Teichmüller geometry of moduli space, II: $\mathcal{M}(S)$ seen from far away

Benson Farb and Howard Masur *

January 5, 2009

1 Introduction

Let $S = S_{g,n}$ be a closed, orientable surface with genus $g \geq 0$ with $n \geq 0$ marked points, and let $\mathrm{Teich}(S)$ be the associated Teichmüller space of marked conformal classes or (fixed area) constant curvature metrics on S. Endow $\mathrm{Teich}(S)$ with the Teichmüller metric $d_{\mathrm{Teich}(S)}(\cdot,\cdot)$. Recall that for marked conformal structures $X_1, X_2 \in \mathrm{Teich}(S)$ we define

$$d_{\mathrm{Teich}(S)}(X_1, X_2) = \frac{1}{2}\log K$$

where $K \geq 1$ is the least number such that there is a K-quasiconformal mapping between the marked structures X_1 and X_2. The mapping class group $\mathrm{Mod}(S)$ acts properly discontinuously and isometrically on $\mathrm{Teich}(S)$, thus inducing a metric $d_{\mathcal{M}(S)}(\cdot,\cdot)$ on the quotient moduli space $\mathcal{M}(S) := \mathrm{Teich}(S)/\mathrm{Mod}(S)$. Let $\pi : \mathrm{Teich}(S) \to \mathcal{M}(S)$ be the natural projection.

The goal of this paper is to build an "almost isometric" simplicial model for $\mathcal{M}(S)$, from which we will determine the tangent cone at infinity of $\mathcal{M}(S)$. In analogy with the case of locally symmetric spaces, this can be viewed as a step in building a "reduction theory" for the action of $\mathrm{Mod}(S)$ on $\mathrm{Teich}(S)$. Other results in this direction can be found in [Le].

Moduli space seen from far away. Gromov formalized the idea of "looking at a metric space (X, d) from far away" by introducing the notion of the *tangent cone at infinity* of (X, d). This metric space, denoted $\mathrm{Cone}(X)$, is defined to be a Gromov-Hausdorff limit of based metric spaces (where basepoint $x \in X$ is fixed once and for all):

$$\mathrm{Cone}(X) := \lim_{\epsilon \to 0} (X, \epsilon d)$$

*Both authors are supported in part by the NSF.

So, for example, any compact Riemannian manifold M has $\text{Cone}(X) = *$, a one point space. Let $M = \Gamma\backslash G/K$ be an arithmetic, locally symmetric manifold (or orbifold); so G is a semisimple algebraic $\mathbf{Q}$-group, K a maximal compact subgroup, and Γ an arithmetic lattice. Hattori, Leuzinger and Ji-MacPherson proved that $\text{Cone}(M)$ is a metric cone over the quotient by Γ of the spherical Tits building $\Delta_{\mathbf{Q}}(G)$ associated to $G_{\mathbf{Q}}$. Here the metric on the cone on a maximal simplex of $\Delta_{\mathbf{Q}}(G)$ makes it isometric to the standard (Euclidean) metric on a Weyl chamber in G/K. In particular they deduce:

$$\mathbf{Q}\text{-rank}(\Gamma) = \dim(\text{Cone}(\Gamma\backslash G/K))$$

Our first result is a determination of the metric space $\text{Cone}(\mathcal{M}(S))$. The role of the rational Tits building will be played by the *complex of curves* $\mathcal{C}(S)$ on S. Recall that, except for some sporadic cases discussed below, the complex $\mathcal{C}(S)$ is defined to be the simplicial complex whose vertices are (isotopy classes of) simple closed curves on S, and whose k-simplices are $(k+1)$-tuples of distinct isotopy classes which can be realized simultaneously as disjoint curves on S. Note that $\mathcal{C}(S)$ is a d-dimensional simplicial complex, where $d = 3g - 4 + n$. While $\mathcal{C}(S)$ is locally infinite, its quotient by the natural action of $\text{Mod}(S)$ is a *finite orbicomplex*, by which we mean a finite simplicial complex where each simplex is quotiented out by the action of a finite group. The quotient can be made a simplicial complex by looking at the action on the barycentric subdivision of $\mathcal{C}(S)$. Denote by P the natural quotient map

$$P : \mathcal{C}(S) \to \mathcal{C}(S)/\operatorname{Mod}(S).$$

We now build a metric space which will serve as a coarse metric model for $\mathcal{M}(S)$. Let $\widetilde{\mathcal{V}}(S)$ denote the topological cone

$$\widetilde{\mathcal{V}}(S) := \frac{[0,\infty) \times \mathcal{C}(S)}{\{0\} \times \mathcal{C}(S)}$$

