

Universal Cannon-Thurston maps and the boundary of the curve complex

Christopher J. Leininger*, Mahan Mj† and Saul Schleimer‡

August 17, 2008

Abstract

The fundamental group of a closed surface of genus at least two admits a natural action on the curve complex of the surface with one puncture. Combining ideas from previous work of Kent-Leininger-Schleimer and Mitra, we construct a Universal Cannon-Thurston map from a subset of the circle at infinity for the closed surface group onto the boundary of the curve complex of the once-punctured surface. Further, we show that the boundary of this curve complex is locally path-connected.

AMS subject classification = 20F67(Primary), 22E40 57M50

Contents

1	Introduction	2
1.1	Statement of Results	2
1.2	Notation and conventions	3
2	Point position	9
2.1	A bundle over \mathbb{H}	9
2.2	An explicit construction of Φ	10
2.3	A further description of $\mathcal{C}(S, z)$	13
2.4	Extending to measured laminations	14
2.5	Φ and Ψ	20
3	Universal Cannon-Thurston Maps	22
3.1	Quasiconvex Sets	22
3.2	Rays and existence of Cannon-Thurston Maps	24
3.3	Separation	28
3.4	Surjectivity	31

*partially supported by NSF grant DMS-0603881

†partially supported by a UGC Major Research Project grant

‡partially supported by NSF grant DMS-0508971

1 Introduction

1.1 Statement of Results

Fix a hyperbolic metric on a closed surface S of genus at least 2 identifying the universal cover with the hyperbolic plane $p : \mathbb{H} \rightarrow S$. Fix a basepoint $z \in S$ and a point $\tilde{z} \subset p^{-1}(z)$ defining an isomorphism between the group $\pi_1(S, z)$ of homotopy classes of loops based at z and $\pi_1(S)$ the group of covering transformations of $p : \mathbb{H} \rightarrow S$.

We will also regard the basepoint $z \in S$ as a marked point on S . As such, we write (S, z) for the surface S with the marked point z (we could also work with the punctured surface $S - \{z\}$, though a marked point is more convenient for us).

Let $\mathcal{C}(S)$ and $\mathcal{C}(S, z)$ denote the curve complexes of S and (S, z) respectively, and let $\Pi : \mathcal{C}(S, z) \rightarrow \mathcal{C}(S)$ denote the forgetful projection. From [KLS06], the fiber over $v \in \mathcal{C}^0(S)$ is $\pi_1(S)$ -equivariantly isomorphic to the Bass-Serre tree T_v corresponding to v . We define a map

$$\Phi : \mathcal{C}(S) \times \mathbb{H} \rightarrow \mathcal{C}(S, z)$$

sending $\{v\} \times \mathbb{H}$ to $T_v \cong \Pi^{-1}(v) \subset \mathcal{C}(S, z)$ in a $\pi_1(S)$ -equivariant way then extending over simplices using barycentric coordinates (see Section 2.2).

Given $v \in \mathcal{C}^0(S)$, let Φ^v denote the restriction to $\mathbb{H} \cong \{v\} \times \mathbb{H}$

$$\Phi^v : \mathbb{H} \rightarrow \mathcal{C}(S, z).$$

As we will see in Section 3, there are certain rays in \mathbb{H} whose image has finite diameter in $\mathcal{C}(S, z)$ (namely those that eventually project to lie in a proper essential subsurface of S). The remaining rays define a subset $\mathbb{A}_\infty \subset \partial_\infty \mathbb{H}$ (of full Lebesgue measure). Our first main theorem is the following.

Theorem 1.1 (Universal Cannon–Thurston map). *For any $v \in \mathcal{C}^0(S)$, the map $\Phi^v : \mathbb{H} \rightarrow \mathcal{C}(S, z)$ has a continuous $\pi_1(S)$ -equivariant extension*

$$\overline{\Phi}^v : \mathbb{H} \cup \mathbb{A}_\infty \rightarrow \overline{\mathcal{C}}(S, z).$$

Moreover, $\partial\Phi^v = \overline{\Phi}^v|_{\mathbb{A}_\infty}$ is a quotient map onto $\partial_\infty \mathcal{C}(S, z)$ obtained by identifying the endpoints of each leaf and vertices of each complementary polygon of the lifts of every ending lamination on S .

We recall that a Cannon–Thurston map was constructed in the case that the Kleinian group is the fiber subgroup of a closed hyperbolic 3-manifold fibering over the circle by Cannon–Thurston [CT07], then extended to simply degenerate, bounded geometry Kleinian closed surface groups by Minsky [Min92], and proven in the general simply degenerate case by the second author [Mj05],[Mj06].

In all these cases, one produces a quotient map from the circle $\partial_\infty\mathbb{H}$ onto the limit set of the Kleinian group Γ . The quotient is formed by identifying the endpoints of each leaf and the vertices of each polygon of the lift of the ending laminations for Γ (this is either one or two ending laminations depending on whether the group is singly or doubly degenerate).

The map $\partial\Phi^v$ is universal in that it simultaneously identifies the endpoints of each leaf and the vertices of each complementary polygon of the lifts of *every ending lamination on S* . We remark that the restriction to \mathbb{A}_∞ is necessary to get a reasonable quotient: the quotient space of the entire circle $\partial_\infty\mathbb{H}$ identifying this same set of points is a non-Hausdorff space.

Theorem 1.1 and the techniques of its proof are ingredients in our second main theorem.

Theorem 1.2. *The Gromov boundary $\partial_\infty\mathcal{C}(S, z)$ is path connected and locally path connected.*

We remark that \mathbb{A}_∞ is noncompact and totally disconnected, so unlike the proof of local connectivity in the Kleinian group setting, Theorem 1.2 does not follow immediately from Theorem 1.1.

This strengthens the work of the first and third author in [LS08] in a special case: in [LS08] it was shown that the boundary of the curve complex is connected for surfaces of genus at least 2 with any nonzero number of punctures and closed surfaces of genus at least 4. The boundary of the complex of curves describes the space of simply degenerate Kleinian groups as explained in [LS08]. These results seem to be the first ones providing some information about the *topology* of the boundary of the curve complex, a general problem posed by Minsky in his 2006 I.C.M. address. Gabai has now given a proof of Theorem 1.2 for all hyperbolic surfaces Σ , except the 1-punctured torus and the 3- and 4-punctured sphere, where it is known not to be true.

Acknowledgements. The authors wish to thank the Mathematical Sciences Research Institute for its hospitality during the Fall of 2007 where this work was begun. We would also like to thank the other participants of the two programs, *Kleinian Groups and Teichmüller Theory* and *Geometric Group Theory*, for providing a mathematically stimulating and lively atmosphere.

1.2 Notation and conventions

1.2.1 Laminations

For a discussion of laminations, we refer the reader to [PH92], [CEG87], [Bon88], [Thu80], [CB87].

A measured lamination on S is a lamination with a transverse measure of full support. The measured laminations on S will be denoted λ with the support—the underlying lamination—written $|\lambda|$. We require that all our laminations be essential, which can be taken to mean that the leaves lift to uniform quasi-geodesics in the universal cover.

If a is an arc or curve in S and λ a measured lamination, we write $\lambda(a) = \int_a d\lambda$ for the total variation of λ along a . We say that a is transverse to λ if a is transverse to every leaf of $|\lambda|$. If v is the isotopy class of a simple closed curve, then we write

$$i(v, \lambda) = \inf_{\alpha \in v} \lambda(\alpha)$$

for the intersection number of v with λ , where α varies over representatives of the isotopy class v .

Two measured laminations λ_0 and λ_1 are measure equivalent if for every isotopy class of simple closed curve v , $i(v, \lambda_0) = i(v, \lambda_1)$. Every measured lamination is equivalent to a unique measured geodesic lamination (with respect to the fixed hyperbolic structure on S), that is a measured lamination λ for which $|\lambda|$ is a geodesic lamination. Given a measured lamination λ , we let $\hat{\lambda}$ denote the measure equivalent measured geodesic lamination. We will describe a preferred choice of representative of the measure class of a measured lamination in Section 2 below.

We similarly define measured laminations on (S, z) as compactly supported measured laminations on $S - \{z\}$. These are generally *not* realized as geodesic laminations for a hyperbolic metric on $S - \{z\}$, though any one is measure equivalent to a measured geodesic lamination for a complete hyperbolic metric on $S - \{z\}$.

The spaces of (measure classes of) measured laminations will be denoted by $\mathcal{ML}(S)$ and $\mathcal{ML}(S, z)$. The topology on \mathcal{ML} is the weakest topology for which $\lambda \mapsto i(v, \lambda)$ is continuous for every simple closed curve v . Scaling the measures we obtain an action of \mathbb{R}^+ on $\mathcal{ML}(S) - \{0\}$ and $\mathcal{ML}(S, z) - \{0\}$, and we denote the quotient spaces $\mathbb{P}\mathcal{ML}(S)$ and $\mathbb{P}\mathcal{ML}(S, z)$, respectively.

A particularly important subspace for us is the space of filling laminations which we denote \mathcal{FL} . These are the measure classes of measured laminations λ for which all complementary regions of $|\lambda|$ are simply connected (in $S - \{z\}$, there is also one region with cyclic fundamental group generated by the peripheral loop). The quotient of \mathcal{FL} by forgetting the measures will be denoted \mathcal{EL} and is the space of *ending laminations*.

Train tracks provide another useful tool for describing measured laminations. See [Thu80] and [PH92] for a detailed discussion of train tracks and their relation to laminations. We recall some of the most relevant information.

A lamination \mathcal{L} is carried by a train track τ if there is a map $f : S \rightarrow S$ homotopic to the identity with $f(\mathcal{L}) \subset \tau$ so that for every leaf ℓ of \mathcal{L} the restriction of f to ℓ is an immersion. If λ is a measured lamination carried by a train track τ , then the transverse measure defines weights on the branches of τ satisfying the switch condition—the sum of the weights on the incoming branches equals the sum on the outgoing branches. Moreover, any assignment of nonnegative weights to the branches of a train track satisfying the switch condition uniquely determines an element of \mathcal{ML} . Given a train track τ carrying λ , we write $\tau(\lambda)$ to denote the train track τ together with the weights defined by λ .

The following gives a useful tool for working with the topology on \mathcal{ML} (see Theorem 2.7.4 of [PH92]). If $\{\lambda_n\} \cup \{\lambda\} \subset \mathcal{ML}$ is any sequence and τ is a train track for which $|\lambda|$ and each $|\lambda_n|$ is carried by τ , then $\lambda_n \rightarrow \lambda$ if and only if the weights on each branch of τ defined by λ_n converge to those defined by λ .

A well known construction of train tracks carrying a given lamination which will be useful for us is the following (see [PH92], Theorem 1.6.5, for example). Let \mathcal{L} be a geodesic lamination on S , and $\epsilon > 0$ very small so that the ϵ -neighborhood $N_\epsilon(\mathcal{L})$ admits a foliation transverse to \mathcal{L} . The leaves of this foliation are called *ties*. Taking the quotient by collapsing each tie to a point produces a train track τ on S ; see Figure 1.

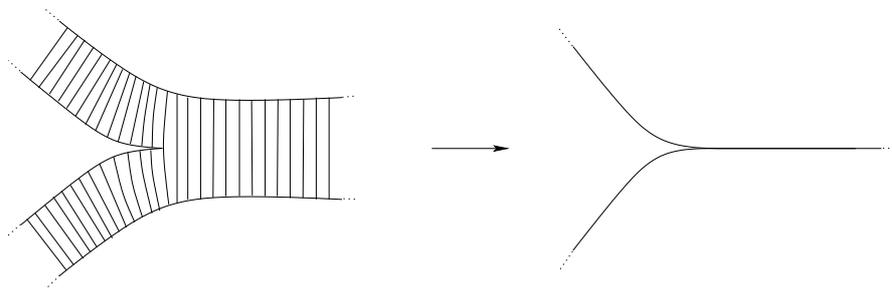


Figure 1: A train track τ constructed from some $N_\epsilon(\mathcal{L})$.

We can view $N_\epsilon(\mathcal{L})$ as being built from finitely many foliated rectangles glued together along arcs of ties in the boundary of the rectangle. In the collapse, each rectangle R projects to a branch β_R of τ . When τ is trivalent, we may assume that $\tau \subset S$ is contained in $N_\epsilon(\mathcal{L})$ with each branch β_R contained in R .

Suppose now that λ is any measured lamination with $|\lambda| \subset N_\epsilon(\mathcal{L})$, and $|\lambda|$ transverse to the ties. If R is a rectangle and a a tie in R , then the weight on the branch β_R associated to R which λ defines is given by $\lambda(a) = \int_a d\lambda$; see Figure 2.

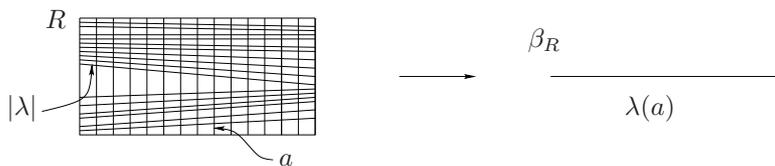


Figure 2: $|\lambda|$ in R and the weight on β_R determined by λ .

1.2.2 Mapping Class Groups

Recall that we have fixed a hyperbolic structure on S as well as a locally isometric universal covering $p : \mathbb{H} \rightarrow S$ and basepoint $\tilde{z} \in p^{-1}(z)$ determining an isomorphism from the covering group $\pi_1(S)$ of $p : \mathbb{H} \rightarrow S$ to $\pi_1(S, z)$, the group of homotopy classes of based loops. All of this is considered fixed for the remainder of the paper.

The **mapping class group** of S is the group $\text{Mod}(S) = \pi_0(\text{Diff}^+(S))$, where $\text{Diff}^+(S)$ is the group of orientation preserving diffeomorphisms of S . We define $\text{Mod}(S, z)$ to be $\pi_0(\text{Diff}^+(S, z))$, where $\text{Diff}^+(S, z)$ is the group of orientation preserving diffeomorphisms of S that fix z .

The evaluation map

$$\text{ev} : \text{Diff}^+(S) \rightarrow S$$

given by $\text{ev}(f) = f(z)$ defines a locally trivial principal fiber bundle

$$\text{Diff}^+(S, z) \rightarrow \text{Diff}^+(S) \rightarrow S.$$

A theorem of Earle and Eells [EE69] says that the component containing the identity $\text{Diff}_0(S)$ is contractible, and so the long exact sequence of a fibration gives rise to J. Birman's exact sequence [Bir69, Bir74]

$$1 \rightarrow \pi_1(S) \rightarrow \text{Mod}(S, z) \rightarrow \text{Mod}(S) \rightarrow 1.$$

We elaborate briefly on the injection $\pi_1(S) \rightarrow \text{Mod}(S, z)$ in Birman's exact sequence. Let

$$\text{Diff}_B(S, z) = \text{Diff}_0(S) \cap \text{Diff}^+(S, z).$$

The long exact sequence of homotopy group identifies $\pi_1(S) \cong \pi_0(\text{Diff}_B(S, z))$. This isomorphism is induced by a homomorphism

$$\text{ev}_* : \text{Diff}_B(S, z) \rightarrow \pi_1(S)$$

given by $\text{ev}_*(h) = [\text{ev}(h_t)]$ where $h_t, t \in [0, 1]$, is a path in $\text{Diff}_0(S)$ from h to Id_S , and $[\text{ev}(h_t)]$ is the based homotopy class of $\text{ev}(h_t) = h_t(z), t \in [0, 1]$. To see that this is a homomorphism, suppose $h, h' \in \text{Diff}_B(S, z)$ and h_t and h'_t are paths from h and h' respectively to Id_S . Write $\gamma(t) = h_t(z)$ and $\gamma'(t) = h'_t(z)$. There is a path H_t from $h \circ h'$ to Id_S given as

$$H_t = \begin{cases} h_{2t} \circ h' & \text{for } t \in [0, 1/2] \\ h'_{2t-1} & \text{for } t \in [1/2, 1] \end{cases}$$

Then $H_t(z)$ is the path obtained by first traversing γ then γ' , while $H_0 = h \circ h'$. So, $\text{ev}_*(h \circ h') = \gamma\gamma'$, and ev_* is the required homomorphism.

Given $h \in \text{Diff}_B(S, z)$, we will write σ_h for a loop (or the homotopy class) representing $\text{ev}_*(h)$. Similarly, we will let h_σ denote the mapping class (or a representative homeomorphism) determined by $\sigma \in \pi_1(S)$. When convenient, we will simply identify $\pi_1(S)$ with a subgroup of $\text{Mod}(S, z)$.

1.2.3 Curve Complexes

A closed curve in S is **essential** if it is homotopically nontrivial in S . We will refer to a closed curve in $S - \{z\}$ simply as a closed curve in (S, z) , and will say it is **essential** if it is homotopically nontrivial *and* nonperipheral in $S - \{z\}$. Essential simple closed curves in (S, z) are isotopic if and only if they are isotopic in $S - \{z\}$.

Let $\mathcal{C}(S)$ and $\mathcal{C}(S, z)$ denote the **curve complexes of S** and (S, z) , respectively; see [Har81] and [MM99]. These are geodesic metric spaces obtained by isometrically gluing regular Euclidean simplices with all edge lengths equal to one. The following is proven in [MM99].

Theorem 1.3 (Masur-Minsky). *The spaces $\mathcal{C}(S)$ and $\mathcal{C}(S, z)$ are δ -hyperbolic for some $\delta > 0$.*

We will refer to a simplex $v \subset \mathcal{C}(S)$ or $u \subset \mathcal{C}(S, z)$ and confuse this with the isotopy class of multicurve it determines. Any simple closed curve u in (S, z) can be viewed as a curve in S which we denote $\Pi(u)$. This well-defines a “forgetful” map

$$\Pi: \mathcal{C}(S, z) \rightarrow \mathcal{C}(S)$$

which is simplicial.

