

The universal Cannon–Thurston map and the boundary of the curve complex

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

Abstract. In genus two and higher, the fundamental group of a closed surface acts naturally 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. Using the techniques we have developed, we also show that the boundary of this curve complex is locally path-connected.

Mathematics Subject Classification (2010). 20F67 Primary; Secondary 22E40, 57M50.

Keywords. Mapping class group, curve complex, ending lamination, Cannon–Thurston map.

Contents

1	Introduction	770
1.1	Statement of results	770
1.2	Notation and conventions	772
2	Point position	778
2.1	A bundle over \mathbb{H}	779
2.2	An explicit construction of Φ	780
2.3	A further description of $\mathcal{C}(S, z)$	783
2.4	Extending to measured laminations	784
2.5	Φ and Ψ	791
3	Universal Cannon–Thurston maps	793
3.1	Quasiconvex sets	793
3.2	Rays and existence of Cannon–Thurston maps	796
3.3	Separation	800
3.4	Surjectivity	803
3.5	Neighborhood bases	809
3.6	$\text{Mod}(S, z)$ -equivariance	812
4	Local path-connectivity	813
	References	814

*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 two. This identifies the universal cover with the hyperbolic plane $p: \mathbb{H} \rightarrow S$. Fix a basepoint $z \in S$ and a point $\tilde{z} \in p^{-1}(z)$. This defines an isomorphism between the group $\pi_1(S, z)$ of homotopy classes of loops based at z and the group $\pi_1(S)$ 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 $S - \{z\}$; however 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, with zero skeleta, $\mathcal{C}^0(S)$ and $\mathcal{C}^0(S, z)$, identified with the sets of isotopy classes of essential simple closed curves in S and (S, z) , respectively. Write $\Pi: \mathcal{C}(S, z) \rightarrow \mathcal{C}(S)$ to denote the forgetful projection. See Section 1.2.4.

From [KLS09], the fiber over $v \in \mathcal{C}^0(S)$ is $\pi_1(S)$ -equivariantly isomorphic to the Bass–Serre tree T_v determined by v . The action of $\pi_1(S)$ on $\mathcal{C}(S, z)$ comes from the inclusion into the mapping class group $\text{Mod}(S, z)$ via the Birman exact sequence; see Section 1.2.3. We define a map

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

by sending $\{v\} \times \mathbb{H}$ to $T_v \cong \Pi^{-1}(v) \subset \mathcal{C}(S, z)$ in a $\pi_1(S)$ -equivariant fashion and 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 $\mathbb{H} \cong \{v\} \times \mathbb{H}$:

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

Suppose that $r \subset \mathbb{H}$ is a geodesic ray that eventually lies in the preimage of some proper essential subsurface of S . We prove in Section 3 that $\Phi_v(r) \subset \mathcal{C}(S, z)$ has finite diameter. The remaining rays define a subset $\mathbb{A} \subset \partial\mathbb{H}$ (of full Lebesgue measure); see Section 3.2. Our first result 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 unique continuous $\pi_1(S)$ -equivariant extension*

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

The map $\partial\Phi = \bar{\Phi}_v|_{\mathbb{A}}$ does not depend on v and is a quotient map onto $\partial\mathcal{C}(S, z)$. Given distinct points $x, y \in \mathbb{A}$, $\partial\Phi(x) = \partial\Phi(y)$ if and only if x and y are ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of the lift of an ending lamination on S .

We recall that a Cannon–Thurston map was constructed by Cannon and Thurston [CT07] for the fiber subgroup of the fundamental group of a closed hyperbolic 3-manifold fibering over the circle. The construction was then extended to simply

degenerate, bounded geometry Kleinian closed surface groups by Minsky [Min92], and to the general simply degenerate case by the second author [Mj06]. In all these cases, one produces a quotient map from the circle $\partial\mathbb{H}$ onto the limit set of the Kleinian group Γ . In the quotient, distinct points are identified if and only if they are ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of the lift of an ending lamination for Γ . This is either one or two ending laminations depending on whether the group is singly or doubly degenerate; see [Mj07].

In a similar fashion, the second author [Mit97] has constructed a Cannon–Thurston map for any δ -hyperbolic extension Γ of a group G by $\pi_1(S)$,

$$1 \rightarrow \pi_1(S) \rightarrow \Gamma \rightarrow G \rightarrow 1$$

(for a discussion of such groups see [Mos97], [FM02]). This is a $\pi_1(S)$ -equivariant quotient map from $\partial\mathbb{H}$ onto the Gromov boundary of Γ . As above, the quotient identifies distinct points if and only if they are ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of the lift of an ending lamination for G .

The map $\partial\Phi$ is universal in that distinct points are identified if and only if they are the ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of the lift of *any* ending lamination on S . We remark that the restriction to \mathbb{A} is necessary to get a reasonable quotient: the same quotient applied to the entire circle $\partial\mathbb{H}$ is a non-Hausdorff space.

It follows from the above description of the various Cannon–Thurston maps that the universal property of $\partial\Phi$ can also be rephrased as follows. If $F: \partial\mathbb{H} \rightarrow \Omega$ is any Cannon–Thurston map as above – so, Ω is either the limit set of a Kleinian group, or the Gromov boundary of a hyperbolic extension Γ – then there exists a map

$$\phi_F: F(\mathbb{A}) \rightarrow \partial\mathcal{C}(S, z)$$

so that $\phi_F \circ F|_{\mathbb{A}} = \partial\Phi$. Moreover, because $\partial\Phi$ identifies *precisely* the required points to make this valid, one sees that any $\pi_1(S)$ -equivariant quotient of \mathbb{A} with this universal property is actually a $\pi_1(S)$ -equivariant quotient of $\partial\mathcal{C}(S, z)$.

It is a classical fact, due to Nielsen, that the action of $\pi_1(S)$ on $\partial\mathbb{H}$ extends to the entire mapping class group $\text{Mod}(S, z)$. It will become apparent from the description of \mathbb{A} given below that this $\text{Mod}(S, z)$ action restricts to an action on \mathbb{A} . In fact, we have

Theorem 1.2. *The quotient map*

$$\partial\Phi: \mathbb{A} \rightarrow \partial\mathcal{C}(S, z)$$

is equivariant with respect to the action of $\text{Mod}(S, z)$.

As an application of the techniques we have developed, we also prove the following.

Theorem 1.3. *The Gromov boundary $\partial\mathcal{C}(S, z)$ 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 1.3 does not follow immediately from Theorem 1.1.

This strengthens the work of the first and third authors in [LS09] in a special case: in [LS09] it was shown that the boundary of the curve complex is connected for surfaces of genus at least 2 with any positive number of punctures and for 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 [LS09]. These results seem to be the first ones providing some information about the *topology* of the boundary of the curve complex. The question of connectivity of the boundary was posed by Storm, and the general problem of understanding its topology was posed by Minsky in his 2006 ICM address. Gabai [Gab09] has now given a proof of Theorem 1.3 for all surfaces Σ for which $\mathcal{C}(\Sigma)$ is nontrivial, except the torus, 1-punctured torus and 4-punctured sphere, where it is known to be false.

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, and the referee for useful suggestions.

1.2. Notation and conventions

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

By a *lamination* on S , we mean a 1-dimensional foliation of a closed subset of S (see e.g. [PH92] and [CC00]). We require that all our laminations be essential, meaning that the leaves lift to quasigeodesics in the universal cover. A *measured lamination* on S is a lamination with a transverse measure of full support. A measured lamination on S will be denoted λ with the support – the underlying lamination – written $|\lambda|$.

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 all 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 metric on S). This is a measured lamination λ for which $|\lambda|$ is a geodesic lamination.

We similarly define measured laminations on (S, z) as compactly supported measured laminations on $S - \{z\}$. In the situations that we will be considering, these will generally not arise 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 defines 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 is the space of *filling laminations* which we denote \mathcal{FL} . These are the measure classes of measured geodesic laminations λ for which all complementary regions of its support $|\lambda|$ are disks (for $S - \{z\}$, there is also a single punctured disk). The quotient of \mathcal{FL} by forgetting the measures will be denoted \mathcal{EL} and is the space of *ending laminations*. For notational simplicity, the image of $\lambda \in \mathcal{FL}$ in \mathcal{EL} will also be denoted by $|\lambda|$.

When there is no confusion, we will let λ simultaneously represent a measured lamination as well as the equivalence class in either \mathcal{ML} or $\mathbb{P}\mathcal{ML}$ which it determines.

Remark 1.4. Because there is a unique measured geodesic lamination in any equivalence class the reader can, unless otherwise stated, consider an element of \mathcal{ML} as a measured geodesic lamination. In Sections 2.4 and 2.5 it will be preferable to choose a slightly different representative for an arbitrary element of $\mathcal{ML}(S)$, though elements of $\mathcal{FL}(S)$ will still be chosen to be measured geodesic laminations. Beginning in Section 3, only $\mathcal{FL}(S)$ and $\mathcal{EL}(S)$ are relevant, and so again, measured geodesic laminations will suffice.

1.2.2. Train tracks. 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 (for (S, z) , we replace S here by $S - \{z\}$). 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 of the weights on the outgoing branches. Conversely, any assignment of nonnegative weights to the branches of a train track satisfying the switch condition uniquely determines an element of \mathcal{ML} ; see [PH92], Theorem 2.7.4. Given a train track τ carrying λ , we write $\tau(\lambda)$ to denote the train track τ together with the weights defined by λ .

Proposition 1.5. *Suppose that $\{\lambda_n\}_{n=1}^\infty \cup \{\lambda\} \subset \mathcal{ML}$ are all carried by the train track 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 .*

Proof. This is an immediate consequence of [PH92], Theorem 2.7.4. □

There is a well-known construction of train tracks carrying a given lamination which will be useful for us. For a careful discussion, see [PH92], Theorem 1.6.5, or Section 4 of [Bro00]. Starting with a geodesic lamination \mathcal{L} one chooses $\epsilon > 0$ very small and constructs a foliation, transverse to \mathcal{L} , of the ϵ -neighborhood $N_\epsilon(\mathcal{L})$. The leaves of this foliation are arcs called *ties*. Taking the quotient by collapsing each tie to a point produces a train track τ on S ; see Figure 1.

We can view $N_\epsilon(\mathcal{L})$ as being built from finitely many rectangles, each foliated by ties, 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 $N_\epsilon(\mathcal{L})$, transverse to the foliation by ties, and the branch β_R is contained in the rectangle 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 , defined by λ , is given by $\lambda(a) = \int_a d\lambda$; see Figure 1.

1.2.3. Mapping class groups. Recall that we have fixed a hyperbolic structure on S as well as a locally isometric universal covering $p: \mathbb{H} \rightarrow S$. We also have a basepoint $\tilde{z} \in p^{-1}(z)$ determining an isomorphism from $\pi_1(S)$, the covering group of p , 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

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




Figure 1. The train track τ is obtained by collapsing the ties of $N_\epsilon(\mathcal{L})$. The lamination λ defines weights on τ : the weight on β_R is $\lambda(a)$, where a is a tie in the rectangle R .

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

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

We elaborate 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 groups 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 an isotopy 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 $\sigma(t) = h_t(z)$ and $\sigma'(t) = h'_t(z)$. 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 σ then σ' , while $H_0 = h \circ h'$ and $H_1 = \text{Id}_S$. So, $\text{ev}_*(h \circ h') = \sigma\sigma'$, 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 $\text{Mod}(S, z)$.