For each maximal simplex σ of $\mathcal{C}(S)$, we will think of the cone over σ as an orthant with coordinates $(x_1, \ldots, x_d)$. We endow this orthant with the standard sup metric:

$$d((x_1, \ldots, x_d), (y_1, \ldots, y_d)) := \frac{1}{2} \max_{1 \le i \le d} |x_i - y_i|.$$

The factor of $\frac{1}{2}$ is designed to be consistent with the definition of the Teichmüller metric.

The metrics on the cones on any two such maximal simplices clearly agree on (the cone on) any common face. We can thus endow $\widetilde{\mathcal{V}}(S)$ with the corresponding path metric. Note that the natural action of $\text{Mod}(S)$ on $\widetilde{\mathcal{V}}(S)$ induces an isometric action of $\text{Mod}(S)$ on $\widetilde{\mathcal{V}}(S)$. The quotient

$$\mathcal{V}(S) := \widetilde{\mathcal{V}}(S)/\operatorname{Mod}(S)$$

thus inherits a well-defined metric. The example $\mathcal{V}(S_{1,2})$ is described in Figure 1. To endow $\mathcal{V}(S)$ with the structure of a simplicial complex instead of an orbicomplex, we can simply replace $\mathcal{C}(S)$ with its barycentric subdivision in the construction above.

Our main result is that $\mathcal{V}(S)$ provides a simple and reasonably accurate geometric model for $\mathcal{M}(S)$.

Theorem 1. *There is a* $(1, D)$-quasi-isometry $\Psi : \mathcal{V}(S) \to \mathcal{M}(S)$. *That is, there is a constant* $D = D(S) \geq 0$ *such that :*

- $|d_{\mathcal{V}(S)}(x, y) - d_{\mathcal{M}(S)}(\Psi(x), \Psi(y))| \leq D$ *for each* $x, y \in \mathcal{V}(S)$, *and*
- *The* D*-neighborhood of* $\Psi(\mathcal{V}(S))$ *in* $\mathcal{M}(S)$ *is all of* $\mathcal{M}(S)$.

The main ingredient in our proof of Theorem 1 is a theorem of Minsky [Mi], which determines up to an additive factor the Teichmüller metric near infinity in Teich(S).

It is clear that Theorem 1 implies that Cone$(\mathcal{M}(S)) =$ Cone$(\mathcal{V}(S))$. Further, it is clear that multiplying the given metric on $\mathcal{V}(S)$ by any fixed constant gives a metric space which is isometric (via the dilatation) to the original metric. In particular, Cone$(\mathcal{V}(S))$ is isometric to $\mathcal{V}(S)$ itself. We thus deduce the following.

Corollary 2. Cone$(\mathcal{M}(S))$ *is isometric to* $\mathcal{V}(S)$.

Using different methods, Leuzinger [Le] has independently proven that $\mathcal{V}(S)$ is bilipschitz homeomorphic to Cone$(\mathcal{M}(S))$. His methods do not seem to yield the isometry type of Cone$(\mathcal{M}(S))$.

Remarks.

1. Corollary 2 has applications to metrics of positive scalar curvature. Namely, it is a key ingredient in the proof by Farb-Weinberger that, while $\mathcal{M}(S)$ admits a metric of positive scalar curvature for most S (e.g. when genus$(S) > 2$), it admits no metric with the same quasi-isometry type as the Teichmüller metric on $\mathcal{M}(S)$. See [FW].