Given a multicurve $v \subset \mathcal{C}(S)$, unless otherwise stated, we assume that v is realized by its geodesic representative in S . Associated to v there is an action of $\pi_1(S)$ on a tree T_v , namely, the Bass-Serre tree for the splitting of $\pi_1(S)$ determined by v . We will make use of the following theorem of [KLS06].

Theorem 1.4 (Kent-Leininger-Schleimer). *The fiber of Π over $v \in \mathcal{C}(S)$ is $\pi_1(S)$ -equivariantly homeomorphic to the tree T_v determined by v .*

1.2.4 Measured laminations and the curve complex.

The curve complex \mathcal{C} naturally injects into $\mathbb{P}\mathcal{ML}$ sending a simplex v to the simplex of measures supported on v . We denote the image subspace $\mathbb{P}\mathcal{ML}_{\mathcal{C}}$. We note that this bijection $\mathbb{P}\mathcal{ML}_{\mathcal{C}} \rightarrow \mathcal{C}$ is not continuous in either direction. We will use the same notation for a point of $\mathbb{P}\mathcal{ML}_{\mathcal{C}}$ and its image in \mathcal{C} .

In [Kla99] Klarreich proved that $\partial_{\infty}\mathcal{C} \cong \mathcal{EL}$. Therefore, if we define

$$\mathbb{P}\mathcal{ML}_{\bar{\mathcal{C}}} = \mathbb{P}\mathcal{ML}_{\mathcal{C}} \cup \mathbb{P}\mathcal{FL}$$

then there is a natural surjective map

$$\mathbb{P}\mathcal{ML}_{\bar{\mathcal{C}}} \rightarrow \bar{\mathcal{C}}$$

extending $\mathbb{P}\mathcal{ML}_{\mathcal{C}} \rightarrow \mathcal{C}$. The following is an immediate consequence of Klarreich’s work [Kla99], stated using our terminology.

Proposition 1.5 (Klarreich). *The natural map $\mathbb{P}\mathcal{ML}_{\bar{\mathcal{C}}} \rightarrow \bar{\mathcal{C}}$ is continuous at every point of $\mathbb{P}\mathcal{FL}$. Moreover, a sequence $\{v_n\}$ in \mathcal{C} converges to $|\lambda|$ if and only if every accumulation point of $\{v_n\}$ in $\mathbb{P}\mathcal{ML}$ has $|\lambda|$ as its support.*

Proof. Theorem 1.4 of [Kla99] implies that if $\{v_n\}$ converges in $\bar{\mathcal{C}}$ to $|\lambda|$, then every accumulation point of $\{v_n\}$ in $\mathbb{P}\mathcal{ML}$ has $|\lambda|$ as its support. We need only verify that if $\lambda \in \mathbb{P}\mathcal{FL}$ and every accumulation point λ' in $\mathbb{P}\mathcal{ML}$ of $\{v_n\}$ has $|\lambda| = |\lambda'|$ then $\{v_n\}$ converges to $|\lambda|$ in $\bar{\mathcal{C}}$.

To see this, let $\{X_n\} \subset \mathcal{T}$ be any sequence in the Teichmüller space \mathcal{T} for which v_n is the shortest curve in X_n , so in particular $\ell_{X_n}(v_n)$ is uniformly bounded. Since every accumulation point of $\{v_n\}$ is in $\mathbb{P}\mathcal{FL}$, it follows that X_n exits every compact set and so accumulates only on $\mathbb{P}\mathcal{ML}$ in the Thurston compactification of \mathcal{T} . Moreover, if λ' is any accumulation point of X_n in $\mathbb{P}\mathcal{ML}$, then $i(\lambda', \lambda) = 0$, and so $|\lambda'| = |\lambda|$ since λ is filling.

Now according to Theorem 1.1 of [Kla99], the map

$$sys : \mathcal{T} \rightarrow \mathcal{C}$$

sending $X \in \mathcal{T}$ to any shortest curve in X extends to

$$\overline{sys} : \mathcal{T} \cup \mathbb{P}\mathcal{FL} \rightarrow \bar{\mathcal{C}}$$

continuously at every point of $\mathbb{P}\mathcal{FL}$. It follows that

$$\lim_{n \rightarrow \infty} v_n = \lim_{n \rightarrow \infty} sys(X_n) = |\lambda|.$$

□

1.2.5 Cannon-Thurston Maps

Definition 1.6. Let X and Y be hyperbolic metric spaces and $i : Y \rightarrow X$ be a continuous map, and $Z \subset \partial_\infty Y$ a subset of the Gromov boundary. A **Z -Cannon-Thurston map** (or just a Cannon-Thurston map, by abuse of notation) is a continuous extension $\bar{i} : Y \cup Z \rightarrow \bar{X} = X \cup \partial_\infty X$ of i . That is, $\bar{i}|_Y = i$.

This definition is more general than that in [Mit98] in the sense that here we require i only to be continuous, whereas in [Mit98] it was demanded that i be an embedding, and we are not requiring \bar{i} to be defined on all of $\bar{Y} = Y \cup \partial_\infty Y$.

To prove the existence of such a Cannon-Thurston map, we shall use the following obvious criterion:

Lemma 1.7. Let X and Y be hyperbolic metric spaces and $i : Y \rightarrow X$ be a continuous map and $Z \subset \partial_\infty Y$. Then a Z -Cannon-Thurston map \bar{i} exists if for every $y \in Z$, there exists a neighborhood basis $\{B_i(y)\}_{i=1}^\infty$ of $y \in Y \cup Z$ and uniformly quasiconvex sets $Q_i(y) \subset X$ with $i(B_i(y) \cap Y) \subset Q_i(y)$ for all i and $d(x, Q_i) \rightarrow \infty$ as $i \rightarrow \infty$ from some basepoint $x \in X$. Moreover, $\bar{i}(y)$ is the unique point of intersection of the sets

$$\bigcap_i \overline{Q_i(y)} = \bigcap_i \partial_\infty Q_i(y) = \{\bar{i}(y)\}.$$

2 Point position

We now describe in more detail the map

$$\Phi : \mathcal{C}(S) \times \mathbb{H} \rightarrow \mathcal{C}(S, z)$$

as promised in the introduction, and explain how this can be extended continuously to $\overline{\mathcal{C}}(S) \times \mathbb{H}$.

2.1 A bundle over \mathbb{H} .

The bundle determining the Birman exact sequence has a subbundle obtained by restricting ev to $\text{Diff}_0(S)$:

$$\text{Diff}_B(S, z) \longrightarrow \text{Diff}_0(S) \xrightarrow{\text{ev}} S.$$

As noted before, Earle and Eells proved that $\text{Diff}_0(S)$ is contractible, and hence there is a unique lift

$$\tilde{\text{ev}} : \text{Diff}_0(S) \rightarrow \mathbb{H}$$

with the property that $\tilde{\text{ev}}(\text{Id}_S) = \tilde{z}$.

The map $\tilde{\text{ev}}$ can also be described as follows. Any diffeomorphism $S \rightarrow S$ has a lift $\mathbb{H} \rightarrow \mathbb{H}$, and the contractibility of $\text{Diff}_0(S)$ allows us to coherently lift diffeomorphisms to obtain an injective homomorphism $\text{Diff}_0(S) \rightarrow \text{Diff}(\mathbb{H})$. Then $\tilde{\text{ev}}$ is the composition of this homomorphism with the evaluation map $\text{Diff}(\mathbb{H}) \rightarrow \mathbb{H}$ determined by \tilde{z} .

Since p is a covering map, $\tilde{\text{ev}}$ is also a fibration. Appealing to the long exact sequence of homotopy groups again, we see that the fiber over \tilde{z} is precisely $\text{Diff}_0(S, z)$. We record this in the following diagram

$$\begin{array}{ccccc}
 & & & & \mathbb{H} \\
 & & & \nearrow \tilde{\text{ev}} & \downarrow p \\
 \text{Diff}_B(S, z) & \longrightarrow & \text{Diff}_0(S) & \xrightarrow{\text{ev}} & S \\
 \uparrow & \nearrow & & & \\
 \text{Diff}_0(S, z) & & & &
 \end{array}$$

The group $\text{Diff}_B(S, z)$ acts on $\text{Diff}_0(S)$ on the left by

$$h \cdot f = f \circ h^{-1}$$

for $h \in \text{Diff}_B(S, z)$ and $f \in \text{Diff}_0(S)$. Also recall from Section 1.2.2 that $\pi_1(S) \cong \pi_0(\text{Diff}_B(S, z))$ with this isomorphism induced by a homomorphism

$$\text{ev}_* : \text{Diff}_B(S, z) \rightarrow \pi_1(S).$$

Lemma 2.1.

$$\tilde{e}v : \text{Diff}_0(S) \rightarrow \mathbb{H}$$

is equivariant with respect to ev_* .

Proof. We need to prove

$$ev_*(h)(\tilde{e}v(f)) = \tilde{e}v(f \circ h^{-1})$$

for all $f \in \text{Diff}_0(S)$ and $h \in \text{Diff}_B(S, z)$. Observe that since $h(z) = z$ for every $h \in \text{Diff}_B(S, z)$, $ev(f) = ev(f \circ h^{-1})$ for every $f \in \text{Diff}_0(S)$. Therefore, since $\tilde{e}v$ is a lift of ev we have

$$p(\tilde{e}v(f)) = ev(f) = ev(f \circ h^{-1}) = p(\tilde{e}v(f \circ h^{-1}))$$

and hence $\tilde{e}v(f)$ differs from $\tilde{e}v(f \circ h^{-1})$ by a covering transformation:

$$\tilde{e}v(f \circ h^{-1}) = \sigma(\tilde{e}v(f))$$

for some $\sigma \in \pi_1(S)$.

The covering transformation σ depends on f and h , however if $h_t, t \in [0, 1]$, is a path in $\text{Diff}_B(S, z)$ then $\tilde{e}v(f \circ h_t^{-1})$ is constant in t as can be seen from the above description of $\tilde{e}v$ as the evaluation map on the lifted diffeomorphism group. It follows that σ depends only on f and the component of $\text{Diff}_B(S, z)$ containing h . In fact, continuity of $\tilde{e}v$ and connectivity of $\text{Diff}_0(S, z)$ implies that σ actually only depends on the component of $\text{Diff}_B(S, z)$ containing h , and not on f at all.

We have

$$\sigma(\tilde{z}) = \sigma(\tilde{e}v(\text{Id}_S)) = \tilde{e}v(\text{Id}_S \circ h^{-1}) = \tilde{e}v(h^{-1}).$$

So if $h_t, t \in [0, 1]$, is a path in $\text{Diff}_0(S)$ from h to Id_S , then since $ev_*(h) = \sigma_h$ where σ_h is represented by the loop $h_t(z), t \in [0, 1]$, it follows that σ_h^{-1} is represented by the loop $h_t^{-1}(z), t \in [0, 1]$.

Now observe that $\tilde{e}v(h_t^{-1}), t \in [0, 1]$, is a lift of the loop $h_t^{-1}(z), t \in [0, 1]$, to a path from $\sigma(\tilde{z})$ to \tilde{z} . Therefore, σ_h^{-1} is σ^{-1} , and hence $\sigma = \sigma_h = ev_*(h)$. \square

2.2 An explicit construction of Φ .

We will define first a map

$$\tilde{\Phi} : \mathcal{C}(S) \times \text{Diff}_0(S) \rightarrow \mathcal{C}(S, z)$$

and show that this descends to a map $\Phi : \mathcal{C}(S) \times \mathbb{H} \rightarrow \mathcal{C}(S, z)$ by composing with $\tilde{e}v$ in the second factor.

Recall that for every $v \in \mathcal{C}^0(S)$, we have realized v by its geodesic representative. We would like to simply define

$$\tilde{\Phi}(v, f) = f^{-1}(v).$$

However, this is not a curve in (S, z) when $f(z)$ lies on the geodesic v . The map we define in the end will agree with this when $f(z)$ is not too close to v , and it is helpful to keep this mind when trying to make sense of the actual definition of $\tilde{\Phi}$.

To carry out the construction of $\tilde{\Phi}$, we first let $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)} \subset \mathbb{R}_+$ be such that

$$N_{\epsilon(v)}(v) \cap N_{\epsilon(v')}(v') = \emptyset \quad \Leftrightarrow \quad i(v, v') = 0 \text{ and } v \neq v'$$

and $N_{\epsilon(v)} \cong S^1 \times [0, 1]$ for all v . We will impose further restrictions on $\{\epsilon(v)\}$ later. We will write $N_{\epsilon(v)}^\circ(v)$ for the interior of $N_{\epsilon(v)}(v)$ and $v_{\epsilon(v)}^\pm$ for the two components of $\partial N_{\epsilon(v)}(v)$.

Given a simplex $v = \{v_0, \dots, v_k\} \subset \mathcal{C}(S)$, we consider the barycentric coordinates for points in v :

$$\left\{ \sum_{j=0}^k s_j v_j \mid \sum_{j=0}^k s_j = 1 \text{ and } s_j \geq 0, \forall j = 0, \dots, k \right\}.$$

To define our map

$$\tilde{\Phi} : \mathcal{C}(S) \times \text{Diff}_0(S) \rightarrow \mathcal{C}(S, z)$$

we first explain how to define it for (v, f) with v a vertex of $\mathcal{C}(S)$. If $f(z) \notin N_{\epsilon(v)}^\circ(v)$, then we set

$$\tilde{\Phi}(v, f) = f^{-1}(v)$$

as suggested above.

If $f(z) \in N_{\epsilon(v)}^\circ(v)$, then $f^{-1}(v_{\epsilon(v)}^+)$ and $f^{-1}(v_{\epsilon(v)}^-)$ are nonisotopic curves in (S, z) . We will define $\tilde{\Phi}(v, f)$ to be a point on the edge between these two vertices of $\mathcal{C}(S, z)$, depending on the distance from $f(z)$ to these boundary components. Specifically, we set

$$\tilde{\Phi}(v, f) = \frac{1}{2\epsilon(v)} \left(d(f(z), v_{\epsilon(v)}^+) f^{-1}(v_{\epsilon(v)}^+) + d(f(z), v_{\epsilon(v)}^-) f^{-1}(v_{\epsilon(v)}^-) \right)$$

in barycentric coordinates on the edge $[f^{-1}(v_{\epsilon(v)}^+), f^{-1}(v_{\epsilon(v)}^-)]$.

In general, for a point $(x, f) \in \mathcal{C}(S) \times \text{Diff}_0(S)$ with $x = \sum_j s_j v_j \subset v = \{v_0, \dots, v_k\}$ we define $\tilde{\Phi}(x, f)$ as follows. First, we shrink all neighborhoods $N_{\epsilon(v_j)}(v_j)$ according to the associated barycentric coordinate s_j to $N_{s_j \epsilon(v_j)}(v_j)$. We let $v_{j, s_j \epsilon(v_j)}^\pm$ denote the two boundary components of $N_{s_j \epsilon(v_j)}(v_j)$.

As before, if $f(z) \notin \bigcup_{j=0}^k N_{s_j \epsilon(v_j)}^\circ(v_j)$, then define

$$\tilde{\Phi}(x, f) = \sum_j s_j f^{-1}(v_j).$$

Otherwise, $f(z) \in N_{s_i \epsilon(v_i)}^\circ(v_i)$ for exactly one $i \in \{0, \dots, k\}$, and we define $\tilde{\Phi}(x, f)$ to be

$$\sum_{j \neq i} s_j f^{-1}(v_j) + \frac{s_i}{2\epsilon(v_i)} \left(\begin{array}{c} d(f(z), v_{i, s_i \epsilon(v_i)}^+) f^{-1}(v_{i, s_i \epsilon(v_i)}^+) \\ + d(f(z), v_{i, s_i \epsilon(v_i)}^-) f^{-1}(v_{i, s_i \epsilon(v_i)}^-) \end{array} \right).$$

The group $\text{Diff}_B(S, z)$ acts on $\mathcal{C}(S) \times \text{Diff}_0(S)$, trivially in the first factor and as described in Section 2.1 in the second factor. Of course, since $\text{Diff}_B(S, z) < \text{Diff}^+(S, z)$ projects into $\text{Mod}(S, z)$ it also acts on $\mathcal{C}(S, z)$. The map $\tilde{\Phi}$ is equivariant: given $h \in \text{Diff}_B(S, z)$, $f \in \text{Diff}_0(S)$ and v a vertex in $\mathcal{C}(S)$, provided $f(z) \notin N_{\epsilon(v)}^\circ(v)$ we have

$$\tilde{\Phi}(h \cdot (v, f)) = \tilde{\Phi}(v, f \circ h^{-1}) = (f \circ h^{-1})^{-1}(v) = h \circ f^{-1}(v) = h \cdot (f^{-1}(v)).$$

The general situation is similar, but notationally more complicated.

Proposition 2.2. *The map $\tilde{\Phi}$ descends to a map Φ making the following diagram commute*

$$\begin{array}{ccc} \mathcal{C}(S) \times \text{Diff}_0(S) & \xrightarrow{\tilde{\Phi}} & \mathcal{C}(S, z) \\ \text{Id}_{\mathcal{C}(S)} \times \tilde{e}v \downarrow & & \nearrow \\ \mathcal{C}(S) \times \mathbb{H} & \xrightarrow{\Phi} & \mathcal{C}(S, z) \end{array}$$

Moreover, Φ is equivariant with respect to the action of $\pi_1(S)$.

Here the action on $\pi_1(S)$ on $\mathcal{C}(S) \times \mathbb{H}$ is trivial on the first factor and the covering group action on the second.

Proof. We suppose that $\tilde{e}v(f_0) = \tilde{e}v(f_1)$ and must show $\tilde{\Phi}(x, f_0) = \tilde{\Phi}(x, f_1)$.