1.2.4. 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 $S - \{z\}$. We write $\mathcal{C}^0(S)$ and $\mathcal{C}^0(S, z)$ to denote the sets of isotopy classes of essential simple closed curves on S and (S, z) , respectively.

The *curve complex* of S is the simplicial complex $\mathcal{C}(S)$ with vertex set $\mathcal{C}^0(S)$ so that $k + 1$ distinct curves v_0, \dots, v_k span a k -simplex if and only if the isotopy classes can be realized disjointly in S ; see [Har81] and [MM99]. The curve complex of (S, z) is denoted $\mathcal{C}(S, z)$ and is defined similarly. A simplex $\{v_0, \dots, v_k\}$ in $\mathcal{C}(S)$ or $\mathcal{C}(S, z)$ naturally determines an isotopy class of *multi-curve*, which is the union of the pairwise disjoint representatives of v_0, \dots, v_k .

We consider $\mathcal{C}(S)$ and $\mathcal{C}(S, z)$ as geodesic metric spaces so that each simplex is isometric to a regular Euclidean simplex with all edge lengths equal to one. The following is proven in [MM99].

Theorem 1.6 (Masur–Minsky). *The spaces $\mathcal{C}(S)$ and $\mathcal{C}(S, z)$ are δ -hyperbolic for some $\delta > 0$.*

Given a simplex $v \subset \mathcal{C}(S)$ or $u \subset \mathcal{C}(S, z)$ we will not distinguish between this simplex and 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 gives a well-defined “forgetful” map

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

which is simplicial.

Unless otherwise stated we assume that a multicurve $v \subset \mathcal{C}(S)$ is realized by its geodesic representative in S (isotopy classes can be realized disjointly if and only if the geodesic representatives are disjoint; see [CB87]). 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 [KLS09].

Theorem 1.7 (Kent–Leininger–Schleimer). *The fiber of Π over a point $x \in \mathcal{C}(S)$ is $\pi_1(S)$ -equivariantly homeomorphic to the tree T_v , where v is the unique simplex containing x in its interior. \square*

1.2.5. 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 $\partial\mathcal{C} \cong \mathcal{EL}$. 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 a consequence of Klarreich’s work [Kla99], stated using our terminology.

Proposition 1.8 (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 $\{v_n\} \subset \mathcal{C}$ converges to $|\lambda| \in \mathcal{EL}$ 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 a sequence $\{v_n\}$ 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 a sequence $\{v_n\}$ has $|\lambda| = |\lambda'|$ then $\{v_n\}$ converges to $|\lambda|$ in $\bar{\mathcal{C}}$. For this it suffices to assume that $\{v_n\}$ converges in $\mathbb{P}\mathcal{ML}$ to λ' with $|\lambda'| = |\lambda|$.

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 . 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 an accumulation point of X_n in $\mathbb{P}\mathcal{ML}$, then $i(\lambda'', \lambda') = 0$ and so $|\lambda''| = |\lambda'| = |\lambda|$ since λ is filling.

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

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

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

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

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

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

in $\bar{\mathcal{C}}$ and we are done. □

1.2.6. Cannon–Thurston maps. Fix X and Y hyperbolic metric spaces, $F: Y \rightarrow X$ a continuous map, and $Z \subset \partial Y$ a subset of the Gromov boundary. A Z -Cannon–Thurston map is a continuous extension $\bar{F}: Y \cup Z \rightarrow X \cup \partial X$ of F . That is, $\bar{F}|_Y = F$. We will simply call \bar{F} a Cannon–Thurston map when the set Z is clear from the context. We sometimes refer to the restriction $\partial F = \bar{F}|_Z$ as a Cannon–Thurston map.

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

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

Lemma 1.9. *Fix X and Y hyperbolic metric spaces, $F: Y \rightarrow X$ a continuous map and $Z \subset \partial Y$ a subset. Fix a basepoint $x \in X$. Then there is a Z -Cannon–Thurston map $\bar{F}: Y \cup Z \rightarrow X \cup \partial X$ if and only if for every $z \in Z$ there is a neighborhood basis $B_i \subset Y \cup Z$ of z and a collection of uniformly quasiconvex sets $Q_i \subset X$ with $F(B_i \cap Y) \subset Q_i$ and $d_X(x, Q_i) \rightarrow \infty$ as $i \rightarrow \infty$. Moreover,*

$$\bigcap_i \bar{Q}_i = \bigcap_i \partial Q_i = \{\bar{F}(z)\}$$

determines $\bar{F}(z)$ uniquely. □

Remark 1.10. For the purposes of the present work, it is more convenient to use the more flexible definition of neighborhood basis. This is a collection of sets $\{B_j(x)\}_j$ associated to each point x in the space with the property that a subset U is open if and only if for every $x \in U$, $B_j(x) \subset U$ for some j . Equivalently, the interior of each $B_j(x)$ is required to contain x , and any open set containing x should contain some $B_j(x)$.

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 $\bar{\mathcal{C}}(S) \times \mathbb{H}$.

2.1. A bundle over \mathbb{H} . Recall from Section 1.2.3 that $\text{Diff}_B(S, z) = \text{Diff}_0(S) \cap \text{Diff}^+(S, z)$ is the subgroup of diffeomorphisms of (S, z) isotopic to the identity on S after forgetting z . The bundle determining the Birman exact sequence has a subbundle obtained by restricting the evaluation map 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

$$\begin{array}{ccccc}
 & & & & \mathbb{H} \\
 & & & \nearrow \tilde{\text{ev}} & \downarrow p \\
 \text{Diff}_B(S, z) & \longrightarrow & \text{Diff}_0(S) & \xrightarrow{\text{ev}} & S \\
 \uparrow & \nearrow & & & \\
 \text{Diff}_0(S, z) & & & &
 \end{array} \tag{1}$$

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. *The lift*

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

is equivariant with respect to ev_ .*

Proof. We need to prove

$$\text{ev}_*(h)(\tilde{\text{ev}}(f)) = \tilde{\text{ev}}(f \circ h^{-1})$$

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

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

and hence $\tilde{\text{ev}}(f)$ differs from $\tilde{\text{ev}}(f \circ h^{-1})$ by a covering transformation $\sigma \in \pi_1(S)$:

$$\tilde{\text{ev}}(f \circ h^{-1}) = \sigma(\tilde{\text{ev}}(f)).$$

The covering transformation σ appears to depend on both f and h . However if $H_t, t \in [0, 1]$, is a path in $\text{Diff}_B(S, z)$ from $h = H_0$ to $h' = H_1$ then $\tilde{\text{ev}}(f \circ H_t^{-1})$ is constant in t : this can be seen from the above description of $\tilde{\text{ev}}$ 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 $\tilde{\text{ev}}$ and connectivity of $\text{Diff}_0(S)$ 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

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

So if $h_t, t \in [0, 1]$, is a path in $\text{Diff}_0(S)$ from $h = h_0$ to $\text{Id}_S = h_1$, then since $\text{ev}_*(h) = \sigma_h$ 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 $h_t^{-1}(z), t \in [0, 1]$; see Section 1.2.3.

Now observe that $\tilde{\text{ev}}(h_t^{-1}), t \in [0, 1]$, 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 $\sigma = \sigma_h = \text{ev}_*(h)$. \square

2.2. An explicit construction of Φ . We are now ready to give an explicit description of the map Φ . We will first define 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 $\Phi: \mathcal{C}(S) \times \mathbb{H}^2 \rightarrow \mathcal{C}(S, z)$ with $\tilde{\Phi} = \Phi \circ \text{Id} \times \tilde{\text{ev}}$.

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 in mind when trying to make sense of the actual definition of $\tilde{\Phi}$.

To carry out the construction of $\tilde{\Phi}$, we now choose $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)} \subset \mathbb{R}_+$ so that $N(v) = N_{\epsilon(v)}(v)$, the $\epsilon(v)$ -neighborhood of v , has the following properties:

- $N(v) \cong \mathcal{S}^1 \times [0, 1]$ and
- $N(v) \cap N(v') = \emptyset$ if $v \cap v' = \emptyset$.

Let $N^\circ(v)$ denote the interior of $N(v)$ and v^\pm denote the boundary components of $N(v)$.

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

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

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 $f(z) \notin N^\circ(v)$, then we set

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

as suggested above.

If $f(z) \in N^\circ(v)$, then $f^{-1}(v^+)$ and $f^{-1}(v^-)$ 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 the two boundary components v^+ and v^- . Specifically, set

$$t = \frac{d(v^+, f(z))}{2\epsilon(v)},$$

where $d(v^+, f(z))$ is the distance inside $N(v)$ from $f(z)$ to v^+ , and define

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

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

In general, for a point $(x, f) \in \mathcal{C}(S) \times \text{Diff}_0(S)$ with $x = \sum_j s_j v_j \in v = \{v_0, \dots, v_k\}$ we define $\tilde{\Phi}(x, f)$ as follows. As before, if $f(z) \notin \cup_j N^\circ(v_j)$, then define

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

Otherwise, $f(z) \in N^\circ(v_i)$ for exactly one $i \in \{0, \dots, k\}$. Set

$$t = \frac{d(v_i^+, f(z))}{2\epsilon(v_i)}$$

as above, and define

$$\tilde{\Phi}(x, f) = s_i (t f^{-1}(v_i^+) + (1 - t) f^{-1}(v_i^-)) + \sum_{j \neq i} s_j f^{-1}(v_j).$$

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 $f(z) \notin N^\circ(v)$ we have

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

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*

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

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

Here the action of $\pi_1(S)$ 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 $\tilde{\text{ev}}(f_0) = \tilde{\text{ev}}(f_1)$ and must show $\tilde{\Phi}(x, f_0) = \tilde{\Phi}(x, f_1)$.

Appealing to diagram (1) 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) .

Recall that the dependence of $\tilde{\Phi}(x, f)$ on f was only via certain isotopy classes $f^{-1}(\alpha)$ and a single distance $d(v^+, f(z))$. Since these data are the same for f_0 and f_1 , 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 $\text{Id}_{\mathcal{C}(S)} \times \tilde{\text{ev}}$ is equivariant with respect to ev_* . Thus, since $\tilde{\Phi}$ is equivariant, so is Φ . □

Proposition 2.3. *Given $x \in \mathcal{C}(S)$, let $v \subset \mathcal{C}(S)$ be the simplex containing x in its interior. Then the restriction*

$$\Phi_x = \Phi|_{\{x\} \times \mathbb{H}} : \mathbb{H} \rightarrow \Pi^{-1}(x)$$

is obtained by first projecting to T_v , then composing with the equivariant homeomorphism $T_v \cong \Pi^{-1}(x)$ from Theorem 1.7.

Proof. Fix $x \in \mathcal{C}(S)$, the simplex $v = \{v_0, \dots, v_k\} \subset \mathcal{C}(S)$ 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 $N(v_i)$ determine a map from \mathbb{H} to the Bass–Serre tree T_v associated to v as follows. We collapse each component U of the preimage $p^{-1}(N(v_i))$ onto an interval, say $[0, 1]$, by the projection defined as the distance to the component of $p^{-1}(v_i^+)$ meeting U , multiplied by $1/(2\epsilon(v_i))$. If we further collapse each component of the complement of

$$p^{-1}(N(v_0) \cup \dots \cup N(v_k))$$

to a point, the quotient space is precisely T_v .

