcal{ML}(S)$ of measured laminations (see e.g. [Thu86b], [Bon86, Prop. 4.6]), and also that they

extend from real-valued functions on Teichmüller space to holomorphic functions on the complex manifold $\mathcal{QF}(S)$ of quasi-Fuchsian representations (see [Bon96, p. 292]) in which $\mathcal{T}(S)$ is a totally real submanifold. The resulting length function $\ell_\lambda : \mathcal{QF}(S) \rightarrow \mathbf{C}$ depends continuously on λ in the locally uniform topology of functions on $\mathcal{QF}(S)$ [Bon98, pp. 20–21].

For holomorphic functions, locally uniform convergence implies locally uniform convergence of derivatives of any fixed order, so we find that the derivatives of ℓ_λ also depend continuously on λ .

Restricting to $\mathcal{T}(S) \subset \mathcal{QF}(S)$, and noting that the length of a nonzero measured lamination does not vanish on $\mathcal{T}(S)$, we see that the 1-form $d\log(\ell_\lambda)$ on $\mathcal{T}(S)$ is real-analytic, and that the map $\lambda \mapsto d\log(\ell_\lambda)$ is continuous from $\mathcal{ML}(S) \setminus \{0\}$ to the C^1 topology of 1-forms on any compact subset of $\mathcal{T}(S)$.

Since the 1-form $d\log \ell_\lambda$ is invariant under scaling λ , it is naturally a function (still C^1 continuous) of $[\lambda] \in \mathcal{PML}(S) = (\mathcal{ML}(S) \setminus \{0\})/\mathbf{R}^+$. Because $\mathcal{PML}(S)$ is compact, this implies that the collection of 1-forms

$$d\log \mathcal{PML} := \{d\log \ell_\lambda : [\lambda] \in \mathcal{PML}(S)\}$$

is locally uniformly bounded in C^1 . In particular it is locally uniformly Lipschitz, and since this collection contains $d\log \mathcal{C}$, we are done. \square

6.2 Shape of the unit sphere

Fix $X \in \mathcal{T}(S)$ for the rest of this section. Let $T_X^1 \mathcal{T}(S)$ denote the unit sphere of Thurston's norm, i.e.

$$T_X^1 \mathcal{T}(S) = \{v \in T_X \mathcal{T}(S) : \|v\|_{\text{Th}} = 1\}.$$

Similarly, let $T_X^{\leq 1} \mathcal{T}(S)$ denote the unit ball of Thurston's norm.

The dual of the convex set $T_X^{\leq 1} \mathcal{T}(S)$ has a convenient description in terms of measured laminations:

Theorem 6.3 (Thurston [Thu86c]). *The map $\mathcal{PML}(S) \rightarrow T_X^* \mathcal{T}(S)$ given by $\mu \mapsto d_X \log \ell_\mu$ embeds $\mathcal{PML}(S)$ as the boundary of a convex neighborhood of the origin. This convex neighborhood is the dual convex set of $T_X^{\leq 1} \mathcal{T}(S)$.*

Unlike this dual set, a typical point in the boundary of $T_X^{\leq 1} \mathcal{T}(S)$ does not have a canonical description in terms of a lamination on S . However, certain points in the sphere arise from stretch paths. Specifically, let \mathcal{CL} denote the set of all complete geodesic laminations on S . We have a map

$$v_X : \mathcal{CL} \rightarrow T_X^1 \mathcal{T}(S)$$

where $v_X(\lambda)$ is the tangent vector at $t = 0$ to the stretch path $t \mapsto \text{stretch}(X, \lambda, t)$. This map is “dual” to the map $d_X \log \ell$, in the weak sense that $d_X \log \ell_\mu(v_X(\lambda)) = 1$ if μ is a measured lamination whose support is contained in λ .

For later use, it will be important to note the continuity of the map v_X , which follows easily from the results of [Bon98]:




Figure 11: At left, the unit sphere $T_X^1\mathcal{T}(S_{1,1})$ of the Thurston norm on the tangent space at the point $X = 0.35 + 1.8i$ in the upper half-plane model of $\mathcal{T}(S_{1,1})$. At right, the unit sphere of the dual norm on the cotangent space.

Lemma 6.4. *The map v_X is continuous with respect to the Hausdorff topology on \mathcal{CL} .*

Proof. Let $\lambda_n \in \mathcal{CL}$ be a sequence that converges in the Hausdorff topology. In [Bon98, pp.20–21], Bonahon shows that the associated shearing embeddings $s_{\lambda_n} : \mathcal{T}(S) \rightarrow \mathbf{R}^N$ converge in the C^k topology¹ to s_λ on any compact subset of $\mathcal{T}(S)$. Since stretch paths are rays from the origin in the shearing coordinates, this shows that the tangent vectors $v_X(\lambda_n)$ to such stretch paths converge to $v_X(\lambda)$. \square

Now we specialize to the punctured torus case. That is, for the rest of this section we assume $S = S_{1,1}$. An example of the Thurston unit sphere (circle) and its dual are shown in Figure 11. We will show that in this case, the shape of the unit sphere determines the hyperbolic structure X up to the action of the mapping class group.