From the discussion in Section 2.1 it follows that $f_0 = f_1 \circ h$ for some $h \in \text{Diff}_0(S, z)$. We suppose that α is a simple closed curve on S and $f_0(z) \notin \alpha$. Then $f_1(z) = f_1(h(z)) = f_0(z) \notin \alpha$ and

$$d(f_0(z), \alpha) = d(f_1(h(z)), \alpha) = d(f_1(z), \alpha).$$

Moreover, $f_0^{-1}(\alpha) = h^{-1}(f_1^{-1}(\alpha))$ and since h^{-1} is isotopic to the identity in (S, z) , it follows that $f_0^{-1}(\alpha)$ and $f_1^{-1}(\alpha)$ are isotopic in (S, z) .

Now because $\tilde{\Phi}(x, f)$ is defined in terms of the isotopy classes of curves of the form $f^{-1}(\alpha)$ and numbers of the form $d(f(z), \alpha)$ it follows that

$$\tilde{\Phi}(x, f_0) = \tilde{\Phi}(x, f_1)$$

and so $\tilde{\Phi}$ descends to $\mathcal{C}(S) \times \mathbb{H}$ as required.

Lemma 2.1 implies that $\text{Id}_{\mathcal{C}(S)} \times \tilde{e}v$ is equivariant with respect to ev_* . Thus, since $\tilde{\Phi}$ is equivariant, so is Φ . \square

Suppose that $x \in \mathcal{C}(S)$ and $\{v_0, \dots, v_k\} = v \subset \mathcal{C}(S)$ is the simplex containing x in its interior and write

$$x = \sum_{i=1}^k s_i v_i$$

in terms of barycentric coordinates.

We note that the neighborhoods $N_{s_i \epsilon(v_i)}^\circ(v_i)$ determine a map to the Bass–Serre tree T_v associated to v as follows. We collapse each component U of the preimage $p^{-1}(N_{s_i \epsilon(v_i)}^\circ(v_i))$ onto an interval, say $(0, 1)$, by the projection defined as the distance to the component of $p^{-1}(v_{i, s_i \epsilon(v_i)}^-)$ meeting U , multiplied by $1/(2s_i \epsilon(v_i))$. If we further collapse each component of the complement of

$$p^{-1}(N_{s_0 \epsilon(v_0)}(v_0) \cup \dots \cup N_{s_k \epsilon(v_k)}(v_k))$$

to a point, the quotient space is precisely T_v .

The map Φ restricted to $\{x\} \times \mathbb{H} \cong \mathbb{H}$, which we denote Φ^x , is constant on the fibers of the projection to T_v . That is, $\Phi^x : \{x\} \times \mathbb{H} \rightarrow \Pi^{-1}(x) \subset \mathcal{C}(S, z)$ factors through the projection to T_v

$$\begin{array}{ccc} \{x\} \times \mathbb{H} & \xrightarrow{\Phi^x} & \Pi^{-1}(x) \\ & \searrow & \nearrow \\ & T_v & \end{array}$$

Moreover, the equivariance of Φ implies that

$$T_v \rightarrow \Pi^{-1}(x)$$

is equivariant. According to [KLS06], the edge and vertex stabilizers in the domain and range agree. This map is thus the homeomorphism given by Theorem 1.4, and Φ^x agrees with the definition given in the introduction.

2.3 A further description of $\mathcal{C}(S, z)$

We pause here to give a combinatorial description of $\mathcal{C}(S, z)$ which will be useful later, but is also of interest in its own right. Given any simplex $v \subset \mathcal{C}(S)$, the preimage of the interior of v admits a $\pi_1(S)$ –equivariant homeomorphism

$$\Pi^{-1}(\text{int}(v)) \cong \text{int}(v) \times T_v$$

as can be seen from Theorem 1.4. As is well known, the edges of T_v can be labeled by the vertices of v . Now, if $\phi : v' \rightarrow v$ is the inclusion of a face, then there is a $\pi_1(S)$ –equivariant quotient map $\phi^* : T_v \rightarrow T_{v'}$ obtained by collapsing all the edges of T_v labeled by vertices *not* in $\phi(v')$ (compare [GL07], for example). This provides a description of $\Pi^{-1}(v)$, the preimage of the closed simplex as a quotient

$$\left(\bigsqcup_{\phi: v' \rightarrow v} v' \times T_{v'} \right) / \sim$$

where the disjoint union is taken over all faces $\phi : v' \rightarrow v$ and the equivalence relation \sim is defined by

$$(\varphi(x), t) \sim (x, \varphi^*(t))$$

for every inclusion of faces $\varphi : v'' \rightarrow v'$ and every $x \in v, t \in T_{v'}$. Said differently, we take the product $v \times T_v$ and for every face $\phi : v' \rightarrow v$, we glue $v \times T_v$ to $v' \times T_{v'}$ along $\phi(v') \times T_v$ by $\phi^{-1} \times \phi^*$.

We can do this for all simplices, then glue them all together, providing the following useful description of $\mathcal{C}(S, z)$.

Theorem 2.3. *The curve complex $\mathcal{C}(S, z)$ is $\pi_1(S)$ -equivariantly homeomorphic to*

$$\left(\bigsqcup_{v \subset \mathcal{C}(S)} v \times T_v \right) / \sim$$

where the disjoint union is taken over all simplices $v \subset \mathcal{C}(S)$, and the equivalence relation is generated by

$$(\phi(x), t) \sim (x, \phi^*(t))$$

for all inclusions of faces $\phi : v' \rightarrow v$ all $x \in v'$ and all $t \in T_{v'}$.

2.4 Extending to measured laminations

The purpose of this section is to modify the above construction of Φ to build a map

$$\Psi : \mathcal{ML}(S) \times \mathbb{H} \rightarrow \mathcal{ML}(S, z),$$

and to prove that this is continuous at every point of $\mathcal{FL}(S) \times \mathbb{H}$; see Corollary 2.8. We do this by defining a map on $\mathcal{ML}(S) \times \text{Diff}_0(S)$, and checking that it descends to $\mathcal{ML}(S) \times \mathbb{H}$.

Before we can begin, we must specify a particular realization for each element of $\mathcal{ML}(S)$ as a measured lamination. We begin by realizing all elements as measured geodesic laminations (recall we denote these with a hat, $\hat{\lambda}$), then replace all simple closed geodesic components of the support with appropriately chosen annuli. We now explain this more precisely and set some notation.

Given a measured geodesic lamination $\hat{\lambda}$, $|\hat{\lambda}|$ can be decomposed into a finite union of pairwise disjoint minimal sublaminations; see [CB87]. Write

$$\hat{\lambda} = \text{Cur}(\hat{\lambda}) + \text{Min}(\hat{\lambda}),$$

where $|\text{Cur}(\hat{\lambda})|$ and $|\text{Min}(\hat{\lambda})|$ are disjoint with $|\text{Cur}(\hat{\lambda})|$ consisting precisely of the union of simple closed geodesics in $|\hat{\lambda}|$. We construct a measured lamination λ measure equivalent to $\hat{\lambda}$ as

$$\lambda = \text{Ann}(\lambda) + \text{Min}(\lambda),$$

where $\text{Min}(\lambda) = \text{Min}(\hat{\lambda})$ and $\text{Ann}(\lambda)$ is defined as follows.

The sublamination $\text{Cur}(\hat{\lambda})$ can be further decomposed as $\text{Cur}(\hat{\lambda}) = \sum_j t_j v_j$, where $t_j v_j$ means t_j times the transverse counting measure on the simple closed geodesic component v_j of $|\text{Cur}(\hat{\lambda})|$. For each component v_j of $|\text{Cur}(\hat{\lambda})|$ with $t_j \leq 1$, there is a component of $|\text{Ann}(\lambda)|$ which is the annulus $N_{t_j \epsilon(v_j)}(v_j)$ given the foliation by equidistant curves to v_j . This is assigned the transverse measure which is $1/2\epsilon(v_j)$ times the distance between leaves. For each component v_j with $t_j > 1$, there is a corresponding component of $|\text{Ann}(\lambda)|$ which is the annulus $N_{\epsilon(v_j)}(v_j)$ foliated again by equidistant curves to v . This is assigned the transverse measure $t_j/2\epsilon(v_j)$ times the distance between leaves; see Figure 3 for a cartoon depiction of $\hat{\lambda}$ and λ . For future use, we define $T(\hat{\lambda})$ and $T(\lambda)$ by

$$T(\hat{\lambda}) = T(\lambda) = \max_j t_j.$$

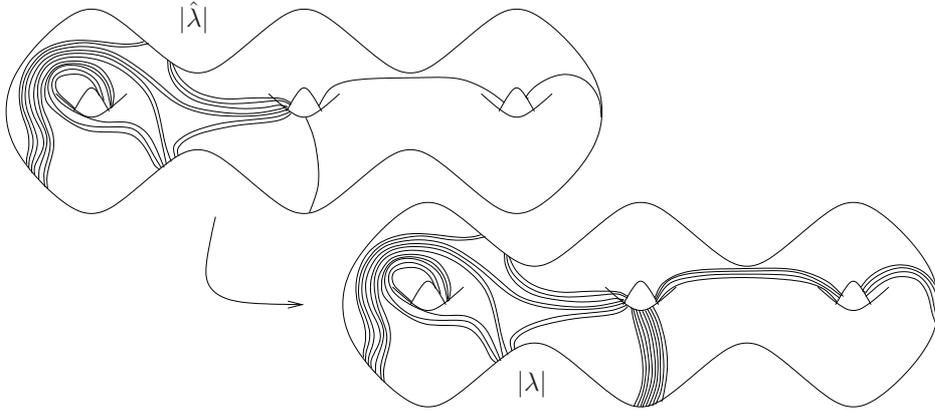


Figure 3: Removing simple closed geodesics and inserting foliated annuli.

Whenever we refer to an element λ of $\mathcal{ML}(S)$ in what follows, we will assume it is realized by such a measured lamination. Of course $|\hat{\lambda}| \subset |\lambda|$, meaning that as subsets of S , $|\hat{\lambda}|$ is a subset of $|\lambda|$, and that each leaf of $|\hat{\lambda}|$ is a leaf of $|\lambda|$. The difference between the total variations assigned an arc by λ and $\hat{\lambda}$ is estimated by the following.

Lemma 2.4. *If a is any arc transverse to $|\lambda|$, then it is also transverse to $|\hat{\lambda}|$ and we have*

$$|\lambda(a) - \hat{\lambda}(a)| \leq T(\hat{\lambda})$$

Proof. The transversality statement is an immediate consequence of $|\hat{\lambda}| \subset |\lambda|$.

Since $\text{Min}(\lambda) = \text{Min}(\hat{\lambda})$, we see that

$$|\lambda(a) - \hat{\lambda}(a)| = |\text{Ann}(\lambda)(a) - \text{Cur}(\hat{\lambda})(a)|$$

The intersection of $|\text{Ann}(\lambda)| \cap a$ is a union of subarcs of a , each containing an intersection point of $|\text{Cur}(\hat{\lambda})| \cap a$, with the possible exception of those arcs which meet the endpoints of a . If $a_0 \subset a$ is one of the subarcs which meets the boundary, then we have $|\text{Ann}(\lambda)(a_0) - \text{Cur}(\hat{\lambda})(a_0)| \leq T(\hat{\lambda})/2$. Since there are at most 2 such arcs, the desired inequality follows. \square

The following is also useful.

Lemma 2.5. *Suppose $\lambda_n \rightarrow \lambda$ in $\mathcal{ML}(S)$ and $\text{Ann}(\lambda) = \emptyset$. Further suppose that $|\lambda_n|$ converges in the Hausdorff topology on closed subsets of S to a set \mathcal{L} . Then \mathcal{L} is a geodesic lamination containing $|\lambda|$.*

Proof. If $|\lambda_n| = |\hat{\lambda}_n|$ is a geodesic lamination for all n , then the fact that \mathcal{L} is a geodesic lamination is well known; see [CB87].

Since $\lambda_n \rightarrow \lambda$ and $\text{Ann}(\lambda) = \emptyset$, it follows that $T(\lambda_n) \rightarrow 0$ as $n \rightarrow \infty$. To see this, suppose this were not the case. Then, up to subsequence, some component $v_n \subset \text{Cur}(\hat{\lambda}_n)$ would have weight $t_n > t > 0$ for some $t > 0$. Convergence in \mathcal{ML} implies that the length of v_n must therefore be bounded above for all n , and so after passing to a further subsequence, we would have all v_n equal to some fixed simple closed geodesic v . But because $t_n > t > 0$, the limit λ would have $v \subset \text{Cur}(\hat{\lambda})$, which is a contradiction.

Suppose now that $x = \lim_{n \rightarrow \infty} x_n$ with $x_n \in |\lambda_n|$. Then there exists $\hat{x}_n \in |\hat{\lambda}_n|$ with $\lim_{n \rightarrow \infty} d(x_n, \hat{x}_n) = 0$, so $x = \lim_{n \rightarrow \infty} \hat{x}_n$. It follows that the Hausdorff limit of $\{|\lambda_n|\}$ is the same as that of $\{|\hat{\lambda}_n|\}$ and hence \mathcal{L} is a geodesic lamination. \square

Now, given any $(\lambda, f) \in \mathcal{ML}(S) \times \text{Diff}_0(S)$, we would like to simply define

$$\tilde{\Psi}(\lambda, f) = f^{-1}(\lambda).$$

As before, this does not make sense when $f(z)$ lies on the supporting lamination $|\lambda|$. This is remedied by first splitting open the lamination along the leaf which $f(z)$ meets to produce a new measured lamination λ' representing the measure class λ (there is no ambiguity about how the measure is split since λ has no atoms). The new lamination $|\lambda'|$ has either a bigon or annular region containing $f(z)$ and $f^{-1}(\lambda)$ is defined to be $f^{-1}(\lambda')$. The support $|f^{-1}(\lambda')|$ is contained in $f^{-1}(|\lambda'|)$, and this containment can be proper since $f^{-1}(|\lambda|)$ may have an isolated leaf. Note that this happens precisely when $f(z)$ lies on a boundary leaf of $|\lambda|$.

Train tracks provide a more concrete description of $\tilde{\Psi}(f, \lambda)$ which will be useful in proving continuity results. Let \mathcal{L} be any geodesic lamination on S and $\epsilon > 0$ sufficiently small so that the quotient of $N_\epsilon(\mathcal{L})$ by collapsing the ties defines a train track τ as in Section 1.2.1. Suppose that λ is a measured lamination on S for which $|\lambda|$ is contained in $N_\epsilon(\mathcal{L})$ and is transverse to the ties. If $f(z) \notin N_\epsilon(\mathcal{L})$, then $\tilde{\Psi}(f, \lambda)$ is the lamination on (S, z) determined by the weighted train track $f^{-1}(\tau(\lambda))$ as described in Section 1.2.1.

If $f(z) \in N_\epsilon(\mathcal{L})$, by small perturbation of ϵ , we may assume that $f(z)$ does not lie on a boundary-tie of any rectangle and that each switch of τ is trivalent. Then either $f(z)$ is outside $N_\epsilon(\mathcal{L})$ and we are in the situation above, or else $f(z)$ is in the interior of some rectangle R . Furthermore, τ can be realized in $N_\epsilon(\mathcal{L})$ with the branch β_R associated to R contained in R .

We modify the train track τ at the branch β_R as follows. Remove an arc in the interior of β_R leaving two subarcs β_R^ℓ and β_R^r of β_R and insert a two branches β_R^u and β_R^d creating a bigon containing $f(z)$; see Figure 4. The result, denoted τ' is not a train track on S , but is a train track on $(S, f(z))$.

If $f_t \in \text{Diff}_0(S)$ is an isotopy with $f = f_0$ and $f_t(z) \subset R$ for every $t \in [0, 1]$, and τ'_t is constructed for f_t as τ is constructed for f (so $\tau' = \tau'_0$), then $f_t^{-1}(\tau'_t)$ is (isotopic to) $f^{-1}(\tau')$ for all t .

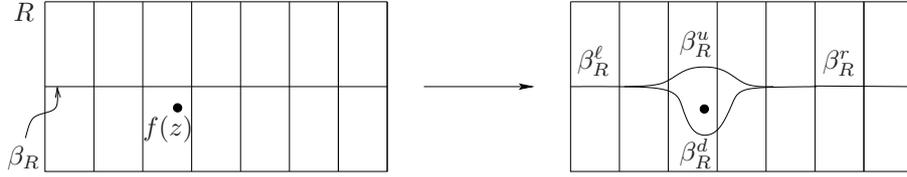


Figure 4: Modifying τ to τ' .

The measured lamination λ makes τ' into a weighted train track $\tau'(\lambda)$ on $(S, f(z))$ as follows. For the branches of τ' that are the same as those of τ , the weights are defined as before. To define the weights on the new branches, we first consider the tie $a \subset R$ that contains $f(z)$, and write it as the union of subarcs $a = a^u \cup a^d$ with $a^u \cap a^d = \{f(z)\}$. We define the weights on the branches β_R^u and β_R^d of the bigon to be $\lambda(a^u)$ and $\lambda(a^d)$, respectively, while the weights on the branches β_R^ℓ and β_R^r are both $\lambda(a) = \lambda(a^u) + \lambda(a^d)$; see Figure 5. The lamination $f^{-1}(\lambda)$ is the lamination determined by the weighted train track $f^{-1}(\tau'(\lambda))$.

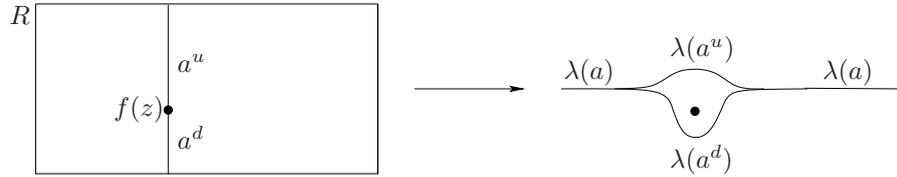


Figure 5: Weights on τ' determined by λ and $f(z)$.