The map 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 [KLS09], the edge and vertex stabilizers in the domain and range agree, and in fact this map is the homeomorphism given by Theorem 1.7, as required. □

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 1.7. 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

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

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

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

for every inclusion of faces $\phi: v'' \rightarrow v'$ 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 2.4. *The curve complex $\mathcal{C}(S, z)$ is $\pi_1(S)$ -equivariantly homeomorphic to*

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

Here 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 2.10. 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. Given any element $\lambda \in \mathcal{ML}(S)$, we suppose λ also denotes the measured geodesic lamination representing it. We then produce an element $\hat{\lambda}$ measure equivalent to λ by replacing all simple closed geodesic components

of the support $|\lambda|$ with appropriately chosen foliated annuli. We now explain this more precisely and set some notation.

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

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

where $|\text{Cur}(\lambda)|$ is the union of all simple closed geodesics in $|\lambda|$ and $|\text{Min}(\lambda)| = |\lambda| - |\text{Cur}(\lambda)|$. We construct a measured lamination $\hat{\lambda}$ measure equivalent to λ by taking

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

where $\text{Min}(\hat{\lambda}) = \text{Min}(\lambda)$ and $\text{Ann}(\hat{\lambda})$ is a measured lamination whose support is a foliation on annular neighborhoods of $|\text{Cur}(\lambda)|$ defined as follows.

The sublamination $\text{Cur}(\lambda)$ can be further decomposed as $\text{Cur}(\lambda) = \sum_j t_j v_j$, where $t_j v_j$ means t_j times the transverse counting measure on the simple closed geodesic component v_j of $|\text{Cur}(\lambda)|$. Then $|\text{Ann}(\hat{\lambda})|$ is the disjoint union $\cup_j N(v_j)$, with each $N(v_j)$ given the foliation by curves equidistant to v_j . This foliation of $N(v_j)$ is assigned the transverse measure which is $t_j / (2\epsilon(v_j))$ times the distance between leaves, and $\text{Ann}(\hat{\lambda})$ is the sum of these measured laminations; see Figure 2 for a cartoon depiction of λ and $\hat{\lambda}$. Choosing $\{\epsilon(v)\}$ sufficiently small it follows that $|\text{Ann}(\hat{\lambda})| \cap |\text{Min}(\hat{\lambda})| = \emptyset$ for all λ .

For future use, if $\text{Cur}(\lambda) = \sum_j t_j v_j$, then we define

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

If $|\text{Cur}(\lambda)| = \emptyset$ we set $T(\lambda) = T(\hat{\lambda}) = 0$.

For any measured geodesic lamination λ , by construction we have $|\lambda| \subset |\hat{\lambda}|$, meaning that as subsets of S , $|\lambda|$ is a subset of $|\hat{\lambda}|$, and that each leaf of $|\lambda|$ is a leaf of $|\hat{\lambda}|$. The difference between the total variations assigned an arc by λ and $\hat{\lambda}$ is estimated by the following.

Lemma 2.5. *If a is any arc transverse to $|\hat{\lambda}|$, then it is also transverse to $|\lambda|$ and we have*

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

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

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

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




Figure 2. Removing simple closed geodesics and inserting foliated annuli.

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

The following is also useful.

Lemma 2.6. *Suppose $\lambda_n \rightarrow \lambda$ in $\mathcal{ML}(S)$ with $\lambda \in \mathcal{FL}(S)$. Further suppose that $|\hat{\lambda}_n|$ converges in the Hausdorff topology on closed subsets of S to a set \mathcal{L} . Then \mathcal{L} is a geodesic lamination containing the geodesic lamination $|\lambda|$.*

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]), as is the fact that any Hausdorff limit of $\{|\lambda_n|\}$ contains $|\lambda|$ (compare [Thu80], Proposition 8.10.3).

Since $\lambda_n \rightarrow \lambda$ and $\lambda \in \mathcal{FL}$, it follows that no simple closed geodesic occurs infinitely often in $\{|\text{Cur}(\lambda_n)|\}$. Further note that if $\{v_n\}$ is any sequence of distinct simple closed geodesics in S , then their lengths tend to infinity and hence $\epsilon(v_n) \rightarrow 0$. Therefore, the Hausdorff distance between $|\lambda_n|$ and $|\hat{\lambda}_n|$ tends to zero, and so the Hausdorff limits of $|\lambda_n|$ and $|\hat{\lambda}_n|$ are the same. As above, we see that \mathcal{L} is a geodesic lamination containing $|\lambda|$. \square

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

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

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

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 1.2.2. Suppose that $\hat{\lambda}$ is a measured lamination on S for which $|\hat{\lambda}|$ 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 $f^{-1}(\tau(\hat{\lambda}))$.

If $f(z) \in N_\epsilon(\mathcal{L})$ then by a small perturbation of ϵ we may assume that $f(z)$ does not lie on a the boundary 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 . Insert two branches β_R^u and β_R^d creating a bigon containing $f(z)$; see Figure 3. The result, denoted τ' , is a train track on $(S, f(z))$.




Figure 3. Modifying τ to τ' .

If $f_t \in \text{Diff}_0(S)$ is an isotopy with $f = f_0$ and $f_t(z) \in \text{int}(R)$ 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 .

The measured lamination $\hat{\lambda}$ 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 $\hat{\lambda}(a) = \hat{\lambda}(a^u) + \hat{\lambda}(a^d)$; see Figure 4. The lamination $f^{-1}(\hat{\lambda})$ is the lamination determined by the weighted train track $f^{-1}(\tau'(\hat{\lambda}))$.




Figure 4. Weights on τ' determined by $\hat{\lambda}$ and $f(z)$.

The proof of the following is similar to that of Proposition 2.2 and we omit it.

Proposition 2.7. $\tilde{\Psi}$ descends to a $\pi_1(S)$ -equivariant map Ψ :

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

□

Remark 2.8. Our construction of Ψ depends on our choice of representation of elements in \mathcal{ML} . An alternative would be to realize each element of \mathcal{ML} by a measured foliation. Indeed, the hyperbolic metric determines a complex structure, and for each element $\lambda \in \mathcal{ML}$ there is a unique quadratic differential for which the vertical foliation realizes λ ; see [HM79]. Our choices are more compatible with the use of hyperbolic geometry in Sections 3 and 4.

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 can break down. In particular, we have the following.

Proposition 2.9. The map $\tilde{\Psi}$ is continuous on $\mathcal{FL}(S) \times \text{Diff}_0(S)$.

Proof. We will show that for any sequence $\{(\lambda_n, f_n)\}$ in $\mathcal{ML}(S) \times \text{Diff}_0(S)$ converging to $(\lambda, f) \in \mathcal{FL}(S) \times \text{Diff}_0(S)$ there is a subsequence for which $\{\tilde{\Psi}(\lambda_{n_k}, f_{n_k})\}$ converges to $\tilde{\Psi}(\lambda, f)$. Since we will find such a subsequence for *any* sequence converging to (λ, f) , continuity of $\tilde{\Psi}$ at (λ, f) will follow.

We begin by passing to a subsequence for which the supports $\{|\hat{\lambda}_n|\}$ converge in the Hausdorff topology to a closed set \mathcal{L} . It follows from Lemma 2.6, that \mathcal{L} is a geodesic lamination containing $|\lambda| = |\hat{\lambda}|$.

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 $n \geq N$, $|\hat{\lambda}_n| \subset N_\epsilon(\mathcal{L})$. Therefore, for all $n \geq N$, $\hat{\lambda}$ and $\hat{\lambda}_n$ determine weighted train tracks $\tau(\hat{\lambda})$ and $\tau(\hat{\lambda}_n)$, respectively. Since $\hat{\lambda}_n \rightarrow \hat{\lambda}$, it follows that $\tau(\hat{\lambda}_n) \rightarrow \tau(\hat{\lambda})$ as $n \rightarrow \infty$.

Since f_n is isotopic to f by an isotopy keeping the image of z in $N_\epsilon(f(z))$, it follows that $f^{-1}(\tau) = f_n^{-1}(\tau)$, up to isotopy. Therefore, $f^{-1}(\tau(\hat{\lambda}_n))$ and $f_n^{-1}(\tau(\hat{\lambda}_n))$ are isotopic and so we have convergence of weights $f^{-1}(\tau(\hat{\lambda}_n)) \rightarrow f^{-1}(\tau(\hat{\lambda}))$ 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 choose $\epsilon > 0$ sufficiently small so that the quotient of $N_\epsilon(\mathcal{L})$ by collapsing ties is a train track τ , so that $f(z)$ lies in the interior of some rectangle R of $N_\epsilon(\mathcal{L})$ and so that τ is trivalent.

Let $N > 0$ be such 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 $\tau(\hat{\lambda}_n)$ converge to $\tau(\hat{\lambda})$. Therefore, to prove that the weighted train tracks $f^{-1}(\tau'_n(\hat{\lambda}_n)) = f_n^{-1}(\tau'_n(\hat{\lambda}_n))$ converge to $f^{-1}(\tau'(\hat{\lambda}))$, it suffices to prove that the weights assigned to $f^{-1}(\beta_R^u)$ and $f^{-1}(\beta_R^d)$ by $\hat{\lambda}_n$ converge to the weights assigned to these branches by $\hat{\lambda}$. This is sufficient 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 the $\hat{\lambda}_n$ converge to the weight defined by $\hat{\lambda}$. 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 of one of these, say, $f^{-1}(\beta_R^u)$.

To define the required weights, first recall that we have the 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 $a_n^u \cap a_n^d = \{f_n(z)\}$. 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 $\hat{\lambda}_n$ and $\hat{\lambda}$ are given by

$$\hat{\lambda}_n(a_n^u) \quad \text{and} \quad \hat{\lambda}(a^u),$$

respectively.

Therefore, we must verify that $\hat{\lambda}_n(a_n^u) \rightarrow \hat{\lambda}(a^u)$. However, since $T(\hat{\lambda}_n) \rightarrow 0$ as $k \rightarrow \infty$, Lemma 2.5 implies that it suffices to prove $\lambda_n(a_n^u) \rightarrow \lambda(a^u)$.

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

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

so that

$$\lambda(a_-^u) \leq \lambda(a^u) \leq \lambda(a_+^u) \quad \text{and} \quad \lambda(a_+^u) - \lambda(a_-^u) < \delta.$$

Since $\lambda_n \rightarrow \lambda$, it follows that we also have

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

and

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

Furthermore, since $a_n \rightarrow a$ and $a_n^u \rightarrow a^u$ in the C^1 -topology, we see that

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

and

$$\liminf_{n \rightarrow \infty} \lambda_n(a_n^u) \geq \lim_{n \rightarrow \infty} \lambda_n(a_-^u) = \lambda(a_-^u).$$

Since $\liminf \lambda_n(a_n^u) \leq \limsup \lambda_n(a_n^u)$, combining all of the above, we obtain

$$\left| \limsup_{n \rightarrow \infty} \lambda_n(a_n^u) - \lambda(a^u) \right| + \left| \liminf_{n \rightarrow \infty} \lambda_n(a_n^u) - \lambda(a^u) \right| < 2\delta.$$

As δ was arbitrary, it follows that

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

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 2.10. *The map Ψ is continuous on $\mathcal{FL}(S) \times \mathbb{H}$.*

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

2.5. Φ and Ψ . The map Ψ descends to a map $\mathbb{P}\mathcal{ML}(S) \times \mathbb{H} \rightarrow \mathbb{P}\mathcal{ML}(S, z)$ in the obvious way. We denote this map by Ψ with the meaning determined by context.

