

Hyperbolic spaces in Teichmüller spaces

Christopher J. Leininger and Saul Schleimer

October 29, 2011

Abstract

We prove, for any n , that there is a closed connected orientable surface S so that the hyperbolic space \mathbb{H}^n almost-isometrically embeds into the Teichmüller space of S , with quasi-convex image lying in the thick part. As a consequence, \mathbb{H}^n quasi-isometrically embeds in the curve complex of S .

1 Introduction

We denote the Teichmüller space of a surface S by $\mathcal{T}(S)$, and the ϵ -thick part by $\mathcal{T}_\epsilon(S)$; see Section 4. An *almost-isometric embedding* of one metric space into another is a $(1, C)$ -quasi-isometric embedding, for some $C \geq 0$; see Section 2. Let \mathbb{H}^n denote hyperbolic n -space. The main result of this paper is the following.

Theorem 1.1. *For any $n \geq 2$, there exists a surface of finite type S and an almost-isometric embedding*

$$\mathbb{H}^n \rightarrow \mathcal{T}(S).$$

Moreover, the image is quasi-convex and lies in $\mathcal{T}_\epsilon(S)$ for some $\epsilon > 0$.

According to Proposition 4.4 below, Theorem 1.1 remains true if we replace “surface of finite type” with “closed surface”. Our work is motivated, in part, by the following open question (see [7] for the case $n = 2$).

Question 1.2. *Does there exist a closed surface S of genus at least 2, a closed hyperbolic n -manifold B with $n \geq 2$, and an S -bundle E over B for which $\pi_1(E)$ is Gromov hyperbolic?*

To explain the relationship with our theorem, suppose that

$$S \rightarrow E \rightarrow B$$

is an S -bundle over $B = \mathbb{H}^n/\Gamma$, for some closed surface S and some torsion free cocompact lattice $\Gamma < \text{Isom}(\mathbb{H}^n)$. The monodromy is a homomorphism to the mapping class group of S , $\rho: \pi_1(B) = \Gamma \rightarrow \text{Mod}(S)$. The mapping class group $\text{Mod}(S)$ acts on $\mathcal{T}(S)$ by isometries with respect to the Teichmüller metric, and according to work of Farb-Mosher [7] and Hamenstädt [12], $\pi_1(E)$ is δ -hyperbolic if and only if we can construct a Γ -equivariant quasi-isometric embedding

$$f: \mathbb{H}^n \rightarrow \mathcal{T}(S)$$

with quasi-convex image lying in $\mathcal{T}_\epsilon(S)$ for some $\epsilon > 0$; see also [25]. (In fact the Γ -equivariance and quasi-isometric embedding assumptions imply that the image lies in $\mathcal{T}_\epsilon(S)$.)

This work is in the public domain.

Our main theorem states that if we drop the assumption of equivariance, then quasi-isometric embeddings with all the remaining properties exist. On the other hand, as was shown in [6], one can find cocompact lattices $\Gamma < \text{Isom}(\mathbb{H}^2)$ and Γ -equivariant quasi-isometries into $\mathcal{T}(S)$ with image in $\mathcal{T}_\epsilon(S)$ —for these examples the image is not quasi-convex.

The main theorem for $n = 2$ also contrasts with the situation of isometrically embedding hyperbolic planes in $\mathcal{T}(S)$. More precisely, every geodesic in $\mathcal{T}(S)$ is contained in an isometrically embedded hyperbolic plane (with the Poincaré metric) called a Teichmüller disk. However, it is well-known that *no* Teichmüller disk lies in any thick part—this follows from [21] which guarantees that along a dense set of geodesic rays in the Teichmüller disk the hyperbolic length of some curve on S tends to zero.

The curve complex of S is a metric simplicial complex $\mathcal{C}(S)$ whose vertices are isotopy classes of essential simple closed curves, and for which $k + 1$ distinct isotopy classes of curves span a k -simplex if they can be realized disjointly. In [23], Masur and Minsky proved that $\mathcal{C}(S)$ is δ -hyperbolic. One of the key ingredients in their proof is the construction of a coarsely Lipschitz map $\mathcal{T}(S) \rightarrow \mathcal{C}(S)$. The restriction of this map to any quasi-convex subset of $\mathcal{T}_\epsilon(S)$ is a quasi-isometry (see for example [27, Lemma 4.4] or [15, Theorem 7.6]). Composing the almost-isometry of Theorem 1.1 with the map $\mathcal{T}(S) \rightarrow \mathcal{C}(S)$ we have the following corollary.

Corollary 1.3. *For every $n \geq 2$, there exists a surface of finite type S and a quasi-isometric embedding*

$$\mathbb{H}^n \rightarrow \mathcal{C}(S).$$

The case of $n = 2$ here can be compared to the result of Bonk and Kleiner [5] in which it is shown that every δ -hyperbolic group which is not virtually free contains a quasi-isometrically embedding hyperbolic plane. The assumption that the group is not virtually free implies the existence of an arc in the boundary. According to [9] (see also [19, 18]) with the exception of a few small surfaces, there are indeed arcs in the boundary of $\mathcal{C}(S)$. In [5] however, essential use is made of the fact that there is an action of the group, and so even in the case $n = 2$, Corollary 1.3 does not follow from [5].

We now explain the idea for the construction in the case $n = 2$. Given a closed Riemann surface Z and a point $z \in Z$, the Teichmüller space $\mathcal{T}(Z, z)$ is naturally a \mathbb{H}^2 -bundle over $\mathcal{T}(Z)$; see Section 4.3. Given a biinfinite geodesic τ in $\mathcal{T}(Z)$, the preimage of τ in $\mathcal{T}(Z, z)$ is a 3-manifold. The parameterization $t \mapsto \tau(t)$ lifts to a flow on the preimage of τ for which the flow lines are geodesics in $\mathcal{T}(Z, z)$. The fiber over $\tau(0)$ admits a pair of transverse 1-dimensional singular foliations—these are naturally associated to the vertical and horizontal foliations of the quadratic differential defining τ . Any two flow lines meeting the same nonsingular leaf of the vertical foliation are forward asymptotic. Therefore, we have a 1-parameter family of forward asymptotic geodesics in $\mathcal{T}(Z, z)$. We use this to define a map from \mathbb{H}^2 to $\mathcal{T}(Z, z)$: we pick a horocycle $C \subset \mathbb{H}^2$ and send the pencil of geodesics perpendicular to C to our set of forward asymptotic geodesics in $\mathcal{T}(Z, z)$.

At the beginning of Section 5.2 we give a brief explanation of how this can be modified to give the construction for $n = 3$. The idea for $n \geq 4$ is then a straightforward inductive construction.

Acknowledgements. We would like to thank Richard Kent for useful conversations as well as having originally asked about the existence of quasi-isometric embeddings of hyperbolic planes into $\mathcal{C}(S)$.

2 Hyperbolic geometry

Suppose that (X, d_X) and (Y, d_Y) are metric spaces.

Definition 2.1. A map $F: X \rightarrow Y$ is a K -almost-isometric embedding if for all $x, x' \in X$ we have

$$|d_X(x, x') - d_Y(F(x), F(x'))| \leq K.$$

We use the *exponential model* for hyperbolic space: $\mathbb{H}^n = \mathbb{R}^{n-1} \times \mathbb{R}$ with length element

$$ds^2 = e^{-2t}(dx_1^2 + \dots + dx_{n-1}^2) + dt^2.$$

For two points $p, q \in \mathbb{H}^n$ we use $d_{\mathbb{H}}(p, q)$ to denote the distance between them. The exponential model of hyperbolic space is related to the upper-half space model $U = \mathbb{R}^{n-1} \times (0, \infty)$ by the map $\mathbb{H}^n \rightarrow U$ given by $(x, t) \mapsto (x, e^t)$. In the exponential model, for every $x \in \mathbb{R}^{n-1}$ the path $\eta_x(t) = (x, t)$ is a *vertical geodesic* and is parameterized by arc-length.

Lemma 2.2. *Suppose (X, d_X) is a geodesic metric space and $\delta, \epsilon, R > 0$ are constants. Suppose $F: \mathbb{H}^n \rightarrow X$ is a function with the following properties.*

1. $F \circ \eta_x$ is a geodesic for all $x \in \mathbb{R}^{n-1}$.
2. For distinct $x, x' \in \mathbb{R}^{n-1}$ the geodesics $F \circ \eta_x$ and $F \circ \eta_{x'}$ are two sides of an ideal δ -slim triangle in (X, d_X) .
3. For any $x, x' \in \mathbb{R}^{n-1}$ if $e^{-t}|x - x'| < \epsilon$ then $d_X(F(x, t), F(x', t)) \leq R$.
4. If $(x_k, t_k), (x'_k, t_k) \in \mathbb{H}^n$ satisfy $\lim_{k \rightarrow \infty} e^{-t_k}|x_k - x'_k| = \infty$, then $\lim_{k \rightarrow \infty} d_X(F(x_k, t_k), F(x'_k, t_k)) = \infty$.

Then there exists a constant K so that F is a K -almost isometric embedding.

A useful consequence of Property 3 is that for any $x, x', t \in \mathbb{R}$ we have

$$d(F(x, t), F(x', t)) \leq \frac{R}{\epsilon} e^{-t}|x - x'| + R. \quad (1)$$

The remainder of this section gives the proof of Lemma 2.2. We begin by controlling how F moves the centers of ideal triangles. To be precise: Suppose that $T = P \cup Q \cup R \subset \mathbb{H}^n$ is an ideal triangle where P and Q are distinct vertical geodesics. Let $r \in R$ denote the point of R with maximal t -coordinate. We call r the *midpoint* of R . Thus r serves as a center for T . Define $x = x(P), x' = x(Q)$.

Observe, say from the upper-half space model, that for all $t \geq t(r)$ we have

$$d_{\mathbb{H}}((x, t), (x', t)) \leq e^{-t}|x - x'| \leq e^{-t(r)}|x - x'| = 2. \quad (2)$$

Thus, by Inequality (1) we have $d_X(F(x, t), F(x', t)) \leq 2R/\epsilon + R$. Define $\Delta = \max\{3\delta, 2R/\epsilon + R\}$ and define the *displaced height* of T to be

$$h_T = h(T) = \min \left\{ t \in \mathbb{R} \mid d_X(F(x, t), F(Q)) \text{ or } d_X(F(P), F(x', t)) \leq \Delta \right\}.$$

It follows that $h(T) \leq t(r)$. Note that for any vertical triangle T , Property 2 implies that $h(T) > -\infty$.

Claim 2.3. *For any vertical triangle $T = P \cup Q \cup R \subset \mathbb{H}^n$,*

$$d_X(F(x, h_T), F(x', h_T)) \leq 3\Delta,$$

where $x = x(P), x' = x(Q)$.

Proof. Breaking symmetry, in this setting, allows us to assume that there is some $s \in \mathbb{R}$ so that $d_X(F(x', s), F(x, h_T)) \leq \Delta$. Let $t' = \max\{s, t(r)\}$. Using the triangle inequality, Inequality 1 and Property 1 we have

$$\begin{aligned} t' - h_T &= d_X(F(x, t'), (x, h_T)) \\ &\leq d_X(F(x, t'), F(x', t')) + d_X(F(x', t'), F(x', s)) + d_X(F(x', s), F(x, h_T)) \\ &\leq (2R/\epsilon + R) + (t' - s) + \Delta \end{aligned}$$

and similarly

$$t' - s \leq 2R/\epsilon + R + t' - h_T + \Delta.$$

Thus $|h_T - s| \leq 2R/\epsilon + R + \Delta$. Another application of the triangle inequality and Property 1 implies that $d_X(F(x, h_T), F(x', h_T)) \leq 2R/\epsilon + R + 2\Delta \leq 3\Delta$, as desired. \square

As mentioned above, for every vertical triangle T we have $h(T) > -\infty$ and hence $t(r) - h(T) < \infty$. We now obtain a uniform bound on this quantity.