The proof of the following is essentially the same as that of Proposition 2.2

and we omit it.

Proposition 2.6. $\tilde{\Psi}$ descends to a map Ψ :

$$\begin{array}{ccc} \mathcal{ML}(S) \times \text{Diff}_0(S) & \xrightarrow{\tilde{\Psi}} & \mathcal{ML}(S, z) \\ \text{Id}_{\mathcal{ML}(S)} \times \tilde{ev} \downarrow & & \\ \mathcal{ML}(S) \times \mathbb{H} & \xrightarrow{\Psi} & \end{array}$$

□

Because of the particular way we have realized our laminations, neither the map $\tilde{\Psi}$ nor the map Ψ need be continuous at measured laminations with nontrivial annular component. However, this is the only place where continuity breaks down.

Proposition 2.7. *The map $\tilde{\Psi}$ is continuous at every (λ, f) with $\text{Ann}(\lambda) = \emptyset$.*

Proof. It suffices to show that for any sequence $\{(\lambda_n, f_n)\}$ in $\mathcal{ML}(S) \times \text{Diff}_0(S)$ converging to (λ, f) , there exists a subsequence (which we also denote by $\{(\lambda_n, f_n)\}$ for convenience) with $\{\tilde{\Psi}(\lambda_n, f_n)\}$ converging to $\tilde{\Psi}(\lambda, f)$.

We begin by passing to a subsequence for which the supports $\{|\lambda_n|\}$ converge in the Hausdorff topology to a closed set \mathcal{L} . It follows from Lemma 2.5, that \mathcal{L} is a geodesic lamination containing $|\lambda|$.

Case 1. Suppose $f(z) \notin \mathcal{L}$.

In this case, there is an $\epsilon > 0$ so that the ϵ -neighborhoods of $f(z)$ and \mathcal{L} are disjoint. Since $f_n \rightarrow f$ as $n \rightarrow \infty$, there exists $N > 0$ so that for all $n \geq N$, $f_n(z) \in N_\epsilon(f(z))$, and moreover, f_n is isotopic to f through an isotopy f_t such that $f_t(z) \in N_\epsilon(f(z))$ for all t . Taking N even larger if necessary, we may assume that for $n \geq N$, $\lambda_n \subset N_\epsilon(\mathcal{L})$. Therefore, for all $n \geq N$, λ and λ_n determine weighted train tracks $\tau(\lambda)$ and $\tau(\lambda_n)$, respectively. Since $\lambda_n \rightarrow \lambda$, it follows that $\tau(\lambda_n) \rightarrow \tau(\lambda)$ as $n \rightarrow \infty$.

Since f_n is isotopic to the f keeping the image of z in $N_\epsilon(f(z))$, it follows that $f^{-1}(\tau) = f_n^{-1}(\tau)$, up to isotopy. Therefore, $f^{-1}(\tau(\lambda_n))$ and $f_n^{-1}(\tau(\lambda_n))$ are isotopic and so we have convergence of weights $f^{-1}(\tau(\lambda_n)) \rightarrow f^{-1}(\tau(\lambda))$ which implies the associated measured laminations converge

$$\tilde{\Psi}(\lambda_n, f_n) \rightarrow \tilde{\Psi}(\lambda, f)$$

as required. This completes the proof for Case 1.

Case 2. Suppose that $f(z) \in \mathcal{L}$.

We let $\epsilon > 0$ be chosen sufficiently small so that the quotient of $N_\epsilon(\mathcal{L})$ by collapsing ties is a train track τ , and so that $f(z)$ lies in the interior of some rectangle R of $N_\epsilon(\mathcal{L})$ and so that τ is trivalent.

Let $N > 0$ so that for all $n \geq N$, $f_n(z)$ also lies in the interior of R and f is isotopic to f_n by an isotopy f_t with $f_t(z)$ contained in R for all t . For each $n \geq N$, the train track τ associated to $N_\epsilon(\mathcal{L})$ and the points $f_n(z)$ and $f(z)$ define tracks τ'_n and τ' , respectively, with bigons as described above. Moreover, $f_n^{-1}(\tau'_n)$ and $f^{-1}(\tau')$ are isotopic, and we simply identify the two as the same train track on (S, z) .

Since λ_n is converging to λ as $n \rightarrow \infty$, it follows that the weighted train tracks converge $\tau(\lambda_n) \rightarrow \tau(\lambda)$. Therefore, to prove that the weighted train tracks $f^{-1}(\tau'_n(\lambda_n)) = f_n^{-1}(\tau'_n(\lambda_n))$ converges to $f^{-1}(\tau'(\lambda))$, it suffices to prove that the weights assigned to $f^{-1}(\beta_R^u)$ and $f^{-1}(\beta_R^d)$ by λ_n converge to the weights assigned to these branches by λ . This is because the weights on the remaining branches agree with weights on the corresponding branches of τ , where we already know convergence. From this it will follow that $\tilde{\Psi}(\lambda_n, f_n) \rightarrow \tilde{\Psi}(\lambda, f)$.

Note that the weights on β_R determined by λ_n converge to the weight defined by λ . So, since the sum of the weights on $f^{-1}(\beta_R^u)$ and $f^{-1}(\beta_R^d)$ is precisely the weight on β_R , it suffices to prove convergence for the weights on one of these, say, $f^{-1}(\beta_R^u)$.

To define the required weights, recall that we first choose a tie $a_n \subset R$ with $f_n(z) \in a_n$, and write a_n as a union of subarcs $a_n = a_n^u \cup a_n^d$ with $a_n^u \cap a_n^d = \{f_n(z)\}$ (and similarly we have a tie $a \subset R$ with $a = a^u \cup a^d$ and $a^u \cap a^d = \{f(z)\}$). Then the weights on $f^{-1}(\beta_R^u)$ determined by λ_n and λ are given by

$$w_n^u = \lambda_n(a_n^u) \quad \text{and} \quad w^u = \lambda(a^u),$$

respectively.

Therefore, we must verify that $\lambda_n(a_n^u) \rightarrow \lambda(a^u)$. However, since $T(\lambda_n) \rightarrow 0$ as $k \rightarrow \infty$, Lemma 2.4 implies that it suffices to prove $\hat{\lambda}_n(a_n^u) \rightarrow \lambda(a^u)$.

Fix any $\delta > 0$. Since $\text{Cur}(\hat{\lambda}) = \emptyset$, $\hat{\lambda}|_a$ has no atoms, and so we can find subarcs a_-^u and a_+^u of a with

$$a_-^u \subsetneq a^u \subsetneq a_+^u \subset a$$

so that

$$\hat{\lambda}(a_-^u) \leq \hat{\lambda}(a^u) \leq \hat{\lambda}(a_+^u)$$

with $\hat{\lambda}(a_+^u) - \hat{\lambda}(a_-^u) < \delta$.

Since $\hat{\lambda}_n \rightarrow \hat{\lambda}$, it follows that we also have

$$\lim_{n \rightarrow \infty} \hat{\lambda}_n(a_+^u) = \hat{\lambda}(a_+^u)$$

and

$$\lim_{n \rightarrow \infty} \hat{\lambda}_n(a_-^u) = \hat{\lambda}(a_-^u).$$

Furthermore, since $a_n \rightarrow a$ and $a_n^u \rightarrow a^u$ in the C^1 -topology, combining this with the previous two equations, we see that

$$\limsup_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) \leq \lim_{n \rightarrow \infty} \hat{\lambda}_n(a_+^u) = \hat{\lambda}(a_+^u)$$

and

$$\liminf_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) \geq \lim_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) = \hat{\lambda}(a^u).$$

Since $\liminf \hat{\lambda}_n(a_n^u) \leq \limsup \hat{\lambda}_n(a_n^u)$, combining all of the above, we see that

$$|\limsup_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) - \hat{\lambda}(a^u)| + |\liminf_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) - \hat{\lambda}(a^u)| < \delta.$$

As δ was arbitrary, it follows that

$$\lim_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) = \limsup_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) = \liminf_{n \rightarrow \infty} \hat{\lambda}_n(a_n^u) = \hat{\lambda}(a^u)$$

and this completes the proof of Case 2. Since Cases 1 and 2 exhaust all possibilities, this also completes the proof of the proposition. \square

Corollary 2.8. *The maps Ψ is continuous at every (λ, x) with $\text{Ann}(\lambda) = \emptyset$. In particular, Ψ is continuous on $\mathcal{FL}(S) \times \mathbb{H}$.*

Proof. The map $\tilde{e}v$ is a quotient map. \square

2.5 Φ and Ψ

We let Ψ_c denote the restriction of Ψ to $\mathbb{PML}_c(S) \times \mathbb{H}$ which maps to $\mathbb{PML}_c(S, z)$.

Lemma 2.9. *The following diagram commutes*

$$\begin{array}{ccc} \mathbb{PML}_c(S) \times \mathbb{H} & \xrightarrow{\Psi_c} & \mathbb{PML}_c(S, z) \\ \downarrow & & \downarrow \\ \mathcal{C}(S) \times \mathbb{H} & \xrightarrow{\Phi} & \mathcal{C}(S, z) \end{array}$$

The vertical arrows here are the natural maps.

Proof. On $\mathbb{PML}_c(S)$, Ψ was defined in the same way as Φ . \square

If we let $\Psi_{\bar{c}}$ be the restriction of the map Ψ to $\mathbb{PML}_{\bar{c}}(S) \times \mathbb{H}$, then we have

Proposition 2.10. *There is a continuous extension $\hat{\Phi} : \bar{\mathcal{C}}(S) \times \mathbb{H} \rightarrow \bar{\mathcal{C}}(S, z)$ which fits into a commutative diagram*

$$\begin{array}{ccc} \mathbb{PML}_{\bar{c}}(S) \times \mathbb{H} & \xrightarrow{\Psi_{\bar{c}}} & \mathbb{PML}_{\bar{c}}(S, z) \\ \downarrow & & \downarrow \\ \bar{\mathcal{C}}(S) \times \mathbb{H} & \xrightarrow{\hat{\Phi}} & \bar{\mathcal{C}}(S, z) \end{array}$$

Proof. From the construction of Ψ and the definition of \mathcal{FL} , we see that $\Psi(\mathcal{FL}(S)) \subset \mathcal{FL}(S, z)$, and that $|\Psi(\lambda, x)| = |\Psi(\lambda', x)|$ for all $\lambda, \lambda' \in \mathbb{PML}(S)$, whenever $|\lambda| = |\lambda'|$. Since $\partial_\infty \mathcal{C}(S) = \mathcal{EL}(S)$ from Klarreich's work as discussed in Section 1.2.4, we obtain the desired extension $\hat{\Phi}$ making the diagram commute. Continuity follows from Proposition 1.5 and Corollary 2.8. \square

We will also need the following

Proposition 2.11. *Suppose $\{v_n\} \subset \mathcal{C}(S)$, $\{x_n\} \subset \mathbb{H}$, and $x_n \rightarrow x \in \mathbb{H}$. If $\{v_n\}$ does not accumulate on $\partial_\infty \mathcal{C}(S)$, then $\{\Phi(v_n, x_n)\}$ does not accumulate on $\partial_\infty \mathcal{C}(S, z)$.*

Proof. Suppose $\{\Phi(v_n, x_n)\}$ accumulates on some lamination $|\mu| \in \partial_\infty \mathcal{C}(S, z)$. Pass to a subsequence which converges to $|\mu|$ in $\overline{\mathcal{C}}(S, z)$. If any curve in the sequence $\{v_n\}$ occurs infinitely often, then passing to a further subsequence, we can assume v_n is constant equal to v and $\Phi(v_n, x_n) = \Phi(v, x_n)$ lies in a compact subset of the tree $T_v \cong \Pi^{-1}(v)$, so does not accumulate on $\partial_\infty \mathcal{C}(S, z)$, as required.

We therefore assume that the v_n are all distinct. As such, it follows that if we let λ_n be the laminations corresponding to v_n under the natural map $\mathbb{P}\mathcal{ML}\mathcal{C}(S) \rightarrow \mathcal{C}(S)$, then $T(\lambda_n) \rightarrow 0$, so as in the proof of Lemma 2.5, we may pass to a subsequence so that $|\lambda_n|$ converges to a geodesic lamination \mathcal{L} . Passing to a further subsequence, we also assume $\lambda_n \rightarrow \lambda$ in $\mathbb{P}\mathcal{ML}(S)$.

It follows from Proposition 1.5, no sublamination of \mathcal{L} lies in $\mathcal{EL}(S)$. In particular, removing the infinite isolated leaves of \mathcal{L} , we obtain a lamination which is disjoint from a simple closed curve v' and contains the support of $\hat{\lambda}$. Choosing $\epsilon > 0$ sufficiently small, we can assume that the train track τ obtained from $N_\epsilon(\mathcal{L})$ contains a subtrack τ_0 so that (1) τ_0 is disjoint from some representative α of v' and (2) $\tau(\lambda)$ has nonzero weights only on the branches of τ_0 .

Now let $f \in \text{Diff}_0(S)$ be such that $\tilde{e}v(f) = x$. After modifying τ and τ_0 to τ' and τ'_0 as in the previous section if necessary (that is, possibly inserting a bigon around $f(z)$), it follows that for sufficiently large n , $f^{-1}(\tau'(\lambda_n))$ determines the lamination $\Psi(\lambda_n, x_n)$. After passing to yet a further subsequence if necessary, we can assume that $f^{-1}(\tau'(\lambda_n))$ converges (projectively) to some $f^{-1}(\tau')(\mu_0)$, also having nonzero weights only on $f^{-1}(\tau'_0)$. It follows that μ_0 , the (projective) limit of $\Psi(\lambda_n, x_n)$, is *not* in $\mathcal{FL}(S, z)$ since its support is disjoint from $f^{-1}(\alpha)$. By Proposition 1.5, $|\mu'| = |\mu|$, which is a contradiction. \square

Just as we restricted Φ to $v \times \mathbb{H}$ to map onto the Bass–Serre tree, we can restrict $\hat{\Phi}$ to $\{|\lambda|\} \times \mathbb{H}$ for any $|\lambda| \in \mathcal{EL}(S)$. Of course, $\hat{\Phi}(|\lambda|, x) = \hat{\Phi}(|\lambda|, x')$ if x, x' lie on the same leaf, or in the closure of the same component of $\mathbb{H} - p^{-1}(|\lambda|)$. We also have the following, which states that the converse is also true.

Lemma 2.12. *For $(|\lambda|, x), (|\lambda'|, x') \in \mathcal{EL}(S) \times \mathbb{H}$, $\hat{\Phi}(|\lambda|, x) = \hat{\Phi}(|\lambda'|, x')$ if and only if $|\lambda| = |\lambda'|$ and x and x' are in the same leaf or the closure of the same complementary region of $\mathbb{H} - p^{-1}(|\lambda|)$.*

Proof. Only one direction requires explanation. Assuming $\hat{\Phi}(|\lambda|, x) = \hat{\Phi}(|\lambda'|, x')$, we must show that $|\lambda| = |\lambda'|$ and x and x' are in the same leaf or closure of the same complementary region of $\mathbb{H} - p^{-1}(|\lambda|)$.

Assume we have applied an isotopy so that the laminations $\hat{\Phi}(|\lambda|, x)$ and $\hat{\Phi}(|\lambda'|, x')$ are equal (not just isotopic). Forgetting z , the laminations remain

the same (though they may have a bigon complementary region, and so are not necessarily geodesic laminations), and hence $|\lambda| = |\lambda'|$.

Proving the statement about x and x' is slightly more subtle. For simplicity, we assume that x and x' lie in components of $\mathbb{H} - p^{-1}(|\lambda|)$ for simplicity (the general case is similar, but the notation is more complicated). Let $f, f' \in \text{Diff}_0(S)$ be such that $\tilde{e}v(f) = x$ and $\tilde{e}v(f') = x'$. Let \tilde{f} and \tilde{f}' be lifts of f and f' with $\tilde{f}(\tilde{z}) = x$ and $\tilde{f}'(\tilde{z}) = x'$. Modifying f by an element of $\text{Diff}_0(S, z)$ if necessary, we may assume that $f^{-1}(|\lambda|) = \hat{\Phi}(|\lambda|, x)$ and $f'^{-1}(|\lambda|) = \hat{\Phi}(|\lambda|, x')$ are equal (again, not just isotopic).

Since $f^{-1}(|\lambda|) = f'^{-1}(|\lambda|)$, it follows that $f' \circ f^{-1}(|\lambda|) = |\lambda|$. Back in \mathbb{H} this means $\tilde{f}' \circ \tilde{f}^{-1}(p^{-1}(|\lambda|)) = p^{-1}(|\lambda|)$. Since $\tilde{f}' \circ \tilde{f}^{-1}(x) = x'$, and $\tilde{f}' \circ \tilde{f}^{-1}$ is the identity on $\partial_\infty \mathbb{H}$, it must be that x and x' lie in the same complementary region of $\mathbb{H} - p^{-1}(|\lambda|)$, as required. \square

3 Universal Cannon-Thurston Maps

3.1 Quasiconvex Sets

Consider a biinfinite geodesic $\gamma \subset \mathbb{H}$ for which $p(\gamma)$ is a *filling* closed geodesic in S , by which we mean that $p(\gamma)$ is a closed geodesic and the complement of $p(\gamma)$ is a union of disks in S . We will consider this as fixed for the remainder of the paper. Let $\delta \in \pi_1(S)$ generate the (infinite cyclic) stabilizer of γ . We will make several statements about γ , though they will also obviously apply to any $\pi_1(S)$ -translate of γ .