2. For locally symmetric M, we know that Cone(M) is nonpositively curved in the CAT(0) sense. In contrast, $\mathcal{V}(S)$ strongly exhibits aspects of *positive* curvature, since even within the cone on a single simplex, any two points $x, y \in \mathcal{V}(S)$ have whole families of distinct geodesics between them, and these geodesics get arbitrarily far apart as $d(x, y) \to \infty$. This is a basic property of the sup metric on a quadrant.

3. Corollary 2 implies that any metric on $\mathcal{M}(S)$ quasi-isometric to the Teichmüller metric must have a cone which is bilipshitz homeomorphic to $\mathcal{V}(S)$.

The authors would like to thank the referee for some extremely helpful comments.

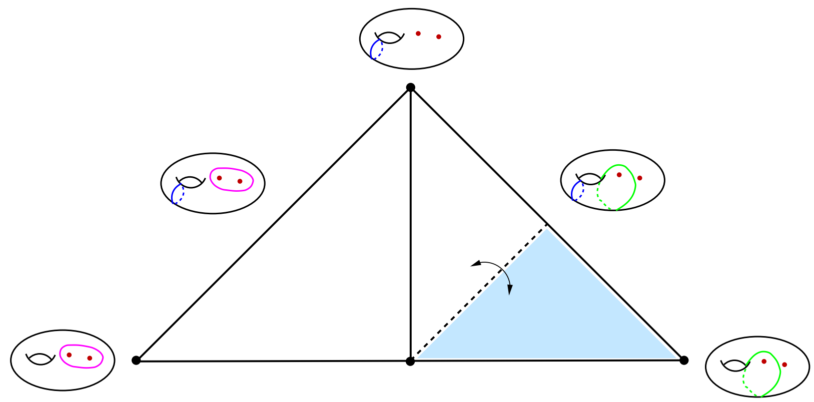

Figure 1: The metric space $\mathcal{V}(S_{1,2})$. The fundamental domain for the action of $\mathrm{Mod}(S)$ on $\mathcal{C}(S)$ is the union of two edges, one corresponding to a separating/nonseparating pair of curves, the other to a nonseparating/nonseparating pair. These are the only combinatorial types. Note that the latter edge has an order two symmetry, corresponding to the mapping class which switches the curves. Thus $\mathcal{V}(S)$ is the union of a Euclidean quadrant and a quotient of a Euclidean quadrant by a reflection along the $y = x$ ray.

2 The proof of Theorem 1

2.1 Minsky's Product Theorem

In this subsection we recall some work of Minsky which will be crucial for what follows.

Let $d = 3g-3+n$. Fix $\epsilon > 0$ smaller than the Margulis constant for hyperbolic surfaces. Let $\mathcal{C} = \{\gamma_1, \ldots, \gamma_p\}$ be a collection of distinct, disjoint, nontrivial homotopy classes of simple closed curves; this is a simplex in $\mathcal{C}(S)$. Let

$$\Omega_{\mathcal{C}}(\epsilon) := \{X \in \mathrm{Teich}(S) : \ell_X(\gamma_i) < \epsilon \text{ for each } i = 1, \ldots, p\}.$$

Extend $\mathcal{C}$ to a maximal collection $\{\gamma_1, \ldots, \gamma_d\}$ of homotopy classes of simple closed curves. Let $\{\theta_i, \ell_i)\}$ denote the corresponding Fenchel-Nielsen coordinates on $\Omega_{\mathcal{C}}(\epsilon)$. Recall that Fenchel-Nielsen coordinates give global coordinates on $\mathrm{Teich}(S)$; henceforth we will identify points in $\mathrm{Teich}(S)$ with their corresponding coordinates.

Consider the Teichmüller space $\mathrm{Teich}(S \setminus \mathcal{C})$, which is the space of complete, finite area hyperbolic metrics on $S \setminus \mathcal{C}$. Note that the coordinates $\{(\theta_i, \ell_i) : i > p\}$ give Fenchel-Nielsen coordinates on $\mathrm{Teich}(S \setminus \mathcal{C})$.

