Shadows of mapping class groups: capturing convex cocompactness

Autumn Kent and Christopher J. Leininger *

February 1, 2007

for Dick and Geri

1 Introduction

A Kleinian group Γ is a discrete subgroup of $\mathrm{PSL}_2(\mathbb{C})$. When non-elementary, such a group possesses a unique non-empty minimal closed invariant subset Λ_Γ of the Riemann sphere, called the limit set. A Kleinian group acts properly discontinuously on the complement Δ_Γ of Λ_Γ and so this set is called the domain of discontinuity.

Such a group is said to be convex cocompact if it acts cocompactly on the convex hull H_Γ in $\mathbb{H}^3$ of its limit set Λ_Γ. This is equivalent to the condition that an orbit of Γ is quasi-convex in $\mathbb{H}^3$—or that the orbit defines a quasi-isometric embedding $\Gamma \to \mathbb{H}^3$. Equivalent to each of these is the property that every limit point of Γ is conical, and still another definition is that Γ has a compact Kleinian manifold—meaning that Γ acts cocompactly on $\mathbb{H}^3 \cup \Delta_\Gamma$. We refer the reader to [10] and the references therein for the history of these notions and the proof of their equivalence (see also [65]).

Let S denote an oriented complete hyperbolic surface of finite area, $\mathrm{Mod}(S) = \pi_0(\mathrm{Homeo}^+(S))$ its group of orientation preserving self–homeomorphisms up to isotopy, and $\mathcal{T}(S)$ the Teichmüller space of S equipped with Teichmüller's metric.

The mapping class group $\mathrm{Mod}(S)$ acts on Teichmüller space $\mathcal{T}(S)$ by isometries, and W. Thurston discovered a $\mathrm{Mod}(S)$–equivariant compactification of $\mathcal{T}(S)$ by an ideal sphere, the sphere of compactly supported projective measured laminations $\mathbb{P}\mathcal{ML}(S)$. J. McCarthy and A. Papadopoulos have shown that a subgroup G of $\mathrm{Mod}(S)$ has a well defined limit set Λ_G, although it need not be unique or minimal, and that there is a certain enlargement $Z\Lambda_G$ of Λ_G on whose complement G acts properly discontinuously [48]. So such a group has a domain of discontinuity $\Delta_G = \mathbb{P}\mathcal{ML}(S) - Z\Lambda_G$.

In general, the limit set of a subgroup of $\mathrm{Mod}(S)$ has no convex hull to speak of, as there are pairs in $\mathbb{P}\mathcal{ML}(S)$ that are joined by no geodesic in $\mathcal{T}(S)$. Nevertheless, if every pair of points in Λ_G are the negative and positive directions of a geodesic in $\mathcal{T}(S)$, one can define the weak hull $\mathfrak{H}_G$ of Λ_G to be the union of all such geodesics. This is precisely what B. Farb and L. Mosher do in [20], where they develop a notion of convex cocompact mapping class groups. They prove the following

*The second author was supported by an N.S.F. postdoctoral fellowship.

Theorem 1.1. **(Farb–Mosher)** *Given a finitely generated subgroup G of* $\mathrm{Mod}(S)$*, the following statements are equivalent:*

- *Some orbit of G is quasi-convex in $\mathcal{T}(S)$.*
- *Every orbit of G is quasi-convex in $\mathcal{T}(S)$.*
- *G is hyperbolic and there is a G–equivariant embedding $\partial f\colon\ \partial G \to \mathbb{PML}(S)$ with image Λ_G such that the weak hull $\mathfrak{H}_G$ of Λ_G is defined; the action of G on $\mathfrak{H}_G$ is cocompact; and, if $f\colon\ G \to \mathfrak{H}_G$ is any G–equivariant map, then f is a quasi-isometry and the following map is continuous:*

$$\overline{f} = f \cup \partial f\colon\ G \cup \partial G \to \mathcal{T}(S) \cup \mathbb{PML}(S).$$

A finitely generated subgroup of $\mathrm{Mod}(S)$ is said to be **convex cocompact** if it satisfies one of these conditions.

The interest in convex cocompact Kleinian groups is due in part to combined work of A. Marden [37] and D. Sullivan [64] that implies that such groups are precisely those that are quasiconformally stable, meaning that small perturbations of the identity representation are induced by quasiconformal conjugacies.

The allure of convex cocompact mapping class groups is of a manifestly different nature: Farb and Mosher have shown that when S is closed, convex cocompactness for a subgroup $G < \mathrm{Mod}(S)$ is implied by the δ–hyperbolicity of the associated $\pi_1(S)$–extension of G—see [20]. Moreover, in very recent work, U. Hamenstädt has shown that these are equivalent [27]. In particular, if there is a finite $\mathrm{K}(G,1)$, and an embedding $G \to \mathrm{Mod}(S)$ whose image is not convex cocompact and yet whose non-identity elements are all pseudo-Anosov, then the associated surface group extension is a group with no Baumslag–Solitar subgroups, a finite Eilenberg–Mac Lane space, and which fails to be hyperbolic. This would provide a counterexample to a question of M. Gromov—see [29] and [20].

For more on the geometry of these extensions and related groups, we refer the reader to [20], [21], [54], [55], [56], [57], and [6].

Our purpose here is to strengthen the analogy between convex cocompact Kleinian groups and their cousins in the mapping class group. Our first main result is the following

Theorem 1.2. *Given a finitely generated subgroup G of* $\mathrm{Mod}(S)$*, the following statements are equivalent:*

- *G is convex cocompact.*
- *The weak hull $\mathfrak{H}_G$ is defined and G acts cocompactly on $\mathfrak{H}_G$.*
- *Every limit point of G is conical.*
- *G acts cocompactly on $\mathcal{T}(S) \cup \Delta_G$.*

Remark. The definition of $\mathfrak{H}_G$ used here is more general than that described above and is defined for any infinite irreducible G, see Section 4.2.

That G need only act cocompactly on $\mathfrak{H}_G$ to be convex cocompact follows quickly from the fact, proven in Section 4.2, that geodesic triangles lying in a thick part of $\mathcal{T}(S)$ are thin in the sense of δ–hyperbolic metric spaces: if G acts cocompactly on $\mathfrak{H}_G$ it is coarsely dense therein and the weak hull lies in a thick part of $\mathcal{T}(S)$; the thin triangle condition on $\mathfrak{H}_G$ implies that it is quasi-convex [Theorem 4.5], and a G–orbit is quasi-convex as a result. That triangles lying in a thick part are thin relies on H. Masur's Asymptotic Rays Theorem [39] and Y. Minsky's Contraction Theorem [51].

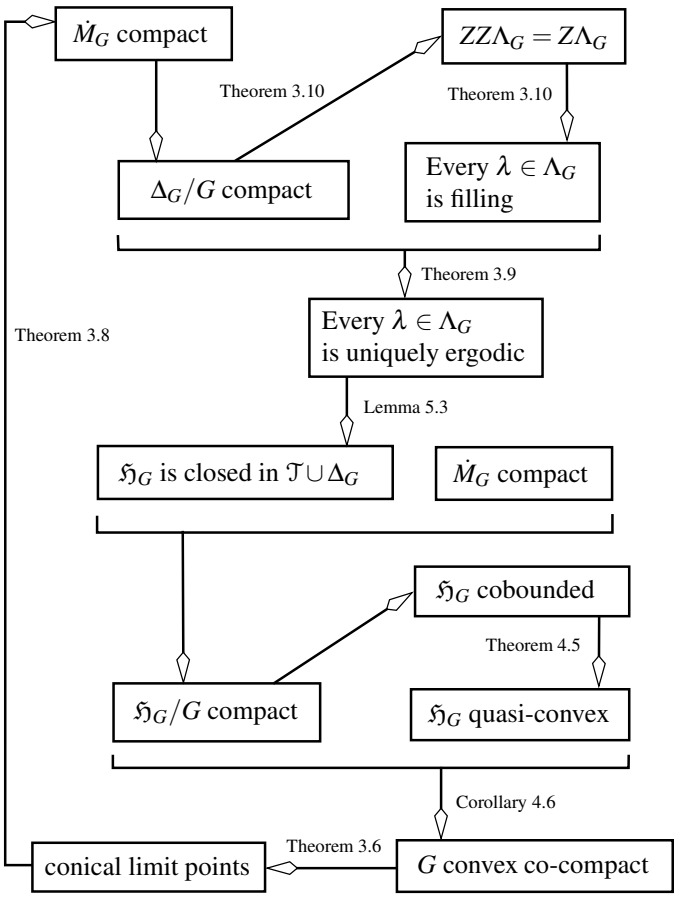

Figure 1: Convex cocompactness, Kleinian manifolds, and conical limit points.

If Γ is a Kleinian group, a limit point for Γ is said to be **conical** if every geodesic in $\mathbb{H}^3$ terminating there has a neighborhood that intersects a Γ–orbit in an infinite set. In the mapping class group, the definition of conical requires some care as there are points in the boundary of Thurston's compactification of $\mathcal{T}(S)$ that are not limits of

Teichmüller geodesic rays. Nonetheless, points exhibiting this behavior are irrelevant by Masur's Two Boundaries Theorem [41], and it is easily seen that convex cocompact groups have all limit points conical [Theorem 3.6].

With the aid of F. Bonahon's work on geodesic currents [8], the arguments given by McCarthy and Papadopoulos to prove that G acts properly discontinuously on Δ_G can be extended to prove proper discontinuity on $\mathcal{T}(S) \cup \Delta_G$. We write $\dot{M}_G = (\mathcal{T}(S) \cup \Delta_G)/G$ and refer to this as the **Kleinian manifold** for G. Along with certain length and intersection number comparisons along Teichmüller geodesic rays, these extended arguments prove that if all limit points are conical, then $\dot{M}_G$ is compact [Theorem 3.8]. The only remaining implication is that having a compact Kleinian manifold implies convex cocompactness.

Minsky's Bounded Geometry Theorem [53] for a doubly degenerate hyperbolic 3–manifold with a type preserving homeomorphism to $S \times \mathbb{R}$ says that the length of the shortest geodesic of such a manifold is bounded below if and only if the Masur–Minsky subsurface projection coefficients of its ending laminations are uniformly bounded above. K. Rafi has proven the analog of this theorem for Teichmüller geodesics [59]: namely, a geodesic lies in a thick part of $\mathcal{T}(S)$ if and only if all of the subsurface projection coefficients of its defining laminations are uniformly bounded.

The set $Z\Lambda_G$ is the set of laminations having zero intersection number with some lamination in Λ_G. The set $ZZ\Lambda_G$ is the set of laminations having zero intersection number with some element of $Z\Lambda_G$. We may continue this procedure to obtain a sequence of sets $Z^{(n)}\Lambda_G$. When a subgroup G of $\mathrm{Mod}(S)$ acts cocompactly on Δ_G, $Z\Lambda_G$ is stable under this operation of taking zero loci and every lamination in Λ_G is filling [Theorem 3.10]. A cocompact action on Δ_G, in conjunction with Rafi's bounded geometry theorem for Teichmüller geodesics, implies that, in fact, every lamination in Λ_G is uniquely ergodic [Theorem 3.9]. Such groups always have weak hulls that are closed in $\mathcal{T}(S) \cup \Delta_G$ [Lemma 5.3] and compactness of $\mathfrak{H}_G/G$ follows from compactness of $\dot{M}_G$. The logic of the proof of Theorem 1.2 is depicted in Figure 1.

Theorem 3.9 provides much stronger information than is needed to prove Theorem 1.2. We state it here as it may be of independent interest.

Theorem 3.9. *Let G be a subgroup of* $\mathrm{Mod}(S)$. *If $\Delta_G \neq \emptyset$ and G acts cocompactly on Δ_G, then every lamination in Λ_G is uniquely ergodic, $Z\Lambda_G = \Lambda_G$, and $\mathfrak{H}_G$ is defined and cobounded. Furthermore, G has a finite index subgroup all of whose non-identity elements are pseudo-Anosov.*

An earlier proof that convex cocompact mapping class groups have compact Kleinian manifolds mirrored the proof in the Kleinian group setting and revealed that weak hulls lying in a thick part of Teichmüller space have closest points projections with contraction properties similar to convex hulls in $\mathbb{H}^3$, generalizing the quasi-projection theorems of Minsky—we have preserved this projection theorem in Section 6.