Claim 2.4. *There is a constant $C_0 = C_0(F)$ so that $t(r) - h(T) \leq C_0$ for all vertical triangles $T \subset \mathbb{H}^n$.*

Proof. Suppose not. Then we are given a sequence of vertical triangles $T_k = P_k \cup Q_k \cup R_k$ where $t(r_k) - h(T_k)$ tends to infinity with k . Here r_k is the midpoint of R_k , the non-vertical side. Define $t_k = t(r_k)$, $h_k = h(T_k)$. Define $x_k = x(P_k)$, $x'_k = x(Q_k)$ to be the horizontal coordinates of the vertical sides of T_k .

Note that by Equation (2)

$$e^{-t_k} |x_k - x'_k| = 2$$

and so

$$e^{-h_k} |x_k - x'_k| = e^{-h_k} \cdot 2e^{t_k} = 2e^{t_k - h_k}.$$

Thus $e^{-h_k} |x_k - x'_k|$ tends to infinity with k . From Property 4 we deduce that the quantity $d_X(F(x_k, h_k), F(x'_k, h_k))$ also tends to infinity with k . This last, however, contradicts Claim 2.3. \square

We give the proof of Lemma 2.2. Fix any $p, q \in \mathbb{H}^n$. If $x(p) = x(q)$ then we are done by Property 1. Suppose instead that $x(p) \neq x(q)$. Let $P \cup Q \cup R$ denote the vertical triangle having vertical sides P and Q so that $x(P) = x(p)$, $x(Q) = x(q)$; let $r \in R$ be the midpoint of the non-vertical side. Define $C_1 = 2C_0 + 5\Delta + 1$. There are now two cases to consider.

Case. *Suppose that $t(p) \geq h(T) - C_1$.*

Let $p' \in P$ and $q' \in Q$ be the points with $t(p') = t(q') = \max\{t(p), t(r)\}$. Then by the triangle inequality and Equation (2) we have

$$\begin{aligned} d_{\mathbb{H}}(p, q') &\leq d_{\mathbb{H}}(p, p') + d_{\mathbb{H}}(p', q') \\ &\leq t(p') - t(p) + 2 \\ &\leq t(r) - h(T) + C_1 + 2 \\ &\leq C_0 + C_1 + 2. \end{aligned}$$

It follows that $d_{\mathbb{H}}(p, q)$ is estimated by $d_{\mathbb{H}}(q', q) = |t(q') - t(q)|$ up to an additive error at most $C_0 + C_1 + 2$. Appealing to Property 1, Inequality (1), and the triangle inequality we similarly have

$$\begin{aligned} d_X(F(p), F(q')) &\leq d_X(F(p), F(p')) + d_X(F(p'), F(q')) \\ &\leq t(p') - t(p) + 2R/\epsilon + R \\ &\leq C_0 + C_1 + 2R/\epsilon + R. \end{aligned}$$

Thus $d_X(F(p), F(q))$ is estimated by $d_X(F(q'), F(q)) = d_{\mathbb{H}}(q', q)$ with an additive error at most $C_0 + C_1 + 2R/\epsilon + R$. This completes the proof in this case.

Case. Suppose that $t(p), t(q) \leq h(T) - C_1$.

In this case, since the triangle $P \cup Q \cup R$ is slim in \mathbb{H}^n , we find that that $d_{\mathbb{H}}(p, q)$ is estimated by $t(r) - t(p) + t(r) - t(q)$ up to an additive error of at most 2. We now show that $d_X(F(p), F(q))$ is also estimated by the latter quantity, with a uniformly bounded error. Using Property 1 and Inequality (1) deduce

$$d_X(F(p), F(q)) \leq t(r) - t(p) + 2R/\epsilon + R + t(r) - t(q).$$

We now give a lower bound for $d_X(F(p), F(q))$. Recall that $F(P)$ and $F(Q)$ are two sides of a δ -slim triangle in X . Let R_X be the third side of this triangle. Since

$$d_X(F(p), F(Q)), d_X(F(P), F(q)) > \Delta \geq \delta$$

it follows that there are points $p_X, q_X \in R_X$ so that $d_X(F(p), p_X), d_X(q_X, F(q)) \leq \delta$. Thus the distance $d_X(p_X, q_X)$ is an estimate for $d_X(F(p), F(q))$ with an additive error at most 2δ .

Define $a = (x, h_T), b = (x', h_T)$. Again, as in the previous paragraph, there are points $a_X, b_X \in R_X$ within distance δ of $F(a), F(b)$. Since $d_{\mathbb{H}}(a, b) \leq 2(t(r) - h(T)) + 2$ we find

$$\begin{aligned} d_X(a_X, b_X) &\leq 2\delta + 2(t(r) - h(T)) + 2R/\epsilon + R \\ &\leq 2\delta + 2C_0 + 2R/\epsilon + R. \end{aligned}$$

Note that the geodesic segments $[p_X, a_X], [b_X, q_X] \subset R_X$ have length at least $h(T) - t(p) - 2\delta$ and $h(T) - t(q) - 2\delta$ respectively. Each of these is greater than $C_1 - 2\delta$.

If $p_X \in [a_X, b_X]$ then $C_1 - 2\delta \leq 2\delta + 2C_0 + 2R/\epsilon + R$ and this is a contradiction. Similarly, deduce $q_X \notin [a_X, b_X]$. If $p_X = q_X$ then $d_X(F(p), F(q)) \leq 2\delta < \Delta$, contradicting our assumption that $t(p) < h(T)$. Finally, if $p_X \in (b_X, q_X)$ then an intermediate value argument using the fact that R_X is a geodesic implies $d_X(F(p), F(Q)) \leq 3\delta$, again a contradiction. Similarly q_X is not in (p_X, a_X) . Thus, $[p_X, a_X] \cap [b_X, q_X]$ is either empty or is equal to $[a_X, b_X]$. We deduce that

$$\begin{aligned} d_X(p_X, q_X) &\geq 2h(T) - t(p) - t(q) - 4\delta - 2\delta - 2C_0 - 2R/\epsilon - R \\ &\geq 2t(r) - t(p) - t(q) - 7\Delta - 4C_0. \end{aligned}$$

The proof of Lemma 2.2 is complete. \square

3 Foliations and projections

Let Z be a closed surface of genus at least 2 and \mathbf{z} a set of marked points. A *measured singular foliation* \mathcal{F} on (Z, \mathbf{z}) is a singular topological foliation so that

- \mathcal{F} has only prong-type singularities,
- all one-prong singularities of \mathcal{F} appear at points of \mathbf{z} , and
- \mathcal{F} is equipped with a transverse measure of full support.

We refer the reader to [8, 20] for a detailed discussion of measured foliations. Two measured (respectively, topological) foliations are *measure equivalent* (respectively, *topologically equivalent*) if they differ by isotopy and Whitehead moves. We will only be concerned with those foliations which appear as the vertical foliation for some meromorphic quadratic differential on Z (see Section 4.1). Every measured singular foliation is measure equivalent to such a foliation for a fixed complex structure on Z ; see [13].

The space of measure classes of measured foliation on (Z, \mathbf{z}) is denoted by $\mathcal{MF}(Z, \mathbf{z})$ and its projectivization by $\mathbb{P}\mathcal{MF}(Z, \mathbf{z})$. A measured foliation $\mathcal{F} \in \mathcal{MF}(Z, \mathbf{z})$ is *arational* if it has no closed leaf cycles. We say that \mathcal{F} is *uniquely ergodic* if whenever $\mathcal{F}' \in \mathcal{MF}(Z, \mathbf{z})$ is topologically equivalent to \mathcal{F} , then \mathcal{F} and \mathcal{F}' project to the same point in $\mathbb{P}\mathcal{MF}(Z, \mathbf{z})$. Both of these notions depend only on the topological classes of the foliations, and not the transverse measures.

If \mathcal{F} is a measured foliation representing an element of $\mathcal{MF}(Z)$, and $\mathbf{z} \subset Z$ is a set of marked points, then \mathcal{F} also determines an element of $\mathcal{MF}(Z, \mathbf{z})$. We note that it is important in this case that \mathcal{F} be a foliation, and not an equivalence class of foliations. If \mathcal{F} is arational as an element of $\mathcal{MF}(Z)$, and if $\mathbf{z} = \{z\}$ is a single point, then \mathcal{F} is also arational as an element of $\mathcal{MF}(Z, z)$; see [19].

By a *strict subsurface* $Y \subset Z - \mathbf{z}$ we mean a properly embedded surface with *nonempty* boundary and a set of punctures, possibly empty, such that every component of ∂Y is an *essential curve* in $Z - \mathbf{z}$; that is, homotopically nontrivial and nonperipheral. We also assume that Y is not a sphere with k punctures and j boundary components where $k + j = 3$. We will only refer to subsurfaces in one context, and that is as follows. Given a pair of arational measured foliation $\mathcal{F}, \mathcal{G} \in \mathcal{MF}(Z, \mathbf{z})$ and a proper subsurface $Y \subset Z - \mathbf{z}$, we have the *projection distance*

$$d_Y(\mathcal{F}, \mathcal{G}) \in \mathbb{Z}_{\geq 0}$$

between \mathcal{F} and \mathcal{G} in Y . This is the distance in the arc-and-curve complex of Y between the the subsurface projections of \mathcal{F} and \mathcal{G} to Y . For a detailed discussion, see [23, 24]. All we use is that d_Y satisfies a triangle inequality

$$d_Y(\mathcal{F}_1, \mathcal{F}_2) \leq d_Y(\mathcal{F}_1, \mathcal{G}) + d_Y(\mathcal{G}, \mathcal{F}_2)$$

for all arational measured foliations $\mathcal{F}_1, \mathcal{F}_2, \mathcal{G} \in \mathcal{MF}(Z, \mathbf{z})$. This relates to Teichmüller geometry by Theorem 4.2 below.

4 Teichmüller spaces

Here we set notation and recall some basic properties of Teichmüller space. For background on Teichmüller space, we refer the reader to any of [2, 10, 1, 14].

4.1 Teichmüller space, quadratic differentials and geodesics

Given a closed Riemann surface Z with a finite (possibly empty) set of marked points $\mathbf{z} \subset Z$, let $\mathcal{T}(Z, \mathbf{z})$ denote the *Teichmüller space* of equivalence classes of marked Riemann surfaces

$$\mathcal{T}(Z, \mathbf{z}) = \left\{ [f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})] \mid \begin{array}{l} f \text{ is an orientation preserving homeo-} \\ \text{morphism to the Riemann surface } X \end{array} \right\}.$$

The equivalence relation is defined by

$$(f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})) \sim (g: (Z, \mathbf{z}) \rightarrow (Y, \mathbf{y}))$$

if $f \circ g^{-1}: (Y, \mathbf{y}) \rightarrow (X, \mathbf{x})$ is isotopic (rel marked points) to a conformal map. If $\mathbf{z} = \emptyset$, then we write $\mathcal{T}(Z) = \{[f: Z \rightarrow X]\}$.