Define

$$\mathfrak{X}(\gamma) = \Phi(\mathcal{C}(S) \times \gamma).$$

Let $H^+(\gamma)$ and $H^-(\gamma)$ denote the two half spaces bounded by γ and define

$$\mathcal{H}^\pm(\gamma) = \Phi(\mathcal{C}(S) \times H^\pm(\gamma)).$$

We assume that we have chosen $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)}$ so that each arc of $\gamma \cap p^{-1}(N_{\epsilon(v)}(v))$ traverses the neighborhood from one boundary component to the other (rather than being allowed to enter and exit the neighborhood through the same boundary component). Since $p(\gamma)$ is a closed geodesic in S , this can be easily arranged by shrinking each of the originally chosen numbers $\epsilon(v)$ as needed.

A subset X of a geodesic metric space is called *weakly convex* if for any two points of the set there exists a geodesic connecting the points contained in the set (in a Gromov hyperbolic space, weakly convex sets are in particular uniformly quasi-convex).

Proposition 3.1. $\mathfrak{X}(\gamma), \mathcal{H}^\pm(\gamma)$ are simplicial subcomplexes of $\mathcal{C}(S, z)$ spanned by their vertex sets and are weakly convex.

To say that a subcomplex $\Omega \subset \mathcal{C}(S, z)$ is *spanned by its vertex set*, we mean that Ω is the largest subcomplex having $\Omega^{(0)}$ as its vertex set.

Proof. We describe the case of $\mathfrak{X}(\gamma)$, with $\mathcal{H}^\pm(\gamma)$ handled by similar arguments. We will construct a simplicial (and hence 1-Lipschitz) projection onto $\mathfrak{X}(\gamma)$. This easily implies the proposition. This is most succinctly done in terms of the description of $\mathcal{C}(S, z)$ given in Theorem 2.3, and so we identify $\mathcal{C}(S, z)$ with the quotient space as described in that theorem.

For any $x \in \text{int}(v)$, $\mathfrak{X}(\gamma) \cap \Pi^{-1}(x) = \Phi(\{x\} \times \gamma)$, which is a biinfinite geodesic in the tree $\Pi^{-1}(x) \cong T_v$. One can also see this as the axis of δ in T_v (since $p(\gamma)$ is filling, δ can never be elliptic in T_v) and we denote this as $\gamma_v \subset T_v$. Recall that an inclusion of faces $\phi : v' \rightarrow v$ induces a quotient of associated trees $\phi^* : T_v \rightarrow T_{v'}$. Since the axis of δ in T_v is sent to the axis of δ in $T_{v'}$ by ϕ^* , we have $\phi^*(\gamma_v) = \gamma_{v'}$. Therefore, with respect to our homeomorphism with the quotient of Theorem 2.3, we have

$$\mathfrak{X}(\gamma) \cong \left(\bigsqcup_{v \in \mathcal{C}(S)} v \times \gamma_v \right) / \sim \quad (1)$$

where, as in Theorem 2.3, the disjoint union is over all simplices $v \in \mathcal{C}(S)$, and the equivalence relation is generated by

$$(\phi(x), t) \sim (x, \phi^*(t))$$

for all faces $\phi : v' \rightarrow v$, all $x \in v'$ and all $t \in \gamma_v$. We also use the homeomorphism in (1) to identify the two spaces.

The simplices of $\mathcal{C}(S, z)$ via the homeomorphism of Theorem 2.3 are precisely the images of cells $v \times \sigma$ in the quotient, where $v \in \mathcal{C}(S)$ is a simplex and $\sigma \subset T_v$ an edge or vertex. Thus, if the image of $v \times \sigma$ is a simplex, and we let v_0, \dots, v_k be the vertices of v and t_0, t_1 the vertices of σ (assuming, for example, that σ is an edge) then the vertices of the simplex determined by $v \times \sigma$ are image of (v_i, t_j) for $i = 0, \dots, k$ and $j = 0, 1$. If these vertices lie in $\mathfrak{X}(\gamma)$, then $t_0, t_1 \in \gamma_v$, hence $\sigma \subset \gamma_v$ and the image of $v \times \sigma$ lies in $\mathfrak{X}(\gamma)$. It follows that $\mathfrak{X}(\gamma)$ is a simplicial subcomplex of $\mathcal{C}(S, z)$ spanned by its vertex set.

Next, we will define a projection

$$\rho : \mathcal{C}(S, z) \rightarrow \mathfrak{X}(\gamma)$$

by first defining it on each of the pieces $\rho_v : v \times T_v \rightarrow v \times \gamma_v$ as the identity on the first factor and the nearest-point projection $\eta_v : T_v \rightarrow \gamma_v$ on the second factor. Observe that if $\phi : v' \rightarrow v$ is a face, then nearest-point projections commute

$$\eta_{v'} \circ \phi^* = \phi^* \circ \eta_v$$

since a geodesic segment in T_v from a point t to γ_v is taken to a geodesic segment from $\phi^*(t)$ to $\gamma_{v'}$. From this it follows that the maps ρ_v glue together to well-define the map ρ .

All that remains is to verify that ρ is simplicial. Given a simplex which is the image of $v \times \sigma$ in the quotient, for some $v \in \mathcal{C}(S)$ and $\sigma \subset T_v$, we have ρ of this simplex is the image of $\rho_v(v \times \sigma) = v \times \eta_v(\sigma)$ in the quotient. Since $\eta_v(\sigma)$ is either an edge or vertex, $v \times \eta_v(\sigma)$ projects to a simplex in the quotient, as required. \square

Proposition 3.2. *We have*

$$\mathcal{H}^+(\gamma) \cup \mathcal{H}^-(\gamma) = \mathcal{C}(S, z)$$

and

$$\mathcal{H}^+(\gamma) \cap \mathcal{H}^-(\gamma) = \mathfrak{X}(\gamma).$$

Proof. The first statement follows from the fact that $H^+(\gamma) \cup H^-(\gamma) = \mathbb{H}$ and that Φ is surjective.

The second can be seen by looking in each tree $\Pi^{-1}(v) \cong T_v$ (we use this homeomorphism to identify the two spaces). For any vertex $v \in \mathcal{C}(S)$, let $H^\pm(\gamma_v) = \mathcal{H}^\pm(\gamma) \cap \Pi^{-1}(v) = \Phi(\{v\} \times H^\pm(\gamma))$. Since $H^\pm(\gamma)$ is connected and $H^\pm(\gamma_v)$ is a subgraph (by Proposition 3.1), so $H^\pm(\gamma_v)$ is a subtree. Therefore, $H^+(\gamma_v) \cap H^-(\gamma_v)$ is a subtree containing γ_v . Let $u \in H^+(\gamma_v) \cap H^-(\gamma_v)$ be any vertex, and we must show that $u \in \gamma_v$.

We can write $u = \Phi(\{v\} \times U)$ where U is a component of $\mathbb{H} - p^{-1}(N_\epsilon(v))$. Therefore, $U \cap H^+(\gamma) \neq \emptyset$ and $U \cap H^-(\gamma) \neq \emptyset$. Since U is connected and γ separates $H^+(\gamma)$ and $H^-(\gamma)$ we have $U \cap \gamma \neq \emptyset$, and hence $\Phi(\{v\} \times U) = u \in \gamma_v$ as required.

Now suppose $u \in \mathcal{H}^+(\gamma) \cap \mathcal{H}^-(\gamma)$ is any vertex. Setting $v = \Pi(u)$, we have

$$u \in (\Pi^{-1}(v) \cap \mathcal{H}^+(\gamma) \cap \mathcal{H}^-(\gamma)) = \gamma_v = \mathfrak{X}(\gamma) \cap \Pi^{-1}(v) \subset \mathfrak{X}(\gamma).$$

Since $\mathfrak{X}(\gamma)$ is a subcomplex spanned by its vertex set, this completes the proof. \square

It will be convenient to keep the terminology in the proofs of these propositions and think of γ_v as “bounding half-trees” $H^\pm(\gamma_v) \subset T_v \cong \Pi^{-1}(v)$.

3.2 Rays and existence of Cannon-Thurston Maps

A ray $r \subset \mathbb{H}$ is said to have the *all tails filling property* (briefly, r is an *ATF ray*) if for any simple closed geodesic v on S and any proper tail $r' \subset r$

$$p(r') \cap v \neq \emptyset.$$

This clearly depends only on the asymptote class of r , and we say that a point $x \in \partial_\infty \mathbb{H}$ is an *ATF point*, if any ray ending at x is an ATF ray. Let $\mathbb{A}_\infty \subset \partial_\infty \mathbb{H}$ denote the set of all ATF points.

Lemma 3.3. *If r is not an ATF ray, then $\Phi(\{v\} \times r)$ has bounded diameter for all $v \in \mathcal{C}^0(S)$.*

Proof. Since r is not an ATF ray, there is some $v' \in \mathcal{C}^0(S)$, and a tail $r' \subset r$ for which $p(r')$ is disjoint from v' . Therefore, $\Phi(\{v'\} \times r')$ is vertex in $\Pi^{-1}(v') \cong T_v$, and hence $\Phi(\{v'\} \times r)$ has finite diameter. Since $\Phi(\{v\} \times r)$ is a finite Hausdorff distance from $\Phi(\{v'\} \times r)$, the latter also has finite diameter. \square

Recall that we have fixed once and for all a geodesic $\gamma \subset \mathbb{H}$ which projects to a nonsimple closed geodesic in S . Consider a set $\{\gamma_n\}$ of pairwise distinct $\pi_1(S)$ -translates of γ , with the property that the half spaces are nested:

$$H^+(\gamma_1) \supset H^+(\gamma_2) \supset \dots$$

Since the γ_n are all distinct, proper discontinuity of the action of $\pi_1(S)$ on \mathbb{H} implies that

$$\bigcap_{n=1}^{\infty} \overline{H^+(\gamma_n)} = \{x\}$$

for some $x \in \partial_{\infty}\mathbb{H}$. Here the bar denotes closure in $\overline{\mathbb{H}} = \mathbb{H} \cup \partial_{\infty}\mathbb{H}$. We say that $\{\gamma_n\}$ *nest down on* x . Note that $\{H^+(\gamma_n)\}$ determines a neighborhood basis for x .

Given any $x \in \partial_{\infty}\mathbb{H}$, if $r \subset \mathbb{H}$ is a geodesic ray ending at x , then since $p(\gamma)$ is filling, $p(r)$ intersects $p(\gamma)$ infinitely often. It follows that there is a sequence $\{\gamma_n\}$ which nest down on x .

Proposition 3.4. *If $\{\gamma_n\}$ is a sequence nesting down on an ATF point $x \in \mathbb{A}_{\infty}$, then for any choice of basepoint $u_0 \in \mathcal{C}(S, z)$,*

$$d(u_0, \mathcal{H}^+(\gamma_n)) \rightarrow \infty$$

as $n \rightarrow \infty$.

Proof. In what follows, all distances are computed in the 1-skeleton (which differs from the distance in the curve complex by a uniformly bounded multiplicative constant). We write $u_0 = \Phi(v_0, y)$ for some vertex $v_0 \in \mathcal{C}(S)$ and $y \in \mathbb{H}$. By discarding a finite number of initial elements of the sequence $\{\gamma_n\}$ we may assume that $y \in H^-(\gamma_n)$ for all n , and so $u_0 \in \mathcal{H}^-(\gamma_n)$ for all n .

Now, fix any $R > 0$. Since

$$\mathcal{H}^+(\gamma_1) \supset \mathcal{H}^+(\gamma_2) \supset \mathcal{H}^+(\gamma_3) \supset \dots$$

we must show that there exists $N > 0$ so that for all $u \in \mathcal{H}^+(\gamma_N)$, $d(u_0, u) \geq R$.

Claim 1. It suffices to prove that there exists $N > 0$, so that for all $u \in (\mathcal{H}^+(\gamma_N) \cap \Pi^{-1}(B(v_0, R)))$, the distance inside $\Pi^{-1}(B(v_0, R))$ from u_0 to u is at least R .

Proof. Observe that any edge path from a point $u \in \mathcal{C}(S, z)$ to u_0 which meets $\mathcal{C}(S, z) - \Pi^{-1}(B(v_0, R))$ projects to a path which meets both $\mathcal{C}(S) - B(v_0, R)$ and v_0 , and therefore has length at least R . Since Π is 1-Lipschitz, the length of the path in $\mathcal{C}(S, z)$ is also at least R . \square

The intersection of $\mathcal{H}^+(\gamma_n)$ with each fiber $\Pi^{-1}(v) \cong T_v$ is a half tree denoted $H^+(\gamma_{n,v})$ bounded by $\gamma_{n,v} = \mathfrak{X}(\gamma_n) \cap \Pi^{-1}(v)$ (see the proof of Proposition

3.2 and comments following it).

Claim 2. For any $k > 0$, there exists positive integers $N_1 < N_2 < N_3 < \dots < N_k$ so that

$$\gamma_{N_j, v} \cap \gamma_{N_{j+1}, v} = \emptyset \quad (2)$$

for all $j = 1, \dots, k-1$ and all $v \in B(v_0, R)$.

Proof. The proof is by induction on k . For $k = 1$, the condition is vacuously satisfied by setting $N_1 = 1$. So, we assume it is true for $k \geq 1$, and prove it true for $k+1$. Thus, by hypothesis, we have found $N_1 < N_2 < \dots < N_k$ so that (2) is true, and we need to find N_{k+1} so that

$$\gamma_{N_k, v} \cap \gamma_{N_{k+1}, v} = \emptyset \quad (3)$$

for all $v \in B(v_0, R)$.

We suppose that no such N_{k+1} exists and arrive at a contradiction. Observe that the nesting

$$H^+(\gamma_{1, v}) \supset H^+(\gamma_{2, v}) \supset \dots$$

means that if $\gamma_{n, v} \cap \gamma_{m, v} = \emptyset$ for some $m > n$, then $\gamma_{n, v} \cap \gamma_{m+j, v} = \emptyset$ for all $j \geq 0$.

Thus, if there is no N_{k+1} , it must be that for every $j > 0$, there exists a curve $v_j \in B(v_0, R)$ so that $\gamma_{N_k, v} \cap \gamma_{N_k+j, v} \neq \emptyset$. Let $u_j \in \gamma_{N_k, v} \cap \gamma_{N_k+j, v}$ be a vertex in common. This vertex is the image of a component $U_j \subset \mathbb{H} - p^{-1}(v_j)$ which meets both γ_{N_k} and γ_{N_k+j} ; see Figure 6.

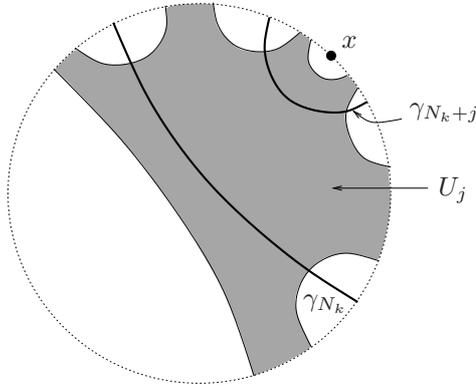


Figure 6: The region U_j and the geodesics γ_{N_k} and γ_{N_k+j} from the sequence nesting down on x .

Let $g_j \subset U_j \subset \mathbb{H}$ be a geodesic segment connecting a point $y_j^- \in \gamma_{N_k}$ to $y_j^+ \in \gamma_{N_k+j}$. Note that as $j \rightarrow \infty$, y_j^+ converges to x . Furthermore, observe

that all y_j^- must lie on a some compact arc $K \subset \gamma_{N_k}$. For, if not then there must exist g_j intersecting γ_{N_k} at arbitrarily small angles. In particular, there must be g_j which run very nearly parallel to γ_{N_k} for any arbitrarily long time. However, because $p(\gamma_{N_k})$ is a filling curve and g_j is contained in U_j which is disjoint from $p^{-1}(v_j)$, this is impossible.

So, after passing to a subsequence if necessary, we may assume that y_j^- converge to some point $y \in \gamma_{N_k}$, and hence the sequence of geodesics g_j converges to a geodesic ray r_∞ from y to x . By passing to a further subsequence, we can assume that v_j limits in the Hausdorff topology to a geodesic lamination \mathcal{L} , and that $p(r_\infty)$ does not transversely intersect \mathcal{L} . Because v_j are all contained in $B(v_0, R)$, \mathcal{L} cannot contain an ending lamination as a sublamination. It follows from [CB87] that \mathcal{L} is obtained from a lamination supported on a proper subsurface Σ by adding a finite number of isolated leaves. Any geodesic ray in S which does not transversely intersect \mathcal{L} can only transversely intersect $\partial\Sigma$ twice (when it possibly exits/enters a crown; see [CB87]), and hence some tail of $p(r_\infty)$ is disjoint from $\partial\Sigma$. Therefore, r_∞ cannot be an ATF ray, which is a contradiction since it ends at x . \square

Now, pick an integer $k > R + 1$ and let $N_1 < N_2 < \dots < N_k$ be as in Claim 2. There can be no vertices $u \in \mathfrak{X}(\gamma_{N_j}) \cap \mathfrak{X}(\gamma_{N_{j+1}}) \cap B(v_0, R)$, and since these are subcomplexes, it must be that

$$\mathfrak{X}(\gamma_{N_j}) \cap \mathfrak{X}(\gamma_{N_{j+1}}) \cap \Pi^{-1}(B(v_0, R)) = \emptyset.$$

Moreover, since

$$\mathcal{H}^+(\gamma_{N_1}) \supset \mathcal{H}^+(\gamma_{N_2}) \supset \dots \supset \mathcal{H}^+(\gamma_{N_k})$$

it follows from Proposition 3.2 that

$$\mathfrak{X}(\gamma_{N_j}) \cap \mathfrak{X}(\gamma_{N_i}) \cap \Pi^{-1}(B(v_0, R)) = \emptyset \quad (4)$$

for all $i \neq j$ between 1 and k .