In [Thu86c], Thurston studies the facets of the unit ball in $T_X\mathcal{T}(S)$, showing in particular that they correspond to simple curves on the surface. We will require a slight extension of the result about these facets given by Thurston in Theorem 10.1 of that paper. While a corresponding result for any surface is suggested by Thurston’s work, here we will use an *ad hoc* argument specific to the punctured torus case.

Let $\mathcal{RL} \subset \mathcal{CL}$ be the set of canonical completions of maximal chain-recurrent geodesic laminations on $S_{1,1}$. Thus for any simple curve α on $S_{1,1}$ we have $\alpha^+, \alpha^- \in \mathcal{RL}$, and any $\lambda \in \mathcal{RL}$ is either of this form or is a completion of a measured lamination without closed leaves.

Theorem 6.5. *Let L be a support line of the unit ball of $\|\cdot\|_{\text{Th}}$. Then either:*

- (i) $L \cap T_X^1\mathcal{T}(S_{1,1})$ is a line segment with endpoints $v_X(\alpha^+)$ and $v_X(\alpha^-)$ for a simple curve α , in which case $L = \{v : (d_X \log \ell_\alpha)(v) = 1\}$, or

¹More precisely, Bonahon shows locally uniform convergence of a sequence of holomorphic embeddings that complexify the shearing coordinates. Locally uniform convergence of holomorphic maps implies local C^k convergence.

- (ii) $L \cap T_X^1 \mathcal{J}(S_{1,1})$ is a point, and is equal to $\{v_X(\tilde{\lambda})\}$ for $\tilde{\lambda}$ the canonical completion of a measured lamination λ with no closed leaves.

Proof. First, note that Theorem 5.1 implies that $v_X(\alpha^+) \neq v_X(\alpha^-)$, so case (i) always yields a (nondegenerate) line segment.

By the duality between the embedding of $\mathcal{PM}\mathcal{L}(S)$ in $T_X^* \mathcal{J}(S_{1,1})$ and the norm ball $T_X^{\leq 1} \mathcal{J}(S_{1,1})$, the support lines of the latter are exactly the sets

$$L_\mu = \{\mu : (d_X \log \ell_\mu)(v) = 1\}$$

for nonzero $\mu \in \mathcal{ML}(S_{1,1})$. Thus it suffices to characterize the set

$$L'_\mu := L_\mu \cap T_X^1 \mathcal{J}(S_{1,1})$$

for such μ . Since L_μ is a support line of $T_X^{\leq 1} \mathcal{J}(S_{1,1})$, we have that L'_μ is a compact convex subset of a line, i.e. either a point or a segment. If L'_μ is a segment, then at any interior point p of this segment the line L_μ is the *unique* support line of $T_X^{\leq 1} \mathcal{J}(S_{1,1})$ through p .

Suppose α is a simple curve. Then $v_X(\alpha^\pm) \in L'_\alpha$ since $\alpha \subset \alpha^\pm$. By convexity of L'_α , the closed segment with endpoints $v_X(\alpha^\pm)$ is also a subset of L' .

If L'_α properly contained this segment, then at least one of $v_X(\alpha^+)$ or $v_X(\alpha^-)$ would be an interior point of L' , and hence there would be a neighborhood of that point in $T_X^1 \mathcal{J}(S_{1,1})$ in which L_α is the unique support line.

To see that this is not the case, choose $\lambda \in \mathcal{RL}$ that does not contain α (such as $\lambda = \beta^+$ for β a simple curve that intersects α). Then the sequence of Dehn twists $\lambda_n = D_\alpha^n(\lambda)$ converges to α^\pm in the Hausdorff topology as $n \rightarrow \pm\infty$, and the stump μ_n of λ_n has $[\mu_n] \neq [\alpha] \in \mathcal{PM}\mathcal{L}(S_{1,1})$ for all n . By Lemma 6.4, the sequence $v_X(\lambda_n)$ converges (again as $n \rightarrow \pm\infty$) to $v_X(\alpha^\pm)$. Also, $v_X(\lambda_n)$ lies on the support line L_{μ_n} . Since $\mathcal{PM}\mathcal{L}(S_{1,1})$ embeds in $T_X^* \mathcal{J}(S_{1,1})$ (Theorem 6.3), the lines L_{μ_n} are all distinct from L_α . This shows that L_α is not the unique support line in any neighborhood of $v_X(\alpha^\pm)$, and that (i) holds in this case.

Now consider L'_μ for μ a measured lamination with no closed leaves. Let $\tilde{\mu} \in \mathcal{RL}$ be the canonical completion of μ . Then $v_X(\tilde{\mu}) \in L'_\mu$. To complete the proof we show $L'_\mu = \{v_X(\tilde{\mu})\}$, so that these support lines give case (ii).