The *Teichmüller distance* on $\mathcal{T}(Z, \mathbf{z})$ is defined by

$$d_{\mathcal{T}}([f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})], [g: (Z, \mathbf{z}) \rightarrow (Y, \mathbf{y})]) = \inf \left\{ \frac{1}{2} \log(K_h) \mid h \simeq f \circ g^{-1} \right\}$$

where K_h is the dilatation of h and where $h: (Y, \mathbf{y}) \rightarrow (X, \mathbf{x})$ ranges over all quasi-conformal maps isotopic (rel marked points) to $f \circ g^{-1}$.

Given $\epsilon > 0$, the ϵ -thick part of Teichmüller space $\mathcal{T}_\epsilon(Z, \mathbf{z}) \subset \mathcal{T}(Z, \mathbf{z})$ is the set of points $[f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})] \in \mathcal{T}(Z, \mathbf{z})$ where the unique complete hyperbolic surface in the conformal class of $X - \mathbf{x}$ has its shortest geodesic of length at least ϵ . When ϵ is understood from context we will simply refer to $\mathcal{T}_\epsilon(Z, \mathbf{z})$ as the *thick part of Teichmüller space*.

Let $\mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{M}(Z, \mathbf{z})$ denote the projection to moduli space obtained by forgetting the marking

$$[f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})] \mapsto [(X, \mathbf{x})]$$

or, equivalently, by taking the quotient by the mapping class group. Mumford's compactness criterion [3] now implies: For any $\epsilon > 0$, the thick part $\mathcal{T}_\epsilon(Z, \mathbf{z})$ projects to a compact subset of $\mathcal{M}(Z, \mathbf{z})$. Conversely, the preimage of any compact subset of $\mathcal{M}(Z, \mathbf{z})$ is contained in $\mathcal{T}_\epsilon(Z, \mathbf{z})$ for some $\epsilon > 0$.

Suppose (X, \mathbf{x}) is a closed Riemann surface with marked points and $q \in \mathcal{Q}(X, \mathbf{x})$ is a unit norm, meromorphic quadratic differential with all poles simple and contained in \mathbf{x} . We also use q to denote the associated Euclidean cone metric on X . We note that $\mathcal{Q}(X) \subset \mathcal{Q}(X, \mathbf{x})$, for any set of marked point $\mathbf{x} \subset X$. Given $q \in \mathcal{Q}(X)$ we view it as an element of $\mathcal{Q}(X, \mathbf{x})$ whenever it is convenient.

Given $q \in \mathcal{Q}(X, \mathbf{x})$ and $t \in \mathbb{R}$, let $g_t: (X, \mathbf{x}) \rightarrow (X_t, g_t(\mathbf{x}))$ denote the e^{2t} -quasi-conformal Teichmüller mapping defined by (q, t) . Let $q_t \in \mathcal{Q}(X_t, g_t(\mathbf{x}))$ denote the terminal quadratic differential. For any point $p \in X$ which is not a zero or pole of q we have a preferred coordinate z_0 for (X, q) near p and preferred coordinate z_t for (X_t, q_t) near $g_t(p)$. In these coordinates $q = dz_0^2$ and $q_t = dz_t^2$, and g_t is given by $(u, v) \mapsto (e^t u, e^{-t} v)$. If we mark (X, \mathbf{x}) by $f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})$, then setting $f_t = g_t \circ f$ we have

$$\tau_q(t) = [f_t: (Z, \mathbf{z}) \rightarrow (X_t, g_t(\mathbf{x}))]$$

being a Teichmüller geodesic through $[f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})]$; note that every Teichmüller geodesic can be described in this way. The Teichmüller geodesic τ is ϵ -thick if the image of τ lies in $\mathcal{T}_\epsilon(Z, \mathbf{z})$. We also simply say a geodesic τ is *thick* if it is ϵ -thick for some $\epsilon > 0$. A collection of geodesics $\{\tau_\alpha\}$ is *uniformly thick* if there is an $\epsilon > 0$ so that each τ_α is ϵ -thick.

Given $q \in \mathcal{Q}(X, \mathbf{x})$ we will let $\mathcal{F}(q), \mathcal{G}(q)$ denote the vertical and horizontal foliations respectively; that is, the preimage in preferred coordinates of the foliations of \mathbb{C} by vertical and horizontal lines. For $q \in \mathcal{Q}(X, \mathbf{x})$ and $t \in \mathbb{R}$ consider the associated Teichmüller mapping $g_t: (X, \mathbf{x}) \rightarrow (X_t, g_t(\mathbf{x}))$ as above. If $c: \mathbb{R} \rightarrow X$ is a nonsingular leaf of $\mathcal{F}(q)$ parameterized by arc-length with respect to the q -metric, then composing with g_t we obtain a nonsingular leaf of the vertical foliation for the terminal quadratic differential $\mathcal{F}(q_t)$,

$$g_t \circ c: \mathbb{R} \rightarrow X_t.$$

From the description of g_t in local coordinates we see that this is parameterized *proportional* to arc-length and, in fact, the q_t -length is given by

$$\ell_{q_t}(g_t \circ c|_{[x, x']}) = e^{-t}|x' - x|. \quad (3)$$

4.2 Properties of Teichmüller geodesics

Suppose $\tau = \tau_q$ is the Teichmüller geodesic determined by $[f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})] \in \mathcal{T}(Z, \mathbf{z})$ and $q \in \mathcal{Q}(X, \mathbf{x})$. The forward asymptotic behavior of τ is reflected in the structure of the vertical foliation $\mathcal{F}(q)$. For us, the most important instance of this is a result of Masur [22].

Theorem 4.1 (Masur). *If there exists $\epsilon > 0$ and $\{t_k\}_{k=1}^\infty$ such that*

- $t_k \rightarrow \infty$ as $k \rightarrow \infty$ and

- $\tau_q(t_k) \in \mathcal{T}_\epsilon(Z, \mathbf{z})$ for all k

then $\mathcal{F}(q)$ is arational and uniquely ergodic.

In particular, if τ_q is thick then both $\mathcal{F}(q)$ and $\mathcal{G}(q)$ are uniquely ergodic. We say a pair of arational foliations \mathcal{F} and \mathcal{G} are K -cobounded if for all strict subsurfaces $Y \subset X - \mathbf{x}$ we have $d_Y(\mathcal{F}, \mathcal{G}) \leq K$. A result of Rafi [26, Theorem 1.5] relates the thickness of a geodesic $\tau_q \subset \mathcal{T}$ to the coboundedness of the associated vertical and horizontal foliations.

Theorem 4.2 (Rafi). *For all $\epsilon > 0$ there exists $K > 0$ so that if $q \in \mathcal{Q}(X, \mathbf{x})$ has τ_q being ϵ -thick then $\mathcal{F}(q)$ and $\mathcal{G}(q)$ are K -cobounded.*

Conversely, for all $K > 0$ there exists $\epsilon > 0$ so that if $q \in \mathcal{Q}(X, \mathbf{x})$ has $\mathcal{F}(q)$ and $\mathcal{G}(q)$ being K -cobounded then τ_q is ϵ -thick.

4.3 Forgetting the marked point: the Bers fibration

Suppose now that Z is a closed surface and $z \in Z$ is a single marked point; we use (Z, z) to denote $(Z, \{z\})$. Let $p: \tilde{Z} \rightarrow Z$ denote the universal covering. Given $[f: (Z, z) \rightarrow (X, f(z))]$ we can forget the marked point to obtain an element $[f: Z \rightarrow X] \in \mathcal{T}(Z)$. This defines a holomorphic map

$$\Pi: \mathcal{T}(Z, z) \rightarrow \mathcal{T}(Z)$$

called the *Bers fibration* [4]. The fiber of this map over $[f: Z \rightarrow X]$ is holomorphically identified with \tilde{X} , the universal covering of X . Moreover, this identification is canonical, up to the action of the covering group on \tilde{X} .

The projection of Teichmüller spaces $\Pi: \mathcal{T}(Z, z) \rightarrow \mathcal{T}(Z)$ descends to a projection of moduli spaces $\hat{\Pi}: \mathcal{M}(Z, z) \rightarrow \mathcal{M}(Z)$. The fiber of $\hat{\Pi}$ over $X \in \mathcal{M}(Z)$ is just $X/\text{Aut}(X)$ and this is compact.

Recall that puncturing a closed surface once increases the hyperbolic systole. (Lift to universal covers and apply the Schwarz-Pick lemma.) It follows that the preimage of $\mathcal{T}_\epsilon(Z)$ by Π^{-1} is contained in $\mathcal{T}_\epsilon(Z, z)$.

By a theorem of Royden [28] the Teichmüller metric agrees with the Kobayashi metric on Teichmüller space. Recall that the inclusion of the universal covering $\tilde{X} \rightarrow \mathcal{T}(Z, z)$ is a holomorphic embedding [4]. Thus, if we give \tilde{X} the Poincaré metric ρ_0 — one-half the hyperbolic metric — then $(\tilde{X}, \rho_0) \rightarrow (\mathcal{T}(Z, z), d_{\mathcal{T}})$ is a contraction [16]. Kra [17] further proved the following.

Theorem 4.3 (Kra). *There exists a homeomorphism $h: [0, \infty) \rightarrow [0, \infty)$ so that for any $[f: Z \rightarrow X] \in \mathcal{T}(Z)$, and any $\tilde{x}_1, \tilde{x}_2 \in \tilde{X} \subset \mathcal{T}(Z, z)$, we have*

$$h(\rho_0(\tilde{x}_1, \tilde{x}_2)) \leq d_{\mathcal{T}}(\tilde{x}_1, \tilde{x}_2) \leq \rho_0(\tilde{x}_1, \tilde{x}_2).$$

The function h can be described concretely in terms of the solution to a certain extremal mapping problem for the hyperbolic plane which was solved by Teichmüller [29] and Gehring [11]. We will extend h to a nondecreasing function, $h: \mathbb{R} \rightarrow [0, \infty)$ by declaring $h(t) = 0$ for all $t \leq 0$.

4.4 Branched covers

Here we use branched covers to induce maps on Teichmüller space.

Suppose $P: \Sigma \rightarrow Z$ is a branched cover, branched over some finite set of points $\mathbf{z} \subset Z$. Then any complex structure on Z pulls back to a complex structure on Σ , and thus induces a map $P^*: \mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{T}(\Sigma)$. Regarding Teichmüller space as the space of marked Riemann surfaces, $\mathcal{T}(Z, \mathbf{z}) = \{[f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})]\}$, the embedding is described as follows. The branched covering $P: \Sigma \rightarrow (Z, \mathbf{z})$ induces a branched covering

$U: \Omega \rightarrow (X, \mathbf{x})$, for some Riemann surface Ω , namely the branched cover induced by the subgroup $(f \circ P)_*(\pi_1(\Sigma - P^{-1}(\mathbf{z}))) < \pi_1(X - \mathbf{x})$. By construction, there is a lift of the marking homeomorphism $\phi: \Sigma \rightarrow \Omega$. This is described by the following commutative diagram.

$$\begin{array}{ccc} \Sigma & \xrightarrow{\phi} & \Omega \\ P \downarrow & & \downarrow U \\ (Z, \mathbf{z}) & \xrightarrow{f} & (X, \mathbf{x}). \end{array}$$

Then, we have

$$P^*([f: (Z, \mathbf{z}) \rightarrow (X, \mathbf{x})]) = [\phi: \Sigma \rightarrow \Omega].$$

We now give a well-known consequence of these definitions.