An obstacle to shining light on $\mathrm{Mod}(S)$ presents itself when one has taken a point of view based on the analogy between $\mathcal{T}(S)$ and $\mathbb{H}^3$: the Teichmüller space with the Teichmüller metric is not hyperbolic in any reasonable sense of the word [38, 47] (nor is it hyperbolic with any reasonable $\mathrm{Mod}(S)$–invariant metric [13]). Indeed, even if the map sending a subgroup G of $\mathrm{Mod}(S)$ to its orbit in $\mathcal{T}(S)$ is a quasi-isometric

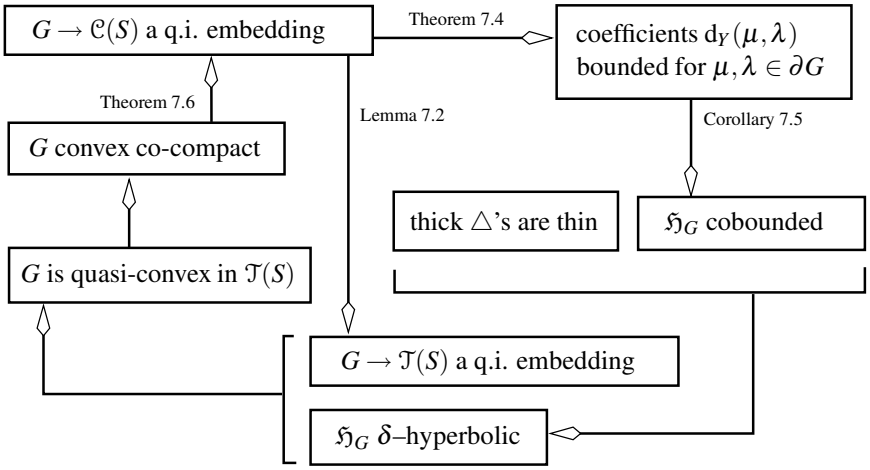

Figure 2: Quasi-isometrically embedding in $\mathcal{C}(S)$ and convex cocompactness.

embedding, G need not be convex cocompact—not even when G is cyclic [46]. On the other hand, $\mathrm{Mod}(S)$ acts by isometries on W. Harvey's complex of curves $\mathcal{C}(S)$, which is δ–hyperbolic by a celebrated theorem of H. Masur and Y. Minsky [44, 11]. Illuminating $\mathrm{Mod}(S)$ from this vantage point has some advantages over the view from $\mathcal{T}(S)$—as well as disadvantages due to the fact that $\mathcal{C}(S)$ fails to be locally compact. Our second main theorem is the following.

Theorem 1.3. *A finitely generated subgroup G of $\mathrm{Mod}(S)$ is convex cocompact if and only if sending G to an orbit in the complex of curves defines a quasi-isometric embedding $G \to \mathcal{C}(S)$.*

Remark. This theorem was independently discovered by U. Hamenstädt [27].

The proof that convex cocompact groups have this property is a novel application of Minsky's Contraction Theorem combined with the fact that $\mathcal{C}(S)$ is quasi-isometric to the electric Teichmüller space $\mathcal{T}_{\mathrm{el}}(S)$.

Given a quasi-isometric embedding $G \to \mathcal{C}(S)$, we obtain from general principles that G is δ–hyperbolic and that the given map extends continuously to a map

$$G \cup \partial G \to \mathcal{C}(S) \cup \partial \mathcal{C}(S)$$

whose restriction to ∂G is an embedding. The boundary of the complex of curves is naturally parameterized by the space $\mathcal{EL}(S)$ of potential ending laminations for geometrically infinite hyperbolic 3–manifolds homeomorphic to $S \times \mathbb{R}$, by a theorem of E. Klarreich [33]. The space $\mathcal{EL}(S)$ sits naturally—as the set of filling laminations—in the quotient of $\mathbb{PML}(S)$ obtained by forgetting transverse measures.

Using hyperbolicity of $\mathcal{C}(S)$ and Masur and Minsky's Bounded Geodesic Image Theorem [45], we are able to uniformly bound the projection coefficients for endpoints

in $\partial\mathcal{C}(S)$ of bi-infinite geodesics in G. Rafi's Bounded Geometry Theorem and a proposition of Klarreich allow us to lift ∂G to the limit set Λ_G and demonstrate that the weak hull $\mathfrak{H}_G$ is defined and cobounded. The fact that triangles in a thick part are thin again tells us that $\mathfrak{H}_G$ is δ–hyperbolic. The quasi-isometric embedding $G \to \mathcal{C}(S)$ yields a quasi-isometric embedding $G \to \mathfrak{H}_G$ and we conclude that G is quasi-convex in $\mathcal{T}(S)$ by hyperbolicity of the hull. See Figure 2.

In [20], Farb and Mosher prove that when S is closed, a free subgroup of $\mathrm{Mod}(S)$ is convex cocompact if and only if the associated surface-by-free group is δ–hyperbolic. Such subgroups are called **Schottky**. In [57], Mosher proves that for any finite set of independent pseudo-Anosov mapping classes $\varphi_1, \dots, \varphi_n$—meaning that the fixed points are pairwise distinct—there is a natural number m such that the surface-by-free group associated to $\varphi_1^m, \dots, \varphi_n^m$ is δ–hyperbolic. This demonstrates the abundance of Schottky groups. Theorem 1.3 yields a new proof of this fact (without the closed hypothesis on S).

Theorem 1.4 (Abundance of Schottky groups [20]). *Given a finite set of independent pseudo-Anosov mapping classes $\varphi_1, \dots, \varphi_n$, there is a number ℓ so that for all natural numbers $m > \ell$, the group generated by $\varphi_1^m, \dots, \varphi_n^m$ is Schottky.*

Proof. We refer the reader to Section 2 for terminology.

Let $\mathfrak{G}_1, \dots, \mathfrak{G}_n$ denote the Cayley graphs of $\langle\varphi_1\rangle, \dots, \langle\varphi_n\rangle$, respectively. Fix $\alpha \in \mathcal{C}(S)$ and $\langle\varphi_i\rangle$–equivariant embeddings $\mathfrak{G}_i \to \mathcal{C}(S)$ by sending each vertex to the associated point of the orbit of α and sending edges to geodesics joining the images of their endpoints. By Theorem 4.6 of [44], these are all (K_0, C_0)–quasi-isometric embeddings for some $K_0 \geq 1$ and $C_0 \geq 0$. Moreover, since $\varphi_1, \dots, \varphi_n$ are independent, all the endpoints of these quasi-geodesic rays in $\partial\mathcal{C}(S)$ are distinct.

For any $m > 0$, we consider the abstract free group $\langle \varphi_1^m, \dots, \varphi_n^m \,|\, — \rangle$ equipped with the metric $m\,\mathrm{d}$, where d is the word metric with respect to the φ_i^m. Note that we naturally obtain a metric on the associated Cayley graph $\mathfrak{G}(m)$. There is a canonical isometric embedding $\mathfrak{G}_i \to \mathfrak{G}(m)$ restricting to the "identity" on $\langle\varphi_i^m\rangle$. Our chosen embeddings of $\mathfrak{G}_i \to \mathcal{C}(S)$ induces a $\langle \varphi_1^m, \dots, \varphi_n^m \,|\, — \rangle$–equivariant map $\mathfrak{G}(m) \to \mathcal{C}(S)$.

Now, given two quasi-geodesic rays τ_1 and τ_2 with a common origin in a δ–hyperbolic metric space $\mathcal{X}$ representing distinct points in $\partial\mathcal{X}$, the union $\tau_1 \cup \tau_2$ (suitably parameterized) is a quasi-geodesic line with constants depending only on the τ_i, their quasi-geodesic constants, and δ.

Fix $K \geq K_0 \geq 1$ and $C \geq C_0 \geq 1$ quasi-geodesic constants for the embeddings of $\mathfrak{G}_i$ into $\mathcal{C}(S)$ and all quasi-geodesic lines obtained as unions of distinct rays in unions $\mathfrak{G}_i \cup \mathfrak{G}_j$ (via the embeddings into $\mathcal{C}(S)$).

Since $\mathcal{C}(S)$ is δ–hyperbolic for some δ, there is an $R \geq 1$ such that any (K, C, R)–local-quasi-geodesic is a quasi-geodesic, see Théorème 1.4 of [18]. If $m \geq R$, then every geodesic segment of length R through $\mathbf{1}$ in $\mathfrak{G}(m)$ is contained in a union $\mathfrak{G}_i \cup \mathfrak{G}_j$. In fact, such a segment is contained in a union of two geodesic rays from $\mathbf{1}$ contained in $\mathfrak{G}_i$ and $\mathfrak{G}_j$ respectively and by the choice of K and C, this segment is sent to a (K, C)–quasi-geodesic segment in $\mathcal{C}(S)$.

As any geodesic segment in $\mathfrak{G}(m)$ may be translated to a segment through $\mathbf{1}$, we conclude that every geodesic in $\mathfrak{G}(m)$ is sent to a (K, C, R)–local-quasi-geodesic in

$\mathcal{C}(S)$, and that $\langle \varphi_1^m, \dots, \varphi_n^m \,|\!-\!-\rangle$ is quasi-isometrically embedded by its orbit in $\mathcal{C}(S)$ provided $m \geq R$. In particular, $\langle \varphi_1^m, \dots, \varphi_n^m \rangle \to \mathrm{Mod}(S)$ is injective, and has convex cocompact image by Theorem 1.3. □

Remark. Although Farb and Mosher work with closed surfaces in considering convex cocompactness, their definitions carry over verbatim for the case of finite area hyperbolic surfaces. Theorem 1.1 also easily extends to this setting. A quick way to see this (that requires no verification of Farb-Mosher's techniques) is to observe that the Teichmüller spaces of punctured surfaces isometrically embed in the Teichmüller spaces of closed surfaces with nice mapping class group equivariance properties (via appropriate branched covers).

In addition, using this method, Farb and Mosher's proof of corollary 1.4 easily implies the corollary for finite area hyperbolic surfaces. Furthermore, we note that this trick allows any example of a convex cocompact subgroups of finite area hyperbolic surfaces to be promoted virtually to examples in closed surface mapping class groups.

Acknowledgments. The authors thank Jeff Brock, Jason DeBlois, Moon Duchin, Benson Farb, Cameron Gordon, Yair Minsky, Lee Mosher, Alan Reid, and Peter Storm for useful conversations. We also thank the referee for several helpful comments and suggestions.

2 Background

2.1 Coarse geometry

Given metric spaces $\mathcal{X}$ and $\mathcal{Y}$ and constants $K \geq 1$ and $C \geq 0$, a map $f\colon \mathcal{X} \to \mathcal{Y}$ is a (K,C)**–quasi-isometric embedding** if

$$\frac{1}{K}\mathrm{d}_{\mathcal{X}}(a,b) - C \leq \mathrm{d}_{\mathcal{Y}}(f(a),f(b)) \leq K\mathrm{d}_{\mathcal{X}}(a,b) + C$$

for all a and b in $\mathcal{X}$, and a (K,C)**–quasi-isometry** if its image is A–dense for some A. Such a map is said to be K**–bi-Lipschitz** if $C = 0$.

A map $f\colon \mathcal{X} \to \mathcal{Y}$ is said to be (K,C)**–coarsely–Lipschitz** if

$$\mathrm{d}_{\mathcal{Y}}(f(a),f(b)) \leq K\mathrm{d}_{\mathcal{X}}(a,b) + C$$

for all a and b in $\mathcal{X}$, and K**–Lipschitz** if $(K,0)$–coarsely–Lipschitz.

A map from an interval in $\mathbb{R}$ or $\mathbb{Z}$ to a metric space $\mathcal{X}$ is a (K,C)**–quasi-geodesic** if it is a (K,C)–quasi-isometric embedding and a (K,C,R)**–local–quasi-geodesic** if its restriction to any interval of length R is a (K,C)–quasi-isometric embedding.