Let

$$\Phi = (\Phi_1, \Phi_2) : \Omega_{\mathcal{C}}(\epsilon) \to \mathrm{Teich}(S \setminus \mathcal{C}) \times \prod_{i=1}^{p} \mathbf{H}^2$$

be defined by

$$\Phi((\theta_1, \ldots, \theta_d, \ell_1, \ldots, \ell_d,)) := (\theta_{p+1}, \ldots, \theta_d, \ell_{p+1}, \ldots, \ell_d,) \times \prod_{i=1}^{p} (\theta_i, 1/\ell_i).$$

Notice that we are changing the last set of length coordinates from ℓ to $1/\ell$ giving coordinates in the upper half-space model of $\mathbf{H}^2$. We give $\mathbf{H}^2$ the metric $ds^2 = \frac{1}{4}(dx^2 + dy^2)/y^2$. Note that the factor of $\frac{1}{4}$ leads to a factor of $\frac{1}{2}$ in the distance, and is consistent with the factor of $\frac{1}{2}$ in the metric on the Euclidean octant. If $S \setminus \mathcal{C}$ is disconnected, then $\mathrm{Teich}(S \setminus \mathcal{C})$ is itself a product of the Teichmuller spaces of the components of $S \setminus \mathcal{C}$; we endow this total product space itself with the sup metric, denoted by d. We remark that Φ is a homeomorphism onto its image, and its image is $\mathrm{Teich}(S \setminus \mathcal{C}) \times \prod_{i=1}^{p} \{(x_i, y_i) \in \mathbf{H}^2 : y_i > 1/\epsilon\}$.

The following was proved in [Mi].

Theorem 3 (Minsky Product Theorem). *With notation as above, there exists D such that for all $X, Y \in \Omega_{\mathcal{C}}(\epsilon)$,*

$$|d(\Phi(X), \Phi(Y)) - d_{\mathrm{Teich}(S)}(X, Y)| \leq D.$$

We will need the following lemma about distances in $\mathcal{M}(S)$.

Lemma 4. *Given constants C, C' there is a constant C'' with the following property. Let $\sigma = \{\alpha_1, \ldots, \alpha_d\}$ be a maximal simplex of $\widetilde{V}(S)$. Let $X, Y \in \text{Teich}(S)$ be such that $\ell_X(\alpha_i) \leq C$ and $\ell_Y(\alpha_i) \leq C$ for each i. Suppose also that $|\log(\ell_X(\alpha_i)/\ell_Y(\alpha_i))| \leq C'$. Then $d_{\mathcal{M}(S)}(\pi(X), \pi(Y)) \leq C''$.*

Proof. This follows from Theorem 3. We can find a point Y' which differs from Y by Dehn twists about curves in σ so that the Fenchel-Nielsen twist coordinates of X, Y' have bounded difference. Now we consider the list of curves shorter than ϵ on both X and Y'. Since the ratios of lengths of these short curves are bounded above, as are the differences in twist coordinates, it follows that the distances in the corresponding $\mathbf{H}^2$ factors are bounded. The complement of these short curves determines a boundary Teichmuller space. The lengths of the remaining curves are bounded above and below, giving that the surfaces have a bounded distance from each other in this boundary Teichmüller space. The existence of C'' now follows from Theorem 3. $\diamond$

2.2 Defining the map Ψ

We will define a map $\widetilde{\Psi} : \widetilde{\mathcal{V}}(S) \to \mathcal{M}(S)$ by giving its value on a representative of each $\text{Mod}(S)$-orbit in $\widetilde{\mathcal{V}}(S)$, and then define $\widetilde{\Psi}$ to be constant on orbits. It will then follow that $\widetilde{\Psi}$ induces a map $\Psi : \mathcal{V}(S) \to \mathcal{M}(S)$. While this map will not be continuous, we will prove that it is a $(1, D)$-quasi-isometry for some $D \geq 0$.