Suppose for contradiction that L'_μ contains a nontrivial segment. Then L_μ is the unique support line of $T_X^1 \mathcal{J}(S_{1,1})$ in the interior of that segment, which has $v_X(\tilde{\mu})$ in its closure.

We can approximate $\tilde{\mu}$ in the Hausdorff topology by completions α_n^+ of simple curves α_n , and can furthermore do so with $[\alpha_n] \in \mathcal{PM}\mathcal{L}(S)$ converging to $[\mu] \in \mathcal{PM}\mathcal{L}(S)$ from either side (recalling that $\mathcal{PM}\mathcal{L}(S) \simeq S^1$, so that removing $[\mu]$ separates a small neighborhood in $\mathcal{PM}\mathcal{L}(S)$ into two sides). Thus the directions of the support lines $L_{\alpha_n^+}$ can be taken to converge to that of L_μ from a given side. As in the previous case, Lemma 6.4 shows that $v_X(\alpha_n^+)$ converges to $v_X(\tilde{\mu})$, and since $v_X(\alpha_n^+)$ lies on L_{α_n} , this convergence can be taken to be from either side of $v_X(\tilde{\mu})$. Since the support lines L_{α_n} are distinct from L_μ , this shows that the support line is not unique in any interval whose closure contains $v_X(\tilde{\lambda})$, which is the desired contradiction. \square

Having established that the maximal line segments in $T_X^1\mathcal{J}(S_{1,1})$ are exactly those with endpoints $v_X(\alpha^\pm)$ for α a simple curve, we now study the geometry of these segments. From now on we refer to these simply as *facets*. Let $|F(X, \alpha)|$ denote the length of the facet corresponding to a curve α with respect to $\|\cdot\|_{\text{Th}}$, i.e.

$$|F(X, \alpha)| = \|v_X(\alpha^+) - v_X(\alpha^-)\|_{\text{Th}}$$

To estimate this length, we first need the following lemma.

Lemma 6.6. *Let $\text{EQ}_\alpha(X, t)$ be the earthquake path along α with $\text{EQ}_\alpha(X, 0) = X$. Let $\text{EQ}_\alpha = \frac{d}{dt} \text{EQ}_\alpha(t) \Big|_{t=0}$. Then we have*

$$\|\text{EQ}_\alpha\|_{\text{Th}} \stackrel{*}{\asymp}_X \ell_\alpha(X).$$

Proof. In fact, this is true for arbitrary measured laminations, for the Teichmüller space of any surface S , and for any norm on $T_X\mathcal{J}(S)$. It is essentially just a rephrasing of [Thu86c, Theorem 5.2] and the subsequent discussion.

The map $\lambda \mapsto \text{EQ}_\lambda$ is a homeomorphism $\mathcal{ML}(S) \rightarrow T_X\mathcal{J}(S)$ (compare [Gar95, Theorem 5.1]), and in particular the tangent vector to the earthquake path of a nonzero lamination is always nonzero. The function $\lambda \mapsto \text{EQ}_\lambda/\ell_\lambda(X)$ is invariant under scaling of λ and hence gives a well-defined continuous map $\mathcal{PM}\mathcal{L}(S) \rightarrow T_X\mathcal{J} \setminus \{0\}$. By compactness of $\mathcal{PM}\mathcal{L}(S)$ the function $\|\text{EQ}_\lambda/\ell_\lambda(X)\|$ is bounded above and below by positive constants, which is equivalent to the claim of the Lemma. \square

Proposition 6.7. *For every curve α , we have*

$$|F(X, \alpha)| \stackrel{*}{\asymp}_X \ell_\alpha(X)^2 e^{-\ell_\alpha(X)}.$$

Proof. Let X_t^+ and X_t^- be as in Section 5.1. These are paths with $X_0^+ = X_0^- = X$ and with tangent vectors $v_X(\alpha^+)$ and $v_X(\alpha^-)$, respectively, at $t = 0$. Note that the length of α is the same in X_t^+ and X_t^- , hence

$$X_t^+ = \text{EQ}_\alpha(X_t^-, \Delta(t)), \tag{24}$$

where as before EQ_α is the earthquake map along α and Δ is the function

$$\Delta(t) = \tau_\alpha(X_t^+) - \tau_\alpha(X_t^-) = 4e^t \log \coth \frac{\ell_\alpha(X)}{2} - 4 \log \coth \frac{e^t \ell_\alpha(X)}{2}.$$

Note that $\Delta(0) = 0$, and define $\dot{\Delta} = \frac{d}{dt} \Delta(t) \Big|_{t=0}$. Differentiating (24) at $t = 0$ we find

$$\frac{d}{dt} X_t^+ \Big|_{t=0} = (D_1 \text{EQ}_\alpha)_{(X,0)} \left(\frac{d}{dt} X_t^- \Big|_{t=0} \right) + (D_2 \text{EQ}_\alpha)_{(X,0)} (\dot{\Delta}) \tag{25}$$

Here, D_1 and D_2 denote the derivatives of EQ_α with respect to its first and second arguments, respectively. Now, as observed above, the left hand side of (25) is $v_X(\alpha^+)$. Also, since $\text{EQ}_\alpha(Y, 0) = Y$ for all Y , we have that $(D_1 \text{EQ}_\alpha)_{(X,0)}$ is the identity map, and the first term

on the right hand side of (25) becomes $v_X(\alpha^-)$. Recalling that $\dot{E}Q_\alpha = \frac{d}{dt}EQ_\alpha(X, t)|_{t=0}$, the second term on the right hand side of (25) is equal to $\dot{\Delta} \dot{E}Q_\alpha$.