If $\mathcal{X}$ is a geodesic metric space, $\mathcal{Y} \subset \mathcal{X}$, then $\mathcal{Y}$ is said to be A**–quasi-convex** if every geodesic joining two points in $\mathcal{Y}$ is contained in the A–neighborhood $\mathcal{N}_A(\mathcal{Y})$ of $\mathcal{Y}$.

Given a finitely generated group G with finite generating set $\mathcal{U}$, let $\mathrm{d}_{\mathcal{U}}$ denote the induced word metric. A **geodesic** $\mathcal{G}$ in $(G,\mathrm{d}_{\mathcal{U}})$ is a $(1,0)$–quasi-geodesic defined on an interval $\mathbb{I} \subset \mathbb{Z}$. We represent $\mathcal{G}$ by a sequence of group elements $\mathcal{G} = \{h_j\}_{j\in\mathbb{I}}$ and emphasize that the defining characteristic of being a $(1,0)$–quasi-geodesic is that

$$\mathrm{d}_{\mathcal{U}}(h_j,h_k) = |j-k|.$$

Note that, given any point $h_i \in \mathcal{G}$, we can translate $\mathcal{G}$ to a geodesic through the identity $\mathbf{1}$ taking h_i to $\mathbf{1}$, namely

$$h_i^{-1}(\mathcal{G}) = \{h_i^{-1}h_j\}_{j\in\mathbb{I}}.$$

2.2 δ–hyperbolic spaces

We refer the reader to [18] and [12] for more on hyperbolic metric spaces and coarse geometry.

A geodesic triangle in a metric space is **δ–thin** if each of its sides is contained in the δ–neighborhood of the union of the other two sides.

A geodesic metric space is **δ–hyperbolic** in the sense of M. Gromov and J. Cannon if every geodesic triangle is δ–thin.

Let $\mathcal{X}$ be a metric space. Given x, y, and z in $\mathcal{X}$, the **Gromov product of y and z** with respect to x is defined to be

$$(y\cdot z)_x = \frac{1}{2}(\mathrm{d}_{\mathcal{X}}(y,x) + \mathrm{d}_{\mathcal{X}}(z,x) - \mathrm{d}_{\mathcal{X}}(y,z)).$$

Fix a basepoint x in $\mathcal{X}$. A sequence $\{x_n\}$ **converges at infinity** in $\mathcal{X}$ if

$$\lim_{m,n\to\infty}(x_m\cdot x_n)_x = \infty$$

and two sequences $\{x_n\}$ and $\{y_m\}$ are **equivalent** if

$$\lim_{m,n\to\infty}(y_m\cdot x_n)_x = \infty.$$

If $\mathcal{X}$ is a δ–hyperbolic geodesic metric space, the **Gromov boundary** $\partial\mathcal{X}$ of $\mathcal{X}$ is the set of equivalence classes of sequences in $\mathcal{X}$ that converge at infinity. The set $\overline{\mathcal{X}} = \mathcal{X}\cup\partial\mathcal{X}$ admits a natural topology in which a sequence $\{x_n\}$ in $\mathcal{X}$ converges to a point $y = \{y_n\}$ in $\partial\mathcal{X}$ if and only if $\{x_n\}$ is equivalent to y.

A geodesic ray based at x uniquely determines a point in $\partial\mathcal{X}$ given by any sequence of points on the ray that converges at infinity.

2.3 Teichmüller theory

We refer the reader to [3, 1, 23] for more on quasiconformal mappings and Teichmüller theory. In what follows, unless otherwise stated, all Riemann surfaces and complex/conformal structures are of finite analytic type, and hyperbolic surfaces and structures are complete with finite area.

Let X be a Riemann surface homeomorphic to S. We view X as either a complex 1–manifold, or an oriented hyperbolic 2–manifold—the Uniformization Theorem permits us to change this view at will. A **marking** of X is an orientation preserving homeomorphism $f\colon S\to X$, and the Teichmüller space $\mathcal{T}(S)$ of S is the set of equivalence classes of marked Riemann surfaces $f\colon S\to X$. The equivalence relation is defined by declaring $f_1\colon S\to X_1$ to be equivalent to $f_2\colon S\to X_2$ if $f_2\circ f_1^{-1}$ is isotopic to an isomorphism of Riemann surfaces.

Abusing notation, we often refer to a Riemann surface X as a point in Teichmüller space, by which we mean the equivalence class of X implicitly marked by some homeomorphism.

We may also think of $\mathcal{T}(S)$ as the space of complex, conformal, or hyperbolic structures on S, up to isotopy, as such a structure is induced on S by pulling back via the marking.

Let X_1 and X_2 be two Riemann surfaces equipped with markings $f_1\colon S \to X_1$ and $f_2\colon S \to X_2$. A homeomorphism $f\colon X_1 \to X_2$ is **K–quasiconformal** if it is absolutely continuous on lines and $|f_{\bar{z}}| \leq k|f_z|$ in every local coordinate z where $k = (K-1)/(K+1) < 1$. The minimal value of K for which f is K–quasiconformal is the **dilatation** of f and is denoted $K[f]$. The **Teichmüller distance** between X_1 and X_2 is defined to be

$$\mathrm{d}_{\mathcal{T}}(X_1, X_2) = \frac{1}{2} \inf \log K[f]$$

where the infimum is taken over all quasiconformal maps f isotopic to $f_2 \circ f_1^{-1}$. There is a unique extremal quasiconformal map $X_1 \to X_2$ realizing the above distance, called the **Teichmüller mapping**.

A **holomorphic quadratic differential** q on X is an assignment of a holomorphic function $\varphi(z)$ to each local coordinate z such that for two coordinates z_1 and z_2,

$$\varphi_1(z_1)(\mathrm{d}z_1/\mathrm{d}z_2)^2 = \varphi_2(z_2).$$

We say that q is **integrable** if $\int_X |q| < \infty$. We equip the vector space of integrable holomorphic quadratic differentials $\mathcal{Q}(X)$ with the norm $\| \ \| = \int_X | \ |$. Varying X over $\mathcal{T}(S)$ and assembling the vector spaces $\mathcal{Q}(X)$ one obtains a vector bundle $\mathcal{Q}(S)$ over $\mathcal{T}(S)$. We let $\mathcal{Q}^1(S)$ denote the associated unit sphere bundle over $\mathcal{T}(S)$ and $\mathcal{Q}^*(S)$ the complement of the zero section of $\mathcal{Q}(S) \to \mathcal{T}(S)$. Throughout the remainder of this paper, we will refer to a point of $\mathcal{Q}^*(S)$ simply as a quadratic differential, with the holomorphic, integrability, and non-zero conditions implicit. We denote a quadratic differential by q or sometimes (X, q) if we wish to emphasize the underlying complex structure $X \in \mathcal{T}(S)$.

A **natural parameter** at a point p associated to a quadratic differential $q = \varphi(z)\,\mathrm{d}z^2$ is given by

$$\zeta(w) = \int_{z_0}^{w} \sqrt{\varphi(z)}\,\mathrm{d}z,$$

where $z_0 = z(p)$. Away from the zeroes of q, ζ is a *bona fide* coordinate for which q takes the particularly simple form $q = \mathrm{d}\zeta^2$.

Pulling back the horizontal and vertical foliations of $\mathbb{C}$ via a natural parameter ζ yields a pair of transverse measured (singular) foliations $\mathcal{F}_-(q)$ and $\mathcal{F}_+(q)$ associated to q, where the transverse measures are the pullbacks of the vertical and horizontal total variations in $\mathbb{C}$, respectively. A branched flat metric, called the **q–metric**, is obtained from ζ by pulling back the Euclidean metric from $\mathbb{C}$. The total area of X with respect to the q-metric is $\|q\|$.

The Teichmüller mapping between Riemann surfaces may be described explicitly in terms of a holomorphic quadratic differential. Fix such a differential q with natural parameter ζ and a number $0 \le k < 1$. The (q,k)**–Teichmüller deformation** of X is the Riemann surface X' with underlying surface X and complex structure given by the local parameter $\zeta' = (\zeta + k\overline{\zeta})/(1-k)$. As X' is determined by X, q, and k, we let (X,q,k) denote X'. The quadratic differential $q' = (d\zeta')^2/\|(d\zeta')^2\|$ is called the **terminal differential** of the deformation. The horizontal and vertical measured foliations of the terminal differential are given by $\mathcal{F}_-(q') = K^{-1/2}\mathcal{F}_-(q)$ and $\mathcal{F}_+(q') = K^{1/2}\mathcal{F}_+(q)$, where $K = (1+k)(1-k)$. The "identity" $X \to X'$ is the Teichmüller mapping in its homotopy class, and $\mathrm{d}_{\mathcal{T}}(S,S') = \frac{1}{2}\log K$. Teichmüller's Theorem asserts that for any $X \in \mathcal{T}(S)$ the map from the unit ball $B_1 \subset \mathcal{Q}(X)$ to $\mathcal{T}(S)$ given by $q \mapsto (X,q,\|q\|)$ is a homeomorphism.

Given a quadratic differential q on X, define the **Teichmüller geodesic** $\tau = \tau_q\colon \mathbb{R} \to \mathcal{T}(S)$ as follows. For $t \in \mathbb{R}$, let s_t be given by $t = \frac{1}{2}\log((1+s_t)/(1-s_t))$, and define

$$\tau(t) = \begin{cases} (X,q,s_t) & \text{for } t > 0 \\ (X,-q,-s_t) & \text{otherwise} \end{cases}$$

The parameter is chosen so that τ is a geodesic parameterized by arc–length with respect to the Teichmüller metric—note that the terminal differential at time t has horizontal and vertical measured foliations $\mathcal{F}_-(q_t) = e^{-t}\mathcal{F}_-(q)$ and $\mathcal{F}_+(q_t) = e^{t}\mathcal{F}_+(q)$. We refer to the projective classes of the horizontal and vertical foliations of q as the horizontal and vertical foliations of τ.

The mapping class group $\mathrm{Mod}(S)$ acts on $\mathcal{T}(S)$ by pulling back hyperbolic structures. This is a properly discontinuous action by isometries of the Teichmüller metric and the quotient $\mathcal{M}(S)$ is the **moduli space** of S.

The ε**–thick part** of $\mathcal{T}(S)$ is the set of hyperbolic structures on S for which the length of the shortest geodesic is greater than $\varepsilon > 0$. A set is said to be ε**–cobounded** if it lies in the ε–thick part of $\mathcal{T}(S)$ and **cobounded** if it is ε–cobounded for some $\varepsilon > 0$. D. Mumford's Compactness Criterion [58] says that a set is ε–cobounded for some ε if and only if it projects to a precompact set in $\mathcal{M}(S)$, see also [15].

2.4 Laminations and foliations

For a nice survey of geodesic laminations, see [9].

Fix a hyperbolic metric on S. A **geodesic lamination** on S is a nonempty compact subset of S that is a union of pairwise disjoint simple complete local geodesics on S, called the **leaves** of the lamination. By a simple complete local geodesic, we mean the image of an injective locally isometric immersion of $\mathbb{R}$. A **transverse measure** on a geodesic lamination is an assignment of a Radon measure to each arc (and so each curve) α transverse to the lamination such that (1) the measure on a subarc α' of an arc α is the restriction to α' of the measure on α and (2) so that two arcs are assigned the same measure if they are isotopic through arcs transverse to the lamination. A **measured lamination** λ is a geodesic lamination $|\lambda|$ called the **support** of λ, along with a transverse measure. We will further always require that our transverse measures have **full support**: if the intersection of a transverse arc α with $|\lambda|$ is non-empty, then

the measure on α is non-zero (equivalently, the support of the measure on α is exactly $\alpha \cap |\lambda|$).

The set of measured laminations admits a natural topology, see [9], and the resulting space is denoted $\mathcal{ML}(S)$. The set of non-zero laminations up to scale, the **projective measured laminations**, is denoted $\mathbb{P}\mathcal{ML}(S)$. We let $\mathcal{UML}(S)$ denote the quotient of the set of measured laminations obtained by forgetting transverse measures. The space $\mathcal{ML}(S)$ depends on the choice of hyperbolic metric, but there is a canonical identification between spaces obtained by different metrics.

