GAFA, Geom. funct. anal.
Vol. 17 (2007) 936 – 959
1016-443X/07/030936-24
DOI 10.1007/s00039-007-0615-x
ONLINE FIRST: July 2007

© Birkhäuser Verlag, Basel 2007

GAFA Geometric And Functional Analysis

A COMBINATORIAL MODEL FOR THE TEICHMÜLLER METRIC

Kasra Rafi

Abstract. We study how the length and the twisting parameter of a curve change along a Teichmüller geodesic. We then use our results to provide a formula for the Teichmüller distance between two hyperbolic metrics on a surface, in terms of the combinatorial complexity of curves of bounded lengths in these two metrics.

1 Introduction

This paper should be considered a sequel to [R]. We continue here to study the geometry of Teichmüller space using combinatorial properties of curves on surfaces. The main result is a formula for the Teichmüller distance between two points in Teichmüller space, in terms of the combinatorial information extracted from short curves of these two points. Let S be a surface of finite type with negative Euler characteristic and let σ_1 and σ_2 be two points in the thick part of Teichmüller space $\mathcal{T}(S)$ of S. Let μ_1 and μ_2 be short markings on σ_1 and σ_2, respectively.

Theorem 1.1. *There exists $k > 0$ such that*

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \asymp \sum_{Y} \big[d_Y(\mu_1, \mu_2)\big]_k + \sum_{\alpha} \log \big[d_\alpha(\mu_1, \mu_2)\big]_k. \tag{1}$$

In the above theorem, the first sum is over all subsurfaces of S that are not annuli and the second sum is over all simple closed curves on S; $d_Y(\mu_1, \mu_2)$ measures the relative complexity of the restrictions of μ_1 and μ_2 to a subsurface Y, and $d_\alpha(\mu_1, \mu_2)$ measures the relative twisting of μ_1 and μ_2 around a curve α; the function $[x]_k$ is equal to zero when $x < k$ and is equal to x when $x \geq k$, that is, we take into account only terms that are large enough; and the function log is a modified logarithm so that, for

Keywords and phrases: Teichmüller geodesics, curve complex, quadratic differentials.
AMS *Mathematics Subject Classification*: 30F60, 32G15, 57M07

$x \in [0,1]$, $\log x = 0$. A general version of this theorem, where σ_1 and σ_2 are not necessarily in the thick part, is stated in §6 (Theorem 6.1).

Other recent results relate the geometry of Teichmüller space to combinatorial spaces. In [MaM1] Masur and Minsky show that the electrified Teichmüller space is quasi-isometric to the complex of curves and therefore is also δ-hyperbolic. Brock has shown ([Br]) that Teichmüller space equipped with the Weil–Petersson metric is quasi-isometric to the pants complex. Most recent developments in studying the Weil–Petersson metric have resulted from this analogy.

To drive our formula, we need to acquire an understanding of how the length and the twisting parameter of a curve change along a Teichmüller geodesic. [R] provides a description of short curves. In this paper, we prove the following *convexity* property for the length of a curve along a Teichmüller geodesic. Let $g\colon \mathbb{R} \to \mathcal{T}(S)$ be a geodesic in the Teichmüller space of S. For a curve α on S, denote the hyperbolic length of the geodesic representative of α at $g(t)$ by l_t.

Theorem 1.2. *Assume α is balanced at t_α and $s \geq t_\alpha$ (respectively, $s \leq t_\alpha$). Then, for any $t \geq s$ ($t \leq s$), we have*

$$\tfrac{1}{l_s} \succ \tfrac{1}{l_t}\,.$$

We also give the following estimate for the twisting parameter along a Teichmüller geodesic. Let ν_+ be the stable foliation of the geodesic g. The twisting parameter around a curve α at $g(t)$ is (roughly) the number of times that ν_+ twists around α relative to a curve perpendicular to α in the hyperbolic metric of $g(t)$, and is denoted by tw_t^+.

Theorem 1.3. *There exists a constant $d_\alpha > 0$ such that*

$$tw_t^+(\alpha) = \frac{d_\alpha\, e^{-2\,(t-t_\alpha)}}{e^{2\,(t-t_\alpha)} + e^{-2\,(t-t_\alpha)}} \pm O(1/l_t)\,.$$

Some notation. To simplify our presentation, we avoid keeping track of constants that depend on the topology of the surface only. Instead, we use the following notation: When two functions f and g are equal up to additive constants, that is, when there exists a C depending on the topology of S, such that $g(x) - C \leq f(x) \leq g(x) + C$, we write $f(x) \stackrel{+}{\asymp} g(x)$. Similarly, $f(x) \stackrel{+}{\succ} g(x)$ and $f(x) \stackrel{+}{\prec} g(x)$ mean that the inequalities are true up to an additive constant. When an inequality is true up to a multiplicative constant, we use symbols $\stackrel{\cdot}{\asymp}$, $\stackrel{\cdot}{\succ}$ and $\stackrel{\cdot}{\prec}$; and, when it is true up to an additive constant and a multiplicative constant, we use symbols $\asymp$, $\prec$ and $\succ$. For example, $f(x) \asymp g(x)$ means that there are constants c and C, depending on the topology of the surface only, such that

$$\frac{1}{c}\, g(x) - C \le f(x) \le c\, g(x) + C\,.$$

Acknowledgments. I would like to thank Yair Minsky for suggesting the statement of Theorem 1.1 and Jason Behrstock for persuading me to prove a more general version and suggesting the statement of Theorem 6.1. I would also like to thank Yair Minsky and Young-Eun Choi for several helpful conversations, as well as the referee for many useful comments.

2 Preliminaries

2.1 Curves and markings. By a *curve* in S we mean a non-trivial, non-peripheral, simple closed curve in S. The free homotopy class of a curve α is denoted by $[\alpha]$. By an *essential arc* ω we mean a simple arc, with endpoints on the boundary of S, that cannot be pushed to the boundary of S. In case S is not an annulus, $[\omega]$ represents the homotopy class of ω relative to the boundary of S. When S is an annulus, $[\omega]$ is defined to be the homotopy class of ω relative to the endpoints of ω.

Define $\mathcal{C}(S)$ to be the set of all homotopy classes of curves and essential arcs on the surface S. To simplify notation, we often write $\alpha \in \mathcal{C}(S)$ instead of $[\alpha] \in \mathcal{C}(S)$. Define a distance on $\mathcal{C}(S)$ as follows: For $\alpha, \beta \in \mathcal{C}(S)$, define $d_S(\alpha, \beta)$ to be equal to one if $\alpha \neq \beta$ and if α and β can be represented by disjoint curves or arcs. Let the metric on $\mathcal{C}(S)$ be the maximal metric having the above property, i.e. $d_S(\alpha, \beta) = n$ if $\alpha = \gamma_0, \gamma_1, \dots, \gamma_n = \beta$ is the shortest sequence of curves or arcs on S such that, for $i = 1, \dots, n$, γ_{i-1} is distance one from γ_i. (See [MaM1].)

Let $\{\alpha_1, \dots, \alpha_m\}$ be a pants decomposition of S. A *marking* on S is a set $\mu = \{(\alpha_1, \beta_1), \dots, (\alpha_m, \beta_m)\}$ such that the curve β_i is disjoint from α_j, for $i \neq j$, and intersects α_i once (twice) if the surface filled by α_i and β_i is a once-punctured torus (four-times-punctured sphere). The α_i are called the base curves of μ. For every i, β_i is called the transverse curve to α_i in μ. When the distinction between the base curves and the transverse curves is not important, we represent a marking as a set of curves $\{\beta_1, \dots, \beta_n\}$ including all the base curves and the transverse curves. Denote the space of all markings on S by $\mathcal{M}(S)$ (see [MaM2].)

2.2 Subsurface intersection and subsurface distance. Let ν be a subset of $\mathcal{C}(S)$ (e.g. curves appearing in a marking) or a singular foliation on S, and let Y be a subsurface of S. We define the *projection* of ν to the subsurface Y as follows: Let

$$f : \bar{S} \to S$$

be a regular covering of S such that $f_*(\pi_1(\bar{S}))$ is conjugate to $\pi_1(Y)$ (the Y-cover of S). Since S admits a hyperbolic metric, $\bar{S}$ has a well-defined boundary at infinity. Let $\bar{\nu}$ be the lift of ν to $\bar{S}$. Components of $\bar{\nu}$ that are essential arcs or curves on $\bar{S}$, if any, form a subset of $\mathcal{C}(\bar{S})$. The surface $\bar{S}$ is homeomorphic to Y. We call the corresponding subset of $\mathcal{C}(Y)$ the projection of ν to Y and will denote it by ν_Y. If there are no essential arcs or curves in $\bar{\nu}$, ν_Y is the empty set; otherwise we say that ν *intersects* Y *essentially*. This projection depends on the homotopy class of elements of ν only.