Thus we have

$$v_X(\alpha^+) = v_X(\alpha^-) + \dot{\Delta} \dot{E}Q_\alpha$$

and hence

$$|F(X, \alpha)| = \|v_X(\alpha^+) - v_X(\alpha^-)\|_{\text{Th}} = |\dot{\Delta}| \left\| \dot{E}Q_\alpha \right\|_{\text{Th}}. \quad (26)$$

Using the formula for $\Delta(t)$ given above we compute

$$\dot{\Delta} = 4 \log \coth \frac{\ell_\alpha(X)}{2} + 4\ell_\alpha(X) \operatorname{csch}(\ell_\alpha(X)) > 0.$$

For large values of x we have

$$\log \coth(x) \sim 2e^{-2x} \quad \text{and} \quad \operatorname{csch}(x) \sim 2e^{-x}$$

Hence for large values of $\ell_\alpha(X)$, we have

$$|\dot{\Delta}| = \dot{\Delta} \sim 8e^{-\ell_\alpha(X)} + 4\ell_\alpha(X)e^{-\ell_\alpha(X)} \stackrel{*}{\asymp} \ell_\alpha(X)e^{-\ell_\alpha(X)},$$

and by Lemma 6.6,

$$\left\| \dot{E}Q_\alpha \right\|_{\text{Th}} \stackrel{*}{\asymp} \ell_\alpha(X).$$

Substituting these estimates for $\dot{\Delta}$ and $\left\| \dot{E}Q_\alpha \right\|_{\text{Th}}$ into (26) gives the proposition. \square

Theorem 6.8. *Let α and β be curves with $i(\alpha, \beta) = 1$. Let $\beta_n = D_\alpha^n(\beta)$. Then*

$$\lim_{n \rightarrow \infty} \frac{|\log |F(X, \beta_n)||}{n} = \ell_\alpha(X).$$

Proof. For large values of n ,

$$\ell_{\beta_n}(X) \stackrel{\pm}{\asymp} n\ell_\alpha(X).$$

The theorem now follows from Proposition 6.7. \square

Using the results above we can now show that the shape of the unit sphere in $T_X^1\mathcal{J}(S_{1,1})$ determines X up to the action of the mapping class group.

Proof of Theorem 1.4. Within the convex curve $T_X^1\mathcal{J}(S_{1,1})$ let U_0 denote an open arc disjoint from $F(X, \alpha)$ which has $v_X(\alpha^+)$ as one endpoint. We use “interval notation” to refer to open arcs within U_0 , where (x, y) refers to the open arc in U_0 with endpoints x, y . Thus for example U_0 itself is $(v_X(\alpha^+), y)$ for some y .

Let $\mathcal{S}(U_0)$ denote the set of curves γ such that $F(X, \gamma) \subset U_0$. Thus $\mathcal{S}(U_0)$ corresponds to the rational points of an interval in $\mathcal{PML}(S_{1,1})$ with α as one of its endpoints. Any sequence of simple closed geodesics in this interval converging to α in $\mathcal{PML}(S_{1,1})$ also converges in the Hausdorff topology, to α_0^+ .




Figure 12: *The train track structure of a neighborhood of α_0^+ (shown in a collar about α).*

Thus for any $\epsilon > 0$, by choosing U_0 small enough we can assume that all of the curves $\gamma \in \mathcal{S}(U_0)$ have geodesic representative in X that is contained in an ϵ -neighborhood of the geodesic lamination α_0^+ . This neighborhood has the structure of a thickened train track τ with three branches (as shown in Figure 12): Along α there is a “thick” branch and a “thin” branch, and there is a third branch κ which connects one side of α to the other. Such a curve γ is therefore determined by a pair of coprime nonnegative integers (p, q) , where p is the weight of the thin branch along α and q is the weight of κ . (By the switch relations, these two weights determine the weight of the thick branch to be $p + q$). We call (p, q) the *coordinates* of γ . In terms of these coordinates, q is the geometric intersection number of the curve with α .

Let ℓ_1 denote the length of the branch κ , which we define to be the minimum length of a path in the rectangle joining its short sides, and let $\ell_{\text{thick}}, \ell_{\text{thin}}$ denote the lengths of the branches along α . Note that ℓ_1 increases without bound as $\epsilon \rightarrow 0$, so we assume from now on that ℓ_1 is much larger than $\ell_\alpha(X)$. Also, since the union of the thick and thin branches gives a small neighborhood of α we have $\ell_{\text{thick}} + \ell_{\text{thin}} = \ell_\alpha(X) + O(\epsilon)$.