Certain definitions for (or properties of) $\lambda \in \mathcal{ML}(S)$ may depend only on its projective class $[\lambda]$ or its support $|\lambda|$. In these cases, we will freely apply such definitions (or refer to such properties) for $[\lambda]$ or $|\lambda|$, as is appropriate. Occasionally, we confuse a measured lamination λ with its projective class $[\lambda]$, and even its support $|\lambda|$, referring to all three simply as λ when the meaning is clear.

Let $\mathcal{S}$ be the set of isotopy classes of essential simple closed curves on S. Essential here means homotopically nontrivial and nonperipheral. Taking geodesic representatives with transverse counting measures, we identify $\mathcal{S}$ with a subset of $\mathcal{ML}(S)$. The geometric intersection number $i(\cdot,\cdot)\colon\ \mathcal{S}\times\mathcal{S}\to\mathbb{R}_+=[0,\infty)$ extends naturally to a continuous function

$$i(\cdot,\cdot)\colon\ \mathcal{ML}(S)\times\mathcal{ML}(S)\to\mathbb{R}_+\,.$$

When λ is a measured lamination and α is a simple closed curve, $i(\lambda,\alpha)$ is simply the total mass of transverse measure on α assigned by λ. The injection $\mathcal{S}\to\mathcal{ML}(S)$ remains injective upon descending to the quotients $\mathcal{S}\to\mathbb{P}\mathcal{ML}(S)$ and $\mathcal{S}\to\mathcal{UML}(S)$.

A measured lamination λ is said to be **filling** if it intersects non-trivially any measured lamination whose support is different from that of λ. This property for λ depends only on $|\lambda|$.

There is a closely related theory of (singular) **measured foliations** on S. There is a space $\mathcal{MF}(S)$ of (measure classes of) measured foliations, an intersection function $i:\mathcal{MF}(S)\times\mathcal{MF}(S)\to\mathbb{R}$, and a space of such foliations up to scale, $\mathbb{P}\mathcal{MF}(S)$. The horizontal and vertical foliations of a holomorphic quadratic differential $q\in\mathcal{Q}(S)$ are examples of transversely measured foliations. In fact, it is a theorem of J. Hubbard and H. Masur [28] and S. Kerckhoff [32] that for any fixed complex structure, each (measure class) of measured foliation appears as the vertical foliation of a unique holomorphic quadratic differential.

As discussed in the previous subsection, for any $X\in\mathcal{T}(S)$, Teichmüller's Theorem provides a homeomorphism from the open unit ball in $\mathcal{Q}(X)$ to $\mathcal{T}(S)$. Therefore, the closed unit ball serves as a visual compactification of $\mathcal{T}(S)$. Identifying a quadratic differential with the projective class of its vertical foliation, we interpret this as a compactification by $\mathbb{P}\mathcal{MF}(S)$. This is **Teichmüller's compactification** of $\mathcal{T}(S)$.

We have chosen to work primarily with measured laminations rather than foliations, and so refer the reader to [22] for a detailed treatment of the latter. However, we need to know that there is a "dictionary" between $\mathcal{ML}(S)$ and $\mathcal{MF}(S)$—see G. Levitt's paper [36] for details. Given a measured foliation F, there is an associated lamination λ_F, obtained from F by straightening the leaves of F. The foliation F can be recovered from λ_F by a certain collapsing procedure applied to the complementary regions of λ_F

(when λ_F has isolated leaves, one must first replace such a leaf by an annulus foliated by parallel copies of the core). We call λ_F the **lamination associated to** F, and F the **foliation associated to** λ_F.

The identification $\mathcal{ML}(S) \cong \mathcal{MF}(S)$ is completely natural. Specifically, we have $i(F_1, F_2) = i(\lambda_{F_1}, \lambda_{F_2})$ and for every $t > 0$ and $F \in \mathcal{MF}(S)$, one has $\lambda_{tF} = t\lambda_F$. Moreover, for any $g \in \mathrm{Mod}(S)$, we have $g\lambda_F = \lambda_{gF}$.

Given a Teichmüller geodesic τ defined by a holomorphic quadratic differential q with horizontal and vertical foliations $\mathcal{F}_-$ and $\mathcal{F}_+$, we call the associated laminations $\lambda_- = \lambda_{\mathcal{F}_-}$ and $\lambda_+ = \lambda_{\mathcal{F}_+}$ (or their projective classes) the **negative and positive directions** of τ, respectively. If τ is a geodesic ray, the lamination λ_+ associated to $\mathcal{F}_+$ is called the **direction** of τ. The boundary of Teichmüller's compactification of $\mathcal{T}(S)$ is thus identified with $\mathbb{PML}(S)$ as the directions of rays emanating from a point $X \in \mathcal{T}(S)$.

If two projective measured geodesic laminations $[\lambda_-], [\lambda_+]$ **bind** S, meaning that every complementary region of the union of their supports is a disk or once-punctured disk, then they are the negative and positive directions of a unique Teichmüller geodesic which we denote $\tau(\lambda_-, \lambda_+)$. We note that the binding condition depends only on $|\lambda_-|, |\lambda_+|$, while $\tau(\lambda_-, \lambda_+)$ depends on the projective classes of measures.

We say that a geodesic lamination is **uniquely ergodic** if it is filling and supports exactly one transverse measure up to scale. By definition, being uniquely ergodic depends only on the support of the lamination.

2.5 Lengths, intersection numbers, and geodesic currents

There is a function

$$\ell\colon \mathcal{T}(S) \times \mathcal{S} \to \mathbb{R}_+$$

that assigns a pair (X, α) the length $\ell_X(\alpha)$ of the geodesic representative of α with respect to the hyperbolic metric X. This function extends continuously and homogeneously (in the second factor) to a function

$$\ell\colon \mathcal{T}(S) \times \mathcal{ML}(S) \to \mathbb{R}_+$$

called the **hyperbolic length** function, see [9].

There is also a function

$$\mathrm{ext}\colon \mathcal{T}(S) \times \mathcal{S} \to \mathbb{R}_+$$

that assigns a pair (X, α) the extremal length $\mathrm{ext}_X(\alpha)$ of the curve α in the Riemann surface X—$\mathrm{ext}_X(\alpha)$ is the infimum of the reciprocals of conformal moduli of embedded annuli about α. Kerckhoff proves in [32] that this function has a continuous square–homogeneous (in the second factor) extension to $\mathcal{MF}(S)$, namely

$$\mathrm{ext}\colon \mathcal{T}(S) \times \mathcal{MF}(S) \to \mathbb{R}.$$

Moreover, if $(X, q) \in \mathcal{Q}^*(S)$, then

$$\mathrm{ext}_X(\mathcal{F}_\pm(q)) = \|q\|$$

This allows the following characterization of the Teichmüller metric in terms of extremal length.

Theorem 2.1 (Kerckhoff [32]). *For any X and Y in $\mathcal{T}(S)$*

$$\mathrm{d}_{\mathcal{T}}(X,Y) = \frac{1}{2}\sup\log\left(\frac{\mathrm{ext}_X(\alpha)}{\mathrm{ext}_Y(\alpha)}\right),$$

where the supremum is taken over all α in $\mathcal{MF}(S)$.

Both $\mathcal{T}(S)$ and $\mathcal{ML}(S)$ admit proper embeddings into F. Bonahon's space $\mathfrak{C}(S)$ of **geodesic currents**: the space of $\pi_1(S)$–invariant Radon measures on the space of geodesics in the universal cover of S—we recommend the beautiful [8] for details of what follows.

There is a natural $\mathbb{R}_+$–action on $\mathfrak{C}(S)$ and the quotient $\mathbb{P}\mathfrak{C}(S)$ of $\mathfrak{C}(S)-\{0\}$ by this action is compact. The above embeddings descend to embeddings of $\mathcal{T}(S)$ and $\mathbb{P}\mathcal{ML}(S)$ into $\mathbb{P}\mathfrak{C}(S)$. There is an "intersection function":

$$\iota\colon\ \mathfrak{C}(S)\times\mathfrak{C}(S)\to\mathbb{R},$$

which is a continuous symmetric homogeneous (in each factor) function with the following properties.

Theorem 2.2 (Bonahon). *Identifying $\mathcal{T}(S)$ and $\mathcal{ML}(S)$ with their images in $\mathfrak{C}(S)$, if $\mu,\nu\in\mathcal{ML}(S)$ and $X\in\mathcal{T}(S)$, then:*

1. $\iota(\mu,\nu) = i(\mu,\nu)$
2. $\iota(X,\mu) = \ell_X(\mu)$
3. $\iota(X,X) = \pi^2|\chi(S)|$

Moreover, $\mathcal{ML}(S)$ consists of precisely those currents $\eta\in\mathfrak{C}(S)$ for which $\iota(\eta,\eta)=0$.

Taking the closure of $\mathcal{T}(S)$ in $\mathbb{P}\mathfrak{C}(S)$ provides a compactification of $\mathcal{T}(S)$, as $\mathbb{P}\mathfrak{C}(S)$ is compact. Properness of the embedding $\mathcal{T}(S)\to\mathfrak{C}(S)$ with part 3 of Theorem 2.2 guarantees that any point of $\overline{\mathcal{T}(S)}-\mathcal{T}(S)$ in $\mathbb{P}\mathfrak{C}(S)$ is represented by an element $\eta\in\mathfrak{C}(S)$ satisfying $\iota(\eta,\eta)=0$. By the final remark in Theorem 2.2, this is a compactification by $\mathbb{P}\mathcal{ML}(S)$. Bonahon proves that this is precisely **Thurston's compactification** [8], which we write as $\mathcal{T}(S)\cup\mathbb{P}\mathcal{ML}(S)$. We comment that this is related to Teichmüller's compactification (see the next section) but the two compactifications are different [32].

2.6 Three theorems of H. Masur

We use the following theorems of H. Masur repeatedly.

Theorem 2.3 (Criterion for unique ergodicity [43]). *If the direction of a Teichmüller geodesic ray τ is not uniquely ergodic, then the projection of τ to the moduli space $\mathcal{M}(S)$ leaves every compact set.*

A topological ray $[0,\infty)\to\mathcal{M}(S)$ **leaves** (or **exits**) a compact set $\mathcal{K}$ if the pullback of $\mathcal{K}$ to $[0,\infty)$ is compact.

Theorem 2.4 (Two Boundaries Theorem [41]). *A Teichmüller geodesic ray τ with direction a uniquely ergodic lamination λ converges in $\mathcal{T}(S) \cup \mathbb{P}\mathcal{ML}(S)$ to λ.*

So Teichmüller's compactification of $\mathcal{T}(S)$ agrees with Thurston's on the set of uniquely ergodic laminations.

Theorem 2.5 (Asymptotic Rays Theorem [39]). *Let X and Y be points in $\mathcal{T}(S)$ and let σ and τ be Teichmüller geodesic rays from X and Y, respectively, with common direction a uniquely ergodic lamination $[\lambda]$ in $\mathbb{P}\mathcal{ML}(S)$. Then σ and τ are asymptotic.*

Here we say that σ and τ are **asymptotic** if

$$\lim_{t \to \infty} \mathrm{d}_{\mathcal{T}}(\sigma(t), \tau(t)) = 0$$

for an appropriate choice of unit speed parameterization.

2.7 Complexes of curves and arcs

Given a surface Y with compact boundary for which the interior $\mathrm{int}(Y)$ is a surface of genus g with p punctures, we let $\xi(Y) = 3g + p$. We assume throughout that $2 \leq \xi(Y) < \infty$.

A simple closed curve in such a surface Y is **essential** if it is essential in $\mathrm{int}(Y)$. A simple (compact) arc is **essential** if it is homotopically essential relative to ∂Y.

Suppose that $\xi(Y) \geq 5$. Harvey's **complex of curves** $\mathcal{C}(Y)$ of Y is the simplicial complex whose k–cells are collections of isotopy classes of $k+1$ disjoint pairwise non-isotopic essential simple closed curves.