We let $\Psi_{\mathcal{C}}$ denote the restriction of Ψ to $\mathbb{P}\mathcal{ML}_{\mathcal{C}}(S) \times \mathbb{H}$. The map $\Psi_{\mathcal{C}}$ has image $\mathbb{P}\mathcal{ML}_{\mathcal{C}}(S, z)$.

Lemma 2.11. *The following diagram commutes:*

$$\begin{array}{ccc} \mathbb{P}\mathcal{ML}_{\mathcal{C}}(S) \times \mathbb{H} & \xrightarrow{\Psi_{\mathcal{C}}} & \mathbb{P}\mathcal{ML}_{\mathcal{C}}(S, z) \\ \downarrow & & \downarrow \\ \mathcal{C}(S) \times \mathbb{H} & \xrightarrow{\Phi} & \mathcal{C}(S, z). \end{array}$$

The vertical arrows here are the natural maps.

Proof. The two maps are defined identically on $\mathcal{ML}_{\mathcal{C}}(S)$ as there $\lambda = \text{Cur}(\lambda)$. \square

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

Proposition 2.12. *There is a continuous $\pi_1(S)$ -equivariant extension $\hat{\Phi}: \bar{\mathcal{C}}(S) \times \mathbb{H} \rightarrow \bar{\mathcal{C}}(S, z)$ which fits into a commutative diagram*

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

Proof. Via Klarreich’s work, as discussed in Section 1.2.5, we identify $\partial\mathcal{C}$ with \mathcal{EL} . Moreover, the vertical maps in the statement of the proposition send $\mathbb{P}\mathcal{FL}(S) \times \mathbb{H}$ and $\mathbb{P}\mathcal{FL}(S, z)$ onto $\mathcal{EL}(S) \times \mathbb{H}$ and $\mathcal{EL}(S, z)$, respectively, using this identification.

From the construction of Ψ and the definition of \mathcal{FL} , one can see that

$$\Psi(\mathcal{FL}(S) \times \mathbb{H}) \subset \mathcal{FL}(S, z).$$

Furthermore, if $\lambda, \lambda' \in \mathcal{ML}(S)$ with $|\lambda| = |\lambda'|$, then $|\Psi(\lambda, x)| = |\Psi(\lambda', x)|$. Thus, Ψ determines a map

$$\mathcal{EL}(S) \times \mathbb{H} \rightarrow \mathcal{EL}(S, z)$$

which extends Φ to the required map

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

Continuity follows from Proposition 1.8 and Corollary 2.10, and equivariance follows from continuity and the equivariance of Φ . \square

We will also need the following

Proposition 2.13. *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 $\partial\mathcal{C}(S)$, then $\{\Phi(v_n, x_n)\}$ does not accumulate on $\partial\mathcal{C}(S, z)$.*

Proof. The proof is by contradiction. Suppose $\{\Phi(v_n, x_n)\}$ accumulates on some lamination $|\mu| \in \partial\mathcal{C}(S, z)$, and 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 and equal to v . Then

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

This is a contradiction since $|\mu| \in \partial\mathcal{C}(S, z)$. So without loss of generality, we may assume that all the v_n are distinct.

Fix elements $\lambda_n \in \mathcal{ML}(S)$ representing the projective classes associated to v_n via the natural bijection $\mathbb{P}\mathcal{ML}_{\mathcal{C}}(S) \rightarrow \mathcal{C}(S)$, and as in Section 2.4, we let $\hat{\lambda}_n$ be our preferred representative. After passing to a further subsequence and rescaling if necessary, we may assume that for some $\lambda \in \mathcal{ML}(S)$, we have convergence $\lambda_n \rightarrow \lambda$ in the space $\mathcal{ML}(S)$. Since v_n are all distinct, $T(\lambda_n) \rightarrow 0$. Thus, as in the proof of Lemma 2.6, we may pass to yet a further subsequence if necessary so that $|\hat{\lambda}_n|$ converges to a geodesic lamination \mathcal{L} .

It follows from Proposition 1.8 that 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 λ . Choosing $\epsilon > 0$ sufficiently small, we can assume that the train track τ obtained from $N_\epsilon(\mathcal{L})$ (as described in Section 1.2.1) 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 $\widetilde{ev}(f) = x$. After modifying τ and τ_0 to τ' and τ'_0 as in Section 2.4 if necessary (that is, after a possible isotopy and insertion of a bigon around $f(z)$), it follows that for sufficiently large n , $f^{-1}(\tau'(\hat{\lambda}_n))$ determines the lamination $\Psi(\lambda_n, x_n)$. After passing to a further subsequence if necessary, we can assume that $f^{-1}(\tau'(\hat{\lambda}_n))$ converges to some $f^{-1}(\tau')(\mu_0)$, also having nonzero weights only on $f^{-1}(\tau'_0)$. It follows that μ_0 , the limit of $\Psi(\lambda_n, x_n)$, is not in $\mathcal{FL}(S, z)$ since its support is disjoint from $f^{-1}(\alpha)$. Since $\mathbb{P}\Psi(\lambda_n, x_n) \in \mathbb{P}\mathcal{ML}_{\mathcal{C}}(S, z)$, Proposition 1.8 implies $|\mu_0| = |\mu|$, which is a contradiction. \square

Lemma 2.14. *For $(|\lambda|, x), (|\lambda'|, x') \in \mathcal{EL}(S) \times \mathbb{H}$, $\widehat{\Phi}(|\lambda|, x) = \widehat{\Phi}(|\lambda'|, x')$ if and only if $|\lambda| = |\lambda'|$ and x and x' are in the same leaf of $p^{-1}(|\lambda|)$ or in the closure of the same complementary region of $\mathbb{H} - p^{-1}(|\lambda|)$.*

Proof. If x, x' lie on the same leaf of $p^{-1}(|\lambda|)$ or in the closure of the same component of $\mathbb{H} - p^{-1}(|\lambda|)$, then it is straightforward to see that $\widehat{\Phi}(|\lambda|, x) = \widehat{\Phi}(|\lambda|, x')$.

Now we prove the forward direction; suppose that $\widehat{\Phi}(|\lambda|, x) = \widehat{\Phi}(|\lambda'|, x')$. We must show that $|\lambda| = |\lambda'|$ and x and x' are in the same leaf of $p^{-1}(|\lambda|)$ or in the closure of the same complementary region of $\mathbb{H} - p^{-1}(|\lambda|)$.

We first apply an isotopy so that the laminations $\widehat{\Phi}(|\lambda|, x)$ and $\widehat{\Phi}(|\lambda'|, x')$ 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 $\mathbb{H} - p^{-1}(|\lambda|)$ (the general case is similar, but the notation is more complicated). Let $f, f' \in \text{Diff}_0(S)$ be such that $\widetilde{e\nu}(f) = x$ and $\widetilde{e\nu}(f') = x'$. Let \tilde{f} and \tilde{f}' be lifts of f and f' with $\tilde{f}(\tilde{z}) = x$ and $\tilde{f}'(\tilde{z}) = x'$ (see Section 2.1). Modifying f and f' by an element of $\text{Diff}_0(S, z)$ if necessary, we may assume that $f^{-1}(|\lambda|) = \widehat{\Phi}(|\lambda|, x)$ and $f'^{-1}(|\lambda|) = \widehat{\Phi}(|\lambda|, x')$ 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 $\partial\mathbb{H}$, 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. For the remainder of the paper, fix a bi-infinite geodesic γ in the universal covering $p: \mathbb{H} \rightarrow S$ 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 . 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

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

where Φ is the map constructed in Section 2.2. Let $H^\pm(\gamma)$ denote the two half spaces bounded by γ and define

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

Recall that $N(v) = N_{\epsilon(v)}(v)$ is a small neighborhood of the geodesic representative of $v \in \mathcal{C}^0(S)$. We may assume that the $\epsilon(v)$ are small enough to ensure that every component $\alpha \subset \gamma \cap p^{-1}(N(v))$ is essential in the strip of $p^{-1}(N(v))$ that α meets. Here, we say that arc is *essential* if it is not homotopic into the boundary keeping the endpoints fixed.

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. Observe that the image of a 1-Lipschitz retraction on any geodesic metric space is weakly convex. 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 of Proposition 3.1. We describe the case of $\mathfrak{X}(\gamma)$, with $\mathcal{H}^\pm(\gamma)$ handled by similar arguments. First we appeal to Proposition 2.3 and Theorem 2.4 to describe the structure of $\mathfrak{X}(\gamma) \subset \mathcal{C}(S, z)$. Next we prove that $\mathfrak{X}(\gamma)$ is spanned by its vertices and finally we construct a simplicial retraction $\rho: \mathcal{C}(S, z) \rightarrow \mathfrak{X}(\gamma)$. The existence of ρ implies the proposition since a simplicial map is 1-Lipschitz.

For any $x \in \text{int}(v)$, $\mathfrak{X}(\gamma) \cap \Pi^{-1}(x) = \Phi(\{x\} \times \gamma)$, which is a bi-infinite geodesic in the tree $\Pi^{-1}(x) \cong T_v$; see Proposition 2.3. One can also see this as the axis of δ in T_v (since $p(\gamma)$ is filling, δ is not elliptic in T_v). We denote this axis by $\gamma_v \subset T_v$. Recall from Section 2.3 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 2.4, we have

$$\mathfrak{X}(\gamma) \cong \left(\bigsqcup_{v \subset \mathcal{C}(S)} v \times \gamma_v \right) / \sim \quad (2)$$

where, as in Theorem 2.4, the disjoint union is over all simplices $v \subset \mathcal{C}(S)$, 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 (2) to identify the two spaces.

We can now show that $\mathfrak{X}(\gamma)$ is spanned by its vertices. The simplices of $\mathcal{C}(S, z)$ via the homeomorphism of Theorem 2.4 are precisely the images of cells $v \times \sigma$ in the quotient, where $v \subset \mathcal{C}(S)$ is a simplex and $\sigma \subset T_v$ is 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 images 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 retraction

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

Let $\eta_v: T_v \rightarrow \gamma_v$ be the nearest point projection map. Extend this to

$$\rho_v = \text{Id} \times \eta_v: v \times T_v \rightarrow v \times \gamma_v.$$

Observe that if $\phi: v' \rightarrow v$ is a face, then nearest-point projections commute

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

This is because 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 give a well-defined 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 $v \subset \mathcal{C}(S)$ and $\sigma \subset T_v$, the ρ -image 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

Throughout what follows we continue to denote the axis of δ in T_v by $\gamma_v \subset T_v$ or, with respect to the homeomorphism $T_v \cong \Pi^{-1}(v)$, by $\gamma_v = \Phi(\{v\} \times \gamma)$.

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.