Let ν and ν' be subsets of $\mathcal{C}(S)$ or singular foliations on S that intersect a subsurface Y essentially. We define the *Y-intersection* (*Y-distance*) between ν and ν' to be the maximum geometric intersection number in Y (maximum distance in $\mathcal{C}(Y)$) between the elements of projections ν_Y and ν'_Y and denote it by

$$i_Y(\nu, \nu') \quad \text{(respectively, } d_Y(\nu, \nu')\text{)}\,.$$

If Y is an annulus whose core is the curve α, then we also denote $i_Y(\nu, \nu')$ and $d_Y(\nu, \nu')$ by $i_\alpha(\nu, \nu')$ and $d_\alpha(\nu, \nu')$, respectively.

LEMMA 2.1 [B, Lem. 1.2]. *Let Y, ν and ν' be as above.*

1. *If Y is not an annulus, then*

$$d_Y(\nu, \nu') \prec \log i_Y(\nu, \nu')\,.$$

2. *For a curve α,*

$$d_\alpha(\nu, \nu') \asymp i_\alpha(\nu, \nu')\,.$$

2.3 Quadratic differentials. Let q be a meromorphic quadratic differential of area one on S. (See [GL] for definition and details.) We assume that q has a discrete set of finite critical points (i.e. critical points of q are either zeroes or poles of order 1). Corresponding to q, there are two singular measured foliations called the horizontal and the vertical foliations, which we denote by ν_- and ν_+. We call the singular Euclidean metric $|q|$ the *q-metric* on S. For a curve α in S, the q-geodesic representative of α exists and is unique except for the case where it is one of the continuous family of closed geodesics in a flat annulus, which we refer to as the flat annulus corresponding to α. (Some difficulties arise when q has poles of order 1. See [R] for precise definitions and discussion.) We denote the q-length of α by $l_q(\alpha)$, the horizontal length of α by $h_q(\alpha)$ and the vertical length of α by

$v_q(\alpha)$. We also denote the q-length, the horizontal length and the vertical length of the q-geodesic representative of α, by $l_q([\alpha])$, $h_q([\alpha])$ and $v_q([\alpha])$, respectively. In general, for any metric τ, $l_\tau(\alpha)$ represents the τ-length of α and $l_\tau([\alpha])$ represents the τ-length of the τ-geodesic representative of α.

The following theorem, which is an analogue of the collar lemma, compares the lengths, in quadratic differential metric, of intersecting curves, assuming one of them has bounded hyperbolic length. Let σ be the hyperbolic metric in conformal class of q.

Theorem 2.2 [R, Th. 1.3]**.** *For every $L > 0$, there exists D_L, $\log D_L \asymp e^L$, such that, if α and β are two simple closed curves in S intersecting non-trivially with $l_\sigma(\beta) \leq L$, then*

$$D_L l_q(\alpha) \geq l_q(\beta)\,.$$

2.4 Regular and primitive annuli in q. Let Y be a subsurface of S and γ be a boundary component of Y. (We always assume that curves are piecewise smooth.) The curvature of γ with respect to Y, $\kappa_Y(\gamma)$, is well defined as a measure with atoms at the corners. We choose the sign to be positive when the acceleration vector points into Y. If γ is curved non-negatively (or non-positively) with respect to Y at every point, we say it is *monotonically curved* with respect to Y. Let A be an open annulus in S with boundaries γ_0 and γ_1. Suppose both boundaries are monotonically curved with respect to A and $\kappa_A(\gamma_0) \leq 0$. Further, suppose that the boundaries are equidistant from each other, and the interior of A contains no zeroes. We call A a *primitive* annulus and write $\kappa(A) = -\kappa_A(\gamma_0)$. If $\kappa(A) > 0$, we call A *expanding* and say that γ_0 is the inner boundary and γ_1 is the outer boundary. When $\kappa(A) = 0$, A is a *flat* annulus and is foliated by closed Euclidean geodesics homotopic to the boundaries. The following lemma is useful for computing the modulus of a primitive annulus.

LEMMA 2.3 [R, Lem. 3.6]. *Let A and γ_0 be as above, and let d be the distance between the boundaries of A. Then*

$$\begin{cases} \kappa \ \mathrm{Mod}(A) \asymp \log\left(\frac{d}{l_q(\gamma_0)}\right) & \text{if } \kappa(A) > 0\,, \\ \mathrm{Mod}(A) l_q(\gamma_0) = d & \text{if } \kappa(A) = 0\,. \end{cases}$$

Every annulus of large modulus contains a primitive annulus with comparable modulus. This is a consequence of Corollary 5.5 in [MaS], and was proven again in the following form in [Mi1, Th. 4.6].

Theorem 2.4 (Masur–Smillie, Minsky)**.** *There exists an $\epsilon_0 > 0$ such that, for a curve α in S, if $l_\sigma([\alpha]) \leq \epsilon_0$, then there exists a primitive annulus A*

such that

$$\frac{1}{l_\sigma([\alpha])} \asymp \mathrm{Mod}(A)\,.$$

Throughout this paper, ϵ_0 is a fixed constant smaller than the Margulis constant, such that the above theorem and Theorem 2.5 are true.

2.5 Product regions in Teichmüller space. The Teichmüller space of S, $\mathcal{T}(S)$, is the space of conformal structures on S up to isotopy. The Teichmüller distance between two points σ_1 and σ_2 is defined as

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) = \tfrac{1}{2} K(\sigma_1, \sigma_2)\,,$$

where $K(\sigma_1, \sigma_2)$ is the smallest quasi-conformal dilatation of a homeomorphism from σ_1 to σ_2. Let Γ be a system of disjoint curves on S, and let $\mathrm{Thin}_\epsilon(\Gamma)$ denote the set of all $\sigma \in \mathcal{T}(S)$ such that, for all $\gamma \in \Gamma$, the length of γ in σ, $l_\sigma(\gamma)$, is less than or equal to ϵ. Let $\mathcal{T}_\Gamma$ denote the product space

$$\mathcal{T}(S \setminus \Gamma) \times \prod_{\gamma \in \Gamma} \mathbb{H}_\gamma\,,$$

where $S \setminus \Gamma$ is considered as a punctured space and each $\mathbb{H}_\gamma$ is a copy of the hyperbolic plane. Endow $\mathcal{T}_\Gamma$ with the sup metric. Minsky has shown, for small enough ϵ, that $\mathrm{Thin}_\epsilon(\Gamma)$ has a product structure.

Theorem 2.5 (Minsky [Mi3]). *The Fenchel–Nielsen coordinates on $\mathcal{T}(S)$ give rise to a natural homeomorphism $\pi\colon \mathcal{T}(S) \to \mathcal{T}_\Gamma$. There exists an $\epsilon_0 > 0$ sufficiently small that this homeomorphism restricted to* $\mathrm{Thin}_{\epsilon_0}(\Gamma)$ *distorts distances by a bounded additive amount.*

Note that $\mathcal{T}(S \setminus \Gamma) = \prod_Y \mathcal{T}(Y)$, where the product is over all connected components Y of $S \setminus \Gamma$. Let π_0 denote the component of π mapping to $\mathcal{T}(S \setminus \Gamma)$, let π_Y denote the component mapping to $\mathcal{T}(Y)$, and, for $\gamma \in \Gamma$, let π_γ denote the component mapping to $\mathbb{H}_\gamma$. For the rest of the paper, we fix $L_0 > 0$ such that, for a hyperbolic metric σ on S, if $l_\sigma(\alpha) \geq \epsilon_0$, then there exists a curve β intersecting α with $l_\sigma(\beta) \leq L_0$.

3 Behavior of a Geodesic in the Thin Part of Teichmüller Space

In this section, we prove Theorem 1.2, restated as Theorem 3.1, and study how the combinatorics of short markings change along a Teichmüller geodesic. We show that, for every curve α in S, there exists a connected interval where α is *short* (Corollary 3.3), and the projections of the short markings to a subsurface can only change while all the boundaries of that subsurface are short (Proposition 3.7). This is an essential component of the proof of the main theorem.

3.1 Teichmüller geodesics. For $t \in \mathbb{R}$, let q_t be the quadratic differential obtained from q by scaling its vertical foliation by a factor of e^t, and its horizontal foliation by a factor of e^{-t}. Define $g(t)$ to be the conformal structure corresponding to q_t. Then $g\colon \mathbb{R} \to \mathcal{T}(S)$ is a geodesic in $\mathcal{T}(S)$ parametrized by arc length. For a curve α in S, the horizontal and vertical lengths of α vary with time as follows:

$$h_{q_t}(\alpha) = h_q(\alpha)e^t \quad \text{and} \quad v_{q_t}(\alpha) = v_q(\alpha)e^{-t}. \tag{2}$$

We say α is *balanced*, *mostly horizontal* or *mostly vertical* at time t if, respectively, $v_t([\alpha]) = h_t([\alpha])$, $v_t([\alpha]) \leq h_t([\alpha])$ or $v_t([\alpha]) \geq h_t([\alpha])$.

3.2 Hyperbolic length along a geodesic. The behavior of the hyperbolic length of a curve along a Teichmüller geodesic is somewhat mysterious. For the Weil–Petersson metric on $\mathcal{T}(S)$, the hyperbolic length of a curve along a geodesic is a convex function of time. In the Teichmüller metric, the quadratic differential length of a curve is also convex. The following result is a weaker but analogous statement. It roughly states that a curve assumes its shortest length when it is balanced and the length is *non-decreasing* as one moves away in either direction. Let σ_t denote the hyperbolic metric on $g(t)$.