If $\xi(Y) = 4$, then $\mathrm{int}(Y)$ is a sphere with four punctures or a punctured torus. In these cases, we define $\mathcal{C}(Y)$ to be the graph whose vertices are essential simple closed curves and that two vertices are joined by an edge if they intersect in as few a number of points as is possible for two such curves in Y.

When $\xi(Y) \leq 3$, we declare that $\mathcal{C}(Y)$ be empty.

Given a surface Y, it is convenient to have a **complex of arcs** $\mathcal{A}(Y)$. If $\xi(Y) \geq 3$, we define $\mathcal{A}(Y)$ to be the simplicial complex whose k cells are collections of isotopy classes of $k+1$ disjoint pairwise non-isotopic essential simple closed curves and arcs—where isotopy classes of arcs are defined relative to ∂Y.

If Y is a compact annulus, we define $\mathcal{A}(Y)$ to be the graph whose vertices are isotopy classes of essential arcs in Y *relative to their endpoints in* ∂Y and that two vertices are joined by an edge if they may be realized disjointly.

For any Y, we metrize $\mathcal{C}(Y)$ and $\mathcal{A}(Y)$ by demanding that any simplex is a regular Euclidean simplex with all side lengths equal to one and taking the induced path metric.

We note that when $\xi(Y) \leq 4$ and has no boundary, or when Y is a noncompact annulus, we are declaring $\mathcal{A}(Y)$ to be empty.

Although the complex $\mathcal{A}(Y)$ is uncountable when Y is a compact annulus, it is nonetheless quasi-isometric to $\mathbb{Z}$ [45].

2.8 Ivanov–Masur–Minsky subsurface projections

A subsurface Y of S is said to be a **proper domain** if it is not equal to S, it is properly embedded (a compact set of S intersects Y in a compact set), and the induced map on fundamental groups is injective. We further assume that every boundary component of Y is essential (thus the only punctures of Y are also punctures of S). We note that under these assumptions, $\mathcal{A}(Y)$ is always nonempty.

Fix a hyperbolic metric on S. The definitions for measured laminations which follow are equally valid for their projective classes as well as their supports. We note that the projections and projection coefficients defined below do not depend on our choice of hyperbolic metric used to realize laminations as geodesic laminations, but the angle θ and the bound given in Lemma 2.6 do.

Given a proper domain Y in S, pass to the cover $\widetilde{Y}$ of S corresponding to the fundamental group of Y. Adding the domain of discontinuity for $\pi_1(Y)$ to $\mathbb{H}^2$ and taking the quotient, we obtain an enlargement $\overline{Y}$ of $\widetilde{Y}$ which is homeomorphic to Y.

Given $\lambda \in \mathcal{ML}(S)$, we may lift $|\lambda|$ to a (noncompact) geodesic lamination $|\widetilde{\lambda}|$ in $\widetilde{Y}$. Compactifying $|\widetilde{\lambda}|$ in $\overline{Y}$ and identifying any two arcs in the same parallelism class (and disregarding inessential arcs) yields a simplex in $\mathcal{A}(Y)$—when Y is an annulus, parallelism is defined relative to the endpoints and instead of a simplex, we obtain a set with diameter at most one. This simplex (or set in the annulus case) is the **projection of λ to** Y, denoted $\pi_Y(\lambda)$—note that $\pi_Y(\lambda)$ is allowed to be empty.

Given μ and λ in $\mathcal{ML}(S)$, the **projection coefficient for μ and λ at** Y is defined to be

$$\mathrm{d}_Y(\mu,\lambda) = \mathrm{diam}_{\mathcal{A}(Y)}(\pi_Y(\mu) \cup \pi_Y(\lambda))$$

provided that $\pi_Y(\mu) \neq \emptyset$ and $\pi_Y(\lambda) \neq \emptyset$. If either of $\pi_Y(\mu)$ or $\pi_Y(\lambda)$ is empty, we define $\mathrm{d}_Y(\mu,\lambda) = \infty$.

We henceforth write $\mathrm{diam}_Y(\cdot)$ to denote $\mathrm{diam}_{\mathcal{A}(Y)}(\cdot)$.

When Y is an annulus, distance in $\mathcal{A}(Y)$ is determined by the intersection number: if α and β are distinct vertices of $\mathcal{A}(Y)$,

$$\mathrm{d}_Y(\alpha,\beta) = 1 + i(\alpha,\beta).$$

When convenient, we refer to the core of an annulus Y as ∂Y.

Fix a hyperbolic structure on S and λ and μ in $\mathcal{ML}(S)$. Define the angle

$$\theta(\mu,\lambda) = \inf_{x \in \mu \cap \lambda} \theta(x,\mu,\lambda)$$

where $\theta(x,\mu,\lambda) \in [0, \frac{\pi}{2}]$ is the smaller of the two angles between tangent lines to $|\mu|$ and $|\lambda|$ at x.

Lemma 2.6. *Suppose $\alpha \in \mathcal{S}$ is a simple closed geodesic in S and $\mu,\lambda \in \mathcal{ML}(S)$ are two laminations with $\theta = \min\{\theta(\mu,\alpha), \theta(\lambda,\alpha)\} > 0$. If we let Y denote the annulus with $\partial Y = \alpha$, then we have*

$$\mathrm{d}_Y(\mu,\lambda) \leq 2\left\lceil 2\cosh^{-1}(1/\sin(\theta))/\ell(\alpha)\right\rceil + 2.$$

Proof. Let $\widetilde{Y}$ denote the annular cover of S associated to Y, $\widetilde{\alpha}$ the lift of α to the core of $\widetilde{Y}$, and $\widetilde{\mu}$ and $\widetilde{\lambda}$ any lifts of leaves of $|\mu|$ and $|\lambda|$, respectively, that meet $\widetilde{\alpha}$.

Let $\widetilde{\beta}$ denote any geodesic that meets $\widetilde{\alpha}$ orthogonally. By symmetry and the triangle inequality for d_Y, it suffices to prove that

$$i(\widetilde{\mu},\widetilde{\beta}) \leq \left\lceil 2\cosh^{-1}(1/\sin(\theta))/\ell(\alpha)\right\rceil .$$

To verify this inequality, further lift the picture to the universal cover $\mathbb{H}^2 \to \widetilde{Y}$. Let $\widetilde{\alpha}_0$ be a geodesic covering $\widetilde{\alpha}$ that is stabilized by an isometry ζ generating the cyclic covering group. Let $\widetilde{\mu}_0$ denote a lift of $\widetilde{\mu}$ intersecting $\widetilde{\alpha}_0$ in a point x. Set $\widetilde{\beta}_0$ to be any lift of $\widetilde{\beta}$ and note that all other lifts of $\widetilde{\beta}$ are of the form $\zeta^n(\widetilde{\beta}_0)$.

Since every point of intersection of $\widetilde{\mu}$ with $\widetilde{\beta}$ has exactly one lift on $\widetilde{\mu}_0$, we see that

$$i(\widetilde{\mu},\widetilde{\beta}) = \sum_{n\in\mathbb{Z}} i(\widetilde{\mu}_0,\zeta^n(\widetilde{\beta}_0)).$$

However, from elementary trigonometric formulae for right triangles we see that a geodesic orthogonal to $\widetilde{\alpha}_0$ will intersect $\widetilde{\mu}_0$ if and only if the distance from this geodesic to x is less than $\cosh^{-1}(1/\sin(\theta(x,\widetilde{\mu}_0,\widetilde{\alpha}_0)))$. The desired inequality follows from the fact that the translation length of ζ is $\ell(\alpha)$ and that $\theta \leq \theta(x,\widetilde{\mu}_0,\widetilde{\alpha}_0)$. □

2.9 Bounded geometry theorems

Minsky's Bounded Geometry Theorem [53] says that a doubly degenerate hyperbolic 3–manifold homeomorphic to $S\times\mathbb{R}$ has the length of its shortest geodesic bounded below if and only if the subsurface projection coefficients of its ending laminations are all uniformly bounded above.

K. Rafi has characterized the short curves in hyperbolic structures on a Teichmüller geodesic in terms of the intersections of the subsurface projections of its defining laminations [59]. With the global connection between intersection numbers and subsurface projection coefficients described in [45], this yields the following bounded geometry theorem for Teichmüller geodesics—the theorem is implicit in the the proof of Theorem 1.5 of [59].

Theorem 2.7 (Rafi). *For every $D>0$, there exists $\varepsilon>0$ such that if $\tau=\tau(\lambda_-,\lambda_+)$ is a Teichmüller geodesic with λ_- and λ_+ in $\mathbb{PML}(S)$ satisfying*

$$\mathrm{d}_Y(\lambda_-,\lambda_+) \leq D$$

for every proper domain $Y\subset S$ with $\xi(Y)\neq 3$, then τ is ε–cobounded.

Conversely, for every $\varepsilon>0$ there exists $D>0$ such that if τ is an ε–cobounded Teichmüller geodesic with negative and positive directions λ_- and λ_+, then

$$\mathrm{d}_Y(\lambda_-,\lambda_+) \leq D$$

for every proper domain $Y\subset S$ with $\xi(Y)\neq 3$.

3 Dynamics on $\mathbb{PML}(S)$

3.1 Limit sets

Let G be a subgroup of $\mathrm{Mod}(S)$. The definition of a limit set for the action of G on $\mathbb{PML}(S)$ requires more care than its analogue in the realm of Kleinian groups. We elaborate here on this notion in our setting. We primarily follow the notation and conventions of [48].

A **weak limit point** for the action of G on $\mathbb{PML}(S)$ is any point $[\lambda] \in \mathbb{PML}(S)$ such that for some $[\mu] \in \mathbb{PML}(S)$ and some infinite sequence of distinct elements $\{g_n\}_{n=1}^{\infty} \subset G$, we have $[\lambda] = \lim g_n([\mu])$. The **canonical limit set** (for G acting on $\mathbb{PML}(S)$) is the closure of the set of all weak limit points. A **limit set** for G is any closed G-invariant subset of the canonical limit set.

We say that G is **dynamically irreducible** if it has a unique non-empty minimal limit set $\mathbb{PML}(S)$. We call this **the limit set** and we denote it Λ_G. We refer to the points of Λ_G as **limit points** of G. Dynamically irreducible groups fall into two types—see Theorem 4.6 of [48]. The first type are those which contain a finite index cyclic subgroup (we say it is *virtually cyclic*). Moreover, the finite index cyclic subgroup of G is generated by a pseudo-Anosov mapping class and there is another element of G interchanging the stable and unstable fixed points in $\mathbb{PML}(S)$. The second (more interesting) type of dynamically irreducible subgroup contains two independent pseudo-Anosov mapping classes. For a dynamically irreducible group G, Λ_G can be defined as the closure of the set $\Lambda_0(G)$ of stable laminations of pseudo-Anosov elements of G. Set

$$Z\Lambda_G = \{[\mu] \in \mathbb{PML}(S) \,|\, i(\mu,\lambda) = 0 \text{ for some } [\lambda] \in \Lambda_G\}.$$

In [48], it is shown that G acts properly discontinuously on the set

$$\Delta_G = \mathbb{PML}(S) - Z\Lambda_G,$$

called the **domain of discontinuity** for G.

If G is not dynamically irreducible, we say that it is **dynamically reducible**. By Theorem 4.6 of [48], in this case G is either finite, virtually cyclic (virtually) generated by a single pseudo-Anosov mapping class (and contains no element interchanging the stable and unstable fixed points), or is infinite and **reducible**, which means that there is a nonempty G–invariant set $\mathcal{R} \subset \mathcal{S}$ such that for any $\alpha_1, \alpha_2 \in \mathcal{R}$ we have $i(\alpha_1, \alpha_2) = 0$. We call such a set a **reduction system** for G. We pause to elaborate on the structure of reducible subgroups—see Chapter 7 of [30].