The point of this train track representation for curve in $\mathcal{S}(U_0)$ is that it gives a simple estimate for hyperbolic length. Specifically, the simple curve carried by τ with coordinates (p, q) breaks by τ into p arcs in κ , q arcs in thin branch, and $p + q$ arcs in the thick branch. Each arc has length equal to that of its branch, up to an error proportional to ϵ . Thus the overall length is

$$\begin{aligned} \ell(p, q) &= q \ell_1 + p \ell_{\text{thin}} + (p + q) \ell_{\text{thick}} + O((p + q)\epsilon) \\ &= q (\ell_1 + \ell_{\text{thick}}) + p (\ell_{\text{thin}} + \ell_{\text{thick}}) + O((p + q)\epsilon) \\ &= q (\ell_1 + \ell_{\text{thick}}) + p \ell_\alpha(X) + O((p + q)\epsilon), \end{aligned}$$

where the factor of ϵ accounts for the difference between the length of the branch and the length of a segment of γ contained in the branch.

The quantity p/q (the slope of the curve, for a suitable homology basis) is an affine coordinate for a neighborhood of α in $\mathcal{PML}(S_{1,1}) \simeq \mathbf{RP}^1$ in which α corresponds to $1/0$, thus $\mathcal{S}(U_0)$ corresponds to curves with coordinates satisfying $p/q > m$ for some constant $m \in \mathbf{R}$. By the length estimate above, after shrinking U_0 so that ϵ is much smaller than $\ell_\alpha(X)$ and ℓ_1 , we find the minimum length of a curve in $\mathcal{S}(U_0)$ is attained for $q = 1$ and the smallest integer p with

that $p > m$. Denote these length minimizing coordinates by $(p_0, 1)$, and the corresponding curve by γ_0 . Note that any other curve in $\mathcal{S}(U_0)$ has length exceeding this minimum by at least a fixed positive multiple of $\ell_\alpha(X)$.

By Proposition 6.7, the length of a facet corresponding to a curve whose hyperbolic length is bounded below is exponentially decreasing in length of the curve, up to a fixed multiplicative error. (Here, assuming a lower bound on the length allows us to ignore the factor $\ell_\gamma(X)^2$ in that Proposition, as it is overwhelmed by the exponential decay.) Therefore, long curves with a sufficiently large difference in hyperbolic length have associated facets whose lengths compare in the opposite way. By taking $\ell_\alpha(X)$ to be large enough and U_0 small enough so that all curves in $\mathcal{S}(U_0)$ are long, the hyperbolic length gap between the minimizer γ_0 and any other curve in $\mathcal{S}(U_0)$ noted above implies that $F(X, \gamma_0)$ is the longest facet in U_0 .

Now, we can shrink U_0 to exclude $F(X, \gamma_0)$, find the new longest facet, and iterate this construction. That is, we apply the argument above to the arc $(v_X(\alpha^+), v_X(\gamma_0^-))$ and find hyperbolic length minimizer and facet length maximizer γ_1 . Taking $v_X(\gamma_0^-)$ as the endpoint means that the coordinates (p, q) of curves whose facets lie in this arc now satisfy $p/q > p_0$, so arguing exactly as above we find that the coordinates of γ_1 are $(p_0 + 1, 1)$. Continuing inductively we obtain a sequence γ_i of curves, each corresponding to the longest facet in $(v_X(\alpha^+), v_X(\gamma_{i-1}^-))$ and having coordinates $(p_0 + i, 1)$. We call this the *sequence of longest facets*.

Recall that the Dehn twist about α acts in these coordinates by adding 1 to the slope of the curve. Thus, in more invariant terms we have shown that the sequence of longest facets to one side of α^+ corresponds to the sequence of all sufficiently large positive powers of the Dehn twist about α applied to a curve intersecting α once. This is the sort of collection considered in Theorem 6.8, which shows that the hyperbolic length of α is determined by the asymptotic behavior of these facet lengths.

An argument very similar to the one above shows that the sequence of longest facets in a small neighborhood of $v_X(\alpha^-)$ corresponds to large *negative* powers of the Dehn twist about α applied to a curve intersecting α once, and that through the asymptotics of their lengths, the geometry of the norm sphere near $v_X(\alpha^-)$ also determines the length of α . As before this applies to any simple curve α that is sufficiently long on X . Collectively, we refer to the arguments above as the *longest facet construction*.

Now for $X, Y \in \mathcal{T}(S_{1,1})$, assume that there is a norm preserving linear map

$$L: T_X \mathcal{T}(S_{1,1}) \rightarrow T_Y \mathcal{T}(S_{1,1}).$$

Since L is linear, it maps the facets in $T_X^1 \mathcal{T}(S_{1,1})$ bijectively to those in $T_Y^1 \mathcal{T}(S_{1,1})$. By Theorem 6.5, this induces some permutation on the simple curves that label the facets: For a simple curve γ we denote by γ^* the simple curve such that $L(F(X, \gamma)) = F(Y, \gamma^*)$.