Proposition 4.4. *If $P: \Sigma \rightarrow Z$ is nontrivially branched at every point of $P^{-1}(\mathbf{z})$, then $P^*: \mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{T}(\Sigma)$ is an isometric embedding. Moreover, for all $\epsilon > 0$ there exists $\epsilon' > 0$ so that $P^*(\mathcal{T}_\epsilon(Z, \mathbf{z})) \subset \mathcal{T}_{\epsilon'}(\Sigma)$.*

Proof. When P is a covering then P^* is an isometric embedding; see [27, Section 7]. The proof is identical in the presence of nontrivial branching, as a one-prong singularity at a point of \mathbf{z} lifts to a regular point or to a three-prong or higher singularity.

Let $\widetilde{\mathcal{M}}(Z, \mathbf{z})$ be the quotient of $\mathcal{T}(Z, \mathbf{z})$ by the group of mapping classes of (Z, \mathbf{z}) that lift to Σ . Note that $\widetilde{\mathcal{M}}(Z, \mathbf{z}) \rightarrow \mathcal{M}(Z, \mathbf{z})$ is a finite sheeted (orbifold) covering. The embedding $P^*: \mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{T}(\Sigma)$ descends to a map $\widetilde{\mathcal{M}}(Z, \mathbf{z}) \rightarrow \mathcal{M}(\Sigma)$, giving a commutative square.

$$\begin{array}{ccc} \mathcal{T}(Z, \mathbf{z}) & \xrightarrow{P^*} & \mathcal{T}(\Sigma) \\ \downarrow & & \downarrow \\ \widetilde{\mathcal{M}}(Z, \mathbf{z}) & \longrightarrow & \mathcal{M}(\Sigma) \end{array}$$

By Mumford's compactness criteria [3], the image of $\mathcal{T}_\epsilon(Z, \mathbf{z})$ in $\widetilde{\mathcal{M}}(Z, \mathbf{z})$ is compact, and hence so is the image in $\mathcal{M}(\Sigma)$. Appealing to Mumford's criteria again (for $\mathcal{M}(\Sigma)$), it follows that for some $\epsilon' > 0$ we have $P^*(\mathcal{T}_\epsilon(Z, \mathbf{z})) \subset \mathcal{T}_{\epsilon'}(\Sigma)$. \square

In general, for any branched cover $P: \Sigma \rightarrow Z$, branched over $\mathbf{z} \subset Z$, consider $\sigma = P^{-1}(\mathbf{z})$ as a set of marked points on Σ . Then again there is an isometric embedding

$$P^*: \mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{T}(\Sigma, \sigma).$$

If $\omega \subset \sigma$ then define $\Pi_\omega: \mathcal{T}(\Sigma, \sigma) \rightarrow \mathcal{T}(\Sigma, \omega)$ by forgetting the points of σ not in ω . When ω is empty we may omit the subscript. In this notation, the composition $\Pi \circ P^*$ gives the map of Proposition 4.4. So, if P is non-trivially branched at all points of σ then $\Pi \circ P^*$ is an isometric embedding. If P is not branched at all points of σ then $\Pi \circ P^*$ fails to be an isometric embedding; however it remains 1-Lipschitz.

Proposition 4.5. *If $P: \Sigma \rightarrow Z$ is branched over \mathbf{z} and if $\omega \subset \sigma = P^{-1}(\mathbf{z})$ is any subset then*

$$\Pi_\omega \circ P^*: \mathcal{T}(Z, \mathbf{z}) \rightarrow \mathcal{T}(\Sigma, \omega)$$

is 1-Lipschitz.

Proof. The Bers fibration is a holomorphic map [4] and, by forgetting the points of $\sigma - \omega$ one at a time, we see that $\Pi_\omega: \mathcal{T}(\Sigma, \sigma) \rightarrow \mathcal{T}(\Sigma, \omega)$ is a composition of holomorphic maps, hence holomorphic. In particular, because the Teichmüller metric agrees with the Kobayashi metric [28], it follows that Π_ω is 1-Lipschitz [16]. Since P^* is an isometric embedding, the composition is 1-Lipschitz. \square

5 An inductive construction

The proof of Theorem 1.1 is constructive, but also appeals to an inductive procedure. We begin by constructing the required embedding of \mathbb{H}^2 into some Teichmüller space as the base case of the induction, then produce an embedding of \mathbb{H}^3 into some other Teichmüller space, then an embedding of \mathbb{H}^4 , and so on. All the main ideas and technical difficulties are present in the construction of the embedding of \mathbb{H}^2 and then the embedding of \mathbb{H}^3 from that of \mathbb{H}^2 . The only further complications which arise to describe the embedding of \mathbb{H}^n from \mathbb{H}^{n-1} for $n \geq 4$ are in the notation, which becomes increasingly messy as n increases. This is due to the fact that the proof for n really depends on the proof for all $2 \leq k < n$ (rather than just $n - 1$). For this reason, we carefully describe the cases $n = 2$ and $n = 3$, and sketch the general inductive step indicating only those things that require modification.

5.1 The hyperbolic plane case

Let Z be a closed hyperbolic surface. Let $q \in \mathcal{Q}(Z)$ be a nonzero holomorphic quadratic differential on Z so that the associated Teichmüller geodesic $[g_t: Z \rightarrow Z_t]$ is thick. Write $\mathcal{F} = \mathcal{F}(q)$ and $\mathcal{G} = \mathcal{G}(q)$ for the vertical and horizontal foliations of q , respectively. Next, let $c: \mathbb{R} \rightarrow Z$ be a nonsingular leaf of \mathcal{F} parameterized by arc-length with respect to q and let $z = c(0)$ be a marked point on Z ; see Section 4.

Our goal is to construct an almost-isometric embedding

$$\mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z).$$

We consider an isotopy $Z \times \mathbb{R} \rightarrow Z$, written $(w, x) \mapsto f^x(w)$, where $f^x: Z \rightarrow Z$ is a homeomorphism for all $x \in \mathbb{R}$, f^0 is the identity and $f^x(z) = c(x)$ for all $x \in \mathbb{R}$. We further assume that f^x preserves \mathcal{F} for all $x \in \mathbb{R}$. To construct such an isotopy we can, for example, take the flow of a complete, time-dependent vector field ξ_x on Z which is tangent to \mathcal{F} for all $x \in \mathbb{R}$ and with $\xi_x(c(x)) = c'(x)$. We think of the isotopy as “pushing z along c ”. This determines the *horocyclic coordinate*

$$\tilde{c}: \mathbb{R} \rightarrow \mathcal{T}(Z, z)$$

given by

$$\tilde{c}(x) = [f^x: (Z, z) \rightarrow (Z, c(x))].$$

The image of \tilde{c} lies in the Bers fiber over the basepoint $[\text{Id}: Z \rightarrow Z] \in \mathcal{T}(Z)$; the fiber is identified with the universal cover \tilde{Z} of Z . As such, we can identify \tilde{c} with a lift of c to \tilde{Z} and write

$$\tilde{c}: \mathbb{R} \rightarrow \tilde{Z} \subset \mathcal{T}(Z, z).$$

Applying the Teichmüller mapping $g_t: Z \rightarrow Z_t$ determined by q and $t \in \mathbb{R}$ gives the *height coordinate*. These coordinates together define $\mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z)$ where

$$\mathbf{Z}(x, t) = [g_t \circ f^x: (Z, z) \rightarrow (Z_t, g_t(c(x)))].$$

Here we are using the coordinates (x, t) on \mathbb{H}^2 described in Section 2.

Since the marking homeomorphisms are determined by x and t , we simplify notation and denote the values in Teichmüller space by

$$\mathbf{Z}(x, t) = \tilde{c}_t(x) = (Z_t, g_t(c(x))). \tag{4}$$

We also write

$$\mathbf{Z}(x, 0) = \tilde{c}(x) = (Z, c(x)).$$

As the notation suggests, $\tilde{c}_t: \mathbb{R} \rightarrow \tilde{Z}_t \subset \mathcal{T}(Z, z)$ is a lift of $g_t \circ c: \mathbb{R} \rightarrow Z_t$ to the universal cover \tilde{Z}_t , thought of as the fiber over $[g_t: Z \rightarrow Z_t]$.

Theorem 5.1. *The map $\mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z)$ is an almost-isometric embedding. Moreover, the image lies in the thick part and is quasi-convex.*

Proof. We verify the hypothesis of Lemma 2.2 to prove that \mathbf{Z} is an almost-isometric embedding and, along the way, prove that the image is quasi-convex and lies in the thick part.

First, fix any $x \in \mathbb{R}$ so that $\eta_x(t) = (x, t)$ is a vertical geodesic in \mathbb{H}^2 . Then $t \mapsto \mathbf{Z} \circ \eta_x(t) = \mathbf{Z}(x, t) = (Z_t, g_t(c(x)))$ is a Teichmüller geodesic, and hence Property 1 of Lemma 2.2 holds. Furthermore, since $t \mapsto Z_t$ is a thick geodesic, we see that $\{\mathbf{Z} \circ \eta_x(t)\}_{x \in \mathbb{R}}$ are uniformly thick geodesics. That is, that the union of these geodesics, over all $x \in \mathbb{R}$, project into a compact subset of $\mathcal{M}(Z, z)$; namely, the preimage of the compact subset of $\mathcal{M}(Z)$ containing the image of $t \mapsto Z_t$ (see Section 4.3). In particular, the image of \mathbf{Z} lies in the thick part of $\mathcal{T}(Z, z)$.

For each $x \in \mathbb{R}$, the geodesic $\mathbf{Z} \circ \eta_x$ is defined by the quadratic differential $q \in \mathcal{Q}(Z)$ viewed as a quadratic differential in $\mathcal{Q}(Z, c(x))$. We denote the vertical and horizontal foliations of $q \in \mathcal{Q}(Z, c(x))$ by \mathcal{F}^x and \mathcal{G}^x , respectively, and consider them as measured foliations in $\mathcal{MF}(Z, z)$ by pulling them back via f^x . Since f^x preserves \mathcal{F} , it follows that $\mathcal{F}^x = \mathcal{F}^0 \in \mathcal{MF}(Z, z)$ for all $x \in \mathbb{R}$.

Now, since $t \mapsto Z_t$ is a thick geodesic, by Theorem 4.1 the foliations \mathcal{F} and \mathcal{G} are arational. Puncturing an arational foliation once gives an arational foliation in the punctured surface. Hence \mathcal{F}^x and \mathcal{G}^x are also arational for all x . Since $\mathcal{F}^x = \mathcal{F}^0$ for all $x \in \mathbb{R}$ and since the geodesics $\{\mathbf{Z} \circ \eta_x\}_{x \in \mathbb{R}}$ are uniformly thick, Theorem 4.2 implies that there exists $K > 0$ so that the pairs $(\mathcal{F}^0, \mathcal{G}^x) = (\mathcal{F}^x, \mathcal{G}^x)$ are K -cobounded for all x . By the triangle inequality (applied to each subsurface Y) we see that for all $x, x' \in \mathbb{R}$ the pair $(\mathcal{G}^x, \mathcal{G}^{x'})$ is $2K$ -cobounded.