Theorem 3.1. *Let g be a geodesic in $\mathcal{T}(S)$ and α be a curve in S. Assume α is balanced at t_α and $s \geq t_\alpha$ (respectively, $s \leq t_\alpha$). Then, for any $t \geq s$ ($t \leq s$), we have*

$$\frac{1}{l_{\sigma_s}([\alpha])} \succ \frac{1}{l_{\sigma_t}([\alpha])}\,. \tag{3}$$

Remark 3.2. The reader should be mindful of the additive error in (3). When $l_{\sigma_s}([\alpha])$ and $l_{\sigma_t}([\alpha])$ are both large, $1/l_{\sigma_s}([\alpha])$ and $1/l_{\sigma_t}([\alpha])$ are small and within additive error of each other, so (3) is automatically true and provide no new information. However, if we only know that $l_{\sigma_s}([\alpha])$ is large, then (3) implies that $l_{\sigma_t}([\alpha])$ is bounded below for all $t \geq s$.

Proof. Let F_t be the flat annulus corresponding to α in q_t. (Note that when the q_t-geodesic representative of α is unique, F_t is degenerate and $\operatorname{Mod}(F_t) = 0$.) The modulus of F_t is maximum at t_α, and, for $t \in \mathbb{R}$,

$$\operatorname{Mod}(F_t) \asymp \operatorname{Mod}(F_{t_0})\, e^{-2\,|t-t_\alpha|}. \tag{4}$$

Let A_t be as in Theorem 2.4 for hyperbolic metric σ_t, quadratic differential q_t, and curve α (if $l_t(\alpha) \geq \epsilon_0$, there is nothing to prove). If A_t is flat, then

$$\begin{aligned}
\frac{1}{l_{\sigma_s}([\alpha])} &\asymp \frac{1}{\mathrm{Ext}_{\sigma_s}([\alpha])} && \text{([M])}\\
&\succ \mathrm{Mod}(F_s) && \text{(by definition of } \mathrm{Ext}_{\sigma_s}(\alpha))\\
&\succ \mathrm{Mod}(F_t) && \text{(Equation (4))}\\
&\geq \mathrm{Mod}(A_t) && (A_t \subset F_t)\\
&\asymp \frac{1}{l_{\sigma_t}([\alpha])}\,. && \text{(Theorem 2.4)}
\end{aligned}$$

Assume A_t is not flat. Let d be the distance between the boundary components of A_t and l be the length of the inner boundary of A_t. Let β be a curve intersecting α whose hyperbolic length at s is less than L, for some L such that $e^L \asymp 1/l_{\sigma_s}([\alpha])$. Using the *collar lemma* (Theorem 2.2), we have

$$\frac{1}{l_{\sigma_s}([\alpha])} \succ \log \frac{l_{q_s}([\beta])}{l_{q_s}([\alpha])}\,. \tag{5}$$

But α is mostly vertical at s; therefore, for $t \geq s$,

$$l_{q_t}([\alpha]) \asymp l_{q_{t_s}}([\alpha])\, e^{t-s}.$$

The quadratic differential length of any curve grows at most exponentially; that is, for $t \geq s$,

$$l_{q_t}([\beta]) \prec l_{q_{t_s}}([\beta])\, e^{t-s}.$$

Therefore,

$$\frac{l_{q_t}([\beta])}{l_{q_t}([\alpha])} \leq \frac{l_{q_s}([\beta])}{l_{q_s}([\alpha])}\,. \tag{6}$$

We also have $l_{q_t}([\beta]) \geq d$ (β has to cross A_t) and $l_{q_t}([\alpha]) \leq l$ (α and the inner boundary of A_t are homotopic). Therefore,

$$\begin{aligned}
\frac{1}{l_{\sigma_s}([\alpha])} &\succ \log \frac{l_{q_s}([\beta])}{l_{q_s}([\alpha])} && \text{(Equation (5))}\\
&\geq \log \frac{l_{q_t}([\beta])}{l_{q_t}([\alpha])} && \text{(Equation (6))}\\
&\geq \log \tfrac{d}{l}\\
&\succ \mathrm{Mod}(A_t) && \text{(Lemma 2.3)}\\
&\asymp \frac{1}{l_{\sigma_t}([\alpha])}\,. && \text{(Theorem 2.4)}
\end{aligned}$$

□

COROLLARY 3.3. *There exists $\epsilon_1 \leq \epsilon_0$ such that, for any geodesic in the Teichmüller space and any curve α in S, there exists a connected (perhaps empty) interval I_α such that*

1. *For* $t \in I_\alpha$, $l_{\sigma_t}([\alpha]) \le \epsilon_0$; *and*
2. *For* $t \notin I_\alpha$, $l_{\sigma_t}([\alpha]) \ge \epsilon_1$.

The intersection of connected intervals is a connected interval (or an empty set). Therefore, a similar statement is also true for subsurfaces.

COROLLARY 3.4. *Let* ϵ_0, ϵ_1 *and* g *be as above. For every subsurface* Y, *there exists a connected interval* I_Y *such that*

1. *For* $t \in I_Y$, *the hyperbolic lengths of all boundary components of* Y *at* σ_t *are less than or equal to* ϵ_0, *and*
2. *For* $t \notin I_Y$, *there exists a boundary component of* Y *whose hyperbolic length at* σ_t *is greater than or equal to* ϵ_1.

3.3 A lower bound for distance in the Teichmüller space. Our main theorem describes how the distance between two points in Teichmüller space can be estimated by measuring the combinatorial complexity of curves of bounded size. Here we show that, if two curves of bounded length in σ_1 and σ_2 intersect each other a large number of times, then σ_1 and σ_2 are far apart in $\mathcal{T}(S)$.

First we recall some properties of the extremal length. Let $\operatorname{Ext}_\sigma(\alpha)$ denote the extremal length of α in σ. Minsky has shown (see [Mi2]) that, for curves α and β in S, and $\sigma \in \mathcal{T}(S)$,

$$\operatorname{Ext}_\sigma(\alpha)\ \operatorname{Ext}_\sigma(\beta) \ge i_S(\alpha,\beta)^2. \tag{7}$$

Kerckhoff's theorem (see [K]) states that, for points σ_1 and σ_2 in $\mathcal{T}(S)$,

$$K(\sigma_1,\sigma_2) = e^{2\, d_{\mathcal{T}}(\sigma_1,\sigma_2)} = \sup_\alpha \frac{\operatorname{Ext}_{\sigma_1}(\alpha)}{\operatorname{Ext}_{\sigma_2}(\alpha)}\,, \tag{8}$$

where the sup is over all curves on S. We also know (see [M]) that, if the hyperbolic length of α is short (say, $l_\sigma(\alpha) \le L_0$), then

$$l_\sigma(\alpha) \asymp \operatorname{Ext}_\sigma(\alpha)\,. \tag{9}$$

PROPOSITION 3.5. *Assume, for some* $\sigma_1, \sigma_2 \in \mathcal{T}(S)$ *and curves* α *and* β *in* S, *that* $l_{\sigma_1}(\alpha) \le L_0$ *and* $l_{\sigma_2}(\beta) \le L_0$. *Then*

$$d_{\mathcal{T}}(\sigma_1,\sigma_2) \succ \log i_S(\alpha,\beta)\,.$$

Proof. We have

$$\begin{aligned} i_S(\alpha,\beta)^2 &\le \operatorname{Ext}_{\sigma_1}(\alpha)\ \operatorname{Ext}_{\sigma_1}(\beta) && \text{(Equation (7))}\\ &\le \operatorname{Ext}_{\sigma_1}(\alpha)\ \operatorname{Ext}_{\sigma_2}(\beta)\, K(\sigma_1,\sigma_2) && \text{(Equation (8))}\\ &\prec L_0^2\, K(\sigma_1,\sigma_2)\,. && \text{(Equation (9))}\end{aligned}$$

Note that L_0 is a fixed constant depending on S only. By taking the logarithm of both sides, we obtain the desired inequality. □

3.4 Combinatorics of short markings along a Teichmüller geodesic. For $t \in \mathbb{R}$, let μ_t be the *shortest marking* in σ_t, constructed as follows. Let α_1 be the shortest curve in S and α_2 be the shortest curve disjoint from α_1, and so on, to form a pants decomposition of S. Then, let the transverse curve β_i be the shortest curve intersecting α_i and disjoint from α_j, $i \neq j$. (There may be finitely many such markings.) Proposition 3.7 states that the projection of these markings to a subsurface Y stays in a bounded neighborhood in $\mathcal{C}(Y)$ while the geodesic is outside of the thin part of $\mathcal{T}(S)$ corresponding to Y. The proof makes an essential use of the following theorem.

Theorem 3.6 [R, Th. 5.5]**.** *Let α be a curve in S, β be a curve intersecting α non-trivially, and Y be a component of $S \setminus \alpha$ (Y is allowed to be an annulus). Assume $l_{\sigma_t}([\beta]) \leq L$. We have*

1. *If α is mostly vertical, then*
$$i_Y(\beta, \nu_-) \prec D_L \,.$$
2. *If α is mostly horizontal, then*
$$i_Y(\beta, \nu_+) \prec D_L \,.$$

Here, D_L is a constant depending on L, with $\log D_L \asymp e^L$.