Choose a simple curve α so that $\ell_X(\alpha)$ and $\ell_Y(\alpha^*)$ are large enough so that the longest facet construction applies to both of them. Then we obtain a sequence of curves $\gamma_i = D_\alpha^i \beta$ which satisfy $i(\gamma_i, \alpha) = 1$, and whose facets $F(X, \gamma_i)$ approach one endpoint of $F(X, \alpha)$ with each being longest in some neighborhood of that endpoint. As L is an isometry, the image facets $F(X, \gamma_i^*)$ approach some endpoint of $F(X, \alpha^*)$ and are locally longest in the same sense. Thus

the curves γ_i^* are also obtained by applying powers (positive or negative) of a Dehn twist about α^* to a fixed curve and they satisfy $i(\gamma_i^*, \alpha^*) = 1$. Since $|F(X, \gamma_i)| = |F(Y, \gamma_i^*)|$ we conclude $\ell_X(\alpha) = \ell_X(\alpha^*)$.

Now choose an integer N so that $\ell_X(\gamma_N)$ and $\ell_Y(\gamma_N^*)$ are large enough to apply the longest facet construction (to γ_N and γ_N^* , respectively). Proceeding as in the previous paragraph, we find $\ell_X(\gamma_N) = \ell_Y(\gamma_N^*)$.

At this point we have two pairs of simple curves intersecting once, (α, γ_N) and (α^*, γ_N^*) , and the lengths of the first pair on X are equal to those of the second pair on Y . This implies that X and Y are in the same orbit of the extended mapping class group: Take a mapping class ϕ with $\phi(\alpha) = \alpha^*$ and $\phi(\gamma_N) = \gamma_N^*$. Then α has the same length on X and $\phi^{-1}(Y)$, so these points differ only in the Fenchel-Nielsen twist parameter (relative to pants decomposition α). Since the length of γ_N is the same as well, either the twist parameters are equal and $X = \phi^{-1}(Y)$ or the twist parameters differ by a sign and $X = r(\phi^{-1}(Y))$ where r is the orientation-reversing mapping class which preserves both α and γ_N while reversing orientation of γ_N . \square

6.3 Local and global isometries

Before proceeding with the proof of Theorem 1.5 we recall some standard properties of the extended mapping class group action on $\mathcal{T}(S_{1,1})$. (For further discussion, see for example [Kee74, Section 2] [FM12, Section 2.2.4].)

The mapping class group $\text{Mod}(S_{1,1}) = \text{Homeo}^+(S_{1,1})/\text{Homeo}_0(S_{1,1})$ of the punctured torus is isomorphic to $\text{SL}(2, \mathbf{Z})$, and identifying $\mathcal{T}(S_{1,1})$ with the upper half-plane \mathbf{H} in the usual way, the action of $\text{Mod}(S_{1,1})$ becomes the action of $\text{SL}(2, \mathbf{Z})$ by linear fractional transformations. Similarly, the extended mapping class group $\text{Mod}^\pm(S_{1,1}) = \text{Homeo}(S_{1,1})/\text{Homeo}_0(S_{1,1})$ can be identified with $\text{GL}(2, \mathbf{Z})$, where an element $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ of determinant -1 acts on \mathbf{H} by the conjugate-linear map $z \mapsto \frac{a\bar{z}+b}{c\bar{z}+d}$. Neither of these groups acts effectively on \mathbf{H} , since in each case the elements $\pm I$ act trivially; thus when considering the action on $\mathcal{T}(S_{1,1})$ it is convenient to work with the quotients $\text{PSL}(2, \mathbf{Z})$ and $\text{PGL}(2, \mathbf{Z})$ which act effectively.

The properly discontinuous action of $\text{PGL}(2, \mathbf{Z})$ on \mathbf{H} preserves the standard $(2, 3, \infty)$ triangle tiling of \mathbf{H} (see Figure 13), with the cells of each dimension in this tiling corresponding to different types of isotropy; specifically, we have:

- A point in the interior of a triangle has trivial stabilizer in $\text{PGL}(2, \mathbf{Z})$.
- A point in the interior of an edge has stabilizer in $\text{PGL}(2, \mathbf{Z})$ isomorphic to $\mathbf{Z}/2$ and generated by a *reflection*, i.e. an element conjugate to a either $z \mapsto -\bar{z}$ or $z \mapsto -\bar{z} + 1$.

While vertices of the tiling have larger stabilizers, the only property of such points we will use is that they form a discrete set.

Proof of Theorem 1.5. Let U be an open connected set in $\mathcal{T}(S_{1,1})$ and let $f : (U, d_{\text{Th}}) \rightarrow (\mathcal{T}(S_{1,1}), d_{\text{Th}})$ be an isometric embedding.