Let $u \in \mathcal{H}^+(\gamma_{N_k}) \cap \Pi^{-1}(B(v_0, R))$ be any point and $\{u_0, u_1, \dots, u_m = u\}$ be the vertices of an edge path from u_0 to u within $\Pi^{-1}(B(v_0, R))$. We have $u_0 \in \mathcal{H}^-(\gamma_{N_j})$ for all j and $u \in \mathcal{H}^+(\gamma_{N_k}) \subset \mathcal{H}^+(\gamma_{N_j})$ for all j . By Proposition 3.2, the edge path must meet $\mathfrak{X}(\gamma_{N_j})$ for each j . That is, for each j , there is some $i(j)$ so that $u_{i(j)} \in \mathfrak{X}(\gamma_{N_j})$. By (4), there must therefore be at least $k > R + 1$ vertices in the path, and hence the length of the path is at least R .

Therefore, setting $N = N_k$, we have for all $u \in (\mathcal{H}^+(\gamma_N) \cap \Pi^{-1}(B(v_0, R)))$, the distance inside $\Pi^{-1}(B(v_0, R))$ from u_0 to u is at least R . By Claim 1, this completes the proof of the proposition. \square

We can now prove the first half of Theorem 1.1.

Theorem 3.5. *For any $v \in \mathcal{C}^0(S)$, the map*

$$\Phi^v : \mathbb{H} \rightarrow \mathcal{C}(S, z)$$

has a continuous $\pi_1(S)$ -equivariant extension to

$$\bar{\Phi}^v : \mathbb{H} \cup \mathbb{A}_\infty \rightarrow \bar{\mathcal{C}}(S, z).$$

Proof. Observe that Φ^v is already defined and continuous. All that remains is to extend it to $\overline{\Phi}^v$ on \mathbb{A}_∞ and show that it is continuous there.

Fix a basepoint $u_0 \in \mathcal{C}(S, z)$. Given any $x \in \mathbb{A}_\infty$, let $\{\gamma_n\}$ be any sequence nesting down on x . According to Proposition 3.4, we have

$$d(u_0, \mathcal{H}^+(\gamma_n)) \rightarrow \infty.$$

Moreover, by Proposition 3.1, $\mathcal{H}^+(\gamma_n)$ is weakly convex and hence uniformly quasi-convex. Finally, observe that $\Phi^v(H^+(\gamma_n)) = \Phi(\{v\} \times H^+(\gamma_n)) \subset \mathcal{H}^+(\gamma_n)$. Since $x \in \mathbb{A}_\infty$ was an arbitrary point, Lemma 1.7 implies the existence of an \mathbb{A}_∞ -Cannon-Thurston map $\overline{\Phi}^v$. \square

We note that given $x \in \mathbb{A}_\infty$, $\overline{\Phi}^v(x)$ depends only on x , and not on v , and is given simply as the unique point of intersection of the sets

$$\bigcap_n \overline{\mathcal{H}^+(\gamma_n)}.$$

We can therefore unambiguously define $\overline{\Phi}(x) = \overline{\Phi}^v(x)$, independent of the choice of $v \in \mathcal{C}^0(S)$.

3.3 Separation

Proposition 3.6. *Given $x, y \in \mathbb{A}_\infty$, let ϵ be the geodesic connecting them. Then there are $\pi_1(S)$ -translates γ_x and γ_y of γ defining half-space neighborhoods $\overline{H^+(\gamma_x)}$ and $\overline{H^+(\gamma_y)}$ of x and y , respectively, with*

$$\partial_\infty \mathcal{H}^+(\gamma_x) \cap \partial_\infty \mathcal{H}^+(\gamma_y) = \emptyset$$

if and only if $p(\epsilon)$ is nonsimple (i.e. $p(\epsilon)$ has at least one transverse self intersections).

Before we can give the proof of Proposition 3.6, we will need the analogue of Proposition 3.2 for the boundaries at infinity. Recall that γ was chosen to be a biinfinite geodesic with stabilizer $\langle \delta \rangle$ and $p(\gamma)$ a filling closed geodesic.

Proposition 3.7. *We have*

$$\partial_\infty \mathcal{H}^+(\gamma) \cap \partial_\infty \mathcal{H}^-(\gamma) = \partial_\infty \mathfrak{X}(\gamma).$$

Proof. This follows from general principles. If $|\mu| \in \partial_\infty \mathcal{H}^+(\gamma) \cap \partial_\infty \mathcal{H}^-(\gamma)$, then let $\{u_n^+\} \subset \mathcal{H}^+(\gamma)$ and $\{u_n^-\} \in \mathcal{H}^-(\gamma)$ be sequences converging to $|\mu|$ in $\mathcal{C}(S, z)$. Let g_n be geodesic segments from u_n^+ to u_n^- . By Proposition 3.2, there is a vertex $u_n \in g_n \cap \mathfrak{X}(\gamma)$. It follows that u_n also converges to $|\mu|$, and so $|\mu| \in \partial_\infty \mathfrak{X}(\gamma)$. \square

If $H^+(\gamma_x)$ and $H^+(\gamma_y)$ are disjoint, then to prove that the intersection of Proposition 3.6 is empty, it suffices to show $\partial_\infty \mathfrak{X}(\gamma_x)$ and $\partial_\infty \mathfrak{X}(\gamma_y)$ are disjoint. For this we need a description of $\partial_\infty \mathfrak{X}(\gamma)$.

Recall that γ was chosen to have $p(\gamma)$ a filling closed geodesic on S and we let δ generate the infinite cyclic stabilizer of γ . Work of Kra [Kra81] implies that δ is pseudo-Anosov. We let $|\mu_+|$ and $|\mu_-|$ be the attracting and repelling fixed points of δ , respectively, in $\partial_\infty \mathcal{C}(S, z)$.

Lemma 3.8.

$$\partial_\infty \mathfrak{X}(\gamma) = \hat{\Phi}(\partial_\infty \mathcal{C}(S) \times \gamma) \cup \{|\mu_\pm|\}$$

Proof. Continuity of $\hat{\Phi}$ implies $\hat{\Phi}(\partial_\infty \mathcal{C}(S) \times \gamma) \subset \partial_\infty \mathfrak{X}(\gamma)$. Invariance of γ by δ implies $\{|\mu_\pm|\} \subset \partial_\infty \mathfrak{X}(\gamma)$, and so

$$\partial_\infty \mathfrak{X}(\gamma) \supset \hat{\Phi}(\partial_\infty \mathcal{C}(S) \times \gamma) \cup \{|\mu_\pm|\}.$$

Suppose $\{u_n\}$ is any sequence in $\mathfrak{X}(\gamma)$ with $u_n \rightarrow |\mu| \in \partial_\infty \mathfrak{X}(\gamma)$. We wish to show that $|\mu| \in \hat{\Phi}(\partial_\infty \mathcal{C}(S) \times \gamma) \cup \{|\mu_\pm|\}$, proving the reverse inclusion. By definition of $\mathfrak{X}(\gamma)$ there exists $\{(v_n, x_n)\} \subset \mathcal{C}(S) \times \gamma$ with $\Phi(v_n, x_n) = u_n$ for all n . There are two cases to consider.

Case 1. $\{x_n\} \subset K$, for some compact arc $K \subset \gamma$.

After passing to a subsequence if necessary $x_n \rightarrow x \in K$. By Proposition 2.11, we can assume that v_n accumulates on $\partial_\infty \mathcal{C}(S)$. So, after passing to yet a further subsequence, we can assume that $v_n \rightarrow |\lambda| \in \partial_\infty \mathcal{C}(S)$. Then by continuity of $\hat{\Phi}$ (Proposition 2.10) we have

$$|\mu| = \lim_{n \rightarrow \infty} \Phi(v_n, x_n) = \hat{\Phi}(|\lambda|, x) \in \hat{\Phi}(\partial_\infty \mathcal{C}(S) \times \gamma).$$

Case 2. After passing to a subsequence $x_n \rightarrow x$, where x is one of the endpoints of γ in $\partial_\infty \mathbb{H}$.

Note that $x \in \mathbb{A}_\infty$ since $p(\gamma)$ is filling. Indeed, x is either the attracting or repelling fixed point of δ . Without loss of generality, we assume it is the attracting fixed point. Now suppose γ_1 is any $\pi_1(S)$ translate which nontrivially intersects γ so that $\{\delta^n(\gamma_1)\}$ nests down on x , and hence

$$\bigcap_{n=1}^{\infty} \overline{\mathcal{H}^+(\delta^n(\gamma_1))} = \bigcap_{n=1}^{\infty} \delta^n(\overline{\mathcal{H}^+(\gamma_1)})$$

consists of the single point $|\mu_+|$, the stable fixed point of the pseudo-Anosov δ . After passing to a further subsequence if necessary, we can assume $x_n \in H^+(\delta^n(\gamma_1))$. Therefore, $\Phi(v_n, x_n) \in \mathcal{H}^+(\delta^n(\gamma_1))$, and hence

$$|\mu| = \lim_{n \rightarrow \infty} \Phi(v_n, x_n) = |\mu_+|.$$

□

Proof of Proposition 3.6. First, suppose $p(\epsilon)$ is simple. The closure of $p(\epsilon)$ is a lamination \mathcal{L} [CB87]. Since x and y are \mathbb{A}_∞ points, \mathcal{L} must contain $|\lambda| \in \mathcal{EL}(S)$

as the sublamination obtained by discarding isolated leaves, and ϵ is either a leaf of $p^{-1}(|\lambda|)$ or a diagonal for some complementary polygon of $p^{-1}(|\lambda|)$. Now, suppose γ_x and γ_y are $\pi_1(S)$ -translates of γ for which $H^+(\gamma_x)$ and $H^+(\gamma_y)$ are half-space neighborhoods of x and y , respectively.

It follows from Lemma 2.12 that if $x' \in \gamma_x \cap \epsilon$ and $y' \in \gamma_y \cap \epsilon$, then $\hat{\Phi}(|\lambda|, x') = \hat{\Phi}(|\lambda|, y')$. Appealing to Lemma 3.8 we have

$$\begin{aligned} \emptyset &\neq \hat{\Phi}(\{|\lambda|\} \times \gamma_x) \cap \hat{\Phi}(\{|\lambda|\} \times \gamma_y) \\ &\subset \partial_\infty \mathfrak{X}(\gamma_x) \cap \partial_\infty \mathfrak{X}(\gamma_y) \\ &\subset \partial_\infty \mathcal{H}^+(\gamma_x) \cap \partial_\infty \mathcal{H}^+(\gamma_y) \end{aligned}$$

as required. In fact, it is worth noting that by Lemma 2.12, $\hat{\Phi}(\{|\lambda|\} \times \epsilon)$ is a single point which lies in all $\partial_\infty \mathcal{H}^+(\gamma_x) \cap \partial_\infty \mathcal{H}^+(\gamma_y)$, and is therefore equal to $\overline{\Phi}^v(x) = \overline{\Phi}^v(y)$.

Before we prove the converse, suppose γ_1 and γ_2 are two translates of γ for which $H^+(\gamma_1) \subset H^-(\gamma_2)$ and $H^+(\gamma_2) \subset H^-(\gamma_1)$. Then we have

$$\partial_\infty \mathcal{H}^+(\gamma_1) \subset \partial_\infty \mathcal{H}^-(\gamma_2) \text{ and } \partial_\infty \mathcal{H}^+(\gamma_2) \subset \partial_\infty \mathcal{H}^-(\gamma_1).$$

Therefore, by Proposition 3.7, it follows that

$$\partial_\infty \mathcal{H}^+(\gamma_1) \cap \partial_\infty \mathcal{H}^+(\gamma_2) \subset \partial_\infty \mathfrak{X}(\gamma_1) \cap \partial_\infty \mathfrak{X}(\gamma_2).$$

Further suppose that $\gamma_1 \neq \gamma_2$, so that fixed points of the elements δ_1 and δ_2 generating the stabilizers of γ_1 and γ_2 , respectively, are disjoint in $\partial_\infty \mathcal{C}(S, z)$. If

$$\partial_\infty \mathcal{H}^+(\gamma_1) \cap \partial_\infty \mathcal{H}^+(\gamma_2) \neq \emptyset$$

then by Proposition 3.8 there exists $x_1 \in \gamma_1$ and $x_2 \in \gamma_2$ and $|\lambda_1|, |\lambda_2| \in \partial_\infty \mathcal{C}(S)$ for which $\hat{\Phi}(|\lambda_1|, x_1) = \hat{\Phi}(|\lambda_2|, x_2)$. According to Lemma 2.12, we have $|\lambda_1| = |\lambda_2|$, and x_1 and x_2 lie on the same leaf, or in the closure of the same complementary region of $|\lambda_1|$. In particular, there is a biinfinite geodesic contained in a complementary region or leaf of $p^{-1}(|\lambda_1|)$ which meets both γ_1 and γ_2 .

We now proceed to the proof of the converse. Let $\{\gamma_{n,x}\}$ and $\{\gamma_{n,y}\}$ be sequences of $\pi_1(S)$ -translates of γ which nest down on x and y , respectively. If $\partial_\infty \mathfrak{X}(\gamma_{n,x}) \cap \partial_\infty \mathfrak{X}(\gamma_{n,y}) = \emptyset$ for some n , then we are done. If not, then we must show that the geodesic ϵ with endpoints x and y has $p(\epsilon)$ simple on S .

Assuming $\partial_\infty \mathfrak{X}(\gamma_{n,x}) \cap \partial_\infty \mathfrak{X}(\gamma_{n,y}) \neq \emptyset$, there exists a sequence of laminations $\{|\lambda_n|\} \subset \partial_\infty \mathcal{C}(S)$ so that $\gamma_{x,n}$ and $\gamma_{y,n}$ both meet a leaf or complementary polygon of $p^{-1}(|\lambda_n|)$. It follows that there is a sequence of geodesics $\{\epsilon_n\}$ in \mathbb{H} for which $p(\epsilon_n)$ is simple on S , and $\epsilon_n \cap \gamma_{x,n} \neq \emptyset \neq \epsilon_n \cap \gamma_{y,n}$. The limit ϵ of $\{\epsilon_n\}$ has endpoints x and y and must also have $p(\epsilon)$ simple, as required (the limit of simple geodesics is simple—see [CB87]). \square

The following is now immediate from Proposition 3.6 and its proof.

Corollary 3.9. *Given $x, y \in \mathbb{A}_\infty$, $\overline{\Phi}^v(x) = \overline{\Phi}^v(y)$ if and only if x and y are endpoints of a leaf or vertices of a complementary polygon of $p^{-1}(|\lambda|)$ for some $|\lambda| \in \partial_\infty \mathcal{C}(S)$.*

3.4 Surjectivity

In this section, we prove that our map $\partial\Phi^v$ is surjective. This will require a few lemmas.

According to Birman-Series [BS85], the set

$$\overline{\bigcup_{v \in \mathcal{C}^0(S)} v}$$

is nowhere dense. We fix an $\epsilon > 0$, and assume that $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)}$ is chosen sufficiently small so that

$$S - \bigcup_{v \in \mathcal{C}^0(S)} N_{\epsilon(v)}(v)$$

is ϵ -dense. We further assume that $\epsilon(v) < \epsilon$ for all $v \in \mathcal{C}^0(S)$.

Lemma 3.10. *Suppose $(v_1, x_1), (v_2, x_2) \in \mathcal{C}^0(S) \times \mathbb{H}$ with $\Phi(v_i, x_i) = u_i$ a vertex in $\mathcal{C}(S, z)$ for $i = 1, 2$. Then there is a path*

$$\tilde{v} = (\tilde{v}_1, \tilde{v}_2) : [a, b] \rightarrow \mathcal{C}(S) \times \mathbb{H}$$

such that $\Phi \circ \tilde{v}$ is a geodesic from u_1 to u_2 and $\tilde{v}_2([a, b])$ connects x_1 to x_2 and is contained in the 2ϵ -neighborhood of a geodesic containing x_1, x_2 .

Proof. For each $i = 1, 2$, we can find x'_i in the same component of $S - N_{\epsilon(v_i)}$ as x_i within ϵ of x_i such that x'_1 and x'_2 are contained in some geodesic γ' which projects to a filling closed geodesic in S (the pairs of endpoints of such geodesics in $\partial_\infty \mathbb{H}$ is dense). Then $\Phi(v_i, x_i) = \Phi(v_i, x'_i)$ for $i = 1, 2$. Moreover, the geodesic from x'_1 to x'_2 is within ϵ of the geodesic from x_1 to x_2 . If we can find $\tilde{v}' = (\tilde{v}'_1, \tilde{v}'_2)$ so that $\Phi \circ \tilde{v}'$ is a geodesic from u_1 to u_2 and $\tilde{v}'_2([a, b])$ connects x'_1 to x'_2 and is within ϵ of the geodesic from x'_1 to x'_2 , then we can take $\tilde{v} = (\tilde{v}_1, \tilde{v}_2)$ to be given by $\tilde{v}_1 = \tilde{v}'_1$ and \tilde{v}_2 to first run from x_1 to x'_1 , then traverse \tilde{v}'_2 , and finally run from x'_2 to x_2 (all appropriately reparametrized). This will then provide the desired path proving the lemma.

To construct \tilde{v}' , we suppose for the moment that $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)}$ have been chosen so that any arc of $\gamma' \cap p^{-1}(N_{\epsilon(v)}(v))$ enters and exists the component it meets through distinct boundary components. With this assumption, we can apply the proof of Proposition 3.1 to show that $\mathfrak{X}(\gamma')$ is weakly convex. Now connect u_1 and u_2 by a geodesic edge path within $\mathfrak{X}(\gamma')$ with vertex set $\{u_1 = w_1, w_2, w_3, \dots, w_k = u_2\}$.