If G is infinite and reducible, then there is a **canonical reduction system** for G characterized as the unique largest reduction system $\mathcal{R}$ with the property that if $\beta \in \mathcal{S}$ is any curve with $i(\alpha,\beta) > 0$ for some $\alpha \in \mathcal{R}$, then there exists $g \in G$ so that $\{g^k(\beta)\}_{k=1}^{\infty}$ is an infinite set. Let $S_1, \ldots, S_n$ denote the components of the complement of the curves of $\mathcal{R}$ in S, and we refer to these as the **components of** G. As G leaves $\mathcal{R}$ invariant, one obtains a homomorphism to the mapping class group of the disjoint union of the components

$$G \to \mathrm{Mod}(S_1 \sqcup \ldots \sqcup S_n)$$

by restriction. The kernel of the permutation action on the components is a finite index subgroup $G' < G$ leaving each S_j invariant, and we let $G'|_{S_j}$ denote the restriction of G' to $\mathrm{Mod}(S_j)$. We call a component S_j a **pseudo-Anosov component** if there is an element $g \in G'$ which is pseudo-Anosov in $G'|_{S_j}$. A component S_j is called a **finite component** if the restriction of $G'|_{S_j}$ is finite. Every component S_j is either pseudo-Anosov or finite—see Theorems 7.11 and 7.14 of [30].

As dynamically reducible groups do not have unique closed invariant sets on which they act minimally, we make the following declarations of what are to be considered "**the limit sets**" and "**the domains of discontinuity**" of such groups.

If G is finite, we set $\Lambda_G = Z\Lambda_G = \emptyset$ and declare that $\Delta_G = \mathbb{P}\mathcal{ML}(S)$.

If G is virtually generated by a single pseudo-Anosov mapping class, we define Λ_G, $Z\Lambda_G$, and Δ_G as in the dynamically irreducible case.

If G is infinite and reducible, then we follow McCarthy and Papadopoulos and define the limit set and domain of discontinuity as follows. We let $\mathcal{R}$ denote the canonical reducing system for G and let $S_1, ..., S_n$ be the components of G, which we number so that for some $m \leq n$, $S_1, ..., S_m$ are precisely all the pseudo-Anosov components. for each $j = 1, ..., m$, let $\Lambda_0^j \subset \mathbb{P}\mathcal{ML}(S)$ denote the set of stable laminations of the pseudo-Anosov elements of $G'|_{S_j}$, considered as laminations in S, and Λ^j the closure of this set. The limit set Λ_G is defined to be

$$\Lambda_G = \mathcal{R} \cup \bigcup_{j=1}^{m} \Lambda^j$$

The zero set $Z\Lambda_G$ is defined as before to be

$$Z\Lambda_G = \{[\mu] \in \mathbb{P}\mathcal{ML}(S) \,|\, i(\mu, \lambda) = 0 \text{ for some } [\lambda] \in \Lambda_G\}.$$

McCarthy and Papadopoulos prove that the set $\Delta_G = \mathbb{P}\mathcal{ML}(S) - Z\Lambda_G$ is again a domain of discontinuity for G.

Remark. The limit set Λ_G, its enlargement $Z\Lambda_G$, and the domain Δ_G have their provenance in Masur's work on the mapping class groups of 3–dimensional handlebodies [42].

3.2 Proper discontinuity revisited

The proof that G acts properly discontinuously on Δ_G given in [48] is easily extended to prove

Theorem 3.1. *The action of G on $\mathcal{T}(S) \cup \Delta_G$ is properly discontinuous.*

The proof follows from a series of lemmata mirroring those in Section 6.2 of [48]. We invite the reader to visit that paper for further discussion of these ideas.

Remark. McCarthy and Papadopoulos prove that every orbit in $\mathcal{T}(S) \cup \Delta_G$ is discrete—see Section 8 of [48]. This also follows from Theorem 3.1.

If G is finite, the theorem is trivial and so we assume that G is infinite for the remainder of this section.

Suppose that G contains a mapping class represented by a pseudo-Anosov homeomorphism f and let $|\mathcal{L}| = \{\mathcal{L}_-, \mathcal{L}_+\} \subset \mathcal{ML}(S)$ be unstable and stable measured laminations for f (note that $|\mathcal{L}|$ is a pair of measured lamination, not the support of a lamination). We comment that these are *measured* laminations, not projective classes—this requires an (arbitrary) choice of representative from the projective classes. We refer to $|\mathcal{L}|$ as a **pseudo-Anosov pair** for f. For any h in G, we let $|h\mathcal{L}|$ denote the pair $\{h\mathcal{L}_-, h\mathcal{L}_+\}$, define

$$i(\,\cdot\,, |\mathcal{L}|) = \max\{i(\,\cdot\,, \mathcal{L}_-), i(\,\cdot\,, \mathcal{L}_+)\},$$

and let

$$\Delta_{|\mathcal{L}|} = \{\mu \in \mathcal{ML}(S) \,|\, \forall g \in G,\ \ i(\mu, |\mathcal{L}|) \leq i(\mu, |g\mathcal{L}|)\}.$$

This set is $\mathbb{R}_+$–invariant and so defines a subset of $\mathbb{PML}(S)$ which we also call $\Delta_{|\mathcal{L}|}$.

McCarthy and Papadopoulos show that $\Delta'_{|\mathcal{L}|} = \Delta_{|\mathcal{L}|} \cap \Delta_G$ is a fundamental domain for the action of G on Δ_G.

Remark. If a group Γ acts on a topological space $\mathcal{X}$, we say that a closed subset $\mathcal{D} \subset \mathcal{X}$ is a **fundamental domain** for the action if $\{\gamma\mathcal{D} \,|\, \gamma \in \Gamma\}$ is a locally finite covering of $\mathcal{X}$.

We extend the function $i(\,\cdot\,, |\mathcal{L}|)$ defined on $\mathcal{ML}(S)$ in the obvious way to a function $\imath(\,\cdot\,, |\mathcal{L}|)$ defined on the union $\mathcal{ML}(S) \cup \mathcal{T}(S)$ in $\mathfrak{C}(S)$. We define

$$\widehat{\Delta}_{|\mathcal{L}|} = \{X \in \mathcal{T}(S) \cup \mathcal{ML}(S) \,|\, \forall g \in G,\ \ \imath(X, |\mathcal{L}|) \leq \imath(X, |g\mathcal{L}|)\}.$$

Again, this set is $\mathbb{R}_+$–invariant and so defines a set in $\mathcal{T}(S) \cup \mathbb{PML}(S)$ that we also call $\widehat{\Delta}_{|\mathcal{L}|}$. As with $\Delta_{|\mathcal{L}|}$, one readily checks that $\widehat{\Delta}_{|\mathcal{L}|}$ is closed.

As in [48], we note that for any $g \in G$, we have

$$g(\widehat{\Delta}_{|\mathcal{L}|}) = \widehat{\Delta}_{|g\mathcal{L}|}.$$

As in the proof of proper discontinuity on Δ_G given in [48], special attention must be paid when G is reducible. In this case, we proceed as follows—see Section 3.1 for notation. For each component S_j of G, let $|\mathcal{L}^j| = \{\mathcal{L}^j_-, \mathcal{L}^j_+\}$ be a pseudo-Anosov pair for some pseudo-Anosov automorphism in $G'|_{S_j}$, viewed as laminations in $\mathcal{ML}(S)$. We let υ denote any curve that non-trivially intersects each component of $\mathcal{R}$, and let $|\mathcal{L}|$ denote the union of υ and each $|\mathcal{L}^j|$ for $j = 1, ..., m$. $|\mathcal{L}|$ is called a **complete system for** G. We define

$$\imath(\cdot, |\mathcal{L}|) = \max\{\imath(\cdot, \mathcal{L}^i_-), \imath(\cdot, \mathcal{L}^i_+), \imath(\cdot, \upsilon) \,|\, 1 \leq i \leq n\}$$

and define the sets $\Delta_{|\mathcal{L}|}$, $\Delta'_{|\mathcal{L}|}$, and $\widehat{\Delta}_{|\mathcal{L}|}$ exactly as before.

The first lemma we need is the following (compare Lemma 6.11 [48]). In the following discussion, points of $\mathcal{T}(S) \cup \mathbb{PML}(S)$ will be enclosed in brackets $[X]$, and we will remove the brackets X to denote a representative of $[X]$ in $\mathcal{T}(S) \cup \mathcal{ML}(S)$. This is only relevant when $[X] \in \mathbb{PML}(S)$, in which case X is a representative of the projective class $[X]$. For $[X] \in \mathcal{T}(S)$, we have $X = [X]$.

Lemma 3.2. *Let $[X] \in \mathcal{T}(S) \cup \Delta_G$, and $\{g_n\}$ be an infinite sequence of distinct mapping classes in G. Then the sequence of numbers $\{\imath(X, |g_n\mathcal{L}|)\}$ is unbounded.*

Proof. Suppose that G is irreducible.

Upon passing to a subsequence, the hypothesis implies that one of the sequences $\{g_n\mathcal{L}_-\}$ or $\{g_n\mathcal{L}_+\}$ is unbounded in $\mathcal{ML}(S)$ (see Lemma 2.6 of [48]). That is, there is a curve α so that one of the sequences $\{i(\alpha, g_n\mathcal{L}_+)\}$ or $\{i(\alpha, g_n\mathcal{L}_-)\}$ is tending to infinity. Without loss of generality, we assume $\{g_n\mathcal{L}_+\}$ is unbounded. So there is a sequence of positive real numbers $\{r_n\}$ tending to 0 such that $r_n g_n\mathcal{L}_+ \to \mu$ in $\mathcal{ML}(S)$. Since $[\mathcal{L}_+] \in \Lambda_G$, so is $[\mu]$. We have

$$\imath(X, r_n g_n\mathcal{L}_+) = r_n \imath(X, g_n\mathcal{L}_+) \leq r_n \imath(X, |g_n\mathcal{L}|).$$

If $\imath(X, |g_n\mathcal{L}|)$ were bounded independent of n, the numbers $\imath(X, r_n g_n\mathcal{L}_+)$ would converge to zero, implying that $\imath(X, \mu) = 0$. For $[X] \in \Delta_G = \mathbb{P}\mathcal{ML}(S) - Z\Lambda_G$ this is an obvious contradiction. If $[X] \in \mathcal{T}(S)$, this would mean that $\ell_X(\mu) = 0$, which is also impossible.

If G is reducible, one of $\{g_n\mathcal{L}_-^i\}$, $\{g_n\mathcal{L}_+^i\}$, and $\{g_n\mathcal{V}\}$ is unbounded in $\mathcal{ML}(S)$ (see Lemma 7.6 of [48]), and the proof continues as in the irreducible case. □

The next fact we need is our version of Proposition 6.13 of [48].

Lemma 3.3. *For every $[X] \in \mathcal{T}(S) \cup \Delta_G$, there exists $g \in G$ such that $[X] \in g\widehat{\Delta}_{|\mathcal{L}|}$.*

Proof. By Lemma 3.2 the set $\{\imath(X, |g\mathcal{L}|)\}_{g \in G}$ has no infinite bounded subsets. It follows that there is a minimum $\imath(X, |g\mathcal{L}|)$ for some $g \in G$ and hence $[X] \in \widehat{\Delta}_{|g\mathcal{L}|} = g\widehat{\Delta}_{|\mathcal{L}|}$. □

We now turn to the analog of Proposition 6.14 of [48].

Lemma 3.4. *Let $\mathcal{K}$ be a compact set in $\mathcal{T}(S) \cup \Delta_G$. Then the set of mapping classes $\{g \in G \,|\, \mathcal{K} \cap g\widehat{\Delta}_{|\mathcal{L}|} \neq \emptyset\}$ is finite.*

Proof. Suppose there is an infinite sequence $\{g_n\}$ of distinct elements of G such that $\mathcal{K} \cap g_n\widehat{\Delta}_{|\mathcal{L}|} = \mathcal{K} \cap \widehat{\Delta}_{|g_n\mathcal{L}|} \neq \emptyset$ for every n and let $[X_n] \in \mathcal{K} \cap \widehat{\Delta}_{|g_n\mathcal{L}|}$. In particular, we have $\imath(X_n, |g_n\mathcal{L}|) \leq \imath(X_n, |\mathcal{L}|)$.

Suppose that G is irreducible.