Fix a (finite) collection of maximal simplices that represent all combinatorial types. We will first define $\widetilde{\Psi}$ on the open cone over this collection. Thus let σ be one of these maximal simplices of $\mathcal{C}(S)$ representing a maximal collection of disjoint simple closed curves $\{\alpha_1, \ldots, \alpha_d\}$. Again we think of the cone on σ, as a subspace of $\widetilde{\mathcal{V}}(S)$, as an octant in $\mathbf{R}^d$ with coordinates $x_1, \ldots, x_d$, endowed with the sup metric. Let $\text{Mod}(S, \sigma)$ be the subgroup of $\text{Mod}(S)$ that fixes σ. It acts on the open cone over σ with finite orbit. Take a sector $\Lambda(\sigma)$ inside this cone which is a fundamental domain for the action of $\text{Mod}(S, \sigma)$. For any $(x_1, \ldots, x_d) \in \Lambda(\sigma)$ (no $x_i = 0$), let

$$\widetilde{\Psi}(x_1, \ldots, x_d) := \pi(X) \tag{1}$$

where $\pi(X)$ is any point of $\pi(\Omega_\sigma(\epsilon))$ such that

$$\ell_X(\alpha_i) = \epsilon e^{-x_i} \quad \text{for each } i.$$

Using the action of $\text{Mod}(S, \sigma)$ we extend $\widetilde{\Psi}$ to the entire open cone on σ. Note that $\widetilde{\Psi}$ is continuous on each open cone. We do this for each maximal cone in the finite collection.

Now use the action of $\mathrm{Mod}(S)$ to extend $\widetilde{\Psi}$ to the open cones on all maximal simplices by having it be constant on orbits.

Next let τ be a simplex which is not maximal. Choose some closed maximal simplex $\sigma = \sigma(\tau)$ containing τ. We call this the maximal simplex *associated* to τ. The cone on τ is given by the coordinates $(x_1, \ldots, x_d)$ for the cone on σ as above. The coordinates x_i corresponding to curves in $\sigma - \tau$ are set to 0. Define $\widetilde{\Psi}$ on τ via the equation (1) above. Thus all curves in $\sigma - \tau$ are assigned the fixed length ϵ while the curves in τ can have arbitrarily small length. We extend $\widetilde{\Psi}$ to all of $\widetilde{\mathcal{V}}(S)$ by declaring $\widetilde{\Psi}$ to be constant on each $\mathrm{Mod}(S)$-orbit in $\widetilde{\mathcal{V}}(S)$. It follows that $\widetilde{\Psi}$ induces a map $\Psi : \mathcal{V}(S) \to \mathcal{M}(S)$. We remark that Ψ will in general not be continuous because of the choices made at a face of a maximal simplex. Nonetheless we want to know that the jump in the function at any face is uniformly bounded. We will argue this below using Lemma 4 together with the following lemma.

Lemma 5. *Let τ be a simplex. Let σ_1 a maximal simplex associated to τ and let σ_2 be any other maximal simplex such that $\tau = \sigma_1 \cap \sigma_2$. Then there exists an element $\phi \in \mathrm{Mod}(S)$, fixing τ, such that for each x in the cone over τ there is a point $X \in \mathrm{Teich}(S)$ with $\pi(X) = \widetilde{\Psi}(x)$ and such that the X-length of any curve in $(\sigma_1 - \tau) \cup (\phi(\sigma_2) - \tau)$ is bounded above by a universal constant, and below by the fixed ϵ.*

Proof. The coordinates for curves in $\sigma_1 - \tau$ on the cone over τ are 0. By definition, each curve $\beta \in \sigma_1 - \tau$ then has fixed length ϵ on some X with $\pi(X) = \widetilde{\Psi}(x)$. The curves in $\sigma_2 - \tau$ may have large intersection with curves in $\sigma_1 - \tau$ and therefore large length on X. However, since there are only finitely many combinatorial types of pants decompositions, we can choose ϕ fixing τ so that any curve in $\phi(\sigma_2) - \tau$ has universally bounded intersection with any curve in $\sigma_1 - \tau$. Since $\ell_X(\beta) = \epsilon$ for each $\beta \in \sigma_1 - \tau$, the collar about β has diameter bounded above. Thus we can further compose ϕ by Dehn twists about β, so that for the new ϕ, the curves in $\phi(\sigma_2) - \tau$ have bounded lengths on X. $\diamond$

2.3 Properties of Ψ

Our goal in this subsection is to prove that Ψ is a $(1, D)$-quasi-isometry. In order to do this we will need the following setup.