Appealing to the other direction in Theorem 4.2 the geodesic $\Gamma^{x, x'}$, determined by \mathcal{G}^x and $\mathcal{G}^{x'}$ for distinct $x, x' \in \mathbb{R}$, is uniformly thick, independent of x and x' . From this and [15, Theorem 4.4] it follows that there is a $\delta > 0$ so that $\mathbf{Z} \circ \eta_x$, $\mathbf{Z} \circ \eta_{x'}$ and $\Gamma^{x, x'}$ are the sides of a δ -slim triangle for every pair of distinct points $x, x' \in \mathbb{R}$, and hence Property 2 of Lemma 2.2 holds. From this, it follows that $\mathbf{Z}(\mathbb{H}^2)$ has Hausdorff distance at most δ from the union of the geodesics

$$\mathbf{Z}(\mathbb{H}^2) \cup \left(\bigcup_{x \neq x' \in \mathbb{R}} \Gamma^{x, x'} \right) = \left(\bigcup_{x \in \mathbb{R}} \mathbf{Z} \circ \eta_x \right) \cup \left(\bigcup_{x \neq x' \in \mathbb{R}} \Gamma^{x, x'} \right).$$

This is precisely the *weak hull* of $\{\mathcal{G}^x\}_{x \in \mathbb{R}} \cup \{\mathcal{F}^0\} \subset \mathbb{P}\mathcal{MF}(Z, z)$, and so according to [15, Theorem 4.5], this set, hence also $\mathbf{Z}(\mathbb{H}^2)$, is quasi-convex (the assumption in [15] that the subset of $\mathbb{P}\mathcal{MF}(Z)$ be closed was not used in the proof).

Finally, we must prove that Properties 3 and 4 of Lemma 2.2 hold. For this we can appeal directly to Theorem 4.3. More precisely, observe that because $\{Z_t\}_{t \in \mathbb{R}}$ lies in the thick part, the pull-back of the flat metric on \tilde{Z}_t (which we also denote q_t) is uniformly quasi-isometric to the Poincaré metric ρ_0 on \tilde{Z}_t . That is, there exist constants $A, B \geq 0$ so that

$$\frac{1}{A} (d_{q_t}(\tilde{z}, \tilde{z}') - B) \leq \rho_0(\tilde{z}, \tilde{z}') \leq A d_{q_t}(\tilde{z}, \tilde{z}') + B \quad (5)$$

for all $t \in \mathbb{R}$ and $\tilde{z}, \tilde{z}' \in Z_t$ (see for example [7, Lemma 2.2]).

Applying (4), the upper bound of Theorem 4.3, (5) and (3), in that order, we find

$$\begin{aligned}
d_{\mathcal{T}}(\mathbf{Z}(x, t), \mathbf{Z}(x', t)) &= d_{\mathcal{T}}(\tilde{c}_t(x), \tilde{c}_t(x')) \\
&\leq \rho_0(\tilde{c}_t(x), \tilde{c}_t(x')) \\
&\leq A d_{q_t}(\tilde{c}_t(x), \tilde{c}_t(x')) + B \\
&= Ae^{-t}|x' - x| + B.
\end{aligned}$$

So, setting $\epsilon = 1$ and $R = A + B$, Property 3 of Lemma 2.2 holds.

On the other hand, (4), the lower bound of Theorem 4.3, monotonicity of h , and (3) gives

$$\begin{aligned}
d_{\mathcal{T}}(\mathbf{Z}(x, t), \mathbf{Z}(x', t)) &= d_{\mathcal{T}}(\tilde{c}_t(x), \tilde{c}_t(x')) \\
&\geq h(\rho_0(\tilde{c}_t(x), \tilde{c}_t(x'))) \\
&\geq h\left(\frac{1}{A}(d_{q_t}(\tilde{c}_t(x), \tilde{c}_t(x')) - B)\right) \\
&= h\left(\frac{1}{A}(e^{-t}|x' - x| - B)\right).
\end{aligned}$$

From this, and because h is a homeomorphism on $[0, \infty)$ and hence proper, Property 4 of Lemma 2.2 also holds. This completes the proof of Theorem 5.1. \square

5.2 Hyperbolic 3-space

Before diving into the construction, we explain the basic idea. Our embedding of the hyperbolic plane in Section 5.1 sends (x, t) to $\mathbf{Z}(x, t) \in \mathcal{T}(Z, z)$ by pushing the marked point z distance x along a leaf of the vertical foliation of a quadratic differential then travelling distance t along the Teichmüller flow. There is a simple extension of this construction which produces a map of hyperbolic 3-space into Teichmüller space $\mathcal{T}(Z, \{z, w\})$. Take z and w to lie on distinct leaves and send (x, y, t) to the point of $\mathcal{T}(Z, \{z, w\})$ obtained by pushing z a distance x along its leaf, pushing w a distance y along its leaf, and applying the Teichmüller flow for time t .

The problem is that whenever z and w move close to each other on Z , the corresponding point in $\mathcal{T}(Z, \{z, w\})$ is in the thin part of Teichmüller space; if z and w are very close to each other then there is a simple closed curve surrounding z and w having an annular neighborhood of large modulus. This also shows that this map $(x, y, t) \mapsto \mathcal{T}(Z, \{z, w\})$ is not a quasi-isometric embedding. In fact the map is not even coarsely Lipschitz.

A more subtle construction is required. We first choose a branched cover $P: \Sigma \rightarrow Z$, nontrivially branched at each point of $P^{-1}(z)$. According to Proposition 4.4, this induces an isometric embedding of $\mathcal{T}(Z, z)$ into $\mathcal{T}(\Sigma)$. Fix a suitably generic point $w \in (Z, z)$ and pick a point $\sigma \in P^{-1}(w)$. Roughly, we map our three parameters (x, y, t) into $\mathcal{T}(\Sigma, \sigma)$ as follows. The coordinates (x, t) determine $\mathbf{Z}(x, t) \in \mathcal{T}(Z, z)$ as in Section 5.1. The map P^* applied to $\mathbf{Z}(x, t)$ gives a point in $\mathcal{T}(\Sigma)$ as in Section 4.4. Finally use y to determine a point $\Sigma(x, y, t) \in \mathcal{T}(\Sigma, \sigma)$, lying in the Bers fiber above $P^* \circ \mathbf{Z}(x, t)$. On its face, this new construction avoids the problem we had before. In (Z, z) we have only one marked point; after taking the branched covering over z we forget all of the branch points over z . The single image of σ can now move freely enough so that we stay in the thick part of $\mathcal{T}(\Sigma, \sigma)$. We now explain this construction in more detail and prove that the resulting map has all the required properties.

5.2.1 The construction

The notation from Section 5.1 carries over to this section without change. Let $P: \Sigma \rightarrow Z$ be a branched cover, branched over the marked point $z \in Z$ so that P is nontrivially branched at every point of $P^{-1}(z)$. This determines an isometric embedding of Teichmüller spaces

$$P^*: \mathcal{T}(Z, z) \rightarrow \mathcal{T}(\Sigma)$$

by Proposition 4.4. We write

$$P^*([g_t \circ f^x: (Z, z) \rightarrow (Z_t, g_t(c(x)))] = [\phi_t^x: \Sigma \rightarrow \Sigma_t^x]$$

so that ϕ_t^x is a lift of the marking $g_t \circ f^x$, and P_t^x is the induced branched cover making the following commute:

$$\begin{array}{ccc} \Sigma & \xrightarrow{\phi_t^x} & \Sigma_t^x \\ P \downarrow & & \downarrow P_t^x \\ (Z, z) & \xrightarrow{g_t \circ f^x} & (Z_t, g_t(c(x))). \end{array}$$

The quadratic differentials q_t pull back to quadratic differentials μ_t^x on Σ_t^x , and g_t lifts to Teichmüller mappings of the covers

$$\psi_t^x: \Sigma_0^x \rightarrow \Sigma_t^x$$

so that $t \mapsto \Sigma_t^x$ is a Teichmüller geodesic for all x . The lifts satisfy $\phi_t^x = \psi_t^x \circ \phi_0^x$. We have another commutative diagram which may be helpful in organizing all the maps:

$$\begin{array}{ccccc} \Sigma & \xrightarrow{\phi_0^x} & \Sigma_0^x & \xrightarrow{\psi_t^x} & \Sigma_t^x \\ P \downarrow & & \downarrow P_0^x & & \downarrow P_t^x \\ (Z, z) & \xrightarrow{f^x} & (Z, c(x)) & \xrightarrow{g_t} & (Z_t, g_t(c(x))). \end{array}$$

Denote the vertical foliation for μ_t^x by Φ_t^x . Each nonsingular leaf of Φ_t^x maps isometrically to a nonsingular leaf of the vertical foliation \mathcal{F}_t for q_t via the branched covering $\Sigma_t^x \rightarrow Z_t$ since μ_t^x is the pull back of q_t . Choose any nonsingular leaf $\gamma_0^0: \mathbb{R} \rightarrow \Sigma_0^0 = \Sigma$, parameterized by arc-length. Observe that γ_0^0 maps isometrically by P to a leaf $\gamma: \mathbb{R} \rightarrow Z$ for \mathcal{F} . Note that c and γ are distinct leaves; the preimage of c in Σ consists entirely of singular leaves, namely the leaves that meet the branch points of P .

As we vary x , we can continuously choose lifts of γ to leaves $\gamma_0^x: \mathbb{R} \rightarrow \Sigma_0^x$ which agrees with our initial leaf γ_0^0 when $x = 0$. Specifically, we define the lift to be

$$\gamma_0^x = \phi_0^x \circ \left(P|_{\gamma_0^0(\mathbb{R})} \right)^{-1} \circ (f^x)^{-1} \circ \gamma.$$

Composing with the lifts ψ_t^x , we obtain leaves $\gamma_t^x = \psi_t^x \circ \gamma_0^x$. Observe that via the branched covering $P_t^x: \Sigma_t^x \rightarrow Z_t$, γ_t^x projects to the leaf $g_t \circ \gamma$, independent of x . Furthermore, this shows that the μ_t^x -length of the arc $\gamma_t^x([y, y'])$ is the q_t -length of $g_t \circ \gamma$ which is $e^{-t}|y - y'|$.

We pick a basepoint $\sigma = \gamma_0^0(0) \in \Sigma$, and consider the surface (Σ, σ) , marked by the identity $\text{Id} = \phi_0^0$ as a point in $\mathcal{T}(\Sigma, \sigma)$. Just as we constructed f^x by pushing along c to $c(x)$, we push σ along γ_t^x to $\gamma_t^x(y)$ to obtain maps

$$\xi_t^{x,y}: (\Sigma, \sigma) \rightarrow (\Sigma_t^x, \gamma_t^x(y)).$$

Specifically, we take $\xi_0^{x,y}$ to be the composition of ϕ_0^x and a map isotopic to the identity on Σ_0^x which preserves the foliation Φ_0^x and pushes $\phi_0^x(\sigma)$ along γ_0^x to $\gamma_0^x(y)$. Then $\xi_t^{x,y} = \psi_t^x \circ \xi_0^{x,y}$ maps the foliation Φ_0^x to Φ_t^x .

We denote the associated point in Teichmüller space $[\xi_t^{x,y}: (\Sigma, \sigma) \rightarrow (\Sigma_t^x, \gamma_t^x(y))] \in \mathcal{T}(\Sigma, \sigma)$ simply by $(\Sigma_t^x, \gamma_t^x(y))$ as this point is uniquely determined in this construction by (x, y, t) .