Let $v_i = \Pi(w_i)$. We observe that for every $i = 1, \dots, k$,

$$\Phi^{-1}(w_i) \cap (\mathcal{C}(S) \times \gamma') = \{v_i\} \times \alpha_i$$

where α_i is an arc of $\gamma' \cap (\mathbb{H} - p^{-1}(N_{\epsilon(v_i)}(v_i)))$ which is in particular connected. It follows from the construction of Φ that the geodesics $[w_i, w_{i+1}]$, for $i = 1, \dots, k-1$ are images of paths in $\mathcal{C}(S) \times \gamma'$ which we denote $a_i = (b_i, c_i)$. Explicitly, if $v_i = v_{i+1}$, then b_i is constant and equal to $v_i = v_{i+1}$, and c_i traverses an arc

of $N_{\epsilon(v_i)}(v_i) \cap \gamma'$. If $v_i \neq v_{i+1}$, then b_i traverses from v_i to v_{i+1} and c_i can be taken to be constant.

We can now define $\tilde{\nu}' = (\tilde{\nu}'_1, \tilde{\nu}'_2)$ as follows.

1. Begin by holding $\tilde{\nu}'_1$ constant and let $\tilde{\nu}'_2$ traverse from x'_1 to the initial point of c_1 inside $\alpha_1 \subset \gamma'$.
2. Next, traverse a_1 .
3. After that, hold $\tilde{\nu}'_1$ constant again and let $\tilde{\nu}'_2$ traverse from the terminal point of c_1 to the initial point of c_2 inside $\alpha_2 \subset \gamma'$.
4. We can continue in this way, for $i = 2, \dots, k-2$ traversing a_i , then holding $\tilde{\nu}'_1$ constant and letting $\tilde{\nu}'_2$ go from the terminal point of c_i to the initial point of c_{i+1} inside $\alpha_i \subset \gamma'$.
5. We complete the path by traversing a_{k-1} , then holding $\tilde{\nu}'_1$ constant and letting $\tilde{\nu}'_2$ traverse the path from the terminal point of c_{k-1} to x'_2 inside $\alpha_k \subset \gamma'$.

By construction, the projection of this path $\Phi \circ \tilde{\nu}'$ onto the first coordinate is the geodesic from u_1 to u_2 that we started with (although it stops and is constant at each of the vertices for some interval in the domain of the parametrization). Moreover, $\tilde{\nu}'_2$ is contained in γ' and connects x'_1 to x'_2 , so therefore stays within a distance 0 of the geodesic from x'_1 to x'_2 as required.

The proof so far was carried out under the assumption that for every $v \in \mathcal{C}^0(S)$, every arc of $\gamma' \cap N_{\epsilon(v)}(v)$ enters and exits the component of $N_{\epsilon(v)}(v)$ which it meets in different boundary components. If this is not true, then first shrink all $\epsilon(v)$ to numbers $\epsilon'(v) < \epsilon(v)$ so that it is true, construct the path as above, and call it $\tilde{\nu}'' = (\tilde{\nu}''_1, \tilde{\nu}''_2)$. Note that the numbers $\{\epsilon'(v)\}_{v \in \mathcal{C}^0(S)}$ determine a new map $\Phi' : \mathcal{C}(S) \times \mathbb{H} \rightarrow \mathcal{C}(S, z)$, and $\Phi' \circ \tilde{\nu}''$ is a geodesic. With respect to the original map Φ , $\tilde{\nu}''$ is almost good enough for our purposes. The only problem is that $\Phi \circ \tilde{\nu}''$ may now no longer be a geodesic: If there is some interval in the domain in which $\tilde{\nu}''_1$ is constant equal to v and $\tilde{\nu}''_2$ enters and exits a component $p^{-1}(N_{\epsilon(v)}(v))$ from the same side, then $\Phi \circ \tilde{\nu}''$ will divert from being a geodesic by running (less than half way) into an edge of $\Pi^{-1}(v)$ and running back out. We modify $\tilde{\nu}''$ to the desired path $\tilde{\nu}'$, by pushing $\tilde{\nu}''_2$ outside of $p^{-1}(N_{\epsilon(v)}(v))$ whenever this happens, thus changing it by at most $\epsilon(v) < \epsilon$. The resulting path $\tilde{\nu}'$ has $\tilde{\nu}'_1 = \tilde{\nu}''_1$ and $\tilde{\nu}'_2$ still connects x'_1 to x'_2 and stays within ϵ of γ' , as required. \square

Surjectivity of $\partial\Phi^v$ says that every point of $\partial_\infty\mathcal{C}(S, z)$ is the limit of $\Phi^v(r)$ for some ray $r \subset \mathbb{H}$ ending at a point of \mathbb{A}_∞ . The following much weaker conclusion is easier to arrive at, and will be used in the proof of surjectivity.

Lemma 3.11. *For any $v \in \mathcal{C}^0(S)$, $\overline{\Phi^v(\mathbb{H})} \cap \partial_\infty\mathcal{C}(S, z) = \partial_\infty\mathcal{C}(S, z)$.*

Proof. First, note that since $\pi_1(S) < \text{Mod}(S, z)$ is a normal, infinite subgroup the limit set in $\mathbb{P}\mathcal{ML}(S, z)$ (in the sense of [MP89]) is all of $\mathbb{P}\mathcal{ML}(S, z)$. In

particular, the closure of any $\pi_1(S)$ -equivariant embedding $\mathbb{H} \subset \mathcal{T}(S, z)$ in the Thurston compactification of Teichmüller space meets $\mathbb{P}\mathcal{ML}(S, z)$ in all of $\mathbb{P}\mathcal{ML}(S, z)$. In particular, for any $\mu \in \mathbb{P}\mathcal{FL}$, there is a sequence of points $x_n \in \mathbb{H}$ limiting $\mu \in \mathbb{P}\mathcal{FL}$.

The systol map $sys : \mathcal{T}(S, z) \rightarrow \mathcal{C}(S, z)$ restricts to a $\pi_1(S)$ -equivariant map from \mathbb{H} to $\mathcal{C}(S, z)$, which is therefore a bounded distance from Φ^v . Again appealing to Klarreich's work [Kla99], it follows that sys extends continuously to $\mathbb{P}\mathcal{FL}(S, z)$, and hence $sys(x_n) \rightarrow |\mu| \in \mathcal{EL}(S, z) \cong \partial_\infty \mathcal{C}(S, z)$. Therefore $\Phi^v(x_n) \rightarrow |\mu|$. Since μ was arbitrary, every point of $\partial_\infty \mathcal{C}(S, z)$ is a limit of a sequence in $\Phi^v(\mathbb{H})$, and we are done. \square

Given an arbitrary sequence $\{x_n\}$ in \mathbb{H} , we need to prove the following.

Proposition 3.12. *If $\lim_{n \rightarrow \infty} x_n = x \in \partial_\infty \mathbb{H} - \mathbb{A}_\infty$, then $\Phi^v(x_n)$ does not converge to a point of $\partial_\infty \mathcal{C}(S, z)$.*

One case of this proposition requires a different proof, and we deal with this now.

Lemma 3.13. *If $\{x_n\}$ and x are as in the proposition and x is the endpoint of a lift of a closed geodesic on S , then $\Phi^v(x_n)$ does not converge to a point of $\partial_\infty \mathcal{C}(S, z)$*

Proof. Under the hypothesis of the lemma, there is an element $\eta \in \pi_1(S)$ with x as the stable fixed point. Moreover, because $x \notin \mathbb{A}_\infty$, the geodesic representative of this element of $\pi_1(S)$ is not filling. Therefore, the associated mapping class is reducible (see [Kra81]).

Let γ_0 be a $\pi_1(S)$ -translate of γ for which the half space $\overline{H^+(\gamma_0)}$ is a neighborhood of x . Then $\{\eta^n(\gamma_0)\}$ nest down on x . It follows that after passing to a subsequence (which we continue to denote $\{x_n\}$) that

$$x_n \in H^+(\eta^n(\gamma_0)) = \eta^n(H^+(\gamma_0)).$$

Appealing to the $\pi_1(S)$ -equivariance of Φ we have

$$\Phi^v(x_n) = \Phi(v, x_n) \in \partial_\infty \mathcal{H}^+(\eta^n(\gamma_0)) = \eta^n(\partial_\infty \mathcal{H}^+(\gamma_0)).$$

Suppose now that $\Phi^v(x_n)$ converges to some element $|\mu| \in \partial_\infty \mathcal{C}(S, z)$. It follows that

$$|\mu| \in \bigcap_{n=1}^{\infty} \eta^n \partial_\infty \mathcal{H}^+(\gamma_0).$$

However, any such $|\mu|$ is invariant under η , and since η is a reducible mapping class, it fixes no point of $\partial_\infty \mathcal{C}(S, z)$. This contradiction implies $\Phi^v(x_n)$ does not converge to any $|\mu| \in \partial_\infty \mathcal{C}(S, z)$, as required. \square

Proof of Proposition 3.12. We switch to the notation $\Phi(v, x_n) = \Phi^v(x_n)$ as this will be more descriptive in what follows. Suppose to the contrary that

$$\lim_{n \rightarrow \infty} \Phi(v, x_n) = |\mu| \in \mathcal{EL}(S, z) \cong \partial_\infty \mathcal{C}(S, z).$$

We begin by finding another sequence which also converges to $|\mu|$ to which we can apply the techniques developed so far. By Lemma 3.13, we may assume that x is not the endpoint of a lift of a closed geodesic on S .

Since $x \notin \mathbb{A}_\infty$, it follows that there is a geodesic multicurve $v_0 \subset \mathcal{C}(S)$ and a ray r in \mathbb{H} ending at x with the property that $r \cap p^{-1}(v_0) = \emptyset$. We choose v_0 so that $p(r)$ satisfies the all tails filling condition in the component $Y \subset S - v_0$ containing $p(r)$. By this we mean that for any simple closed geodesic v' which nontrivially intersects the interior of Y , all tails of $p(r)$ nontrivially intersect v' . We note that it is always possible to find such a v_0 unless r is asymptotic to a component of $p^{-1}(v_1)$ for some geodesic $v_1 \in \mathcal{C}^0(S)$. This would imply that x is the endpoint on $\partial_\infty \widetilde{\mathbb{H}}$ of a lift of a closed geodesic which we have assumed not to be the case. Let \widetilde{Y} be the component of $p^{-1}(Y)$ containing r .

We now pass to a subsequence (which we continue to denote $\{x_n\}$) with the property that for every $k > 0$, the geodesic segment β_k connecting x_{2k} to x_{2k+1} passes within some fixed distance, say distance 1, of r and so that furthermore

$$\beta_k \cap \widetilde{Y} \neq \emptyset.$$

Now fix any $k > 0$ and let $\nu^k : [a_k, b_k] \rightarrow \mathcal{C}(S, z)$ be a geodesic from $\Phi(v, x_{2k})$ to $\Phi(v, x_{2k+1})$ of the form

$$\nu^k = \Phi \circ \widetilde{\nu}^k$$

where $\widetilde{\nu}^k = (\widetilde{\nu}_1^k, \widetilde{\nu}_2^k)$ is given by Lemma 3.10. The path $\widetilde{\nu}_2^k$ lies within 2ϵ of a geodesic passing through x_{2k} and x_{2k+1} and so must also pass within a uniformly bounded distance of r (in fact, it passes within $1 + 2\epsilon$).

Choosing the subsequence $\{x_n\}$ sufficiently spread apart, we may assume that β_k spends a very long time in \widetilde{Y} . Doing this we may also arrange to have $\widetilde{\nu}_2^k$ intersecting \widetilde{Y} . Let $t_k \in [a_k, b_k]$ be such that

$$y_k = \widetilde{\nu}_2^k(t_k) \in \widetilde{Y}$$

and let $v_k = \widetilde{\nu}_1^k(t_k) \in \mathcal{C}^0(S)$ (recall that $\widetilde{\nu}_2^k$ is constant when $\widetilde{\nu}_1^k$ is not, so we can assume that t_k is chosen so that $\widetilde{\nu}^k$ is indeed a vertex).

Now observe that since $\nu^k([a_k, b_k])$ is a geodesic from $\Phi(v, x_{2k})$ to $\Phi(v, x_{2k+1})$, the sequence $\{\nu^k(t_k)\} = \{\Phi(v_k, y_k)\}$ also converges to $|\mu|$. Let us write $u_k = \Phi(v_k, y_k)$.

Next, for each $k > 0$ let $f_k \in \text{Diff}_0(S)$ be such that $\widetilde{e\nu}(f_k) = y_k \in \widetilde{Y}$. Since \widetilde{Y} is a single component of $p^{-1}(Y)$, we may assume that any two f_j and f_k differ by an isotopy fixing the complement of the interior of Y . Said differently, they all differ by such an isotopy of f_1 , and so can write $f_t \in \text{Diff}_0(S)$ for $t \in [1, \infty)$ with f_t constant outside of the interior of Y and $y_k = \widetilde{e\nu}(f_k)$ for all integers $k \geq 1$.

We consider the subsurface $(X, z) = (f_1^{-1}(Y), z) \subset (S, z)$ and look at the subsurface projections

$$\pi_{(X, z)}(u_k) \in \mathcal{C}'(X, z)$$

into the arc complex $\mathcal{C}'(X, z)$ of (X, z) . We consider the incomplete metric on (X, z) (which we view now as the punctured surface $X - \{z\}$) for which

$f_1 : (X, z) \rightarrow (Y, f_1(z))$ is an isometry where $Y - \{f_1(z)\}$ is given the induced path metric inside of S .

Claim. *The length of some arc of $\pi_{(X,z)}(u_k)$ tends to infinity.*

We complete the proof modulo the Claim. If this happens, then there are infinitely many projections in the set $\{\pi_{(X,z)}(u_k)\}$ which is impossible if $u_k \rightarrow |\mu|$. Here, length means infimum of lengths over the isotopy class of an arc.

Proof of Claim: So, to prove that the length of some arc tends to infinity, first suppose that $\{\pi_Y(v_k)\}$ contains an infinite set. Then there are arcs $\alpha_k \subset \pi_Y(v_k)$ with $\ell_Y(\alpha_k) \rightarrow \infty$. Now $f_k^{-1}(\alpha_k)$ is an arc of $\pi_{(X,z)}(u_k)$ and $\ell_{(X,z)}(f_k^{-1}(\alpha_k)) = \ell_{(Y,f_1(z))}(f_1 f_k^{-1}(\alpha_k))$. However, $f_1 f_k^{-1}$ is the identity outside the interior of Y , in particular it is the identity on the boundary of Y and isotopic (forgetting z) to the identity in Y . So, we have

$$\ell_{(Y,f_1(z))}(f_1 f_k^{-1}(\alpha_k)) \geq \ell_Y(\alpha_k) \rightarrow \infty$$

and hence there is an arc of $\pi_{(X,z)}(u_k)$ with length tending to infinity as required.

We may now suppose that there are only finitely many arcs in the set $\{\pi_Y(v_k)\}$. By passing to a further subsequence if necessary, we may assume that $\pi_Y(v_k)$ is constant equal to the union of arcs $\alpha_1, \dots, \alpha_m \in \mathcal{C}'(Y)$. We fix attention on one arc $\alpha = \alpha_1$. Again, we see that $f_k^{-1}(\alpha)$ is an arc of $\pi_{(X,z)}(u_k)$ and $\ell_{(X,z)}(f_k^{-1}(\alpha)) = \ell_{(Y,f_1(z))}(f_1 f_k^{-1}(\alpha))$ with $f_1 f_k^{-1}$ equal to the identity outside the interior of Y for all t .

Writing $h_t = f_t f_1^{-1}$, we are required to prove that $\ell_{(Y,f_1(z))}(h_k^{-1}(\alpha))$ tends to infinity as $k \rightarrow \infty$. Observe that h_1 is the identity on S and h_t is the identity outside the interior of Y for all $t \in [1, \infty)$. We can lift h_t to \tilde{h}_t so that \tilde{h}_1 is the identity in \mathbb{H} . It follows from the definition of $\tilde{e}\tilde{v}$ that $\tilde{h}_k(\tilde{e}\tilde{v}(f_1)) = y_k$. Thus, \tilde{h}_t is essentially pushing the point $y = \tilde{e}\tilde{v}(f_1) \in \tilde{Y}$ along the ray r (at least, $\tilde{h}_k(y) = y_k$ comes back to within a uniformly bounded distance to r for every positive integer k , though it is not hard to see that we can choose f_t so that \tilde{h}_t always stays a bounded distance from r).

Now $h_t^{-1}(\alpha)$ can be described as applying the isotopy h_t backward to α . Therefore, if we let $\tilde{\alpha}^k$ be the last arc of $p^{-1}(\alpha)$ intersected by the path $\tilde{h}_t(y)$ for $t \in [1, k]$, then we can push $\tilde{\alpha}^k$ backward along the isotopy \tilde{h}_t as t runs from k back to 1, and the result $\tilde{h}_k^{-1}(\tilde{\alpha}^k)$ projects down by p to $h_k^{-1}(\alpha)$; see Figure 7. Moreover, observe that $\ell_{(Y,f_1(z))}(h_k^{-1}(\alpha))$ is at least the sum of the distances from y to the two boundary components of \tilde{Y} containing the end points of $\tilde{\alpha}^k$.

Finally, since r has all tails filling in Y , the distance from y to the boundary components of \tilde{Y} containing the endpoints of $\tilde{\alpha}^k$ must be tending to infinity as $k \rightarrow \infty$ (otherwise, we would find that r is asymptotic to one of the boundary components of \tilde{Y} which contradicts all tails filling in Y). This implies $\ell_{(Y,f_1(z))}(h_k^{-1}(\alpha))$ tends to infinity as $k \rightarrow \infty$.