Let σ a maximal simplex. Recall P is the quotient map from $\mathcal{C}(S)$ to $\mathcal{C}(S)/\mathrm{Mod}(S)$. Let $d_{P(\sigma)}$ be the path metric on the cone over $P(\sigma)$ and let $d^{\mathcal{M}(S)}_{P(\sigma)}$ be the path metric on the (connected) Ψ image of the cone over $P(\sigma)$ in $\mathcal{M}(S)$ induced from the Teichmüller metric on $\mathcal{M}(S)$. That is, the distance between two points in the image is the infimum of the lengths of paths joining the points that stays in the image of the cone over $P(\sigma)$.

Lemma 6. *There is a constant D_0 such that if x_1, x_2 lie in the cone over $P(\sigma)$, then*

$$|d_{P(\sigma)}(x_1, x_2) - d_{P(\sigma)}^{\mathcal{M}(S)}(\Psi(x_1), \Psi(x_2))| \leq D_0.$$

Proof. We may find a lift X_i of $\Psi(x_i)$ to Teich(S) such that the difference of the twist coordinates of X_1 and X_2 with respect to the Fenchel-Nielsen coordinates defined by σ are bounded and such that

$$d_{\mathrm{Teich}(S)}(X_1, X_2) = d_{P(\sigma)}^{\mathcal{M}(S)}(\Psi(x_1), \Psi(x_2)).$$

If x_1 and x_2 lie in the open cone over $P(\sigma)$, then the lemma follows from Theorem 3 and the definition of the metric $d_{P(\sigma)}$. If not, then one must further quote Lemma 5 and Lemma 4. $\diamond$

Ψ is almost onto: By a theorem of Bers, there is a constant $C = C(g, n)$ such that each $X \in$ Teich(S) has a pants decomposition corresponding to a maximal simplex σ such that each curve of σ has length at most C on X. With respect to these pants curves, each of the twist coordinates is bounded, modulo the action of Dehn twists about the curves in σ, by $2\pi C$. Let τ be the possibly empty face of σ such that the set of curves in $\sigma - \tau$ have lengths on X between ϵ and C. The curves in τ have length at most ϵ. By Lemma 5, there is a point $Y \in$ Teich(S) such that $\pi(Y)$ is in the Ψ-image of the cone on τ, and such that the lengths on Y of the curves in τ are the same as the lengths on X of those curves, and the curves in $\sigma - \tau$ have bounded length on Y. Thus their ratios to the lengths on X are bounded. Applying Lemma 4, we are done.

Ψ is an almost isometry: We need the following lemma.

Lemma 7 (Path Lemma)**.** *The following statements are true.*

1. *Any two points in $\mathcal{V}(S)$ can be joined by a geodesic that enters the cone over each $P(\sigma)$, where σ is a maximal simplex of $\widetilde{V(S)}$, at most once.*

2. *There is a constant C' such that any two points of $\Psi(\mathcal{V}(S))$ can be joined by a $(1, C')$ quasi-geodesic in the metric $d_{\mathcal{M}(S)}$ that enters the cone over each $P(\sigma)$ at most once.*

A first step in proving Lemma 7 is the following.

Lemma 8. *The following statements are true.*

1. *Suppose x, y are points in the cone over $P(\sigma)$ where σ is a maximal simplex. Then there is a geodesic joining x and y that stays in the cone over that $P(\sigma)$.*

2. There is a constant C'' such that if $\Psi(x), \Psi(y)$ lie in the cone over $P(\sigma)$ then

$$d^{\mathcal{M}(S)}_{P(\sigma)}(\Psi(x), \Psi(y)) \leq d_{\mathcal{M}(S)}(\Psi(x), \Psi(y)) + C''.$$

We note that the opposite inequality

$$d_{\mathcal{M}(S)}(\Psi(x), \Psi(y)) \leq d^{\mathcal{M}(S)}_{P(\sigma)}(\Psi(x), \Psi(y))$$

is clearly true.