We define

$$\Sigma: \mathbb{H}^3 \rightarrow \mathcal{T}(\Sigma, \sigma)$$

in the coordinates (x, y, t) for \mathbb{H}^3 from Section 2 by

$$\Sigma(x, y, t) = (\Sigma_t^x, \gamma_t^x(y)).$$

5.2.2 Fibration over \mathbb{H}^2 case

We also require a slightly different description of the map Σ to take advantage of the construction in the \mathbb{H}^2 case. Observe that $P^* \circ \mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z) \rightarrow \mathcal{T}(\Sigma)$ is an almost-isometric embedding, and is given by

$$P^*(\mathbf{Z}(x, t)) = \Sigma_t^x,$$

where Σ_t^x denotes the point $[\phi_t^x: \Sigma \rightarrow \Sigma_t^x]$. Recall that

$$\Pi: \mathcal{T}(\Sigma, \sigma) \rightarrow \mathcal{T}(\Sigma).$$

is the Bers fibration. If we fix $(x, t) \in \mathbb{H}^2$, then for every y we see that $(\Sigma_t^x, \gamma_t^x(y))$ is contained in the fiber $\Pi^{-1}(\Sigma_t^x)$. Since $\Pi^{-1}(\Sigma_t^x)$ is identified with the universal covering $\tilde{\Sigma}_t^x$ of Σ_t^x , just as in the case of \mathbb{H}^2 we see that $t \mapsto (\Sigma_t^x, \gamma_t^x(y))$ is a lift of γ_t^x to $\tilde{\Sigma}_t^x \subset \mathcal{T}(\Sigma, \sigma)$. As such, we use the alternative notation

$$\tilde{\gamma}_t^x: \mathbb{R} \rightarrow \tilde{\Sigma}_t^x \subset \mathcal{T}(\Sigma, \sigma)$$

with

$$\tilde{\gamma}_t^x(y) = (\Sigma_t^x, \gamma_t^x(y))$$

when it is convenient to do so.

Finally we record the equation

$$\Pi \circ \Sigma(x, y, t) = P^* \circ \mathbf{Z}(x, t) \tag{6}$$

which holds for all $(x, y, t) \in \mathbb{H}^3$. The fact that Π is 1-Lipschitz and $P^* \circ \mathbf{Z}$ is an almost-isometric embedding provides us with useful metric information about Σ .

Theorem 5.2. *The map $\Sigma: \mathbb{H}^3 \rightarrow \mathcal{T}(\Sigma, \sigma)$ is an almost-isometric embedding. Moreover, the image lies in the thick part and is quasi-convex.*

Proof. As before, we will verify the hypothesis of Lemma 2.2 to prove that Σ is an almost-isometry and, along the way, prove that the image is quasi-convex and lies in the thick part.

For all $(x, y) \in \mathbb{R}^2$, the geodesic $\eta_{(x,y)}(t)$ in \mathbb{H}^3 is sent to

$$\Sigma \circ \eta_{(x,y)}(t) = (\Sigma_t^x, \gamma_t^x(y)) = (\psi_t^x(\Sigma_0^x), \psi_t^x(\gamma_0^x(y))).$$

This is a geodesic in $\mathcal{T}(\Sigma, \sigma)$ because $\psi_t^x: \Sigma_0^x \rightarrow \Sigma_t^x$ is a Teichmüller mapping; thus Property 1 follows. Furthermore, note that $\Sigma \circ \eta_{(x,y)}(t)$ lies over $P^* \circ \mathbf{Z} \circ \eta_x(t)$ for all (x, y, t) . Since P is nontrivially branched over every point, the uniform thickness of the set of geodesics $\{\mathbf{Z} \circ \eta_x(t)\}_{x \in \mathbb{R}}$ implies the same for $\{P^* \circ \mathbf{Z} \circ \eta_x(t)\}_{x \in \mathbb{R}}$ by Proposition 4.4,

and hence also for $\{\Sigma \circ \eta_{(x,y)}(t) \mid (x,y) \in \mathbb{R}^2\}$ by (6) as discussed in Section 4.3. That is, $\Sigma(\mathbb{H}^3)$ lies in the thick part.

By our choice of maps $\xi_0^{x,y}$, if we pull back the vertical foliation Φ_0^x of μ_0^x to a foliation $\Phi_0^{x,y} \in \mathcal{MF}(\Sigma, \sigma)$ the result is independent of x and y . Furthermore, Theorem 4.1 implies that these foliations, as well as the pull backs of the horizontal foliations, are arational. Thus all strict subsurface projection distances are defined. Theorem 4.2 and the results of [15] can be applied as in the \mathbb{H}^2 case to prove that Property 2 of Lemma 2.2 is satisfied for some $\delta > 0$. Furthermore, $\Sigma(\mathbb{H}^3)$ is quasi-convex.

We now come to the subtle point of the proof, which is verifying Properties 3 and 4 of Lemma 2.2. We start with Property 3.

Claim. *There exists $\epsilon > 0$ and $R > 0$ so that if $e^{-t}|(x,y) - (x',y')| < \epsilon$ then*

$$d_{\mathcal{T}}(\Sigma(x,y,t), \Sigma(x',y',t)) < R.$$

Before we give the proof, we briefly explain the core technical difficulty. Fix t and define $C_x = g_t(c(x))$ and $\Gamma_y = g_t(\gamma(y))$. Observe that, as before, when we vary y we are simply point pushing; thus the change in Teichmüller distance is controlled by Theorem 4.3. On the other hand, varying x means that we are varying the conformal structure on the closed surface Σ_t^x . This is obtained by varying x in (Z_t, C_x) (which is also point pushing) then taking a branched cover. However, while we vary C_x in Z_t we must also keep track of our y coordinate, which means we should also project $\gamma_t^x(y)$ down to Z_t —this is precisely the point Γ_y . Now if C_x and Γ_y are close together and we vary x so as to push these points apart, then this can result in a large distance in the “auxiliary” Teichmüller space $\mathcal{T}(Z, \{z, w\})$, even for small variation of x . The idea is therefore to first vary y , if necessary, to move $\gamma_t^x(y)$ in Σ_t^x and so guaranteeing that Γ_y is not too close to C_x . We can then vary x as required, then vary y back to its original value. Since the variation of y can be carried out independent of x , this will result in a uniformly bounded change in Teichmüller distance.

Proof of Claim. Since the surfaces $\{\Sigma_t^x\}_{t,x \in \mathbb{R}}$ lie in the thick part, the (pulled back) metrics μ_t^x and the Poincaré metric(s) ρ_0 on the universal cover $\tilde{\Sigma}_t^x$ are uniformly comparable. That is, there exist constants A and B so that for all $\tilde{\sigma}, \tilde{\sigma}' \in \tilde{\Sigma}_t^x$

$$\frac{1}{A}(d_{\mu_t^x}(\tilde{\sigma}, \tilde{\sigma}') - B) \leq \rho_0(\tilde{\sigma}, \tilde{\sigma}') \leq A d_{\mu_t^x}(\tilde{\sigma}, \tilde{\sigma}') + B. \quad (7)$$

Applying Theorem 4.3, Equations (7) and (3) we have

$$\begin{aligned} d_{\mathcal{T}}(\Sigma(x,y,t), \Sigma(x',y',t)) &= d_{\mathcal{T}}(\tilde{\gamma}_t^x(y), \tilde{\gamma}_t^x(y')) \\ &\leq \rho_0(\tilde{\gamma}_t^x(y), \tilde{\gamma}_t^x(y')) \\ &\leq A d_{\mu_t^x}(\tilde{\gamma}_t^x(y), \tilde{\gamma}_t^x(y')) + B \\ &= A(e^{-t}|y - y'|) + B. \end{aligned} \quad (8)$$

We now fix t and the notation $C_x = g_t(c(x))$, $\Gamma_y = g_t(\gamma(y))$. To understand the effect of varying x we must consider the branched covering $P_t^x: \Sigma_t^x \rightarrow (Z_t, C_x)$, but also keep track of the image of our marked point $\gamma_t^x(y) = \psi_t^x(\gamma_0^x(y))$ down in (Z_t, C_x) ; that is, the point Γ_y . This results in the surface Z_t with *two marked points*:

$$(Z_t, \{C_x, \Gamma_y\}).$$

Appealing to Proposition 4.5 we have

$$d_{\mathcal{T}}(\Sigma(x,y,t), \Sigma(x',y',t)) \leq d_{\mathcal{T}}((Z_t, \{C_x, \Gamma_y\}), (Z_t, \{C_{x'}, \Gamma_{y'}\})). \quad (9)$$

This is because we are taking a branched covering, $\Sigma_t^x \rightarrow Z_t$, and then forgetting all but one of the marked points in Σ_t^x .

Since Z_t lies in some fixed thick part of $\mathcal{T}(Z)$ for all $t \in \mathbb{R}$, there exists $\epsilon > 0$ so that the 2ϵ -ball about C_x in the q_t metric, $B_{q_t}(C_x, 2\epsilon)$ is a disk for all $t, x \in \mathbb{R}$ (that is, we have a lower bound on the q_t -injectivity radius of Z_t , independent of t). Now suppose Γ_y lies outside this ball

$$\Gamma_y \notin B_{q_t}(C_x, 2\epsilon).$$

Using again the fact that Z_t lies in some thick part of $\mathcal{T}(Z)$ for all $t \in \mathbb{R}$, it follows that there is some $R' > 0$ with the property that for any point $z' \in B_{q_t}(C_x, \epsilon)$ we have

$$d_{\mathcal{T}}((Z_t, \{C_x, \Gamma_y\}), (Z_t, \{z', \Gamma_y\})) < R'.$$

Here the marking homeomorphism for $(Z_t, \{z', \Gamma_y\})$ is assumed to differ from that of $(Z_t, \{C_x, \Gamma_y\})$ by composition with a homeomorphism of Z_t that is the identity outside $B_{q_t}(C_x, 2\epsilon)$. In particular, if $e^{-t}|x - x'| < \epsilon$ and, crucially, $\Gamma_y \notin B_{q_t}(C_x, 2\epsilon)$ then deduce that $C_{x'} \in B_{q_t}(C_x, \epsilon)$ and, from Equation (9), that

$$d_{\mathcal{T}}(\Sigma(x, y, t), \Sigma(x', y, t)) \leq d_{\mathcal{T}}((Z_t, \{C_x, \Gamma_y\}), (Z_t, \{C_{x'}, \Gamma_y\})) < R'. \quad (10)$$

On the other hand, because the leaves of \mathcal{F} are geodesics for q_t and because $B_{q_t}(C_x, 2\epsilon)$ is a disk, if $\Gamma_y \in B_{q_t}(C_x, 2\epsilon)$ then there exists $y' \in \mathbb{R}$ so that $e^{-t}|y' - y| \leq 2\epsilon$ and

$$\Gamma_{y'} \notin B_{q_t}(C_x, 2\epsilon).$$

Then, from (10) we have

$$d_{\mathcal{T}}(\Sigma(x, y', t), \Sigma(x', y', t)) \leq d_{\mathcal{T}}((Z_t, \{C_x, \Gamma_{y'}\}), (Z_t, \{C_{x'}, \Gamma_{y'}\})) < R'.$$

Combining this, inequalities (8) and (10), and the triangle inequality, it follows that for any x, y, x', t with $e^{-t}|x - x'| \leq \epsilon$ there is some $y' \in \mathbb{R}$ with $e^{-t}|y' - y| \leq 2\epsilon$ such that

$$\begin{aligned} d_{\mathcal{T}}(\Sigma(x, y, t), \Sigma(x', y, t)) &\leq d_{\mathcal{T}}(\Sigma(x, y, t), \Sigma(x, y', t)) + d_{\mathcal{T}}(\Sigma(x, y', t), \Sigma(x', y', t)) \\ &\quad + d_{\mathcal{T}}(\Sigma(x', y', t), \Sigma(x', y, t)) \\ &\leq 2(A(e^{-t}|y - y'|) + B) \\ &\quad + d_{\mathcal{T}}((Z_t, \{C_x, \Gamma_{y'}\}), (Z_t, \{C_{x'}, \Gamma_{y'}\})) \\ &< 2(A2\epsilon + B) + R' \\ &\leq 4(A\epsilon + B) + R'. \end{aligned} \quad (11)$$