By Theorem 6.1, the Thurston norm is locally Lipschitz (locally $C_{loc}^{0,1}$). By [MT17, Theorem A] an isometry of such Finsler spaces is $C_{loc}^{1,1}$ and its differential is norm-preserving. Therefore,




Figure 13: The standard $(2,3,\infty)$ triangle tiling of the upper half plane. The marked point is the imaginary unit i .

for each $X \in U$ the differential

$$d_X f : T_X \mathcal{J}(S_{1,1}) \rightarrow T_{f(X)} \mathcal{J}(S_{1,1})$$

is an isometry for the Thurston norm and by Theorem 1.4 there exists $\Phi(X) \in \mathrm{PGL}(2, \mathbf{Z})$ such that

$$f(X) = \Phi(X) \cdot X \quad (27)$$

This property may not determine $\Phi(X) \in \mathrm{PGL}(2, \mathbf{Z})$ uniquely, however, choosing one such element for each point of U we obtain a map $\Phi : U \rightarrow \mathrm{PGL}(2, \mathbf{Z})$.

Let $X_0 \in U$ be a point with trivial stabilizer in $\mathrm{PGL}(2, \mathbf{Z})$. Using proper discontinuity of the $\mathrm{PGL}(2, \mathbf{Z})$ action, we can select neighborhoods V of X_0 and W of $f(X_0)$ so that

$$\{\phi \in \mathrm{PGL}(2, \mathbf{Z}) : \phi \cdot V \cap W \neq \emptyset\} = \{\Phi(X_0)\}$$

However, by continuity of f and (27) we find that the $\Phi(X)$ is an element of this set for all X near X_0 . That is, the map Φ is locally constant at X_0 . More generally, this shows Φ is constant on any connected set consisting of points with trivial stabilizer.

Now we consider the behavior of Φ and f in a small neighborhood V of a point X_1 with $\mathbf{Z}/2$ stabilizer—that is, a point in the interior of an edge e of the $(2,3,\infty)$ triangle tiling. Taking V to be a sufficiently small disk, we can assume $V \setminus e$ has two components, which we label by V_{\pm} , and that each component consists of points with trivial stabilizer (equivalently, V does not contain any vertices of the tiling). By the discussion above Φ is constant on V_+ and on V_- , and we denote the respective values by ϕ_+ and ϕ_- . By continuity of f , the element $\phi_+^{-1} \phi_- \in \mathrm{PGL}(2, \mathbf{Z})$ fixes $e \cap V$ pointwise and is therefore either the identity or a reflection. In the latter case f would map both sides of e (locally, near X_1) to the same side of the edge $f(e)$, and hence it would not be an immersion at X_1 . This is a contradiction, for we have seen that the differential of f is an isomorphism at each point. We conclude $\phi_+ = \phi_-$, and f agrees with this extended mapping class on $V \setminus e$. By continuity of f the same equality extends over the edge e .

Let $U' \subset U$ denote the subset of points with trivial or $\mathbf{Z}/2$ stabilizer. We have now shown that for each $X \in U'$ there exists a neighborhood of X on which f is equal to an element of $\mathrm{PGL}(2, \mathbf{Z})$. An element of $\mathrm{PGL}(2, \mathbf{Z})$ is uniquely determined by its action on any open set, so this local representation of f by a mapping class is uniquely determined and locally constant. Thus on any connected component of U' we have that f is equal to a mapping class. However U' is connected, since U is connected and open and the set of points in $\mathcal{T}(S_{1,1})$ with larger stabilizer (i.e. the vertex set of the tiling) is discrete.

We have therefore shown $f = \phi$ on U' , for some $\phi \in \mathrm{PGL}(2, \mathbf{Z})$. Finally, both f and ϕ are continuous, and U' is dense in U , equality extends to U , as required. \square

References

- [Bon86] F. Bonahon. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math. (2)*, 124(1):71–158, 1986.
- [Bon96] F. Bonahon. Shearing hyperbolic surfaces, bending pleated surfaces and Thurston’s symplectic form. *Ann. Fac. Sci. Toulouse Math. (6)*, 5(2):233–297, 1996.
- [Bon98] F. Bonahon. Variations of the boundary geometry of 3-dimensional hyperbolic convex cores. *J. Differential Geom.*, 50(1):1–24, 1998.
- [BZ04] F. Bonahon and X. Zhu. The metric space of geodesic laminations on a surface. II. Small surfaces. In *Proceedings of the Casson Fest*, volume 7 of *Geom. Topol. Monogr.*, pages 509–547. Geom. Topol. Publ., Coventry, 2004.
- [CR07] Y.-E. Choi and K. Rafi. Comparison between Teichmüller and Lipschitz metrics. *J. Lond. Math. Soc. (2)*, 76(3):739–756, 2007.
- [CRS08] Y.-E. Choi, K. Rafi, and C. Series. Lines of minima and Teichmüller geodesics. *Geom. Funct. Anal.*, 18(3):698–754, 2008.
- [DGK16] J. Danciger, F. Guéritaud, and F. Kassel. Geometry and topology of complete Lorentz spacetimes of constant curvature. *Ann. Sci. Éc. Norm. Supér. (4)*, 49(1):1–56, 2016.
- [Dum13] D. Dumas. CP1: Complex projective structures toolkit, 2013. <https://github.com/daviddumas/cp1>.
- [FM12] B. Farb and D. Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [Gar95] F. P. Gardiner. Infinitesimal bending and twisting in one-dimensional dynamics. *Trans. Amer. Math. Soc.*, 347(3):915–937, 1995.
- [GK17] F. Guéritaud and F. Kassel. Maximally stretched laminations on geometrically finite hyperbolic manifolds. *Geom. Topol.*, 21(2):693–840, 2017.
- [Gué15] F. Guéritaud. Lengthening deformations of singular hyperbolic tori. *Annales de la faculté des sciences de Toulouse*, 24(5):1239–1260, 2015.
- [Gué16] F. Guéritaud. Personal communication, 2016.