PROPOSITION 3.7. *If $[r,s] \cap I_Y = \emptyset$, then*
$$d_Y(\mu_r, \mu_s) = O(1) \,.$$

REMARK 3.8. Here $O(1)$ depends on the choice of ϵ_1. But ϵ_1 is a universal constant depending on the topology of S only. Therefore, we consider $O(1)$ to be a universal constant depending only on the topology of S as well.

Proof. Let L_1 be such that every curve of length larger than ϵ_1 in a hyperbolic surface with geodesic boundary has a transverse curve of length less than L_1. For $t \in [r,s]$, there exists a boundary component γ_t of Y whose σ_t-length is larger than ϵ_1. Therefore, the marking μ_t contains a curve α_t with $l_{\sigma_t}(\alpha_t) \leq L_1$ that intersects Y nontrivially. The projection of μ_t to Y has bounded diameter. Therefore it is sufficient to prove $d_Y(\alpha_r, \alpha_s) = O(1)$.

The curve γ_t is either mostly horizontal or mostly vertical at time t. The set of times at which Y has a boundary component of length larger than or equal to ϵ_1 which is mostly horizontal (or mostly vertical) is closed. Therefore, either

1. γ_r and γ_s are both mostly horizontal or both mostly vertical; or

2. for some $t \in [r,s]$, there are two curves γ_t and γ'_t whose lengths at σ_t are larger than or equal to ϵ_1, and one is mostly horizontal and the other is mostly vertical (possibly $\gamma_t = \gamma'_t$ and γ_t is balanced).

Case 1: If γ_r and γ_s are mostly vertical, Theorem 3.6 implies that

$$i_Y(\alpha_r, \nu_-) \prec D_{L_1} \quad \text{and} \quad i_Y(\alpha_s, \nu_-) \prec D_{L_1}\,.$$

Therefore, using Lemma 2.1,

$$d_Y(\alpha_s, \nu_-) \prec \log i_Y(\alpha_s, \nu_-) \prec \log D_{L_1}.$$

Similarly, $d_Y(\alpha_r, \nu_+) = O(1)$. This implies that $d_Y(\alpha_r, \alpha_s) = O(1)$. The proof is similar if γ_r and γ_s are both mostly horizontal.

Case 2: Assume (without loss of generality) that γ_t is mostly horizontal and γ'_t is mostly vertical. Let α_t and α'_t be the corresponding transverse curves in μ_t of length less than L_1. By the above argument,

$$d_Y(\alpha_t, \nu_+) = O(1) \quad \text{and} \quad d_Y(\alpha'_t, \nu_-) = O(1)\,.$$

But the extremal lengths of α_t and α'_t are bounded by a constant depending on L_1. Equation (7) implies that $i_S(\alpha_t, \alpha'_t) = O(1)$, and, by Lemma 2.1, $d_Y(\alpha_t, \alpha'_t) = O(1)$. Therefore,

$$d_Y(\nu_+, \nu_-) = O(1)\,. \tag{10}$$

Again, as above, the projection of each of α_s and α_r to Y is close to the projection of either ν_+ or ν_- to Y. Thus, (10) and the triangle inequality for d_Y imply that

$$d_Y(\alpha_r, \alpha_s) = O(1)\,.$$

□

Corollary 3.9. *If $I_Y = [c,d] \subset [a,b]$, then*

$$d_Y(\mu_a, \mu_b) \asymp d_Y(\mu_c, \mu_d)\,.$$

4 Twisting in the Hyperbolic Metric vs. Twisting in the Quadratic Differential Metric

Let α be a curve in S. Having a metric in S enables us to define a twisting parameter for curves that cross α. This, roughly speaking, is the number of times that a given curve twists around α in comparison with an arc that is perpendicular to the geodesic representative of α. In this section we define a twisting parameter for ν_+ and ν_- using metrics given by q and σ, and we study how these two quantities are related. We use this to prove Theorem 1.3 at the end of this section.

Let $\bar{S}$ be the annular cover of S with respect to α. Let $\bar{q}$, $\bar{\nu}_+$ and $\bar{\nu}_-$ be the lifts of q, ν_+ and ν_- to $\bar{S}$, respectively, and $\bar{\beta}_q$ be a geodesic arc connecting the boundaries of $\bar{S}$ that is perpendicular (in $\bar{q}$) to the geodesic representative of the core of $\bar{S}$, $\bar{\alpha}$. We define the twisting parameter of ν_+ around α in q to be the maximum intersection number of a leaf of $\bar{\nu}_+$ and an arc $\bar{\beta}_q$, and we denote it by $tw_q(\nu_+, \alpha)$. When it is clear what α is, we denote this by tw_q^+. The twisting parameter tw_q^- of ν_- around α in q is defined similarly. Note that the maximum intersection number is at least one, that is, $tw_q^{\pm}$ are positive integers.

Let F be the flat annulus in q corresponding to α and let β_q be an arc connecting the boundaries of F that is perpendicular to the boundaries of F. The intersection number of the lift of a leaf of ν_+ with $\bar{\beta}_q$ is (up to small additive error) equal to the intersection number of the restriction of this leaf to F with β_q. Therefore, to compute $tw_q^{\pm}$, it is sufficient to understand the picture in F. Consider an isometric embedding of the universal cover of F in $\mathbb{R}^2$ such that the leaves of horizontal foliations are parallel to the x-axis and the leaves of vertical foliations are parallel to the y-axis (see Figure 1).

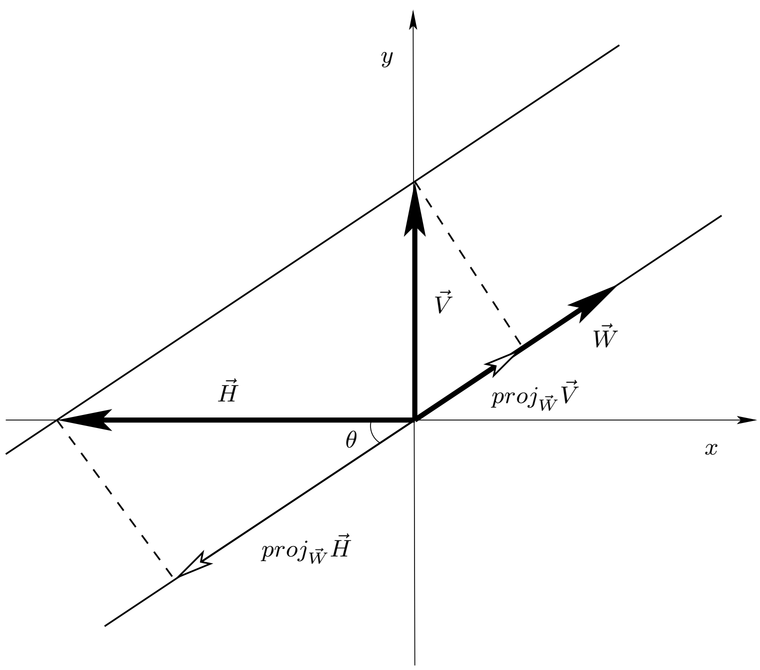

Figure 1: The universal cover of F

Let $\vec{W}$ be the vector representing the translation that generates the deck translation group. Let $\vec{H}$ be the lift of a leaf of ν_- passing through the origin and $\vec{V}$ be the same for ν_+. From the above discussion, we have

$$tw_q(\nu_-, \alpha) \stackrel{+}{\asymp} \frac{\|\mathrm{Proj}_{\vec{W}} \vec{H}\|}{\|\vec{W}\|} \quad \text{and} \quad tw_q(\nu_+, \alpha) \stackrel{+}{\asymp} \frac{\|\mathrm{Proj}_{\vec{W}} \vec{V}\|}{\|\vec{W}\|}.$$

Let θ be the angle between $\vec{W}$ and the x-axis. It is easy to see, using similar triangles, that

$$\frac{\|\mathrm{Proj}_{\vec{W}} \vec{H}\|}{\|\mathrm{Proj}_{\vec{W}} \vec{V}\|} = \frac{\sin^2\theta}{\cos^2\theta}.$$

We also have $\frac{h_q([\alpha])}{v_q([\alpha])} = \frac{\sin\theta}{\cos\theta}$. Therefore,

$$\frac{tw_q^-}{tw_q^+} \stackrel{+}{\asymp} \frac{h_q([\alpha])^2}{v_q([\alpha])^2}. \tag{11}$$

This is a very useful equation that allows us to compute the q-twisting parameter of horizontal and vertical foliations around α along a Teichmüller geodesic (see equation (16)).

We define the twisting parameter for a hyperbolic metric as follows. Let β_σ be the shortest transverse curve to α in the hyperbolic metric σ. Define

$$tw_\sigma^+ = i(\nu_+, \beta_\sigma) \quad \text{and} \quad tw_\sigma^- = i(\nu_-, \beta_\sigma).$$

We would like to prove a statement similar to equation (11) for σ-twisting parameters. However, giving good estimates for $tw_\sigma^\pm$ is difficult when α is very short. The errors in our estimates get larger as $l_\sigma(\alpha)$ gets smaller.