Now, let $\epsilon > 0$ be as above and set $R = 5(A\epsilon + B) + R'$. Given $(x, y, t), (x', y', t) \in \mathbb{H}^3$ with $e^{-t}|(x, y) - (x', y')| < \epsilon$, then we have $e^{-t}|x - x'|, e^{-t}|y - y'| \leq e^{-t}|(x, y) - (x', y')| < \epsilon$. Applying Equations (8) and (11) and the triangle inequality we obtain

$$\begin{aligned} d_{\mathcal{T}}(\Sigma(x, y, t), \Sigma(x', y', t)) &\leq d_{\mathcal{T}}(\Sigma(x, y, t), \Sigma(x', y, t)) + d_{\mathcal{T}}(\Sigma(x', y, t), \Sigma(x', y', t)) \\ &\leq 4(A\epsilon + B) + R' + A\epsilon + B \\ &< 5(A\epsilon + B) + R' = R. \end{aligned}$$

This completes the proof of the claim, and so verifies Property 3 of Lemma 2.2. \square

All that remains to show is Property 4 of Lemma 2.2. Suppose we have a sequence of pairs $\{(x_n, y_n, t_n), (x'_n, y'_n, t_n)\}_{n=1}^{\infty}$ with $e^{t_n}|(x_n, y_n) - (x'_n, y'_n)| \rightarrow \infty$ as $n \rightarrow \infty$. Then, up to subsequence, we must be in one of two cases.

Case. $e^{t_n}|x_n - x'_n| \rightarrow \infty$ as $n \rightarrow \infty$.

Forgetting the marked point is 1-Lipschitz, and so we have

$$\begin{aligned} d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) &\geq d_{\mathcal{T}}(\boldsymbol{\Sigma}_{t_n}^{x_n}, \boldsymbol{\Sigma}_{t_n}^{x'_n}) \\ &= d_{\mathcal{T}}((Z_{t_n}, g_{t_n}(x_n)), (Z_{t_n}, g_{t_n}(x'_n))) \\ &= d_{\mathcal{T}}(\mathbf{Z}(x_n, t_n), \mathbf{Z}(x'_n, t_n)). \end{aligned}$$

However, we have already verified that $\mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z)$ satisfies Lemma 2.2. Therefore the last expression tends to infinity, and hence

$$\lim_{n \rightarrow \infty} d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) = \infty$$

as required.

Case. $e^{t_n} |y_n - y'_n| \rightarrow \infty$ as $n \rightarrow \infty$.

If we also have $e^{t_n} |x_n - x'_n| \rightarrow \infty$, then we can appeal to the previous case and we are done. So we assume, as we may, that $e^{t_n} |x_n - x'_n| < M$, for some constant $M > 0$. Since we have already shown that there are $\epsilon, R > 0$ so that part 3 from Lemma 2.2 holds, it follows from (1) that

$$d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y'_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) \leq \frac{R}{\epsilon} (e^{-t_n} |x_n - x'_n|) + R \leq \frac{R}{\epsilon} M + R.$$

Now, by the triangle inequality we have

$$\begin{aligned} d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) &\geq d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y_n, t_n), \boldsymbol{\Sigma}(x_n, y'_n, t_n)) \\ &\quad - d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y'_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) \\ &\geq d_{\mathcal{T}}(\tilde{\gamma}_{t_n}^{x_n}(y_n), \tilde{\gamma}_{t_n}^{x_n}(y'_n)) - \frac{R}{\epsilon} M - R \end{aligned} \tag{12}$$

We can now appeal to Theorem 4.3 as in our proof for $\mathbf{Z}: \mathbb{H}^2 \rightarrow \mathcal{T}(Z, z)$ to find A, B so that

$$d_{\mathcal{T}}(\tilde{\gamma}_{t_n}^{x_n}(y_n), \tilde{\gamma}_{t_n}^{x_n}(y'_n)) \geq h \left(\frac{1}{A} e^{-t_n} |y'_n - y_n| - B \right)$$

The right-hand side tends to infinity by the properness of h , so we can combine this with (12) to obtain

$$\lim_{n \rightarrow \infty} d_{\mathcal{T}}(\boldsymbol{\Sigma}(x_n, y_n, t_n), \boldsymbol{\Sigma}(x'_n, y'_n, t_n)) = \infty$$

as required. Therefore, Property 4 from Lemma 2.2 holds, and the proof of Theorem 5.2 is complete. \square

5.3 The general case

The previous arguments set up an inductive scheme for producing almost-isometric embeddings of \mathbb{H}^n into Teichmüller spaces. The idea is as follows.

For $n - 1 \geq 3$, induction gives us an almost-isometric embedding $\mathbf{W}: \mathbb{H}^{n-1} \rightarrow \mathcal{T}(W, w)$ satisfying all the hypotheses of Lemma 2.2 for some closed surface W with a marked point w . We again take a branched cover

$$P: \Omega \rightarrow W$$

nontrivially branched over each point in $P^{-1}(w) \subset \Omega$. This determines a map

$$P^* \circ \mathbf{W}: \mathbb{H}^{n-1} \rightarrow \mathcal{T}(\Omega).$$

Using the coordinates $(x, t) = (x_1, x_2, \dots, x_{n-2}, t) \in \mathbb{H}^{n-1}$ we write

$$P^* \circ \mathbf{W}(x, t) = \Omega_t^x.$$

Inductively, we assume that the foliation of \mathbb{H}^{n-1} by asymptotic geodesics $\{\eta_x(t)\}_{x \in \mathbb{R}^{n-2}}$ are all mapped by \mathbf{W} to uniformly thick geodesics in $\mathcal{T}(W, w)$, so the same is true for $P^* \circ W$. These geodesics are obtained by applying the Teichmüller mapping $\psi_t^x: \Omega_0^x \rightarrow \Omega_t^x$ giving

$$P^* \circ \mathbf{W}(x, t) = \psi_t^x \circ P^* \circ \mathbf{W}(x, 0)$$

for all $x \in \mathbb{R}^{n-2}$ and $t \in \mathbb{R}$. Furthermore, the defining quadratic differentials all have the same vertical foliation.

We pick a leaf of this foliation $\gamma: \mathbb{R} \rightarrow \Omega$, and arguing as before, this determines a leaf in each surface $\gamma_t^x: \mathbb{R} \rightarrow \Omega_t^x$ with $\gamma_0^0 = \gamma: \mathbb{R} \rightarrow \Omega_0^0 = \Omega$. Now, pick $\omega = \gamma_0^0(0)$ to be our marked point, add a factor of \mathbb{R} to \mathbb{H}^{n-1} with coordinate $y = x_{n-1}$ to obtain \mathbb{H}^n with coordinates $(x, y, t) = (x_1, \dots, x_{n-2}, x_{n-1}, t)$, and define

$$\mathbf{\Omega}: \mathbb{H}^n \rightarrow \mathcal{T}(\Omega, \omega)$$

by

$$\mathbf{\Omega}(x, y, t) = (\Omega_t^x, \gamma_t^x(y)).$$

So we are again pushing a point along a leaf of the vertical foliation.

Theorem 5.3. *The map $\mathbf{\Omega}: \mathbb{H}^n \rightarrow \mathcal{T}(\Omega, \omega)$ is an almost-isometric embedding. Moreover, the image lies in the thick part and is quasi-convex.*

Sketch of proof. Again, we must verify the hypotheses of Lemma 2.2 and prove that the image of $\mathbf{\Omega}$ is quasi-convex in the thick part, assuming that this is true in all previous steps of the construction.

We can argue exactly as in the case of \mathbb{H}^3 to prove Properties 1 and 2 of Lemma 2.2 as well as the fact that the image of $\mathbf{\Omega}$ is quasi-convex in the thick part. Property 3 requires more care. However, once established, Property 4 follows formally, just as in the case of \mathbb{H}^3 .

We elaborate on the proof that Property 3 holds for some ϵ and R . For this, we must give a more precise description of the construction. Write $\Omega_{n-1} = \Omega$, $\Omega_{n-2} = W$ and

$$P_{n-2} = P: \Omega_{n-1} \rightarrow \Omega_{n-2}$$

for the branched cover used in the construction. Inductively, we have a tower of branched covers

$$\Omega_{n-1} \xrightarrow{P_{n-2}} \Omega_{n-2} \xrightarrow{P_{n-3}} \dots \longrightarrow \Omega_2 \xrightarrow{P_1} \Omega_1$$

In this tower, P_j is nontrivially branched at every point $P_j^{-1}(\omega_j)$ where $\omega_j \in \Omega_j$ is the marked point. To clarify, we note that $\Omega_1 = Z$, $\omega_1 = z$, $\Omega_2 = \Sigma$ and $\omega_2 = \sigma$ from the preceding discussion.

We also have a quadratic differential ν_1 on Ω_1 (this is $\nu_1 = q$ from before), which pulls back via all the branched covers to quadratic differentials $\nu_i = P_{i-1}^*(\nu_{i-1}) \in \mathcal{Q}(\Omega_i)$. On Ω_1 , we have chosen $n-1$ distinct nonsingular leaves from the vertical foliation of ν_1 which we denote $\{\zeta_i: \mathbb{R} \rightarrow \Omega_1\}_{i=1}^{n-1}$. These leaves are parametrized by arc-length so that $\zeta_j(0) = P_1 \circ P_2 \circ \dots \circ P_{j-1}(\omega_j)$.