For the second statement, first observe that since $\gamma \subset H^\pm(\gamma)$, it follows that

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

To prove the other inclusion, look in each of the trees $\Pi^{-1}(v) \cong T_v$. For each vertex $v \in \mathcal{C}(S)$, we define the *half-tree*

$$H^\pm(\gamma_v) := \mathcal{H}^\pm(\gamma) \cap \Pi^{-1}(v) = \Phi(\{v\} \times H^\pm(\gamma)).$$

Let $u \in H^+(\gamma_v) \cap H^-(\gamma_v)$ be any vertex; we will show that $u \in \gamma_v$. We can write $u = \Phi(\{v\} \times U)$ where U is a component of $\mathbb{H} - p^{-1}(N(v))$. Therefore, $U \cap H^+(\gamma) \neq \emptyset$ and $U \cap H^-(\gamma) \neq \emptyset$. Since U is connected and γ separates $H^+(\gamma)$ from $H^-(\gamma)$ we have $U \cap \gamma \neq \emptyset$. Hence $\Phi(\{v\} \times U) = u \in \gamma_v$ as required.

Given any simplex $u = \{u_0, \dots, u_k\} \subset \mathcal{H}^+(\gamma) \cap \mathcal{H}^-(\gamma)$, by the previous paragraph we have $u_j \in \mathfrak{X}(\gamma)$. Since $\mathfrak{X}(\gamma)$ is a subcomplex spanned by its vertex set, we have $u \subset \mathfrak{X}(\gamma)$ and hence

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

which completes the proof. □

It will be convenient to keep the terminology in the proof of this proposition as well. We therefore think of γ_v as “bounding the half-trees” $H^\pm(\gamma_v) \subset T_v \cong \Pi^{-1}(v)$.

3.2. Rays and existence of Cannon–Thurston maps. An essential subsurface of S is either a component of the complement of a geodesic multicurve in S , the annular neighborhood $N(v)$ of some geodesic $v \in \mathcal{C}^0(S)$, or else the entire surface S .

A point $x \in \partial\mathbb{H}$ is a *filling point* for an essential subsurface Y (or simply, x *fills* Y) if

- for every geodesic ray $r \subset \mathbb{H}$ ending at x and for every $v \in \mathcal{C}^0(S)$ which nontrivially intersects Y , we have $p(r) \cap v \neq \emptyset$ and
- there is a geodesic ray $r \subset \mathbb{H}$ ending at x so that $p(r) \subset Y$.

For any r, v as in the first item observe that the ray r meets infinitely many components of $p^{-1}(v)$. Observe also that every point $x \in \partial\mathbb{H}$ fills exactly one essential subsurface of S .

Definition 3.3. Let $\mathbb{A} \subset \partial\mathbb{H}$ be the set of points that fill S .

Lemma 3.4. *If $x \notin \mathbb{A}$ and r is a ray ending at x then $\Phi(\{v\} \times r)$ has bounded diameter for all $v \in \mathcal{C}^0(S)$.*

Proof. Since x does not fill S there is a simple closed geodesic $v' \subset S$ so that $p(r) \cap v'$ is finite. It follows that $\Phi(\{v'\} \times r)$ has bounded diameter in $\Pi^{-1}(v') \subset \mathcal{C}(S, z)$. Since $\Phi(\{v'\} \times r)$ and $\Phi(\{v\} \times r)$ have bounded Hausdorff distance, we are done. □

Recall that we have fixed once and for all a geodesic $\gamma \subset \mathbb{H}$ which projects to a non-simple closed filling geodesic in S . Consider a set $\{\gamma_n\}$ of pairwise disjoint $\pi_1(S)$ -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 disjoint, 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 $x \in \partial\mathbb{H}$. Here the bar denotes closure in $\overline{\mathbb{H}} = \mathbb{H} \cup \partial\mathbb{H}$. We say that $\{\gamma_n\}$ nests down to x . Note that $\{\overline{H^+(\gamma_n)}\}$ is a neighborhood basis for x .

Given any $x \in \partial\mathbb{H}$, 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 to x .

Proposition 3.5. *If $\{\gamma_n\}$ is a sequence nesting down to a point $x \in \mathbb{A}$, 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. Recall that the curve complex and its one-skeleton are quasi-isometric [MM99]. Thus, in what follows all distances will be computed in the 1-skeleton. 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 length at least R . Since Π is simplicial, the length of the path in $\mathcal{C}(S, z)$ is also at least R . □

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

$$\gamma_{N_j,v} \cap \gamma_{N_{j+1},v} = \emptyset \tag{3}$$

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 (3) is true, and we need to find N_{k+1} so that

$$\gamma_{N_k,v} \cap \gamma_{N_{k+1},v} = \emptyset \tag{4}$$

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, since no such N_{k+1} exists, it must be that for every $j > 0$, there exists a curve $v_j \in B(v_0, R)$ so that

$$\gamma_{N_k,v_j} \cap \gamma_{N_{k+j},v_j} \neq \emptyset.$$

Let $u_j \in \gamma_{N_k,v_j} \cap \gamma_{N_{k+j},v_j}$ be a vertex in the intersection. This vertex is the image under the map Φ_{v_j} of a component $U_j \subset \mathbb{H} - p^{-1}(v_j)$ which meets both γ_{N_k} and $\gamma_{N_{k+j}}$; see Figure 5.




Figure 5. The region U_j and the geodesics γ_{N_k} and $\gamma_{N_{k+j}}$ 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 $y_j^+ \in \gamma_{N_k+j}$. Furthermore, we may pass to a subsequence so that the y_j^- converge to some point y (possibly in $\partial\mathbb{H}$) of γ_{N_k} . It follows that the sequence of geodesics g_j converge to a geodesic ray or line r_∞ connecting y and 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 the v_j are all contained in $B(v_0, R)$, \mathcal{L} cannot contain an ending lamination as a sublamination by Proposition 1.8. 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 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]). Since $p(r_\infty)$ meets $\partial\Sigma$ at most twice the point x does not fill S , a contradiction. \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 in $\mathfrak{X}(\gamma_{N_j}) \cap \mathfrak{X}(\gamma_{N_{j+1}}) \cap \Pi^{-1}(B(v_0, R))$, and hence

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

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

$$\mathfrak{X}(\gamma_{N_j}) \cap \mathfrak{X}(\gamma_{N_i}) \cap \Pi^{-1}(B(v_0, R)) = \emptyset \tag{5}$$

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 $i = i(j)$ so that $u_{i(j)} \in \mathfrak{X}(\gamma_{N_j})$. By (5), 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 3.6. *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

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

Proof. Observe that Φ_v is already defined, continuous and equivariant. All that remains is to extend it to $\bar{\Phi}_v$ on \mathbb{A} by checking the criterion of Lemma 1.9 from which continuity follows. Equivariance is a consequence of equivariance of Φ_v and continuity of $\bar{\Phi}_v$.

Fix a basepoint $u_0 \in \mathcal{C}(S, z)$. Given any $x \in \mathbb{A}$, let $\{\gamma_n\}$ be any sequence nesting down on x . According to Proposition 3.5, 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 $x \in \mathbb{A}$ was an arbitrary point, Lemma 1.9 implies the existence of an \mathbb{A} -Cannon–Thurston map $\bar{\Phi}_v$. □

We note that, given $x \in \mathbb{A}$, the image $\bar{\Phi}_v(x)$ depends only on x , not on v , and is the unique point of intersection of the sets

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

We can therefore unambiguously define $\partial\Phi: \mathbb{A} \rightarrow \partial\mathcal{C}(S, z)$ by $\partial\Phi(x) = \bar{\Phi}_v(x)$ for any $x \in \mathbb{A}$, independent of the choice of $v \in \mathcal{C}^0(S)$.

3.3. Separation

Proposition 3.7. *Given distinct $x, y \in \mathbb{A}$, 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 $H^+(\gamma_x)$ and $H^+(\gamma_y)$ of x and y , respectively, with*

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

if and only if $p(\epsilon)$ is non-simple.

Before we can give the proof of Proposition 3.7, we will need the analogue of Proposition 3.2 for the boundaries at infinity. Recall that γ was chosen to be a bi-infinite geodesic with stabilizer $\langle \delta \rangle$ and $p(\gamma)$ a filling closed geodesic.

Proposition 3.8. *We have*

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

and

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

Proof. This first statement is an immediate consequence of Proposition 3.2. The second also follows from this proposition, but requires some additional argument. Since $\mathfrak{X}(\gamma) = \mathcal{H}^+(\gamma) \cap \mathcal{H}^-(\gamma)$, it easily follows that

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

If $|\mu| \in \partial\mathcal{H}^+(\gamma) \cap \partial\mathcal{H}^-(\gamma)$, then let $\{u_n^+\} \subset \mathcal{H}^+(\gamma)$ and $\{u_n^-\} \in \mathcal{H}^-(\gamma)$ be sequences converging to $|\mu|$ in $\mathcal{C}(S, z)$. For each n let g_n be a geodesic segment from u_n^+ to u_n^- . By Proposition 3.2, there is a vertex $u_n \in g_n \cap \mathfrak{X}(\gamma)$. Therefore u_n also converges to $|\mu|$, and so $|\mu| \in \partial\mathfrak{X}(\gamma)$, proving

$$\partial\mathcal{H}^+(\gamma) \cap \partial\mathcal{H}^-(\gamma) \subset \partial\mathfrak{X}(\gamma). \quad \square$$

A theorem of Kra [Kra81] states that, since $p(\gamma)$ is filling on S , δ is pseudo-Anosov as an element of $\text{Mod}(S, z)$. We let $|\mu_+|$ and $|\mu_-|$ be the attracting and repelling fixed points of δ , respectively, in $\partial\mathcal{C}(S, z)$.

Recall from Proposition 2.12 that Φ has a continuous $\pi_1(S)$ -equivariant extension $\widehat{\Phi}: \overline{\mathcal{C}}(S) \times \mathbb{H} \rightarrow \overline{\mathcal{C}}(S, z)$.

Lemma 3.9.

$$\partial\mathfrak{X}(\gamma) = \widehat{\Phi}(\partial\mathcal{C}(S) \times \gamma) \cup \{|\mu_\pm|\}$$

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

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

We are left to prove the reverse inclusion. Suppose $\{u_n\}$ is any sequence in $\mathfrak{X}(\gamma)$ with $u_n \rightarrow |\mu| \in \partial\mathfrak{X}(\gamma)$. We wish to show that $|\mu| \in \widehat{\Phi}(\partial\mathcal{C}(S) \times \gamma) \cup \{|\mu_\pm|\}$. 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 2.13, we can assume that v_n accumulates on $\partial\mathcal{C}(S)$. So, after passing to a further subsequence if necessary, we can assume that $v_n \rightarrow |\lambda| \in \partial\mathcal{C}(S)$. Then by continuity of $\widehat{\Phi}$ (Proposition 2.12) we have

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

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

Note that $x \in \mathbb{A}$ 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 . Thus $\{\delta^n(\gamma_1)\}$ 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 attracting 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

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

completing the proof of Lemma 3.9. □

Proof of Proposition 3.7. We fix $x, y \in \mathbb{A}$ and ϵ the geodesic between them. We write γ_x and γ_y to denote $\pi_1(S)$ -translates of γ for which $\overline{H^+(\gamma_x)}$ and $\overline{H^+(\gamma_y)}$ define disjoint neighborhoods of x and y , respectively. We must show that $p(\epsilon)$ is simple if and only if $\partial\mathcal{H}^+(\gamma_x) \cap \partial\mathcal{H}^+(\gamma_y) \neq \emptyset$ for all such γ_x and γ_y .