As above (and in [48]), one of the sequences $\{g_n\mathcal{L}_+\}$ or $\{g_n\mathcal{L}_-\}$ is unbounded, and we assume without loss of generality that it is the first. After passing to subsequences, there is a pair of sequences $\{r_n\}$ and $\{t_n\}$ of positive real numbers, the first tending to 0, so that

$$r_n g_n\mathcal{L}_+ \to \mu \in \mathcal{ML}(S) \quad \text{and} \quad t_n X_n \to X \in \mathfrak{C}(S).$$

As in [48], it follows from continuity of ι that

$$\begin{aligned}\iota(X,\mu) &= \lim \iota(t_nX_n, r_ng_n\mathcal{L}_+)\\ &\leq \lim r_n\iota(t_nX_n, |g_n\mathcal{L}|)\\ &\leq \lim r_n\iota(t_nX_n, |\mathcal{L}|)\\ &= (\lim r_n)(\lim \iota(t_nX_n, |\mathcal{L}|))\\ &= (0)(\iota(X, |\mathcal{L}|))\\ &= 0.\end{aligned}$$

As in the proof of Lemma 3.2, this contradicts the fact that $\mu \in \Lambda_G$ and $[X] \in \mathcal{K} \subset \mathcal{T}(S) \cup \Delta_G$.

In the reducible case, one of $\{g_n\mathcal{L}_-^i\}$, $\{g_n\mathcal{L}_+^i\}$, and $\{g_n\upsilon\}$ is unbounded and again the proof is formally identical to the irreducible case. □

We may now prove Theorem 3.1 (compare to the proof of Theorem 6.16 of [48]).

Proof. Let $\mathcal{K} \subset \mathcal{T}(S) \cup \Delta_G$ be compact. We show that the set $\{g \in G \,|\, g\mathcal{K} \cap \mathcal{K} \neq \emptyset\}$ is finite.

By Lemma 3.4 the set

$$\{g \in G \,|\, \mathcal{K} \cap g\widehat{\Delta}_{|\mathcal{L}|} \neq \emptyset\}$$

is finite, and we name its elements $g_1, \ldots, g_N$. With Lemma 3.3, we see that

$$\mathcal{K} \subset \bigcup_{j=1}^{N} g_j\widehat{\Delta}_{|\mathcal{L}|}.$$

Now, if $g\mathcal{K} \cap \mathcal{K} \neq \emptyset$, then $g(g_j\widehat{\Delta}_{|\mathcal{L}|}) \cap \mathcal{K} \neq \emptyset$ for some $j \in \{1, \ldots, N\}$. Since

$$g(g_j\widehat{\Delta}_{|\mathcal{L}|}) = (gg_j)\widehat{\Delta}_{|\mathcal{L}|},$$

it follows that $gg_j = g_i$ for some $i \in \{1, \ldots, N\}$. In particular,

$$\{g \in G \,|\, g\mathcal{K} \cap \mathcal{K} \neq \emptyset\} \subset \{g_ig_j^{-1}\}_{i,j=1}^{N}.$$

Since the set on the right is finite, so is the one on the left. □

We have also established

Proposition 3.5. *Let G be an infinite subgroup of* $\mathrm{Mod}(S)$. *Then*

$$\widehat{\Delta}'_{|\mathcal{L}|} = \widehat{\Delta}_{|\mathcal{L}|} \cap (\mathcal{T}(S) \cup \Delta_G)$$

is a fundamental domain for the action of G on $\mathcal{T}(S) \cup \Delta_G$. □

3.3 Conical limit points and compact fundamental domains

A point $[\lambda]$ in Λ_G is a **conical limit point** if for any Teichmüller geodesic ray τ with direction $[\lambda]$, there is a number $R > 0$ such that some G–orbit intersects the R–neighborhood of τ in an infinite set. Note that a conical limit point is uniquely ergodic by Masur's criterion—as the projection of the geodesic τ to the moduli space $\mathcal{M}(S)$ must return to a bounded neighborhood of a point infinitely often and so cannot leave every compact set. In particular, any geodesic ray τ whose direction is a conical limit point in fact terminates at that point in Thurston's compactification of $\mathcal{T}(S)$, by Masur's Two Boundaries Theorem.

If τ is a Teichmüller geodesic ray emanating from a point X with direction $[\lambda]$ whose R–neighborhood contains infinitely many points of a G–orbit, then any geodesic ray σ terminating at $[\lambda]$ has an R'–neighborhood containing infinitely many points from that orbit. To see this, note that σ and τ are asymptotic by Masur's Asymptotic Rays Theorem. In particular, σ and τ are at a finite Hausdorff distance A from each other and it suffices to take $R' = R + A$. So, to verify that a limit point is conical, we need only consider a single ray.

Theorem 3.6. *If G is a convex cocompact subgroup of* $\mathrm{Mod}(S)$*, then every one of its limit points is conical.*

Proof. Let $[\lambda] \in \Lambda_G$ and let τ be a geodesic ray in $\mathfrak{H}_G$ with direction $[\lambda]$. Since G acts cocompactly on $\mathfrak{H}_G$, there is a positive number R such that every point of the image of τ in $\mathfrak{H}_G/G$ is a distance at most R from any fixed point X_0 in $\mathfrak{H}_G/G$. So, if X is a point in the preimage of X_0, τ stays within R of GX. □

For the remainder of this subsection, it is convenient to switch points of view and work primarily with measured foliations instead of measured laminations. As such, we let $|\mathcal{F}| = \{\mathcal{F}_-, \mathcal{F}_+\}$ denote the foliations associated to a pseudo-Anosov pair $\{\mathcal{L}_-, \mathcal{L}_+\}$. If $\lambda \in \mathcal{ML}(S)$ is associated to $F \in \mathcal{MF}(S)$, then we have $i(F, |\mathcal{F}|) = i(\lambda, |\mathcal{L}|)$. Also, if $g \in G$, then we write $|g\mathcal{F}| = \{g\mathcal{F}_-, g\mathcal{F}_+\}$.

The pair $|\mathcal{F}|$ determines a unique point $(X, q) = (X_{|\mathcal{F}|}, q_{|\mathcal{F}|}) \in \mathcal{Q}^*(S)$ with the property that $\mathcal{F}_- = c\mathcal{F}_-(q_{|\mathcal{F}|})$ and $\mathcal{F}_+ = c\mathcal{F}_+(q_{|\mathcal{F}|})$ for some $c > 0$. Scaling $|\mathcal{F}|$, we may assume that $(X, q) \in \mathcal{Q}^1(S)$ (and so $c = 1$).

As in Section 3 of [51], we see that for any $Y \in \mathcal{T}(S)$ and any $\alpha, \beta \in \mathcal{MF}(S)$

$$\mathrm{ext}_Y(\alpha)\,\mathrm{ext}_Y(\beta) \geq i(\alpha, \beta)^2.$$

Since $\mathcal{F}_-$ and $\mathcal{F}_+$ are the horizontal and vertical foliation of q, and since $\|q\| = 1$, we see that

$$\mathrm{ext}_X(\mathcal{F}_-) = \mathrm{ext}_X(\mathcal{F}_+) = 1.$$

Thus, for any $F \in \mathcal{MF}(S)$ with $\lambda \in \mathcal{ML}(S)$ the associated lamination, we have

$$\mathrm{ext}_X(F) = \mathrm{ext}_X(F)\,\mathrm{ext}_X(\mathcal{F}_\pm) \geq i(F, |\mathcal{F}|)^2 = i(\lambda, |\mathcal{L}|)^2. \tag{1}$$

Theorem 3.7. *Let $[\lambda]$ be a conical limit point of G. Then $[\lambda] \notin \widehat{\Delta}_{|\mathcal{L}|}$.*

Proof. Let $|\mathcal{F}|$ and $(X,q) = (X_{|\mathcal{F}|}, q_{|\mathcal{F}|}) \in \mathcal{Q}^1(S)$ be as above. Let $[\lambda]$ be a conical limit point of G and F the measured foliation associated to λ (an arbitrary choice from the projective class of $[\lambda]$). Let τ be the Teichmüller geodesic emanating from X and terminating at $[F]$ defined by a unit norm quadratic differential ω at X. Thus, choosing the representative of the projective class appropriately, we may assume that F is the vertical foliation of ω. Note also that for g in G, $g(X_{|\mathcal{F}|}) = X_{|g\mathcal{F}|}$.

Now, by the conical hypothesis there is a positive number K and an infinite set $\{g_n\}_{n=0}^{\infty} \subset G$ such that $g_n(X)$ is a distance at most $\frac{1}{2}\log K$ from a point Y_n on τ. Since F is the vertical foliation of ω, we have

$$\lim_{n\to\infty} \mathrm{ext}_{Y_n}(F) = 0.$$

Since $\mathrm{d}_{\mathcal{T}}(Y_n, g_n(X))$ is no more than $\frac{1}{2}\log K$,

$$\mathrm{ext}_{g_n(X)}(\alpha) \leq K\,\mathrm{ext}_{Y_n}(\alpha)$$

for all α in $\mathcal{MF}(S)$ by Theorem 2.1 and so

$$\lim_{n\to\infty} \mathrm{ext}_{g_n(X)}(F) = 0.$$

We conclude that

$$\lim_{n\to\infty} i(\lambda, |g_n\mathcal{L}|) = \lim_{n\to\infty} i(F, |g_n\mathcal{F}|) = 0$$

by (1) and so $[\lambda] \notin \widehat{\Delta}_{|\mathcal{L}|}$. □

Theorem 3.8. *Let G be a subgroup of* $\mathrm{Mod}(S)$ *such that every point in Λ_G is conical. Then there is a compact fundamental domain for the action of G on $\mathcal{T}(S) \cup \Delta_G$. In particular, convex cocompact groups act cocompactly on $\mathcal{T}(S) \cup \Delta_G$.*

Proof. By Proposition 3.5, the set $\widehat{\Delta}'_{|\mathcal{L}|}$ is a fundamental domain for the action of G on $\mathcal{T}(S) \cup \Delta_G$. Since conical limit points are uniquely ergodic, every lamination in Λ_G is uniquely ergodic and so $Z\Lambda_G = \Lambda_G$. In particular, $\Delta_G = \mathbb{PML}(S) - \Lambda_G$. By Theorem 3.7, $\widehat{\Delta}'_{|\mathcal{L}|} = \widehat{\Delta}_{|\mathcal{L}|}$. But the set $\widehat{\Delta}_{|\mathcal{L}|}$ is a closed subset of $\mathcal{T}(S) \cup \mathbb{PML}(S)$, and is thus compact. □

3.4 The weak hull

Let $\mathfrak{A}$ be a closed subset of $\mathbb{PML}(S)$. If $\mathfrak{A}$ has the property that for every $[\lambda_-] \in \mathfrak{A}$, there exists a $[\lambda_+] \in \mathfrak{A}$ such that $[\lambda_-]$ and $[\lambda_+]$ bind S, then we define the **weak hull** $\mathfrak{H}_{\mathfrak{A}}$ of $\mathfrak{A}$ to be the union of all geodesics $\tau(\lambda_-, \lambda_+)$ in $\mathcal{T}(S)$ with $[\lambda_-]$ and $[\lambda_+]$ elements of $\mathfrak{A}$ that bind. If $\mathfrak{A}$ does not have this property then we say that the weak hull is not defined. A set $\mathfrak{H}$ is a **weak hull** if it is $\mathfrak{H}_{\mathfrak{A}}$ for some closed $\mathfrak{A} \subset \mathbb{PML}(S)$ with the aforementioned property.

Note that if G is an infinite irreducible subgroup of $\mathrm{Mod}(S)$, then Λ_G possesses a nonempty weak hull $\mathfrak{H}_G = \mathfrak{H}_{\Lambda_G}$: by a theorem of Ivanov ([30], Corollary 7.14), there exists a pseudo-Anosov automorphism in G with stable lamination $[\lambda]$, and any other lamination in Λ_G will bind with $[\lambda]$.

3.5 Compact fundamental domains in Δ_G cobound the hull

Having a compact fundamental domain for the action on the domain of discontinuity is a restrictive condition in itself, and in particular suffices to cobound the hull.