Recall that $y = x_{n-1}$. We can now describe $\Omega(x, y, t) = \Omega(x_1, \dots, x_{n-2}, x_{n-1}, t)$ for any $(x, y, t) \in \mathbb{H}^n$. At the bottom of the tower we push ω_1 along ζ_1 to $\zeta_1(x_1)$, then take the branched cover $\Omega_2^{x_1} \rightarrow (\Omega_1, \zeta_1(x_1))$ induced by P_1 (it is the *induced* branched cover since it branches over $\zeta_1(x_1)$ rather than over $\zeta_1(0) = \omega_1$; see Section 4.4). Next, the lifted marking identifies ω_2 with a point in the preimage of $\zeta_2(0)$, and we push this along an appropriate lift

$\zeta_2^{x_1}$ of ζ_2 to a point $\zeta_2^{x_1}(x_2)$ in the preimage of $\zeta_2(x_2)$. At the next level, there is an branched cover $\Omega_3^{x_1, x_2} \rightarrow (\Omega_2^{x_1}, \zeta_2^{x_1}(x_2))$ induced by P_2 . The lifted marking identifies ω_3 with a point in the preimage of $\zeta_3(0)$ in the composition of branched covers $\Omega_3^{x_1, x_2} \rightarrow \Omega_2^{x_1} \rightarrow \Omega_1$ and we push this along an appropriate lift $\zeta_3^{x_1, x_2}$ of ζ_3 to a point $\zeta_3^{x_1, x_2}(x_3)$ in the preimage of $\zeta_3(x_3)$. We continue in this way to produce a tower of branched covers induced by $P_1, P_2, \dots, P_{n-3}, P_{n-2}$:

$$\Omega_{n-1}^{x_1, \dots, x_{n-2}} \longrightarrow \Omega_{n-2}^{x_1, \dots, x_{n-3}} \longrightarrow \dots \longrightarrow \Omega_3^{x_1, x_2} \longrightarrow \Omega_2^{x_1} \longrightarrow \Omega_1.$$

The point ω_{n-1} is identified with a marked point in $\Omega_{n-1}^{x_1, \dots, x_{n-2}}$ in the preimage of $\zeta_{n-1}(0)$, and then we push this point along an appropriate lift $\zeta_{n-1}^{x_1, \dots, x_{n-2}}$ of ζ_{n-1} to the point $\zeta_{n-1}^{x_1, \dots, x_{n-2}}(y) = \zeta_{n-1}^{x_1, \dots, x_{n-2}}(x_{n-1})$. With this notation

$$\Omega(x, y, 0) = \Omega(x_1, \dots, x_{n-2}, x_{n-1}, 0) = (\Omega_{n-1}^{x_1, \dots, x_{n-2}}, \zeta_{n-1}^{x_1, \dots, x_{n-2}}(x_{n-1})).$$

To find $\Omega(x, y, t)$ for any t , we apply the appropriate Teichmüller deformation to $\Omega(x, y, 0)$. This is the Teichmüller deformation determined by t and the pull back of ν_1 (via the composition of branched covers). We can pull back ν_1 by any of the branched covers, and since the resulting quadratic differential depends only on the surface in this construction, we will simply write Φ_t for the associated Teichmüller deformation on any of the surfaces $\Omega_j^{x_1, \dots, x_{j-1}}$. In particular, we have

$$\Omega(x, y, t) = \Phi_t(\Omega(x, y, 0)).$$

Set $x' = (x'_1, x'_2, \dots, x'_{n-2})$. We now must find an ϵ and R so that if

$$e^{-t}|(x, y) - (x', y')| \leq \epsilon$$

then

$$d_{\mathcal{T}}(\Omega(x, y, t), \Omega(x', y', t)) \leq R.$$

As in the case of \mathbb{H}^3 , appealing to the triangle inequality it suffices to find an ϵ and R' so that if (x_1, \dots, x_{n-2}, y) and $(x'_1, \dots, x'_{n-2}, y')$ agree in all but one coordinate, and in that coordinate differ by at most ϵ , then

$$d_{\mathcal{T}}(\Omega(x, y, t), \Omega(x', y', t)) \leq R'.$$

If (x, y) and (x', y') differ only in the last coordinate, then we can apply Theorem 4.3 just as before to produce $\epsilon = 1$ and $R' = A + B$. Suppose instead that $y = y'$ and x differs from x' in the $n-2$ -coordinate only. We start at the highest coordinate, $y = x_{n-1}$ and work two steps down to x_{n-2} . The idea is similar to what was done in varying x in $(x, y, t) \in \mathbb{H}^3$. We look on $\Phi_t(\Omega_{n-2}^{x_1, \dots, x_{n-3}})$ as an “auxiliary” surface when it is equipped with the *two* marked points $\Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x_{n-2}))$ and the image of $\Phi_t(\zeta_{n-1}^{x_1, \dots, x_{n-2}}(x_{n-1}))$ via the branched covering

$$\Phi_t(\Omega_{n-1}^{x_1, \dots, x_{n-2}}) \rightarrow \Phi_t(\Omega_{n-2}^{x_1, \dots, x_{n-3}}).$$

If these two points are not too close, then we can move from $\Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x_{n-2}))$ to $\Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x'_{n-2}))$ keeping the other marked point fixed, and the distance between these two points in the Teichmüller space of the auxiliary surface with two marked points is uniformly bounded. Since the branched cover induces a 1-Lipschitz map (compare (9)), this means that

$$d_{\mathcal{T}}(\Omega(x_1, \dots, x_{n-3}, x_{n-2}, x_{n-1}, t), \Omega(x_1, \dots, x_{n-3}, x'_{n-2}, x_{n-1}, t))$$

is uniformly bounded.

On the other hand, if the two marked points in $\Phi_t(\Omega_{n-2}^{x_1, \dots, x_{n-3}})$ are close, we

$$\text{move } \Phi_t(\zeta_{n-1}^{x_1, \dots, x_{n-2}}(x_{n-1})) \text{ to } \Phi_t(\zeta_{n-1}^{x_1, \dots, x_{n-2}}(x'_{n-1})),$$

$$\text{move } \Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x_{n-2})) \text{ to } \Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x'_{n-2})),$$

and then

$$\text{move } \Phi_t(\zeta_{n-1}^{x_1, \dots, x'_{n-2}}(x'_{n-1})) \text{ back to } \Phi_t(\zeta_{n-1}^{x_1, \dots, x'_{n-2}}(x_{n-1})).$$

By the triangle inequality, we obtain the desired uniform bound on

$$d_{\mathcal{T}}(\Omega(x_1, \dots, x_{n-3}, x_{n-2}, x_{n-1}, t), \Omega(x_1, \dots, x_{n-3}, x'_{n-2}, x_{n-1}, t)).$$

Note that this required three point pushes in two different auxiliary surfaces. We varied the $(n-1)^{\text{st}}$ coordinate twice, in the highest surface, and varied the $(n-2)^{\text{nd}}$ coordinate once.

Now suppose that x differs from x' in the $(n-3)^{\text{rd}}$ coordinate only. We view $\Phi_t(\Omega_{n-3}^{x_1, \dots, x_{n-4}})$ as an auxiliary surface with *three marked points*: the images of the points $\Phi_t(\zeta_{n-1}^{x_1, \dots, x_{n-2}}(x_{n-1}))$ and $\Phi_t(\zeta_{n-2}^{x_1, \dots, x_{n-3}}(x_{n-2}))$ under the respective branched covers and the point $\Phi_t(\zeta_{n-3}^{x_1, \dots, x_{n-4}}(x_{n-3}))$. We can move this last point a small amount, changing the Teichmüller distance a bounded amount, provided the other two points, higher in the tower, are not too close to it. If they are too close, we first move them out of the way (as in the first two pushes above), move the third point, then move the two higher points back. The triangle inequality together with the 1-Lipschitz property of the branched cover map applied as before, implies a uniform bound on the change in Teichmüller distance

$$d_{\mathcal{T}}(\Omega(x_1, \dots, x_{n-3}, x_{n-2}, x_{n-1}, t), \Omega(x_1, \dots, x'_{n-3}, x_{n-2}, x_{n-1}, t)).$$

It follows that varying x_{n-3} requires at most five point pushes in the three highest auxiliary surfaces.

In general, varying x_{n-k} in this way requires $2k-1$ point pushes in the k highest auxiliary surfaces. Thus we can change any coordinate by a small amount ϵ and change the Teichmüller distance by a bounded amount R' , as required. This completes the sketch of the proof of Theorem 5.3. \square

References

- [1] William Abikoff. *The real analytic theory of Teichmüller space*, volume 820 of *Lecture Notes in Mathematics*. Springer, Berlin, 1980. [6]
- [2] Lars V. Ahlfors. *Lectures on quasiconformal mappings*. Manuscript prepared with the assistance of Clifford J. Earle, Jr. Van Nostrand Mathematical Studies, No. 10. D. Van Nostrand Co., Inc., Toronto, Ont.-New York-London, 1966. [6]
- [3] Lipman Bers. A remark on Mumford's compactness theorem. *Israel J. Math.*, 12:400–407, 1972. [7, 9]
- [4] Lipman Bers. Fiber spaces over Teichmüller spaces. *Acta. Math.*, 130:89–126, 1973. [8, 9]
- [5] Mario Bonk and Bruce Kleiner. Quasi-hyperbolic planes in hyperbolic groups. *Proc. Amer. Math. Soc.*, 133(9):2491–2494 (electronic), 2005. [2]
- [6] Matt Clay, Christopher J Leininger, and Johanna Mangahas. The geometry of right-angled Artin subgroups of mapping class groups. To appear in *Groups Geom. Dyn.*, [arXiv:1007.1129](https://arxiv.org/abs/1007.1129). [2]

- [7] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152 (electronic), 2002. [1, 11]
- [8] A. Fathi, F. Laudenbach, and V. Poénaru. *Travaux de Thurston sur les surfaces*. Société Mathématique de France, Paris, 1991. Séminaire Orsay, Reprint of *Travaux de Thurston sur les surfaces*, Soc. Math. France, Paris, 1979 Astérisque No. 66-67 (1991). [5]
- [9] David Gabai. Almost filling laminations and the connectivity of ending lamination space. *Geom. Topol.*, 13(2):1017–1041, 2009. [2]
- [10] Frederick P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York, 1987. A Wiley-Interscience Publication. [6]
- [11] F. W. Gehring. Quasiconformal mappings which hold the real axis pointwise fixed. In *Mathematical Essays Dedicated to A. J. Macintyre*, pages 145–148. Ohio Univ. Press, Athens, Ohio, 1970. [8]
- [12] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, [arXiv:math.GT/0505244](https://arxiv.org/abs/math/0505244). [1]
- [13] John Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Math.*, 142(3-4):221–274, 1979. [5]
- [14] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. Translated and revised from the Japanese by the authors. [6]
- [15] Richard P. Kent, IV and Christopher J. Leininger. Shadows of mapping class groups: capturing convex cocompactness. *Geom. Funct. Anal.*, 18(4):1270–1325, 2008. [2, 11, 15]
- [16] Shoshichi Kobayashi. Invariant distances on complex manifolds and holomorphic mappings. *J. Math. Soc. Japan*, 19:460–480, 1967. [8, 9]
- [17] Irwin Kra. On the Nielsen-Thurston-Bers type of some self-maps of Riemann surfaces. *Acta Math.*, 146(3-4):231–270, 1981. [8]
- [18] Christopher J. Leininger, Mahan Mj, and Saul Schleimer. The universal Cannon–Thurston map of the curve complex. To appear in *Comm. Math. Helv.*, [arXiv:0808.3521](https://arxiv.org/abs/0808.3521). [2]
- [19] Christopher J. Leininger and Saul Schleimer. Connectivity of the space of ending laminations. *Duke Math. J.*, 150(3):533–575, 2009. [2, 6]
- [20] Howard Masur. Interval exchange transformations and measured foliations. *Ann. of Math. (2)*, 115(1):169–200, 1982. [5]
- [21] Howard Masur. Closed trajectories for quadratic differentials with an application to billiards. *Duke Math. J.*, 53(2):307–314, 1986. [2]
- [22] Howard Masur. Hausdorff dimension of the set of nonergodic foliations of a quadratic differential. *Duke Math. J.*, 66(3):387–442, 1992. [7]
- [23] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999. [2, 6]
- [24] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000. [6]
- [25] Mahan Mj and Pranab Sardar. A combination theorem for metric bundles. Preprint, [arXiv:0912.2715](https://arxiv.org/abs/0912.2715). [1]
- [26] Kasra Rafi. A characterization of short curves of a Teichmüller geodesic. *Geom. Topol.*, 9:179–202, 2005. [8]

- [27] Kasra Rafi and Saul Schleimer. Covers and the curve complex. *Geom. Topol.*, 13(4):2141–2162, 2009. [2, 9]
- [28] 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. [8, 9]
- [29] Oswald Teichmüller. Ein Verschiebungssatz der quasikonformen Abbildung. *Deutsche Math.*, 7:336–343, 1944. [8]