First, suppose $p(\epsilon)$ is simple. The closure of $p(\epsilon)$ is a lamination \mathcal{L} [CB87]. Since $x, y \in \mathbb{A}$, \mathcal{L} must contain some $|\lambda| \in \mathcal{EL}(S)$ as the sublamination obtained by discarding isolated leaves. Therefore ϵ is either a leaf of $p^{-1}(|\lambda|)$ or a diagonal for some complementary polygon of $p^{-1}(|\lambda|)$.

It follows from Lemma 2.14 that if $x' \in \gamma_x \cap \epsilon$ and $y' \in \gamma_y \cap \epsilon$, then $\widehat{\Phi}(|\lambda|, x') = \widehat{\Phi}(|\lambda|, y')$. Appealing to Lemma 3.9 we have

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

as required. In fact, it is worth noting that by Lemma 2.14, $\widehat{\Phi}(\{|\lambda|\} \times \epsilon)$ is a single point which lies in $\partial\mathcal{H}^+(\gamma_x) \cap \partial\mathcal{H}^+(\gamma_y)$ for all allowed choice of γ_x and γ_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

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

Therefore, by Proposition 3.8, it follows that

$$\partial\mathcal{H}^+(\gamma_1) \cap \partial\mathcal{H}^+(\gamma_2) = \partial\mathcal{X}(\gamma_1) \cap \partial\mathcal{X}(\gamma_2).$$

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

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

then by Proposition 3.9 there exists $x_1 \in \gamma_1$ and $x_2 \in \gamma_2$ and $|\lambda_1|, |\lambda_2| \in \partial\mathcal{C}(S)$ for which $\widehat{\Phi}(|\lambda_1|, x_1) = \widehat{\Phi}(|\lambda_2|, x_2)$. According to Lemma 2.14, 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 bi-infinite 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. We suppose that

$$\partial\mathcal{H}^+(\gamma_{n,x}) \cap \partial\mathcal{H}^-(\gamma_{n,y}) \neq \emptyset$$

for all $n \geq 0$, and prove that $p(\epsilon)$ is simple on S . By the discussion in the preceding two paragraphs there exists a sequence of laminations $\{|\lambda_n|\} \subset \partial\mathcal{C}(S)$ 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 . Also $p(\epsilon)$ is simple as it is the limit of simple geodesics [CB87]. □

The following is now immediate from Proposition 3.7 and its proof.

Corollary 3.10. *Given distinct $x, y \in \mathbb{A}$ then $\partial\Phi(x) = \partial\Phi(y)$ if and only if x and y are ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of $p^{-1}(|\lambda|)$ for some $|\lambda| \in \partial\mathcal{C}(S)$.* □

3.4. Surjectivity. In this section, we prove that our map $\partial\Phi$ is surjective.

Birman–Series [BS85] proved that the closure of the union of simple closed geodesics

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

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

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

is ϵ -dense. It follows that $\epsilon(v) \leq \epsilon$ for all $v \in \mathcal{C}^0(S)$.

Lemma 3.11. *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*

$$\mathcal{G} = (\mathbf{v}, \mathbf{x}): [a, b] \rightarrow \mathcal{C}(S) \times \mathbb{H}$$

such that $\Phi \circ \mathcal{G}$ is a geodesic from u_1 to u_2 and \mathbf{x} connects x_1 to x_2 with image contained in the 2ϵ -neighborhood of a geodesic in \mathbb{H} .

Proof. For each $i = 1, 2$ we can find x'_i in the same component of $\mathbb{H} - p^{-1}(N^\circ(v_i))$ 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 are dense in $\partial\mathbb{H} \times \partial\mathbb{H}$). 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 x_1 and x_2 . Suppose we can find $\mathcal{G}' = (\mathbf{v}', \mathbf{x}')$ so that $\Phi \circ \mathcal{G}'$ is a geodesic from u_1 to u_2 and \mathbf{x}' connects x'_1 to x'_2 with image contained in the ϵ -neighborhood of a geodesic containing x'_1 and x'_2 . Then we can take $\mathcal{G} = (\mathbf{v}, \mathbf{x})$ to be such that $\mathbf{v} = \mathbf{v}'$ and \mathbf{x} first runs from x_1 to x'_1 , then traverses \mathbf{x}' , and finally runs from x'_2 to x_2 (all appropriately reparameterized). This will then provide the desired path proving the lemma.

To construct \mathcal{G}' , we suppose for the moment that $\{\epsilon(v)\}_{v \in \mathcal{C}^0(S)}$ have been chosen so that any arc of $\gamma' \cap p^{-1}(N(v))$ is essential. With this assumption, Proposition 3.1 applied to γ' implies 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 $\gamma' \cap (\mathbb{H} - p^{-1}(N(v_i)))$ and is in particular connected. It follows from the construction of Φ that the edges $[w_i, w_{i+1}]$, 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 $\gamma' \cap p^{-1}(N(v_i))$. If $v_i \neq v_{i+1}$, then b_i traverses the edge $[v_i, v_{i+1}]$ and c_i is constant.

We can now define $\mathcal{G}' = (\mathbf{v}', \mathbf{x}')$ as follows.

- (1) Begin by holding \mathbf{v}' constant equal to $u_1 = w_1$ and let \mathbf{x}' 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 \mathbf{v}' constant again and let \mathbf{x}' 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 \mathbf{v}' constant and letting \mathbf{x}' go from the terminal point of c_i to the initial point of c_{i+1} inside $\alpha_{i+1} \subset \gamma'$.

- (5) We complete the path by traversing a_{k-1} , then holding v' constant and letting x' 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 $\Phi \circ \mathcal{G}'$ 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, x' is contained in γ' and connects x'_1 to x'_2 , so therefore stays within a distance zero 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 $\gamma' \cap N(v)$ enters and exits the component of $N(v)$ 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 $\mathcal{G}'' = (v'', x'')$. 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 $\Phi' \circ \mathcal{G}''$ is a geodesic. With respect to the original map Φ , \tilde{v}'' is almost good enough for our purposes. The only problem is that $\Phi \circ \mathcal{G}''$ may now no longer be a geodesic: If there is some interval in the domain in which v'' is constant equal to v and x'' enters and exits a component $p^{-1}(N(v))$ from the same side, then $\Phi \circ \mathcal{G}''$ 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 \mathcal{G}'' to the desired path \mathcal{G}' , by pushing x'' outside of $p^{-1}(N(v))$ whenever this happens, thus changing it by at most $\epsilon(v) \leq \epsilon$. The resulting path \mathcal{G}' has $v' = v''$ and x' still connects x'_1 to x'_2 and stays within ϵ of γ' , as required. □

Surjectivity of $\partial\Phi$ requires that every point of $\partial\mathcal{C}(S, z)$ is the limit of $\Phi_v(r)$ for some $v \in \mathcal{C}^0(S)$ and 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 3.12. *For any $v \in \mathcal{C}^0(S)$, $\partial\mathcal{C}(S, z) \subset \overline{\Phi_v(\mathbb{H})}$.*

Proof. First, note that since $\pi_1(S) < \text{Mod}(S, z)$ is a normal, infinite subgroup, its 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 the boundary $\mathbb{P}\mathcal{ML}(S, z)$ 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 to μ .

The systole map $\text{sys}: \mathcal{T}(S, z) \rightarrow \mathcal{C}(S, z)$ 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 $\text{sys}(x_n) \rightarrow |\mu| \in \mathcal{EL}(S, z) \cong \partial\mathcal{C}(S, z)$. Therefore $\Phi_v(x_n) \rightarrow |\mu|$. Since μ was arbitrary, every point of $\partial\mathcal{C}(S, z)$ is a limit of a sequence in $\Phi_v(\mathbb{H})$, and we are done. □

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

Proposition 3.13. *If $\lim_{n \rightarrow \infty} x_n = x \in \partial\mathbb{H} - \mathbb{A}$, then $\Phi_v(x_n)$ does not converge to a point of $\partial\mathcal{C}(S, z)$.*

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

Lemma 3.14. *If $\{x_n\}$ and x are as in Proposition 3.13 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 $\partial\mathcal{C}(S, z)$*

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

Let γ_0 be a $\pi_1(S)$ -translate of γ such that γ_0 separates x from the repelling fixed point of η . Then $\{\eta^n(\gamma_0)\}$ nests down on x . It follows that after passing to a subsequence (which we continue to denote $\{x_n\}$) we have

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

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

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

Suppose now that $\Phi_v(x_n)$ converges to some element $|\mu| \in \partial\mathcal{C}(S, z)$. It follows that

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

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

Proof of Proposition 3.13. Recall that $\Phi_v(x) = \Phi(v, x)$. Suppose, contrary to the conclusion of the proposition, that

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

We begin by finding another sequence which also converges to $|\mu|$ to which we can apply the techniques developed so far. Since $x \notin \mathbb{A}$ the surface Y filled by x is strictly contained in S . By Lemma 3.14 we may assume that Y is not an annulus. Let $r \subset \mathbb{H}$ be a ray ending at x so that r is contained in a component \tilde{Y} of $p^{-1}(Y)$ and so that $p(r)$ fills Y .

We 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

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

Now fix any $k > 0$ and let $\Phi \circ \mathcal{G}^k : [a_k, b_k] \rightarrow \mathcal{C}(S, z)$ be a geodesic from $\Phi(v, x_{2k})$ to $\Phi(v, x_{2k+1})$ where $\mathcal{G}^k = (\mathbf{v}^k, \mathbf{x}^k)$ is given by Lemma 3.11. The path \mathbf{x}^k connects x_{2k} to x_{2k+1} and has image within 2ϵ of a geodesic in \mathbb{H} which must also pass within a uniformly bounded distance of r (in fact, it passes within a distance $1 + 2\epsilon$).

Choosing the subsequence $\{x_n\}$ carefully, we may assume that β_k spends a very long time in \tilde{Y} . Doing this ensures that the image of \mathbf{x}^k nontrivially intersects \tilde{Y} . Let $t_k \in [a_k, b_k]$ be any time where \mathbf{x}^k meets \tilde{Y} . Then set

$$y_k = \mathbf{x}^k(t_k) \in \tilde{Y} \quad \text{and} \quad v_k = \mathbf{v}^k(t_k) \in \mathcal{C}^0(S)$$

(recall that \mathbf{x}^k is constant when \mathbf{v}^k is not, so we can assume that t_k is chosen so that v_k is indeed a vertex).

Observe that since $\Phi \circ \mathcal{G}^k([a_k, b_k])$ is a geodesic from $\Phi(v, x_{2k})$ to $\Phi(v, x_{2k+1})$, the sequence $\{\Phi \circ \mathcal{G}^k(t_k)\} = \{\Phi(v_k, y_k)\}$ 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 $\tilde{e}\mathbf{v}(f_k) = y_k \in \tilde{Y}$. Since \tilde{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 . That is, there is a path $f_t \in \text{Diff}_0(S)$ for $t \in [1, \infty)$ such that $y_k = \tilde{e}\mathbf{v}(f_k)$ for all positive integers k , and so that

$$f_1|_{S-Y} = f_t|_{S-Y}$$

for all $t \in [1, \infty)$.