Proof. [of Lemma 8] We prove the first statement. Lift to $\widetilde{\mathcal{V}}(S)$ and consider again x, y with the same names such that the distance in the cone over σ realizes the distance between x and y in the cone over $P(\sigma)$. Let the coordinates of x, y be given by $(x_1, \ldots, x_d), (y_1, \ldots, y_d)$. Suppose σ is defined by the curves $\alpha_1, \ldots, \alpha_d$ of a pants decomposition. Without loss of generality assume that $d_\sigma(x, y) = \frac{1}{2}(y_1 - x_1)$. We must show that, for every $\phi \in \mathrm{Mod}(S)$, that does not fix σ, there is no shorter path ρ in $\widetilde{\mathcal{V}}(S)$ from $\phi(x)$ to y.

Suppose first that α_1 is not a vertex in the simplex $\phi(\sigma)$. Then the path from x to y for a last time must enter the cone over a simplex for which α_1 is a vertex at a point z. At z the coordinate corresponding to α_1 is 0, and so

$$d_{\mathcal{V}(S)}(y, z) \geq y_1/2 \geq d_\sigma(x, y).$$

Thus we may assume that the path ρ joining $\phi(x)$ to y lies completely in the cones over simplices for which α_1 is a vertex. Break up this path into segments $\rho = \rho_1 * \rho_2 * \ldots * \rho_N$, where each ρ_i lies in the cone over a single simplex. Let z_1^i (resp. z_1^{i+1}) be the coordinate of α_1 at the beginning (resp. end) of ρ_i, where $z_1^1 = x_1$ and $z_1^{N+1} = y_1$. Then $|\rho_i| \geq \frac{1}{2}|z_1^{i+1} - z_1^i|$. Thus

$$|\rho| = \sum_{i=1}^{N} |\rho_i| \geq \sum \frac{1}{2}|z_1^{i+1} - z_1^i| \geq \frac{1}{2}(y_1 - x_1) = d_{\widetilde{V}(S)}(x, y).$$

We conclude that a shortest path can be found by a geodesic that lies entirely in the cone over σ

We prove the second statement. First lift $\Psi(x), \Psi(y)$ to elements $X, Y \in \mathrm{Teich}(S)$ which lie in $\Omega_\sigma(\epsilon)$, and such that

$$d^{\mathcal{M}(S)}_{P(\sigma)}(\Psi(x), \Psi(y)) = d_{\mathrm{Teich}(S)}(X, Y)$$

and whose twist coordinates are bounded by $2\pi\epsilon$. By Theorem 3, there exists a simple closed curve $\alpha_1 \in \sigma$ such that

$$|d_{\mathrm{Teich}(S)}(X, Y) - \frac{1}{2}\log(\ell_Y(\alpha_1)/\ell_X(\alpha_1))| \leq D'$$

where D' depends on ϵ and on the constant D from Theorem 3. Thus

$$|d^{\mathcal{M}(S)}_{P(\sigma)}(\Psi(x),\Psi(y)) - \frac{1}{2}\log(\ell_Y(\alpha_1)/\ell_X(\alpha_1))| \leq D'. \tag{2}$$

Now let ϕ be a mapping class group element. If α_1 is not a vertex of $\phi(\sigma)$ then any path ρ from $\phi(Y)$ to X must enter a set $\Omega_{\mathcal{C}}(\epsilon)$ for some $\mathcal{C}$ containing α_1 a last time. At that time the length of α_1 is ϵ. By Theorem 3 and Equation (2) we then have

$$|\rho| \geq \frac{1}{2}\log(\epsilon/\ell_X(\alpha_1)) - D \geq d^{\mathcal{M}(S)}_{P(\sigma)}(\Psi(y),\Psi(x)) + \frac{1}{2}\log(\epsilon/\ell_Y(\alpha_1)) - D - D'.$$

Since $\ell_Y(\alpha_1)$ is bounded above, the term $\frac{1}{2}\log(\epsilon/\ell_Y(\alpha_1)) - D - D'$ is bounded below by some constant, and we set $-C''$ to be this constant.