- [Iva97] N. V. Ivanov. Automorphism of complexes of curves and of Teichmüller spaces. *Internat. Math. Res. Notices*, 14:651–666, 1997.
- [Kas09] F. Kassel. Quotients compacts d’espaces homogènes réels ou p-adiques. Ph.D. Thesis, Université Paris-Sud 11, 2009.
- [Kee74] L. Keen. On fundamental domains and the Teichmüller modular group. In *Contributions to analysis (a collection of papers dedicated to Lipman Bers)*, pages 185–194. Academic Press, New York, 1974.
- [LRT12] A. Lenzhen, K. Rafi, and J. Tao. Bounded combinatorics and the Lipschitz metric on Teichmüller space. *Geom. Dedicata*, 159:353–371, 2012.
- [LRT15] A. Lenzhen, K. Rafi, and J. Tao. The shadow of a Thurston geodesic to the curve graph. *J. Topol.*, 8(4):1085–1118, 2015.
- [Min96] Y. N. Minsky. Extremal length estimates and product regions in Teichmüller space. *Duke Math. J.*, 83(2):249–286, 1996.
- [Min99] Y. N. Minsky. The classification of punctured-torus groups. *Ann. of Math. (2)*, 149(2):559–626, 1999.
- [MM00] H.A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [MT17] V. S. Matveev and M. Troyanov. The Myers-Steenrod theorem for Finsler manifolds of low regularity. *Proc. Amer. Math. Soc.*, 145(6):2699–2712, 2017.
- [Pap91] A. Papadopoulos. On Thurston’s boundary of Teichmüller space and the extension of earthquakes. *Topology Appl.*, 41(3):147–177, 1991.
- [PT07] A. Papadopoulos and G. Théret. On Teichmüller’s metric and Thurston’s asymmetric metric on Teichmüller space. In *Handbook of Teichmüller theory. Vol. I*, volume 11 of *IRMA Lect. Math. Theor. Phys.*, pages 111–204. Eur. Math. Soc., Zürich, 2007.
- [Raf05] K. Rafi. A characterization of short curves of a Teichmüller geodesic. *Geometry and Topology*, 9:179–202, 2005.
- [Raf14] K. Rafi. Hyperbolicity in Teichmüller space. *Geom. Topol.*, 18(5):3025–3053, 2014.
- [Roy71] H. L. Royden. Automorphisms and isometries of Teichmüller space. In *Advances in the Theory of Riemann Surfaces (Proc. Conf., Stony Brook, N.Y., 1969)*, pages 369–383. Ann. of Math. Studies, No. 66. Princeton Univ. Press, Princeton, N.J., 1971.
- [Thé07] G. Théret. On the negative convergence of Thurston’s stretch lines towards the boundary of Teichmüller space. *Ann. Acad. Sci. Fenn. Math.*, 32(2):381–408, 2007.
- [Thu86a] W. P. Thurston. Geometry and topology of 3-manifolds. Princeton University Lecture Notes, 1986. Available at www.msri.org/publications/books/gt3m.
- [Thu86b] W. P. Thurston. Hyperbolic Structures on 3-manifolds, II: Surface groups and 3-manifolds which fiber over the circle. Preprint, [arxiv:math/9801045](https://arxiv.org/abs/math/9801045), 1986.

- [Thu86c] W. P. Thurston. Minimal stretch maps between hyperbolic surfaces. Preprint, 1986.
[arXiv:math.GT/9801039](https://arxiv.org/abs/math/9801039).
- [Wal14] C. Walsh. The horoboundary and isometry group of Thurston's Lipschitz metric. In *Handbook of Teichmüller theory. Vol. IV*, volume 19 of *IRMA Lect. Math. Theor. Phys.*, pages 327–353. Eur. Math. Soc., Zürich, 2014.

Department of Mathematics, Statistics, and Computer Science
University of Illinois at Chicago
Chicago, IL, USA
`david@dumas.io`

Department of Mathematics
University of Rennes
Rennes, FR
`anna.lenzhen@univ-rennes1.fr`

Department of Mathematics
University of Toronto
Toronto, CA
`rafi@math.toronto.edu`

Department of Mathematics
University of Oklahoma
Norman, OK, USA
`jing@ou.edu`