Let $X = f_1^{-1}(Y)$ and consider the punctured surfaces

$$Y^\circ = Y - \{f_1(z)\} \quad \text{and} \quad X^\circ = X - \{z\} = f_1^{-1}(Y^\circ).$$

We will be interested in the set of subsurface projections

$$\{\pi_{X^\circ}(u_k)\} \subset \mathcal{C}'(X^\circ)$$

where $\mathcal{C}'(X^\circ)$ is the arc complex of X° ; see [MM00]. We consider the incomplete metric on X° for which $f_1 : X^\circ \rightarrow Y^\circ$ is an isometry where Y° is given the induced path metric inside of S .

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

Here, length means infimum of lengths over the isotopy class of an arc. The claim implies that there are infinitely many arcs in the set $\{\pi_{X^\circ}(u_k)\}$ which is impossible if $u_k \rightarrow |\mu|$. Thus, to complete the proof of the proposition, it suffices to prove the claim.

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 $\pi_{X^\circ}(u_k)$ and $\ell_{X^\circ}(f_k^{-1}(\alpha_k)) = \ell_{Y^\circ}(f_1 f_k^{-1}(\alpha_k))$. 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

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

and hence there is an arc of $\pi_{X^\circ}(u_k)$ 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 and equal to a union of finitely many arcs in Y . We fix attention on one arc, call it α . Again, we see that $f_k^{-1}(\alpha)$ is an arc of $\pi_{X^\circ}(u_k)$ and $\ell_{X^\circ}(f_k^{-1}(\alpha)) = \ell_{Y^\circ}(f_1 f_k^{-1}(\alpha))$ with $f_1 f_t^{-1}$ 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 $\ell_{Y^\circ}(h_k^{-1}(\alpha))$ 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 $\tilde{e}v$ that $\tilde{h}_k(\tilde{e}v(f_1)) = y_k$. Thus, \tilde{h}_t is essentially pushing the point $y = \tilde{e}v(f_1) \in \tilde{Y}$ 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 drag $\tilde{\alpha}^k$ backward using 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 6. Moreover, observe that $\ell_{Y^\circ}(h_k^{-1}(\alpha))$ 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 x fills 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 because x fills Y would imply Y is an annulus, and this is a contradiction. This implies $\ell_{Y^\circ}(h_k^{-1}(\alpha))$ tends to infinity as $k \rightarrow \infty$. This proves the claim, and so completes the proof of the proposition. \square

We can now prove one of the main technical pieces of Theorem 1.1.




Figure 6. 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: dragging $\tilde{\alpha}^k$ back by \tilde{h}_k^{-1} .

Theorem 3.15. *The map*

$$\partial\Phi: \mathbb{A} \rightarrow \partial\mathcal{C}(S, z)$$

is surjective.

Proof. Let $|\mu| \in \partial\mathcal{C}(S, z)$ be an arbitrary point. According to Lemma 3.12 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 $x \in \partial\mathbb{H}$. It follows from Proposition 3.13 that $x \in \mathbb{A}$. Then, by Theorem 3.6

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

Since $|\mu| \in \partial\mathcal{C}(S, z)$ was an arbitrary point, it follows that $\partial\Phi(\mathbb{A}) = \partial\mathcal{C}(S, z)$, and $\partial\Phi$ is surjective. □

3.5. Neighborhood bases. In this section we find neighborhood bases for points of $\partial\mathcal{C}(S, z)$. To do this, we must distinguish between two types of points of \mathbb{A} . We say a point $x \in \mathbb{A}$ is *simple* if there exists a ray r in \mathbb{H} ending at x for which $p(r)$ is simple. Otherwise x is not simple. Equivalently, a point $x \in \mathbb{A}$ is simple if and only if there is a lamination $|\lambda| \in \mathcal{EL}(S)$ such that x is the ideal endpoint of a leaf (or ideal vertex of a complementary polygon) of $p^{-1}(|\lambda|)$.

Lemma 3.16. *If $x \in \mathbb{A}$ is not simple and $\{\gamma_n\}$ are $\pi_1(S)$ -translates of γ which nest down to x , then $\{\partial\mathcal{H}^+(\gamma_n)\}$ is a neighborhood basis for $\partial\Phi(x)$.*

Proof. By Proposition 3.5 the distance of $\mathcal{H}^+(\gamma_n)$ to any fixed basepoint in $\mathcal{C}(S, z)$ tends to infinity as $n \rightarrow \infty$. It follows that the visual diameter of $\overline{\mathcal{H}^+(\gamma_n)}$ measured from any base point also tends to zero. Thus, for any neighborhood U of $\partial\Phi(x)$ in $\partial\mathcal{C}(S, z)$, there exists $N > 0$ so that for all $n \geq N$, $\partial\mathcal{H}^+(\gamma_n) \subset U$.

We must prove $\partial\Phi(x)$ is in $\text{int}(\partial\mathcal{H}^+(\gamma_n))$ for all n . We already know that

$$\bigcap \partial\mathcal{H}^+(\gamma_n) = \{\partial\Phi(x)\}$$

and in particular, $\partial\Phi(x) \in \partial\mathcal{H}^+(\gamma_n)$ for all n . Therefore, it suffices to prove that for any n , there exists $m > n$ so that

$$\partial\mathcal{H}^+(\gamma_m) \subset \text{int}(\partial\mathcal{H}^+(\gamma_n)).$$

It follows from Proposition 3.8 and the fact that $\partial\mathcal{H}^-(\gamma_n)$ is a closed subset of $\partial\mathcal{C}(S, z)$ that

$$\partial\mathcal{H}^+(\gamma_n) - \partial\mathcal{X}(\gamma_n) = \partial\mathcal{C}(S, z) - \partial\mathcal{H}^-(\gamma_n) \subset \text{int}(\partial\mathcal{H}^+(\gamma_n)).$$

For any $m > n$, we also know

$$\partial\mathcal{H}^+(\gamma_m) \subset \partial\mathcal{H}^+(\gamma_n).$$

Thus, if we can find $m > n$ so that

$$\partial\mathcal{X}(\gamma_n) \cap \partial\mathcal{X}(\gamma_m) = \emptyset,$$

then appealing to Proposition 3.8 again, it will follow that

$$\partial\mathcal{H}^+(\gamma_m) \subset \partial\mathcal{H}^+(\gamma_n) - \partial\mathcal{X}(\gamma_n) \subset \text{int}(\partial\mathcal{H}^+(\gamma_n)),$$

as required.

If for all $m > n$ we have $\partial\mathcal{X}(\gamma_m) \cap \partial\mathcal{X}(\gamma_n) \neq \emptyset$, then a similar proof to that given for Proposition 3.7 shows that x is a simple point which is a contradiction. \square

The above lemma gives a neighborhood basis for $\partial\Phi(x)$ when $x \in \mathbb{A}$ not a simple point. The next lemma describes a neighborhood basis $\partial\Phi(x)$, 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 . From Corollary 3.10, $\partial\Phi(x_1) = \dots = \partial\Phi(x_k)$, and the $\partial\Phi$ -image of any simple point has this form.

Lemma 3.17. *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 to x_j for each $j = 1, \dots, k$, then*

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

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

Proof. Let $|\mu| = \partial\Phi(x_1) = \dots = \partial\Phi(x_k)$. As in the proof of the previous lemma, the sets in the proposed neighborhood basis have visual diameter tending to zero as $n \rightarrow \infty$.

Since $|\mu| \in \partial\mathcal{H}^+(\gamma_{j,n})$ for all j and n , we clearly have

$$|\mu| \in \partial\mathcal{H}^+(\gamma_{1,n}) \cup \dots \cup \partial\mathcal{H}^+(\gamma_{k,n})$$

for all n . Thus, we are required to show that $|\mu|$ is an interior point of this set.

This is equivalent to saying that for any sequence $\{|\mu_m|\} \subset \partial\mathcal{C}(S, z)$ converging to $|\mu|$, and every positive integer n , there exists $M > 0$ so that for all $m \geq M$,

$$|\mu_m| \in \partial\mathcal{H}^+(\gamma_{1,n}) \cup \dots \cup \partial\mathcal{H}^+(\gamma_{k,n}). \tag{6}$$

So, let $\{|\mu_m|\} \subset \partial\mathcal{C}(S, z)$ be a sequence converging to $|\mu|$ and n a positive integer. Choose any sequence $\{y_m\} \subset \mathbb{A}$ so that $\partial\Phi(y_m) = |\mu_m|$ (such a sequence exists by surjectivity of $\partial\Phi$). We wish to show that any accumulation point of $\{y_m\}$ is one of the points x_1, \dots, x_k . For then, we can find an $M > 0$ so that for all $m \geq M$

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

and hence (6) holds.

To this end, we pass to a subsequence so that $y_m \rightarrow x \in \partial\mathbb{H}$. Choosing sequences converging to y_m for all m and applying a diagonal argument, we see that there is a sequence $\{q_m\} \subset \mathbb{H}$ with $\lim_{m \rightarrow \infty} q_m = x$ and $\lim_{m \rightarrow \infty} \Phi_v(q_m) = |\mu|$. From Proposition 3.13 we deduce that $x \in \mathbb{A}$.

Now, if $x \in \{x_1, \dots, x_k\}$ then we are done. Suppose not. Then the geodesic ϵ_j from x to x_j has $p(\epsilon_j)$ non-simple for all j . Proposition 3.7 guarantees $\pi_1(S)$ -translates $\gamma_x, \gamma_{1,n}, \dots, \gamma_{k,n}$ of γ defining neighborhoods

$$\overline{H^+(\gamma_x)}, \overline{H^+(\gamma_{1,n})}, \dots, \overline{H^+(\gamma_{k,n})}$$

of x, x_1, \dots, x_k , respectively for which

$$\partial\mathcal{H}^+(\gamma_x) \cap \partial\mathcal{H}^+(\gamma_{j,n}) = \emptyset$$

for all $j = 1, \dots, k$. Since $\partial\Phi$ is continuous, we have

$$|\mu| = \lim_{m \rightarrow \infty} |\mu_m| = \lim_{m \rightarrow \infty} \partial\Phi(y_m) = \partial\Phi(x) \in \partial\mathcal{H}^+(\gamma_x).$$

This is impossible since $|\mu| \in \partial\mathcal{H}^+(\gamma_{j,n})$ for all $j = 1, \dots, k$. Therefore, $x = x_j$ for some j , and the proof is complete. □

We are now ready to prove

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 unique continuous $\pi_1(S)$ -equivariant extension*

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

The map $\partial\Phi = \bar{\Phi}_v|_{\mathbb{A}}$ does not depend on v and is a quotient map onto $\partial\mathcal{C}(S, z)$. Given distinct points $x, y \in \mathbb{A}$, $\partial\Phi(x) = \partial\Phi(y)$ if and only if x and y are ideal endpoints of a leaf (or ideal vertices of a complementary polygon) of the lift of an ending lamination on S .

Proof. By Theorem 3.6, Corollary 3.10 and Theorem 3.15 all that remains is to prove that $\partial\Phi$ is a quotient map. To see this, we need only show that $E \subset \partial\mathcal{C}(S, z)$ is closed if and only if $F = \partial\Phi^{-1}(E)$ is closed. Since $\partial\Phi$ is continuous, it follows that if E is closed, then F is closed.