Thus again we can assume that ρ lies completely in $\Omega_{\mathcal{C}}(\epsilon)$ for a set $\mathcal{C}$ containing α_1. But now the conclusion again follows from Theorem 3. $\diamond$

Proof. [of Lemma 7] Suppose x is in the cone over $P(\sigma_1)$ and that y is in the cone over $P(\sigma_2)$. If $P(\sigma_1) = P(\sigma_2)$ then we are done by Lemma 8. Thus we can assume that $P(\sigma_2) \neq P(\sigma_1)$. Suppose ρ is a geodesic from x to y. Suppose ρ leaves the cone over $P(\sigma_1)$ and returns to it for a last time at some z in the cone over $P\sigma_1) \cap P(\sigma_3)$ for some maximal simplex σ_3. Then by the first part of Lemma 8 we can replace ρ by a geodesic that stays in the cone over $P(\sigma_1)$ from x to z and then follows ρ from z to y never returning to the cone over $P(\sigma_1)$. We now find the last point w that lies in the cone over $P(\sigma_3)$ and replace a segment of ρ with one that stays in the cone over $P(\sigma_3)$ and never returns again to the cone over $P(\sigma_3)$. Since there are only a finite number of simplices in $\mathcal{C}(S)/\operatorname{Mod}(S)$, continuing to apply Lemma 8, we are done. This proves the first statement.

The proof of the second statement is similar, where we now use the second part of Lemma 8. $\diamond$

We now continue with the final step in the proof of Theorem 1, that the map Ψ is an almost isometry. We first prove that

$$d_{\mathcal{M}(S)}(\Psi(x),\Psi(y)) \leq d_{\mathcal{V}(S)}(x,y) + R$$

for some constant R. To prove this, consider a geodesic path $\gamma \subset \mathcal{V}(S)$ connecting x to y. By the first statement of Lemma 7, there exists $c = c(S)$ so that γ can be written as a concatenation $\gamma = \gamma_1 * \cdots * \gamma_c$ with each γ_i a geodesic in the cone over $P(\sigma)$ for σ a maximal simplex σ_i of $\mathcal{V}(S)$. By Lemma 6 each $\Psi(\gamma_i)$ is a $(1, D_0)$-quasigeodesic in the metric $d_{\mathcal{M}(S)}$. It follows that $\Psi(\gamma)$ is a $(1, cD')$-quasigeodesic.

The proof of the opposite inequality

$$d_{\mathcal{V}(S)}(x,y) \leq d_{\mathcal{M}(S)}(\Psi(x),\Psi(y)) + R'$$

for some R' uses the second conclusion of Lemma 7. Any two points can be joined by $(1, C')$ quasi-geodesic in the metric $d_{\mathcal{M}(S)}$ and which intersects a fixed number of cones over image simplices $P(\sigma)$. We now apply Lemma 6 to conclude that $d_{\mathcal{V}(S)}(x, y)$ is only larger by an additive constant.

References

[FW] B. Farb and S. Weinberger, Positive scalar curvature metrics on the moduli space of Riemann surfaces, in preparation.

[Ha1] T. Hattori, Collapsing of quotient spaces of $\mathrm{SO}(n)\backslash\mathrm{SL}(n, R)$ at infinity, *J. Math. Soc. Japan* 47 (1995), no. 2, 193–225.

[Ha2] T. Hattori, Asymptotic geometry of arithmetic quotients of symmetric spaces, *Math. Zeit.* 222 (1996), no. 2, 247–277.

[JM] L. Ji and R. MacPherson, Geometry of compactifications of locally symmetric spaces, *Ann. Inst. Fourier, Grenoble*, Vol. 52, No. 2 (2002), 457–559.

[Le] E. Leuzinger, Reduction theory for mapping class groups, preprint, Jan. 2008.

[Mi] Y. Minsky, Extremal length estimates and product regions in Teichmüller space, *Duke Math. Jour.* 83 (1996), no. 2, 249–286.

Dept. of Mathematics, University of Chicago
5734 University Ave.
Chicago, Il 60637
E-mail: farb@math.uchicago.edu, masur@math.uchicago.edu