Let $\bar{\beta}_\sigma$ be the lift of β_σ to $\bar{S}$ whose end points are in different boundary components of $\bar{S}$. Our strategy is to relate q- and σ-twisting parameters by providing an upper bound for $i(\bar{\beta}_q, \bar{\beta}_\sigma)$.

LEMMA 4.1. *If $i(\bar{\beta}_q, \bar{\beta}_\sigma) = n$, then*

$$\mathrm{Ext}_\sigma(\beta_\sigma) \stackrel{\cdot}{\succ} n^2\, l_\sigma(\alpha).$$

Proof. By definition of the extremal length, for any metric τ on S in the conformal class of σ,

$$\mathrm{Ext}_\sigma(\beta_\sigma) \geq \frac{l_\tau(\beta_\sigma)^2}{\mathrm{area}_\tau(S)}.$$

To find a lower bound for $\mathrm{Ext}_\sigma(\beta_\sigma)$, we need to find an appropriate metric τ. First we establish some notation. Let R be the largest regular neighborhood of F that is still an annulus. Denote the boundary components of R by α_0 and α_c, where c is the q-distance between the boundaries of R. For $t \in (0, c)$, let α_t be a curve in R that is equidistant from a q-geodesic representative of α and whose q-distance from α_0 is t. These curves give a foliation of R into curves in the homotopy class of α. There is a subinterval $[a,b]$ of $[0,c]$ such that, for $t \in [a,b]$, α_t is a q-geodesic representative of α. This gives a division of R into three pieces, the flat annulus F containing all α_t, $t \in [a,b]$, and two expanding annuli R_1 and R_2 on the sides. Theorem 2.4 implies that $\mathrm{Mod}(R) \asymp 1/l_\sigma(\alpha)$. Using Lemma 2.3, we have

$$\frac{1}{l_\sigma(\alpha)} \stackrel{\cdot}{\asymp} \log \frac{a}{l_q([\alpha])} + \frac{(b-a)}{l_q([\alpha])} + \log \frac{(c-b)}{l_q([\alpha])}\,. \tag{12}$$

As t changes in the interval $[b,c]$, the length of α_t increases. The rate of change is equal to the curvature of α_t, which is bounded above and below by constants depending on the topology of S only. A similar statement is true for R_1 as well. Therefore,

$$l_q(\alpha_t) \stackrel{\cdot}{\asymp} \begin{cases} l_q([\alpha]) + (a-t) & \text{if} \quad t \in [0,a]\,, \\ l_q([\alpha]) & \text{if} \quad t \in [a,b]\,, \\ l_q([\alpha]) + (t-b) & \text{if} \quad t \in [b,c]\,. \end{cases} \tag{13}$$

Denote $l_q(\alpha_t)$ by λ_t.

Let Z be the union of R; the λ_0-neighborhood, N_0, of α_0; and the λ_c-neighborhood, N_c, of α_c. Define the metric τ in S in the conformal class of q as follows: if x lies on a curve α_t in R, then we scale the q-metric at x by a factor of $1/\lambda_t$; if x is outside of R and in N_0, then we scale the q-metric at x by a factor of $1/\lambda_0$; if x is outside of R and in N_c, then we scale the q-metric at x by a factor of $1/\lambda_c$ (if x is in both N_0 and N_c, then we scale the q-metric by a factor of $\max(1/\lambda_0, 1/\lambda_c)$); and, if x is outside of Z, then we scale the q-metric at x by a small enough factor so that the τ-area of S is comparable with the τ-area of Z. Note that $\mathrm{area}_q N_0 \stackrel{\cdot}{\prec} \lambda_0^2$ and $\mathrm{area}_q N_c \stackrel{\cdot}{\prec} \lambda_c^2$. We have

$$\begin{aligned}
\mathrm{area}_\tau(S) &\stackrel{\cdot}{\asymp} \mathrm{area}_\tau(Z) \\
&\le \mathrm{area}_\tau N_0 + \mathrm{area}_\tau N_c + \mathrm{area}_\tau A \\
&\stackrel{\cdot}{\asymp} 1 + 1 + \int_0^c 1 \,.\, \frac{dt}{\lambda_t} \\
&\stackrel{\cdot}{\asymp} 2 + \log \frac{a}{l_q([\alpha])} + \frac{(b-a)}{l_q([\alpha])} + \log \frac{(c-b)}{l_q([\alpha])} && \text{(Equation (13))} \\
&\stackrel{\cdot}{\asymp} \frac{1}{l_\sigma(\alpha)}\,. && \text{(Equation (12))}
\end{aligned}$$

Let $\bar{R}$ be the lift of R to $\bar{S}$ that is an annulus, and let $\bar{\alpha}_t$ be the lift of α_t that is in $\bar{R}$ (this is to ensure that $\bar{\alpha}_t$ is a closed curves not an infinite line). Let $\bar{\omega}$ be a sub-arc of $\bar{\beta}_\sigma$ with end points in $\bar{\beta}_q$ that goes around $\bar{S}$

once, that is, if $\bar{\omega}'$ is the sub-arc of $\bar{\beta}_q$ connecting the end points of $\bar{\omega}$, then $\bar{\gamma} = \bar{\omega} \cup \bar{\omega}'$ is a curve in the homotopy class of the core of $\bar{S}$. Let γ be the projection of $\bar{\gamma}$ to S. Then γ is in the homotopy class of α and therefore must intersect R (otherwise, R would not be maximal). Hence, $\bar{\gamma}$ must intersect $\bar{R}$. But $\bar{\beta}_q$ is perpendicular to $\bar{\alpha}_t$, and, once it exits $\bar{R}$, it never returns. Therefore, $\bar{\omega}$ must intersect $\bar{R}$ as well.

Let $\bar{\alpha}_s$ be an equidistant curve in $\bar{R}$ intersecting $\bar{\omega}$ that has the shortest $\bar{q}$-length. We claim that

$$l_{\bar{q}}(\bar{\omega}) \geq l_{\bar{q}}(\bar{\alpha}_t) = \lambda_t\,.$$

Assume $s > b$. The curve $\bar{\alpha}_s$ divides $\bar{S}$ into two annuli. Let B be the annulus that contains $\bar{\alpha}_c$. For $t \in [b, s)$, the $\bar{q}$-length of $\bar{\alpha}_t$ is less than the $\bar{q}$-length of α_s. By assumption $\bar{\alpha}_s$ is the shortest equidistant curve intersecting $\bar{\omega}$, therefore, $\bar{\omega} \subset B$.

The curvature of $\bar{\alpha}_t$ with respect to B is non-positive at all points. Therefore, the closest-point projection from B to $\bar{\alpha}_t$ is length-decreasing. But the end points of $\bar{\omega}$ project to the same point in $\bar{\alpha}_t$ (because $\bar{\beta}_q$ is perpendicular to $\bar{\alpha}_t$), and the projection covers $\bar{\alpha}_t$ completely. Therefore, $l_{\bar{q}}(\bar{\omega}) \geq l_{\bar{q}}(\bar{\alpha}_t)$ in this case.

A similar argument holds if $t < a$. If $t \in [a, b]$, then $\bar{\omega}$ could intersect $\bar{\alpha}_t$ transversally, but, in this case, $\bar{\alpha}_t$ is a $\bar{q}$-geodesic and the curvature of $\bar{\alpha}_t$ is non-positive with respect to both annuli in $\bar{S} \setminus \alpha_t$. Therefore, the claim is true in all cases.

Let ω be the projection of $\bar{\omega}$ to S. If ω exits Z, then its τ-length is larger than the τ-distance between R and ∂Z, which is equal to 1. Otherwise, $\omega \subset Z$. Then, at each point in ω, τ is obtained from q by scaling by a factor of at least $1/\lambda_t$. Therefore,

$$l_\tau(\omega) \geq \tfrac{1}{\lambda_t}\, l_q(\omega) \geq 1\,.$$

There are $(n-1)$ arcs like $\bar{\omega}$, and they all project down to different sub-arcs of β_σ. Therefore,

$$l_\tau(\beta_\sigma) \geq n\,.$$

This implies that

$$\mathrm{Ext}_\sigma(\beta_\sigma) \geq \frac{l_\tau(\beta_\sigma)^2}{\mathrm{area}_\tau S} \stackrel{\cdot}{\succ} \frac{n^2}{1/l_\sigma(\alpha)} = n^2\, l_\sigma(\alpha)\,. \qquad \square$$

COROLLARY 4.2. *For $\bar{\beta}_\sigma$ and $\bar{\beta}_q$ as before, we have*

$$i(\bar{\beta}_\sigma, \bar{\beta}_q) \prec \frac{1}{l_\sigma(\alpha)}\,.$$

Proof. The curve β_σ is the shortest (in σ) transverse curve to α. Therefore, $\mathrm{Ext}_\sigma(\beta_\sigma) \asymp 1/l_\sigma(\alpha)$. Applying Lemma 4.1 we get

$$\frac{1}{l_\sigma(\alpha)} \succ i(\bar\beta_\sigma, \bar\beta_q)^2\, l_\sigma(\alpha)\,,$$

Taking the square root of both side we obtain the desired inequality. □

The following theorem is an immediate consequence of the definitions of the twisting parameters and of Corollary 4.2.