Now, suppose that F is closed. To show that E is closed, we let $|\mu_n| \rightarrow |\mu|$ with $\{|\mu_n|\} \subset E$ and we must check that $|\mu| \in E$. By Lemmas 3.16 and 3.17, after passing to a subsequence if necessary, there is a sequence $\{\gamma_n\}$ nesting down on some point $x \in \partial\Phi^{-1}(|\mu|)$ with $|\mu_n| \in \partial\mathcal{H}^+(\gamma_n)$. Let $x_n \in \partial\Phi^{-1}(|\mu_n|) \subset F$ be such that $x_n \in \partial H^+(\gamma_n)$. It follows that $x_n \rightarrow x$, so since F is closed, $x \in F$. Therefore, $|\mu| = \partial\Phi(x) \in E$, as required. Thus, E is closed, and $\partial\Phi$ is a quotient map. \square

3.6. Mod(S, z)-equivariance. We now prove

Theorem 1.2. *The quotient map*

$$\partial\Phi: \mathbb{A} \rightarrow \partial\mathcal{C}(S, z)$$

constructed in Theorem 1.1 is equivariant with respect to the action of $\text{Mod}(S, z)$.

Proof. It suffices to prove

$$\partial\Phi(\phi(x)) = \phi(\partial\Phi(x)).$$

for every $\phi \in \text{Mod}(S, z)$ and a dense set of points $x \in \mathbb{A}$.

Let $\gamma' \subset \mathbb{H}$ be a geodesic for which $p(\gamma')$ is a filling closed geodesic in S and let $\delta' \in \pi_1(S)$ be the generator of the infinite cyclic stabilizer of γ' . Let $x \in \mathbb{A}$ denote the attracting fixed point of δ' . As previously discussed, according to Kra [Kra81], δ' represents a pseudo-Anosov mapping class in $\text{Mod}(S, z)$, and the $\pi_1(S)$ -equivariance of $\partial\Phi$ implies $\partial\Phi(x)$ is the attracting fixed point for δ' in $\partial\mathcal{C}(S, z)$.

Now, given any $\phi \in \text{Mod}(S, z)$, note that $\phi(x)$ is the attracting fixed point of $\phi \circ \delta' \circ \phi^{-1}$ in \mathbb{A} , and $\phi(\partial\Phi(x))$ is the attracting fixed point for $\phi \circ \delta' \circ \phi^{-1}$ in

$\partial\mathcal{C}(S, z)$. Appealing to the $\pi_1(S)$ -equivariance again, we see that $\partial\Phi$ must take $\phi(x)$ to $\phi(\partial\Phi(x))$. That is

$$\partial\Phi(\phi(x)) = \phi(\partial\Phi(x)).$$

Since the set of endpoints of such geodesics is dense in \mathbb{A} , this completes the proof. \square

4. Local path-connectivity

The following, together with Lemma 3.17 will prove Theorem 1.3.

Lemma 4.1. *$\partial\mathcal{H}^+(\gamma)$ is path-connected.*

Proof. Fix any $|\lambda| \in \mathcal{EL}(S)$. According to Proposition 2.12, $\widehat{\Phi}$ is continuous, so we have a path-connected subset

$$\widehat{\Phi}(\{|\lambda|\} \times H^+(\gamma)) \subset \partial\mathcal{H}^+(\gamma).$$

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

According to Theorem 1.1 there exists $x \in \mathbb{A}$ so that $\partial\Phi(x) = |\mu|$. 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 2.12

$$\widehat{\Phi}(\{|\lambda|\} \times r([t_n, 1)) \subset \widehat{\Phi}(\{|\lambda|\} \times H^+(\gamma_n)) \subset \partial\mathcal{H}^+(\gamma_n).$$

Recall that, by definition, $\partial\Phi(x)$ is the unique point of intersection

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

and hence

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

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

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

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

We now prove

Theorem 1.3. *The Gromov boundary $\partial\mathcal{C}(S, z)$ is path-connected and locally path-connected.*

Proof. From Lemma 4.1, we see that every set of the form $\partial\mathcal{H}^+(\gamma_0)$ is path-connected for any $\pi_1(S)$ -translate γ_0 of γ . According to Lemmas 3.16 and 3.17 there is a basis for the topology consisting of these sets (and finite unions of these sets which all share a point); this proves local path-connectivity. Path-connectivity follows from Lemma 4.1 and Proposition 3.8. \square

References

- [Bir69] J. S. Birman, Mapping class groups and their relationship to braid groups. *Comm. Pure Appl. Math.* **22** (1969), 213–238. [Zbl 0167.21503](#) [MR 0243519](#)
- [Bir74] J. S. Birman, *Braids, links, and mapping class groups*. Ann. of Math. Stud. 82, Princeton University Press, Princeton, N.J., 1974. [Zbl 0305.57013](#) [MR 0375281](#)
- [BS85] J. S. Birman and C. Series, Geodesics with bounded intersection number on surfaces are sparsely distributed. *Topology* **24** (1985), no. 2, 217–225. [Zbl 0568.57006](#) [MR 0793185](#)
- [Bon88] F. Bonahon, The geometry of Teichmüller space via geodesic currents. *Invent. Math.* **92** (1988), no. 1, 139–162. [Zbl 0653.32022](#) [MR 0931208](#)
- [Bro00] J. F. Brock, Continuity of Thurston’s length function. *Geom. Funct. Anal.* **10** (2000), no. 4, 741–797. [Zbl 0968.57011](#) [MR 1791139](#)
- [CB87] A. Casson and S. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*. London Math. Soc. Student Texts 9, Cambridge University Press, Cambridge, 1988. [Zbl 0649.57008](#) [MR 0964685](#)
- [CC00] A. Candel and L. Conlon, *Foliations I*. Grad. Stud. Math. 23, Amer. Math. Soc., Providence, RI, 2000. [Zbl 0936.57001](#) [MR 1732868](#)
- [CEG87] R. D. Canary, D. B. A. Epstein, and P. Green, Notes on notes of Thurston. In *Analytical and geometric aspects of hyperbolic spaces*, London Math. Soc. Lecture Note Ser. 111, Cambridge University Press, Cambridge 1987, 3–92. [Zbl 0612.57009](#) [MR 0903850](#)
- [CT07] J. W. Cannon and W. P. Thurston, Group invariant Peano curves. *Geom. Topol.* **11** (2007), 1315–1355. [Zbl 1136.57009](#) [MR 2326947](#)
- [EE69] C. J. Earle and J. Eells, A fibre bundle description of Teichmüller theory. *J. Differential Geom.* **3** (1969), 19–43. [Zbl 0185.32901](#) [MR 0276999](#)

- [FM02] B. Farb and L. Mosher, Convex cocompact subgroups of mapping class groups. *Geom. Topol.* **6** (2002), 91–152 (electronic). [Zbl 1021.20034](#) [MR 1914566](#)
- [Gab09] D. Gabai, Almost filling laminations and the connectivity of ending lamination space. *Geom. Topol.* **13** (2009), 1017–1041. [Zbl 1165.57015](#) [MR 2470969](#)
- [GL07] V. Guirardel and G. Levitt, Deformation spaces of trees. *Groups Geom. Dyn.* **1** (2007), no. 2, 135–181. [Zbl 1134.20026](#) [MR 2319455](#)
- [Har81] W. J. Harvey, Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference*, Ann. of Math. Stud. 97, Princeton University Press, Princeton, N.J., 1981, 245–251. [Zbl 0461.30036](#) [MR 0624817](#)
- [HM79] J. Hubbard and H. Masur, Quadratic differentials and foliations. *Acta Math.* **142** (1979), no. 3–4, 221–274. [Zbl 0415.30038](#) [MR 0523212](#)
- [Kla99] E. Klarreich, The boundary at infinity of the curve complex and the relative Teichmüller space. Preprint 1999. <http://nasw.org/users/klarreich/research.htm>
- [KLS09] R. P. Kent IV, C. J. Leininger, and S. Schleimer, Trees and mapping class groups. *J. Reine Angew. Math.* **637** (2009), 1–21. [Zbl 1190.57014](#) [MR 2599078](#)
- [Kra81] I. Kra, On the Nielsen-Thurston-Bers type of some self-maps of Riemann surfaces. *Acta Math.* **146** (1981), no. 3–4, 231–270. [Zbl 0477.32024](#) [MR 0611385](#)
- [LS09] C. J. Leininger and S. Schleimer, Connectivity of the space of ending laminations. *Duke Math. J.* **150** (2009), no. 3, 533–575. [Zbl 1190.57013](#) [MR 2582104](#)
- [Min92] Y. N. Minsky, Teichmüller geodesics and ends of hyperbolic 3-manifolds. *Topology* **32** (1993), no. 3, 625–647. [Zbl 0793.58010](#) [MR 1231968](#)
- [Mit97] M. Mitra, Ending laminations for hyperbolic group extensions. *Geom. Funct. Anal.* **7** (1997), no. 2, 379–402. [Zbl 0880.57001](#) [MR 1445392](#)
- [Mit98] M. Mitra, Cannon-Thurston maps for hyperbolic group extensions. *Topology* **37** (1998), 527–538. [Zbl 0907.20038](#) [MR 1604882](#)
- [Mj06] M. Mj, Cannon-Thurston maps for surface groups. Preprint 2006. [arXiv:math/0607509](https://arxiv.org/abs/math/0607509) [math.GT]
- [Mj07] M. Mj, Ending laminations and Cannon-Thurston maps. Preprint 2007. [arXiv:math/0701725](https://arxiv.org/abs/math/0701725) [math.GT]
- [MM99] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves I: hyperbolicity. *Invent. Math.* **138** (1999), 103–139. [Zbl 0941.32012](#) [MR 1714338](#)
- [MM00] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves II: hierarchical structure. *Geom. Funct. Anal.* **10** (2000), 902–974. [Zbl 0972.32011](#) [MR 1791145](#)
- [Mos97] L. Mosher, A hyperbolic-by-hyperbolic hyperbolic group. *Proc. Amer. Math. Soc.*, **125** (1997), no. 12, 3447–3455. [Zbl 0895.20028](#) [MR 1443845](#)
- [MP89] J. D. McCarthy and A. Papadopoulos, Dynamics on Thurston’s sphere of projective measured foliations. *Comment. Math. Helv.* **64** (1989), 133–166. [Zbl 0681.57002](#) [MR 0982564](#)
- [PH92] R. Penner and J. Harer, *Combinatorics of train tracks*. Ann. Math. Stud. 125, Princeton University Press, 1992. [Zbl 0765.57001](#) [MR 1144770](#)

[Thu80] W. P. Thurston, *The geometry and topology of 3-manifolds*. Princeton University Notes, 1980.

Received March 13, 2009

Christopher J. Leininger, Department of Mathematics, University of Illinois, Urbana-Champaign, IL 61801, U.S.A.

E-mail: clein@math.uiuc.edu

Mahan Mj, School of Mathematical Sciences, RKM Vivekananda University, P.O. Belur Math, WB-711202, India

E-mail: mahan@rkmvu.ac.in

Saul Schleimer, Department of Mathematics, University of Warwick, Coventry CV4 7AL, UK

E-mail: s.schleimer@warwick.ac.uk