This is the desired contradiction, and so we cannot have $\lim_{k \rightarrow \infty} u_k = |\mu|$. Therefore, the original sequence $\{x_n\}$ cannot converge to $|\mu|$ and the proof is complete. \square

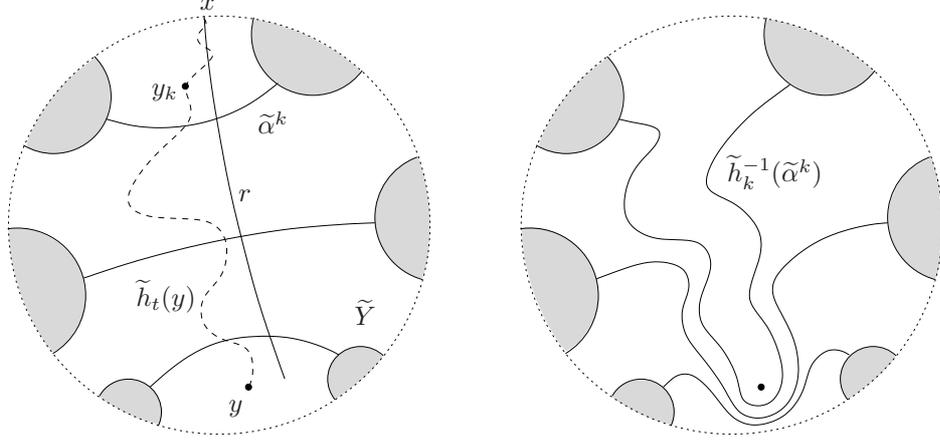


Figure 7: On the left: r inside \tilde{Y} (the complement of the shaded region), the path $\tilde{h}_t(y)$ as it goes through $y_k = \tilde{h}_k(y)$ and the arc $\tilde{\alpha}^k$. On the right: pushing $\tilde{\alpha}^k$ back by \tilde{h}_k^{-1} .

We are now in a position to complete the proof of:

Theorem 1.1 (Universal Cannon–Thurston map). *For any $v \in \mathcal{C}^0(S)$, the map $\Phi^v : \mathbb{H} \rightarrow \mathcal{C}(S, z)$ has a continuous $\pi_1(S)$ -equivariant extension*

$$\bar{\Phi}^v : \mathbb{H} \cup \mathbb{A}_\infty \rightarrow \bar{\mathcal{C}}(S, z).$$

Moreover, $\partial\Phi^v = \bar{\Phi}^v|_{\mathbb{A}_\infty}$ is a quotient map onto $\partial_\infty\mathcal{C}(S, z)$ obtained by identifying the endpoints of each leaf and vertices of each complementary polygon of the lifts of every ending lamination on S .

Proof. The extension and its continuity are given by Theorem 3.5. Next, we need to prove that $\partial\Phi^v$ is surjective. To this end, let $|\mu| \in \partial_\infty\mathcal{C}(S, z)$ be an arbitrary point. According to Lemma 3.11 there exists a sequence $\{x_n\} \subset \mathbb{H}$ with

$$\lim_{n \rightarrow \infty} \Phi^v(x_n) = |\mu|.$$

By passing to a subsequence, we may assume that $\{x_n\}$ converges to a point $x \in \partial_\infty\mathbb{H}$. It follows from Proposition 3.12 that $x \in \mathbb{A}_\infty$. Then, by Theorem 3.5

$$|\mu| = \lim_{n \rightarrow \infty} \Phi^v(x_n) = \bar{\Phi}^v(x) = \partial\Phi^v(x).$$

Since $|\mu|$ was arbitrary, $\partial\Phi^v$ is surjective.

To prove that $\partial\Phi^v$ has the required description as a quotient map, we first observe that Corollary 3.9 shows that two points $x, y \in \mathbb{A}_\infty$ are identified if and only if they are endpoints of a leaf or vertices of a complementary polygon of $p^{-1}(|\lambda|)$ for some $|\lambda| \in \partial_\infty\mathcal{C}(S)$. Therefore, the fibers of $\partial\Phi^v$ are as required for the given description as a quotient map.

The next two lemmas provides a basis for the topology of $\partial_\infty\mathcal{C}(S, z)$ and proves that $\partial\Phi^v$ is indeed a quotient map. \square

To find neighborhood bases for points of $\mathcal{EL}(S, z)$, we must distinguish between two types of points of \mathbb{A}_∞ .

We say a point $x \in \mathbb{A}_\infty$ is *simple* if there exists a ray r in \mathbb{H} ending at x for which $p(r)$ is simple. Otherwise, we say that x is not simple. Equivalently, a point $x \in \mathbb{A}_\infty$ is simple if and only if there is a lamination $|\lambda| \in \mathcal{EL}(S)$ such that x is the endpoint of a leaf of $p^{-1}(|\lambda|)$.

Lemma 3.14. *If $x \in \mathbb{A}_\infty$ is not simple and $\{\gamma_n\}$ are $\pi_1(S)$ -translates of γ which nest down on x , $\{\partial_\infty\mathcal{H}^+(\gamma_n)\}$ is a neighborhood bases for $\bar{\Phi}^v(x)$.*

Proof. We need only show that infinitely many of the $\partial_\infty\mathcal{X}(\gamma_n)$ are pairwise disjoint. If not, then we follow a similar proof to that of Proposition 3.6. \square

This lemma gives neighborhood bases for $\bar{\Phi}^v(x)$ where $x \in \mathbb{A}_\infty$ not a simple point. The next lemma describes a neighborhood basis $\bar{\Phi}^v(x)$, where x is a simple point.

Suppose x_1, x_2 are endpoints of a nonboundary leaf of $p^{-1}(|\lambda|)$ or x_1, \dots, x_k are points of a complementary polygon of some $p^{-1}(|\lambda|)$ for some $|\lambda| \in \mathcal{EL}(S)$. We treat both cases simultaneously referring to these points as x_1, \dots, x_k . Note that $\bar{\Phi}^v(x_1) = \dots = \bar{\Phi}^v(x_k)$, and the $\bar{\Phi}^v$ -image of any simple point has this form.

Lemma 3.15. *If x_1, \dots, x_k are as above, and $\{\gamma_{1,n}\}, \dots, \{\gamma_{k,n}\}$ are sequences of $\pi_1(S)$ -translates of γ with $\{\gamma_{j,n}\}$ nesting down on x_j for each $j = 1, \dots, k$, then*

$$\{\partial_\infty\mathcal{H}^+(\gamma_{1,n}) \cup \dots \cup \partial_\infty\mathcal{H}^+(\gamma_{k,n})\}$$

is a neighborhood basis for $\partial\Phi^v(x_1) = \dots = \partial\Phi^v(x_k)$.

Proof. Let $|\mu| = \partial\Phi^v(x_1) = \dots = \partial\Phi^v(x_k)$. We suppose we have a sequence $\{|\mu_n|\} \subset \mathcal{EL}(S, z)$ converging to $|\mu|$, and we must prove that for every M , there exists $N > 0$ so that for all $n \geq N$,

$$|\mu_n| \in \partial_\infty\mathcal{H}^+(\gamma_{1,M}) \cup \dots \cup \partial_\infty\mathcal{H}^+(\gamma_{k,M}). \quad (5)$$

That is, for every one of these neighborhoods, some tail of a sequence converging to $|\mu|$ lies in this neighborhood. Since $\{\gamma_{1,n}\}, \dots, \{\gamma_{k,n}\}$ nests down on x_1, \dots, x_k , respectively, the intersection of the neighborhoods they define is $|\mu|$, so this will imply that they form a neighborhood basis.

We can find a sequence $\{y_n\} \subset \mathbb{A}_\infty$ so that $\partial\Phi^v(y_n) = |\mu_n|$. We wish to show that any accumulation point of $\{y_n\}$ is one of the points x_1, \dots, x_k . For then, given any M , we can find an $N > 0$ so that for all $n \geq N$

$$y_n \in \overline{H^+(\gamma_{1,M})} \cup \dots \cup \overline{H^+(\gamma_{k,M})}$$

and hence (5) holds.

To this end, we pass to a subsequence so that $y_n \rightarrow x \in \partial_\infty \mathbb{H}$. Choosing sequences converging to y_n for all n and applying a diagonal argument, we see that there is a sequence $\{q_n\} \subset \mathbb{H}$ with $\lim_{n \rightarrow \infty} q_n = x$ and $\lim_{n \rightarrow \infty} \Phi^v(q_n) = |\mu|$. By Proposition 3.12, $x \in \mathbb{A}_\infty$.

Now, either x is one of the points x_1, \dots, x_k or else the geodesic ϵ_i from x to x_j has $p(\epsilon_j)$ nonsimple for all j . In the latter situation, Proposition 3.6 guarantees $\pi_1(S)$ -translates $\gamma_x, \gamma_{1,M}, \dots, \gamma_{k,M}$ of γ defining neighborhoods $\overline{H^+(\gamma_x)}, \overline{H^+(\gamma_{1,M})}, \dots, \overline{H^+(\gamma_{k,M})}$ of x, x_1, \dots, x_k , respectively for which

$$\partial_\infty \mathcal{H}^+(\gamma_x) \cap \partial_\infty \mathcal{H}^+(\gamma_{j,M}) = \emptyset$$

for all $j = 1, \dots, k$. Since $\partial\Phi^v$ is continuous, we have $\Phi^v(y_n) = |\mu_n| \rightarrow |\mu| = \Phi^v(x) \in \partial_\infty \mathcal{H}^+(\gamma_x)$, which is impossible since $|\mu| \in \partial_\infty \mathcal{H}^+(\gamma_{j,M})$ for all $j = 1, \dots, k$. Therefore, $x = x_j$ for some j , and we are done. \square

As noted, these lemmas imply that $\overline{\Phi^v}$ is a quotient map. To see this, we need only show that $E \subset \partial_\infty \mathcal{C}(S, z)$ is closed if and only if $F = (\overline{\Phi^v})^{-1}(E)$ is closed. Since $\overline{\Phi^v}$ is continuous, it follows that E closed implies F closed. Suppose that F is closed. To show that E is closed, we let $|\mu_n| \rightarrow |\mu|$ and must check that $|\mu| \in E$. By Lemmas 3.14 and 3.15, there is a sequence $\{\gamma_n\}$ nesting down on some point $x \in (\overline{\Phi^v})^{-1}(|\mu|)$ with $|\mu_n| \in \partial_\infty \mathcal{H}^+(\gamma_n)$. Let $x_n \in (\overline{\Phi^v})^{-1}(|\mu_n|) \subset F$ be such that $x_n \in H^+(\gamma_n)$. It follows that $x_n \rightarrow x$, so since F is closed, $x \in F$. Therefore, $\overline{\Phi^v}(x) = |\mu| \in E$, as required.

4 Local path connectivity

The following, together with Lemma 3.15 will easily prove Theorem 1.2.

Lemma 4.1. $\partial_\infty \mathcal{H}^+(\gamma)$ is path connected.

Proof. Fix any $|\lambda| \in \mathcal{E}\mathcal{L}(S)$. According to Proposition 2.10, $\hat{\Phi}$ is continuous, so we have a path connected subset

$$\hat{\Phi}(\{|\lambda|\} \times H^+(\gamma)) \subset \partial_\infty \mathcal{H}^+(\gamma).$$

Now let $|\mu| \in \partial_\infty \mathcal{H}^+(\gamma)$ be any point. We will construct a path in $\partial_\infty \mathcal{H}^+(\gamma)$ connecting some point of $\hat{\Phi}(\{|\lambda|\} \times H^+(\gamma))$ to $|\mu|$. This will suffice to prove the lemma.

According to Theorem 1.1 there exists $x \in \mathbb{A}_\infty$ so that for any $v \in \mathcal{C}(S)$, $\overline{\Phi^v}(x) = |\mu|$. Let $r : [0, 1) \rightarrow H^+(\gamma)$ be a ray with

$$\lim_{t \rightarrow 1} r(t) = x.$$

Let $\{\gamma_n\}$ be a sequence of $\pi_1(S)$ -translates of γ which nest down on x . We assume, as we may, that $\gamma_1 = \gamma$. Therefore, there is a sequence $t_1 < t_2 < \dots$ with $\lim_{n \rightarrow \infty} t_n = 1$ and

$$r([t_n, 1]) \subset H^+(\gamma_n)$$

and hence again by Proposition 2.10

$$\hat{\Phi}(\{|\lambda|\} \times r([t_n, 1])) \subset \hat{\Phi}(\{|\lambda|\} \times H^+(\gamma_n)) \subset \partial_\infty \mathcal{H}^+(\gamma_n)$$

Recall that by definition, $\bar{\Phi}^v(x)$ is the unique point of intersection

$$\bigcap_{n=1}^{\infty} \partial_\infty \mathcal{H}^+(\gamma_n),$$

and hence

$$\lim_{t \rightarrow 1} \hat{\Phi}(|\lambda|, r(t)) = |\mu|.$$

Therefore, we can extend $R_{|\lambda|}(t) = \hat{\Phi}(|\lambda|, r(t))$ to a continuous map

$$R_{|\lambda|} : [0, 1] \rightarrow \partial_\infty \mathcal{H}^+(\gamma)$$

with $R_{|\lambda|}(0) \in \hat{\Phi}(\{|\lambda|\} \times H^+(\gamma))$ and $R_{|\lambda|}(1) = |\mu|$. This is the required path completing the proof. \square

We now prove

Theorem 1.2. *The Gromov boundary $\partial_\infty \mathcal{C}(S, z)$ is path connected and locally path connected.*

Proof. From Lemma 4.1, we see that every set of the form $\partial_\infty \mathcal{H}^+(\gamma_0)$ is path connected for any $\pi_1(S)$ -translate γ_0 of γ . Since there is a bases for the topology consisting of these sets, and finite unions of these sets which all share a point by Lemma 3.15, this proves local path connectivity. Path connectivity follows from Lemma 4.1, the fact that $\partial_\infty \mathcal{C}(S, z) = \partial_\infty \mathcal{H}^+(\gamma) \cup \partial_\infty \mathcal{H}^-(\gamma)$, and Proposition 3.7. \square

References

- [Bir69] Joan S. Birman. Mapping class groups and their relationship to braid groups. *Comm. Pure Appl. Math.*, 22, pages 213–238, 1969.
- [Bir74] Joan S. Birman. Braids, links, and mapping class groups. *Princeton University Press, Princeton, N.J., Annals of Mathematics Studies, No. 82.*, 1974.
- [Bon88] Francis Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.

- [BS85] Joan S. Birman and Caroline Series. Geodesics with bounded intersection number on surfaces are sparsely distributed. *Topology*, 24(2):217–225, 1985.
- [CB87] A. Casson and S. Bleiler. Automorphisms of Surfaces after Nielsen and Thurston. *London Math. Soc. Student Texts*, 1987.
- [CEG87] R. D. Canary, D. B. A. Epstein, and P. Green. Notes on Notes of Thurston. in *Analytical and Geometric Aspects of Hyperbolic Spaces*, pages 3–92, 1987.
- [CT07] James W. Cannon and William P. Thurston. Group invariant Peano curves. *Geom. Topol.*, 11:1315–1355, 2007.
- [EE69] Clifford J. Earle and James Eells. A fibre bundle description of Teichmüller theory. *J. Differential Geometry*, 3:19–43, 1969.
- [GL07] Vincent Guirardel and Gilbert Levitt. Deformation spaces of trees. *Groups Geom. Dyn.*, 1(2):135–181, 2007.
- [Har81] W. J. Harvey. Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume 97 of *Ann. of Math. Stud.*, pages 245–251, Princeton, N.J., 1981. Princeton Univ. Press.
- [Kla99] E. Klarreich. The boundary at infinity of the curve complex and the relative teichmüller space. *preprint*, <http://nasw.org/users/klarreich/research.htm>, 1999.
- [KLS06] R. Kent, C. Leininger, and S. Schleimer. Trees and Mapping Class Groups. *arXiv:math/0611241*, preprint 2006.
- [Kra81] Irwin Kra. On the Nielsen-Thurston-Bers type of some self-maps of Riemann surfaces. *Acta Math.*, 146(3-4):231–270, 1981.
- [LS08] C.J. Leininger and S. Schleimer. Connectivity of the space of ending laminations. *arXiv:0801.3058*, preprint 2008.
- [Min92] Y. N. Minsky. Teichmüller Geodesics and Ends of 3-Manifolds. *Topology*, pages 1–25, 1992.
- [Mit98] Mahan Mitra. Cannon-Thurston Maps for Hyperbolic Group Extensions. *Topology* 37, pages 527–538, 1998.
- [Mj05] Mahan Mj. Cannon-Thurston Maps for Surface Groups I: Amalgamation Geometry and Split Geometry. *preprint*, *arXiv:math.GT/0512539*, 2005.

- [Mj06] Mahan Mj. Cannon-Thurston Maps for Surface Groups II: Split Geometry and the Minsky Model. *preprint, arXiv:math.GT/0607509*, 2006.
- [MM99] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves I: Hyperbolicity. *Invent. Math.* 138, pages 103–139, 1999.
- [MP89] J. D. McCarthy and A. Papadopoulos. Dynamics on Thurston’s Sphere of Projective Measured Foliations. *Comment. Math. Helv.* 64, pages 133–166, 1989.
- [PH92] R. Penner and J. Harer. Combinatorics of train tracks. *Ann. Math. Studies* 125, Princeton University Press, 1992.
- [Thu80] W. P. Thurston. The Geometry and Topology of 3-Manifolds. *Princeton University Notes*, 1980.

Department of Mathematics, University of Illinois, Urbana–Champaign, IL 61801
clein@math.uiuc.edu

School of Mathematical Sciences, RKM Vivekananda University, Belur Math,
WB 711202, India
mahan@rkmvu.ac.in

Department of Mathematics, University of Warwick, Coventry CV4 7AL UK
s.schleimer@warwick.ac.uk