Theorem 4.3. *The two twisting parameters are the same up to an additive error comparable to $1/l_\sigma(\alpha)$. That is,*

$$tw_\sigma^\pm = tw_q^\pm \pm O\left(\frac{1}{l_\sigma(\alpha)}\right).$$

4.1 The twisting parameter along a Teichmüller geodesic. In this section, we give estimates for the twisting parameters of $\nu_\pm$ around a curve α in σ_t. Let $d_\alpha = d_\alpha(\nu_+, \nu_-)$ (which is equal to $i_\alpha(\nu_+, \nu_-)$). Note that ν_+ and ν_- twist around α in different directions. Therefore,

$$d_\alpha \overset{+}{\asymp} tw_q^+ + tw_q^-\,. \tag{14}$$

If α is not very short in σ_t, say $l_{\sigma_t}(\alpha) \geq \epsilon_0$, then there exists a curve intersecting α non-trivially whose σ_t-length is not greater than L_0. Theorem 3.6 and (14) imply that

$$\begin{cases} \text{if } \alpha \text{ is mostly horizontal,} & tw_{\sigma_t}^- \overset{+}{\asymp} d_\alpha \quad \text{and} \quad tw_{\sigma_t}^+ \overset{+}{\asymp} 0\,, \\ \text{if } \alpha \text{ is mostly vertical,} & tw_{\sigma_t}^- \overset{+}{\asymp} 0 \quad \text{and} \quad tw_{\sigma_t}^+ \overset{+}{\asymp} d_\alpha\,. \end{cases} \tag{15}$$

Assume α is balanced at t_α. Using equations (11) and (2), we get

$$\frac{tw_{q_t}^-}{tw_{q_t}^+} = \frac{e^{2(t-t_\alpha)}\, h_{q_{t_\alpha}}(\alpha)^2}{e^{-2(t-t_\alpha)}\, v_{q_{t_\alpha}}(\alpha)^2}\,.$$

But $h_{q_{t_\alpha}}(\alpha) = v_{q_{t_\alpha}}(\alpha)$. Therefore,

$$tw_{q_t}^- \overset{+}{\asymp} \frac{d_\alpha e^{2(t-t_\alpha)}}{e^{2(t-t_\alpha)} + e^{-2(t-t_\alpha)}} \quad \text{and} \quad tw_{q_t}^+ \overset{+}{\asymp} \frac{d_\alpha e^{-2(t-t-\alpha)}}{e^{2(t-t-\alpha)} + e^{-2(t-t_\alpha)}}\,. \tag{16}$$

This and Theorem 4.3 prove Theorem 1.3. The following proposition is a different statement for the same basic fact.

Proposition 4.4. *Let $\sigma_t \in \mathcal{T}(S)$ and α be a curve in S with $l_{\sigma_t}(\alpha) \leq \epsilon_0$. Let σ_t' be the point in $\mathcal{T}(S)$ obtained from σ_t by twisting along α such that*

$$tw_{\sigma_t'}^- = \frac{d_\alpha e^{-2(t-t_\alpha)}}{e^{2(t-t_\alpha)} + e^{-2(t-t_\alpha)}}\,.$$

Then $d_{\mathcal{T}}(\sigma_t, \sigma_t') = O(1)$.

Proof. Consider $\pi\colon \mathcal{T}(S) \to \mathcal{T}(S \setminus \alpha) \times \mathbb{H}_\alpha$. Recall that the map π is defined by Fenchel–Nielsen coordinates and $\mathbb{H}_\alpha$ parametrizes the length and the twisting around α. In fact, considering $\mathbb{H}_\alpha$ as the upper half plane, the y-coordinate of $\pi_\alpha(\sigma)$ equals $1/l_\sigma(\alpha)$ and its x-coordinate is the twisting parameter of σ around α compared with some fixed marking. (See [Mi3] for the precise definition). Since $\pi_0(\sigma_t) = \pi_0(\sigma'_t)$, Theorem 2.5 implies that

$$d_{\mathcal{T}}(\sigma_t, \sigma'_t) \stackrel{+}{\asymp} d_{\mathbb{H}_\alpha}\big(\pi_\alpha(\sigma_t)\pi_\alpha(\sigma'_t)\big)\,,$$

and from the geometry of the hyperbolic plane we have

$$d_{\mathbb{H}_\alpha}\big(\pi_\alpha(\sigma_t), \pi_\alpha(\sigma'_t)\big) \stackrel{+}{\asymp} \log\big(l_{\sigma_t}(\alpha)(tw^-_{\sigma_t} - tw^-_{\sigma'_t})\big)\,.$$

Theorem 4.3 states that the σ_t-twisting and the q_t-twisting parameters of ν_- are equal up to an additive error that is comparable with $1/l_{\sigma_t}(\alpha)$. Therefore, the right-hand side of the above equation is uniformly bounded. Thus

$$d_{\mathcal{T}}(\sigma_t, \sigma'_t) = O(1)\,. \qquad \square$$

5 Proof of the Main Theorem

In this section we prove Theorem 1.1. In §5.1, we show how a lower bound for the Teichmüller distance between two points in $\mathcal{T}(S)$ can be obtained by the combinatorial complexity between their short markings. In §5.2, we give an upper bound for the distance between two points in the Teichmüller space by constructing a path in $\mathcal{T}(S)$ of length comparable with the estimate given in Theorem 1.1.

5.1 Lower estimate. Let $g\colon [a,b] \to \mathcal{T}(S)$ be the geodesic segment in the Teichmüller space connecting σ_a to σ_b. Recall that σ_t is the hyperbolic metric of $g(t)$, and μ_t is the short-marking on S corresponding to σ_t.

LEMMA 5.1. *Let Y be a subsurface that is not an annulus and $I = I_Y \cap [a,b]$. Then*

$$|I| \succ d_Y(\mu_a, \mu_b)\,.$$

Proof. Let $I = [c,d]$, $\tau_c = \pi_Y(\sigma_c)$ and $\tau_d = \pi_Y(\sigma_d)$ (see Theorem 2.5). Let η_c and η_d be the short-markings on Y corresponding to τ_c and τ_d, respectively. In fact, $\eta_c \subset \mu_c$ and $\eta_d \subset \mu_d$. We have

$$\begin{aligned} |I| &\succ d_{\mathcal{T}(Y)}(\tau_c, \tau_d)\,, && \text{(Theorem 2.5)}\\ &\succ \log i_Y(\eta_c, \eta_d)\,, && \text{(Proposition 3.5)}\\ &\succ d_Y(\eta_c, \eta_d)\,. && \text{(Lemma 2.1)} \end{aligned}$$

But $d_Y(\eta_c, \eta_d) \asymp d_Y(\mu_c, \mu_d)$ (because they have the same projections to Y). Also, by Proposition 3.7, we have

$$d_Y(\mu_a, \mu_c) = O(1) \quad \text{and} \quad d_Y(\mu_d, \mu_b) = O(1)\,.$$

This proves the lemma. □

A similar lemma is true when the subsurface is an annulus. The difference is that, in Lemma 5.1, there is no restriction on the lengths of the boundaries of Y; but, for the next lemma to be true, we have to assume that α is not very short in σ_a and σ_b. the proofs are almost identical.

Lemma 5.2. *Let α be a curve in S such that $l_{\sigma_a}(\alpha) \geq \epsilon_0$ and $l_{\sigma_b}(\alpha) \geq \epsilon_0$, and let $I = I_\alpha \cap [a,b]$. Then*

$$|I| \succ \log d_\alpha(\mu_a, \mu_b)\,.$$

Proof. Since α is not short at either end, either I_α is disjoint from $[a,b]$ or it is a subset of $[a,b]$. If $I_\alpha \cap [a,b] = \emptyset$, then Proposition 3.7 implies the lemma. If $I_\alpha = [c,d] \subset [a,b]$, then, by Corollary 3.9,

$$d_\alpha(\mu_a, \mu_b) \asymp d_\alpha(\mu_c, \mu_d)\,.$$

Let β_c and β_d be curves transverse to α in markings μ_c and μ_d, respectively. We have

$$i(\beta_c, \beta_d) = d_\alpha(\mu_c, \mu_d)\,.$$

As in the previous lemma, using Theorem 2.5 and Proposition 3.5, we have

$$|I_\alpha| \succ \log i(\beta_c, \beta_d)\,.$$

The combination of the last three equations proves the lemma. □

The following proposition provides a lower bound for the Teichmüller distance between two points in the thick part of $\mathcal{T}(S)$.

Proposition 5.3. *Let σ_1, σ_2 be in the ϵ_0-thick part of $\mathcal{T}(S)$ and μ_1 and μ_2 be the short-markings in σ_1 and σ_2, respectively. There exists a $k_0 > 0$ such that*

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \sum_Y \big[d_Y(\mu_1, \mu_2)\big]_{k_0} + \sum_\alpha \log \big[d_\alpha(\mu_1, \mu_2)\big]_{k_0}.$$

Proof. Let $g\colon [a,b] \to \mathcal{T}(S)$ be the geodesic segment connecting σ_1 and σ_2. Since the end points are in the thick part of $\mathcal{T}(S)$, for every subsurface Y, I_Y either is disjoint from $[a,b]$ or is a subset of $[a,b]$. Let k_0 be a constant such that, if $d_Y(\mu, \eta) \geq k_0$, then $I_Y \subset [a,b]$ (see Proposition 3.5). For $t \in I_Y$, the length of each boundary component of Y is less than ϵ_0. Therefore, there exists a constant C, depending on the topology of S, such that the number of subsurfaces with this property at each given time is at most C. Therefore,

$$d_{\mathcal{T}}(\sigma,\tau) \geq \tfrac{1}{C} \sum |I_Y| \,.$$

Lemmas 5.1 and 5.2 imply the desired inequality. □

5.2 Upper estimate. In [MaM2], Masur and Minsky show how to change one marking to another through elementary moves (described below) efficiently. Their estimate for the number of necessary elementary moves closely resembles the estimate in Theorem 1.1. We use this sequence of elementary moves to construct an efficient path connecting two points in $\mathcal{T}(S)$.

There are two types of elementary moves that transform a marking $\mu = \{(\alpha_1,\beta_1),\dots,(\alpha_m,\beta_m)\}$ to a new marking.

1. Twist: Replace β_i by β_i', where β_i' is obtained from β_i by a Dehn twist or a half twist around α_i.
2. Flip: Replace the pair (α_i,β_i) with (β_i,α_i) and, for $j \neq i$, replace β_j with a curve β_j' that does not intersect β_i, which is now a base curve, in such a way that $d_{\alpha_j}(\beta_j,\beta_j')$ is as small as possible (see [MaM2] for details).

In the first move, a twist can be positive or negative. A half twist is possible when α_i and β_i intersect twice. The following is a consequence of work done in [MaM2]) and [Mi4].

Proposition 5.4. *There exists a large enough k such that: for markings μ and μ', there exists a sequence of markings,*

$$\mu = \mu_1,\dots,\mu_n = \mu',$$

where μ_i and μ_{i+1} differ by an elementary move except, for each α where $d_\alpha(\mu,\mu') \geq k$, there is an index i_α so that

$$\mu_{i_\alpha+1} = D_\alpha^p \mu_{i_\alpha}, \quad \text{and} \quad |p| \asymp d_\alpha(\mu,\mu')\,. \tag{17}$$

Furthermore,

$$n \asymp \sum_{Y\subset S} \big[d_Y(\mu,\mu')\big]_k\,, \tag{18}$$

where the sum is over all subsurfaces Y that are not annuli.

Proof. We use the definitions and notation used in [MaM2] and [Mi4]. [MaM2, 4.6 and 4.20] state that there exists a complete hierarchy H whose initial marking is μ and whose terminal marking is μ'. Any complete marking has a resolution ([MaM2, 5.4]), that is, there is a sequence of markings

$$\mu = \eta_1,\dots,\eta_N = \mu'$$

where η_i and η_{i+1} differ by an elementary move. For k large enough, if $d_\alpha(\mu, \mu') \geq k$, the collar of α appears as a domain in H ([MaM2, 6.2]) exactly once ([Mi4, 5.15]), and the length of the corresponding geodesic in H is comparable to $d_\alpha(\mu, \mu')$ ([MaM2, 6.2]). That is, the number of twist moves around α used in the resolution is comparable to $d_\alpha(\mu, \mu')$. The number of the remaining elementary moves is comparable to the sum of the lengths of geodesics in H whose domains are not annuli, which is comparable to ([MaM2, Lem. 6.2 and eq. (6.4)])

$$\sum_Y \left[d_Y(\mu, \mu')\right]_k .$$

Our goal is, for any α where $d_\alpha(\mu, \mu') \geq k$, to rearrange the elementary moves in the resolution so that all the twist moves around α are applied consecutively. Then we replace the sequence of consecutive twists around α with one large step, which is applying D_α^p, for some $p \asymp d_\alpha(\mu, \mu')$. This will result in the sequence described in the statement of the theorem and has the desired length condition.

We know ([Mi4, 5.16]) that for every curve α, the set J_α of indices i such that α is a base curve in η_i is an interval in $\mathbb{Z}$. Observe that when α is a base curve of a marking, a twist move around α and a twist move around any other curve can be rearranged without any complication. The trouble with the flip moves is that the outcome is not unique. Therefore, after rearranging a flip move and a twist move, we have to make sure the outcomes of two flip moves differ by just a twist around α. For example, assume η_{i-1}, η_i and η_{i+1} all contain α as a base curve, η_i is obtained from η_{i-1} by a flip move and $\eta_{i+1} = D_\alpha \eta_i$. Then, replacing η_i with $\eta_i' = D_\alpha \, \eta_{i-1}$ in our sequence will result in a sequence that is still a resolution of H. Because η_i is obtained from η_{i-1} by applying a flip move, $D_\alpha \, \eta_i$ is also obtained from $D_\alpha \, \eta_{i-1}$ by a flip move (D_α is a homeomorphism). Therefore, we can rearrange the elementary moves in J_α so that all the twist moves around α are done consecutively. □

REMARK 5.5. The constant k can be chosen as large as necessary, and the constants involved in (18) depend on k and the topology of S (see [MaM1, Th. 6.12]). Therefore, we can assume $k \geq k_0$, where k_0 is as chosen in Proposition 5.3.

For a marking μ, let short(μ) be the set of points in $\mathcal{T}(S)$ where all curves in μ have hyperbolic length less than L_0 (L_0 as on page 941). This is a compact subset of $\mathcal{T}(S)$. We define $f(\mu, \mu')$ to be the maximum distance between an element in short(μ) and an element in short(μ').

LEMMA 5.6. *If $i = i_\alpha$, where α is a curve with $d_\alpha(\mu, \mu') \geq k$, then*

$$f(\mu_i, \mu_{i+1}) \asymp \log d_\alpha(\mu, \mu') .$$

Otherwise,

$$f(\mu_i, \mu_{i+1}) = O(1) .$$

Proof. Since short(μ) is compact, it is enough to bound the minimum distance between short(μ_i) and short(μ_{i+1}).

Assume $i = i_\alpha$, for α as above, and let σ be a point in short(μ_i). Then, for some $|p| \asymp d_\alpha(\mu, \eta)$, $\tau = D_\alpha^p \sigma$ is a point in short(μ_{i+1}). The lengths of α in σ and τ are less than L_0, therefore, σ and τ are bounded distance from points σ' and $\tau' = D_\alpha^p(\sigma')$, where the lengths of α in σ' and τ' are less than ϵ_0. Taking $\Gamma = \{\alpha\}$ and π as in Theorem 2.5, the following holds: the distance between σ' and τ' equals, up to additive error, the distance in $\mathbb{H}_\alpha$ between $\pi_\alpha(\sigma')$ and $\pi_\alpha(\tau')$, which, up to multiplicative error, equals $\log |p|$. Therefore, the distance between σ and τ is comparable to $\log |p|$.

Otherwise, μ_i and μ_{i+1} differ by an elementary move. Note that there are only finitely many such pairs of markings up to homeomorphism. Therefore, there exists a uniform upper bound for the minimum distance between short(μ_i) and short(μ_{i+1}), depending on the topology of S only. □

PROPOSITION 5.7. *Let σ_1, σ_2 be in the ϵ_o-thick part of $\mathcal{T}(S)$ and μ_1 and μ_2 be the short-markings in σ_1 and σ_2, respectively. Then*

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \prec \sum_Y \left[d_Y(\mu_1, \mu_2)\right]_k + \sum_\alpha \log \left[d_\alpha(\mu_1, \mu_2)\right]_k .$$

Proof. Let $\mu_1 = \bar{\mu}_1, \dots, \bar{\mu}_n = \mu_2$ be the path in $\mathcal{M}(S)$ described in Proposition 5.4. For each i, let σ_i be a point in short($\bar{\mu}_i$) and let g_i be the geodesic segment connecting σ_i to σ_{i+1}. The distance in $\mathcal{T}(S)$ between σ_1 and σ_2 is less than the sum of the lengths of the g_i. Lemma 5.6 states that the lengths of the g_i are uniformly bounded except when $i = i_\alpha$ and $d_\alpha(\mu_1, \mu_2) \geq k$, in which case the length of g_i is comparable with $\log d_\alpha(\mu_1, \mu_2)$. Therefore,

$$d(\sigma, \tau) \prec n\, O(1) + \sum_\alpha \log \left[d_\alpha(\mu_1, \mu_2)\right]_k .$$

Proposition 5.4 finishes the proof. □

Proof of Theorem 1.1. Propositions 5.3 and 5.7 provide a lower estimate and an upper estimate for the distance between σ and τ. Since $k \geq k_0$ (see Remark 5.5), the estimate given in Proposition 5.7 is smaller than the one given in Proposition 5.3. Therefore $d_{\mathcal{T}}(\sigma_1, \sigma_2)$ is comparable to

$$\sum_Y \left[d_Y(\mu_1, \mu_2)\right]_k + \sum_\alpha \log \left[d_\alpha(\mu_1, \mu_2)\right]_k . \qquad \square$$

6 The General Case

In this section we give an estimate for the distance between two arbitrary points in the Teichmüller space. Let σ_1 and σ_2 be two points in $\mathcal{T}(S)$ and $g\colon [a,b] \to \mathcal{T}(S)$ be the geodesic arc connecting them. If σ_1 and σ_2 are not in the thick part of $\mathcal{T}(S)$, then the set of short curves in σ_1 and σ_2 does not contain enough information to allow us to estimate the distance between σ_1 and σ_2; we also need to know how short these curves are. Therefore, our estimate for the distance contains terms measuring the distance between σ_1 and σ_2 and the thick part of Teichmüller space. An additional complication arises from the case where a curve is short in both σ_1 and σ_2 and remains short along the geodesic. However, the basic idea behind both Theorem 1.1 and Theorem 6.1 is that efficient paths in the space of markings are closely related to geodesics in Teichmüller space.

Let ϵ_0 be as before. Define Γ to be the set of curves that are short in both σ_1 and σ_2, and, for $i = 1, 2$, define Γ_i to be the set of curves that are short in σ_i but not in σ_{3-i}. Let μ_1 and μ_2 be short-markings on σ_1 and σ_2, respectively.

Theorem 6.1. *The distance in $T(S)$ between σ_1 and σ_2 is given by the following formula:*

$$d_{\mathcal{T}}(\sigma_1,\sigma_2) \asymp \sum_{Y} \big[d_Y(\mu_1,\mu_2)\big]_k + \sum_{\alpha\notin\Gamma} \log\big[d_\alpha(\mu_1,\mu_2)\big]_k + \max_{\alpha\in\Gamma} d_{\mathbb{H}_\alpha}(\sigma_1,\sigma_2) + \max_{\substack{\alpha\in\Gamma_i\\ i=1,2}} \log\frac{1}{l_{\sigma_i}(\alpha)}\,. \qquad (19)$$

Proof. Theorem 2.5 implies that

$$d_{\mathcal{T}}(\sigma_1,\sigma_2) \asymp \max_{\alpha\in\Gamma} d_{\mathbb{H}_\alpha}(\sigma_1,\sigma_2) + d_{\mathcal{T}(S\setminus\Gamma)}\big(\pi_0(\sigma_1),\pi_0(\sigma_2)\big)\,.$$

This accounts for the third term on the right-hand side of equation (19). Therefore, without loss of generality, we can assume $\Gamma = \emptyset$.

Let σ_1' and σ_2' be points in the thick part of the Teichmüller space that have the same short-markings as σ_1 and σ_2. We have

$$d_{\mathcal{T}}(\sigma_1,\sigma_2) \le d_{\mathcal{T}}(\sigma_1,\sigma_1') + d_{\mathcal{T}}(\sigma_1',\sigma_2') + d_{\mathcal{T}}(\sigma_2',\sigma_2)\,.$$

The sum of the first two terms in (19) is comparable with $d_{\mathcal{T}}(\sigma_1',\sigma_2')$. Also,

$$d_{\mathcal{T}}(\sigma_1,\sigma_1') \asymp \max_{\beta\in\Gamma_1} \log\frac{1}{l_{\sigma_1}(\beta)} \quad \text{and} \quad d_{\mathcal{T}}(\sigma_2,\sigma_2') \asymp \max_{\gamma\in\Gamma_2}\log\frac{1}{l_{\sigma_2}(\gamma)}\,.$$

Therefore, the right side of (19) is an upper bound for $d_{\mathcal{T}}(\sigma_1,\sigma_2)$ (up to additive and multiplicative constants).

To show that the right side of (19) is also a lower bound for $d_{\mathcal{T}}(\sigma_1, \sigma_2)$, we follow the same argument as in §5.1. However, we can not use Lemma 5.2 when α is short in either σ_1 or σ_2 and using the previous argument we can conclude only that

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \sum_Y \big[d_Y(\mu_1, \mu_2)\big]_k + \sum_{\alpha \notin \Gamma \cup \Gamma_1 \cup \Gamma_2} \log \big[d_\alpha(\mu_1, \mu_2)\big]_k . \tag{20}$$

For every $\alpha \in \Gamma_1$, we have

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) = \log K(\sigma_1, \sigma_2) \geq \log \left| \frac{\mathrm{Ext}_{\sigma_2}(\alpha)}{\mathrm{Ext}_{\sigma_1}(\alpha)} \right| \succ \log \frac{1}{l_{\sigma_1}(\alpha)} .$$

A similar statement is true for $\alpha \in \Gamma_2$. Hence

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \max_{\substack{\alpha \in \Gamma_i \\ i=1,2}} \log \frac{1}{l_{\sigma_i}(\alpha)} . \tag{21}$$

It remains to show, for $\alpha \in \Gamma_1 \cup \Gamma_2$, that $d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log d_\alpha(\mu_1, \mu_2)$. Let β_1 and β_2 be the transverse curves to α in μ_1 and μ_2. We know

$$|tw^+_{\sigma_1} - tw^+_{\sigma_2}| = \big| i_\alpha(\nu_+, \beta_1) - i_\alpha(\nu_+, \beta_2) \big| \overset{+}{\asymp} i_\alpha(\beta_1, \beta_2) = d_\alpha(\mu_1, \mu_2) .$$

Therefore, it is sufficient to show that $d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log |tw^+_{\sigma_1} - tw^+_{\sigma_2}|$. Theorem 4.3 implies that

$$|tw^+_{q_1} - tw^+_{q_2}| \succ \left| tw^+_{\sigma_1} - tw^+_{\sigma_2} - O\left(\frac{1}{l_{\sigma_1}(\alpha)} + \frac{1}{l_{\sigma_2}(\alpha)} \right) \right| ;$$

therefore,

$$|tw^+_{q_1} - tw^+_{q_2}| + \frac{1}{l_{\sigma_1}(\alpha)} + \frac{1}{l_{\sigma_2}(\alpha)} \succ |tw^+_{\sigma_1} - tw^+_{\sigma_2}| ,$$

and equation (11) implies that the q_t-twisting parameter changes at most exponentially fast; hence,

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log(tw^+_{q_1} - tw^+_{q_2}) .$$

We also know that

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log \frac{1}{l_{\sigma_1}(\alpha)} \quad \text{and} \quad d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log \frac{1}{l_{\sigma_2}(\alpha)} .$$

From the last three equations, we can conclude

$$d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log |tw^+_{\sigma_1} - tw^+_{\sigma_2}| .$$

Therefore,

$$\forall \alpha \in \Gamma_1 \cup \Gamma_2 , \quad d_{\mathcal{T}}(\sigma_1, \sigma_2) \succ \log d_\alpha(\mu_1, \mu_2) . \tag{22}$$

The combination of equations (20), (21) and (22) provides the desired lower bound and finishes the proof. □

References

[B] B.H. Bowditch, Intersection numbers and the hyperbolicity of the curve complex, J. reine angew. Math., to appear.

[Br] J.F. Brock, The Weil–Petersson metric and volumes of 3-dimensional hyperbolic convex cores, J. Amer. Math. Soc. 16:3 (2003), 495–535 (electronic).

[GL] F.P. Gardiner, N. Lakic, Quasiconformal Teichmüller theory, Mathematical Surveys and Monographs 76, American Mathematical Society, Providence, RI, 2000.

[K] S.P. Kerckhoff, The asymptotic geometry of Teichmüller space, Topology 19:1 (1980), 23–41.

[M] B. Maskit, Comparison of hyperbolic and extremal lengths, Ann. Acad. Sci. Fenn. Ser. A I Math. 10 (1985), 381–386.

[MaM1] H.A. Masur, Y.N. Minsky, Geometry of the complex of curves, I Hyperbolicity, Invent. Math. 138:1 (1999), 103–149.

[MaM2] H.A. Masur, Y.N. Minsky, Geometry of the complex of curves, II Hierarchical structure, GAFA, Geom. funct. anal. 10:4 (2000), 902–974.

[MaS] H. Masur, J. Smillie, Hausdorff dimension of sets of nonergodic measured foliations, Ann. of Math. 134:3 (1991), 455–543.

[Mi1] Y.N. Minsky, Harmonic maps, length, and energy in Teichmüller space, J. Differential Geom. 35:1 (1992), 151–217.

[Mi2] Y.N. Minsky, Teichmüller geodesics and ends of hyperbolic 3-manifolds, Topology 32:3 (1993), 625–647.

[Mi3] Y.N. Minsky, Extremal length estimates and product regions in Teichmüller space, Duke Math. J. 83:2 (1996), 249–286.

[Mi4] Y.N. Minsky, The classification of kleinian surface groups, i: Models and bounds, arXiv:math.GT/0302208, 2003.

[R] K. Rafi, A characterization of short curves of a Teichmüller geodesic, Geometry and Topology 9 (2005), 179–202.

Kasra Rafi, Department of Mathematics, University of Chicago, 5734 University Avenue, Chicago, Illinois 60637, USA rafi@math.uchicago.edu
http://www.math.uchicago.edu/∼rafi/

Received: October 2005
Revision: April 2006
Accepted: May 2006