Theorem 3.9. *Let G be a subgroup of* $\mathrm{Mod}(S)$. *If $\Delta_G \neq \emptyset$ and G acts cocompactly on Δ_G, then every lamination in Λ_G is uniquely ergodic, $Z\Lambda_G = \Lambda_G$, and $\mathfrak{H}_G$ is defined and cobounded. Furthermore, G has a finite index subgroup all of whose non-identity elements are pseudo-Anosov.*

The first step to prove this theorem is to prove that every lamination in Λ_G is filling. To this end, we define

$$ZZ\Lambda_G = \{[\beta] \in \mathbb{P}\mathcal{ML}(S) \,|\, i(\beta,\mu) = 0 \text{ for some } [\mu] \in Z\Lambda_G\}.$$

and prove

Theorem 3.10 (Insomnia). *Let G be a subgroup of* $\mathrm{Mod}(S)$. *If $\Delta_G \neq \emptyset$, and G acts cocompactly on Δ_G, then*

$$ZZ\Lambda_G = Z\Lambda_G.$$

Moreover, every lamination in Λ_G is filling.

Proof. If G is finite, both $Z\Lambda_G$ and $ZZ\Lambda_G$ are empty, and so we assume that G is infinite. We begin by proving the last statement, assuming the first.

If there is a non-filling lamination in Λ_G, then there is a (projective class of laminations supported on a) simple closed curve in $Z\Lambda_G$. Let $V(Z\Lambda_G) \subset \mathcal{C}(S)$ denote the set of all simple closed curves in $\mathcal{C}(S)$ that lie in $Z\Lambda_G$. Since Δ_G is an open set, it contains a simple closed curve that is not in $Z\Lambda_G$, and hence $\mathcal{C}(S) \neq V(Z\Lambda_G) \neq \emptyset$. Since $\mathcal{C}(S)$ is connected, there is a simple closed curve $[\alpha]$ at a distance 1 from $V(Z\Lambda_G)$. The curve $[\alpha]$ is thus disjoint from some element of $Z\Lambda_G$, and is not in $Z\Lambda_G$ (as it is a positive distance from $V(Z\Lambda_G)$). That is, $[\alpha] \in ZZ\Lambda_G - Z\Lambda_G$, contradicting the first part of the theorem.

We now proceed to the proof of the first statement. If G is irreducible, let $|\mathcal{L}| = \{\mathcal{L}_-, \mathcal{L}_+\}$ be a pseudo-Anosov pair for G. If G is reducible, let $|\mathcal{L}|$ be a complete system for G (see Section 3.2). Let $\mathcal{K}$ be a compact fundamental domain for the action of G on Δ_G. By Propositions 6.14 and 7.10 of [48], the set of mapping classes

$$\{g \in G \,|\, \mathcal{K} \cap \Delta'_{|g\mathcal{L}|} \neq \emptyset\}$$

is finite. Since

$$\begin{aligned}
|\{g \in G \,|\, \mathcal{K} \cap \Delta'_{|g\mathcal{L}|} \neq \emptyset\}| &= |\{g \in G \,|\, g^{-1}\mathcal{K} \cap g^{-1}\Delta'_{|g\mathcal{L}|} \neq \emptyset\}| \\
&= |\{g \in G \,|\, g^{-1}\mathcal{K} \cap \Delta'_{|\mathcal{L}|} \neq \emptyset\}| \\
&= |\{g \in G \,|\, g\mathcal{K} \cap \Delta'_{|\mathcal{L}|} \neq \emptyset\}|
\end{aligned}$$

and $\mathcal{K}$ is a fundamental domain, we conclude that $\Delta'_{|\mathcal{L}|}$ is compact.

Suppose to the contrary that there is a lamination $[\beta]$ in $ZZ\Lambda_G - Z\Lambda_G$. So $i(\beta,\lambda) \neq 0$ for all $[\lambda]$ in Λ_G, and there is a $[\mu]$ in $Z\Lambda_G$ such that $i(\beta,\mu) = 0$. Note that $[\mu]$ cannot be filling, lest $[\beta]$ be an element of $Z\Lambda_G$. So, in fact, there is a simple closed curve $[\alpha]$ in $Z\Lambda_G$ such that $i(\beta,\alpha) = 0$. To see this, take $[\alpha]$ be a component of the boundary of the smallest π_1–injective subsurface containing $|\mu|$. If λ is any lamination with $i(\lambda,\mu) = 0$, then notice that one also has $i(\lambda,\mu) = 0$.

For $t \in [0,1]$, let $\nu_t = (1-t)\alpha + t\beta$. That is, for $t \in (0,1)$, ν_t is supported on $|\alpha| \cup |\beta|$ and assigns to each arc transverse to $|\alpha| \cup |\beta|$ the sum of the transverse measures for α and β weighted by $(1-t)$ and t, respectively. We also have $\nu_0 = \alpha$ and $\nu_1 = \beta$. For any interval $\mathbb{J} \subset [0,1]$, write

$$\nu_{\mathbb{J}} = \{\nu_t \,|\, t \in \mathbb{J}\}$$

and let $[\nu_{\mathbb{J}}]$ denote the image in $\mathbb{P}\mathcal{ML}(S)$. Since $i(\beta,\lambda) \neq 0$ for every $[\lambda] \in \Lambda_G$, the entire interval $[\nu_{(0,1]}]$ is contained in Δ_G.

Now, for any number C, the set

$$\{g \in G \,|\, i(\beta, |g\mathcal{L}|) \leq C\}$$

is finite, by Lemmata 6.11 and 7.7 of [48].

When G is irreducible, the laminations $\mathcal{L}_-$ and $\mathcal{L}_+$ are associated to measured foliations $\mathcal{F}_-$ and $\mathcal{F}_+$. As in Section 3.3, these determine a point $(X,q) = (X_{|\mathcal{F}|}, q_{|\mathcal{F}|}) \in \mathcal{Q}^1(S)$ with $\mathcal{F}_-(q) = \mathcal{F}_-$ and $\mathcal{F}_+(q) = \mathcal{F}_+$.

For any simple closed curve γ, we have

$$\frac{1}{\sqrt{2}}\ell_q(\gamma) \leq i(\gamma, |\mathcal{F}|) = i(\gamma, |\mathcal{L}|) \leq \ell_q(\gamma)$$

where $\ell_q \colon \mathcal{S} \to \mathbb{R}_+$ is the function that assigns a curve its q–length. Now, for any constant C, the set

$$\{g\alpha \,|\, i(g\alpha, |\mathcal{L}|) \leq C\}$$

is finite, since the length spectrum of the q–metric is discrete.

When G is reducible, the set

$$\{g\alpha \,|\, i(g\alpha, |\mathcal{L}|) \leq C\}$$

is again finite. To see this, consider the analogous quadratic differentials q_j on S_j determined by $\mathcal{L}_-^j$ and $\mathcal{L}_+^j$. The same comparison of length and intersection number shows that the set is finite up to Dehn twisting along the components of $\mathcal{R}$. However, an infinite collection of curves that differ only by twists in the components of $\mathcal{R}$ will have unbounded intersection numbers with υ, and therefore the intersection numbers with $|\mathcal{L}|$ will be unbounded.

Note that since $\Delta'_{|\mathcal{L}|}$ is a compact fundamental domain and $[\nu_{(0,1]}]$ is closed in Δ_G and non-compact, the set

$$\{g \in G \,|\, [\nu_{(0,1]}] \cap \Delta'_{|g\mathcal{L}|} \neq \emptyset\}$$

is infinite.

For any $T \in (0,1]$, the set $[\nu_{[T,1]}]$ is compact and contained in Δ_G. By Propositions 6.14 and 7.10 of [48], $[\nu_{[T,1]}]$ only intersects finitely many translates of $\Delta'_{|\mathcal{L}|}$. For any $t \in (0,1]$, $[\nu_t] \in h\Delta'_{|\mathcal{L}|}$ for some h in G, since $\Delta'_{|\mathcal{L}|}$ is a fundamental domain for the action of G on Δ_G. So, we may choose a sequence t_n tending to zero such that $[\nu_{t_n}] \in g_n\Delta'_{|\mathcal{L}|}$ and $\{g_n\}$ is an infinite set—and we do so. By the definition of $\Delta_{|g_n\mathcal{L}|}$, we have

$$\begin{aligned} i(\nu_{t_n}, |g_n\mathcal{L}|) &\leq i(\nu_{t_n}, |\mathcal{L}|) \\ &= (1-t_n)\, i(\alpha, |\mathcal{L}|) + t_n\, i(\beta, |\mathcal{L}|) \end{aligned}$$

and so the $i(\nu_{t_n}, |g_n\mathcal{L}|)$ are uniformly bounded by some number r. In particular,

$$\begin{aligned} i(\nu_{t_n}, |g_n\mathcal{L}|) &= (1-t_n)\, i(\alpha, |g_n\mathcal{L}|) + t_n\, i(\beta, |g_n\mathcal{L}|) \\ &\leq r \end{aligned}$$

and so

$$\begin{aligned} i(g_n^{-1}\alpha, |\mathcal{L}|) &= i(\alpha, |g_n\mathcal{L}|) \\ &\leq r/(1-t_n) \\ &\leq 2r \end{aligned}$$

when n is large. We conclude that the numbers $i(g_n^{-1}\alpha, |\mathcal{L}|)$ are all bounded by some number R. Since the set

$$\{g\alpha \,|\, i(g\alpha, |\mathcal{L}|) \leq R\}$$

is finite, we may pass to a subsequence so that

$$\{g_n^{-1}\alpha\}_{n\in\mathbb{N}} = \{\alpha'\}$$

for some simple closed curve α', and we do so.

Again, by the definition of $\Delta_{|g_n\mathcal{L}|}$, we have

$$i(\nu_{t_n}, |g_n\mathcal{L}|) \leq i(\nu_{t_n}, |g\mathcal{L}|) \tag{2}$$

for all g in G. Since $\{g_n\}$ is infinite and $[\beta] \in \Delta_G$, the sequence of numbers

$$i(g_n^{-1}\beta, |\mathcal{L}|) = i(\beta, |g_n\mathcal{L}|)$$

is unbounded, by Lemmata 6.11 and 7.7 of [48]. So, for some n, we have

$$\begin{aligned} i(\nu_{t_n}, |g_n\mathcal{L}|) &= (1-t_n)\, i(g_n^{-1}\alpha, |\mathcal{L}|) + t_n\, i(g_n^{-1}\beta, |\mathcal{L}|) \\ &= (1-t_n)\, i(\alpha', |\mathcal{L}|) + t_n\, i(g_n^{-1}\beta, |\mathcal{L}|) \\ &> (1-t_n)\, i(\alpha', |\mathcal{L}|) + t_n\, i(g_1^{-1}\beta, |\mathcal{L}|) \\ &= (1-t_n)\, i(g_1^{-1}\alpha, |\mathcal{L}|) + t_n\, i(g_1^{-1}\beta, |\mathcal{L}|) \\ &= i(\nu_{t_n}, |g_1\mathcal{L}|) \end{aligned}$$

which contradicts (2). □

With Theorem 3.10 in hand, we prove Theorem 3.9.

Proof of Theorem 3.9. We begin by proving that there exists a constant $D > 0$ such that for any two distinct points $[\lambda_-], [\lambda_+] \in \Lambda_G$ and any proper domain $Y \subset S$ with $\xi(Y) \neq 3$, we have

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq D.$$

We then appeal to Rafi's Theorem and Masur's criterion for unique ergodicity to see that every lamination in Λ_G is uniquely ergodic, $Z\Lambda_G = \Lambda_G$, and the weak hull $\mathfrak{H}_G$ of Λ_G is defined and cobounded. Our proof is a modification of the proof given in [52] in the case when G is cyclic.

Let $\mathcal{K} \subset \Delta_G$ be a compact set. Fix a hyperbolic metric on S and use this to realize $\mathcal{ML}(S)$.

For $[\lambda]$ in Λ_G and $[\kappa]$ in $\mathcal{K}$, let $L([\lambda],[\kappa])$ denote the supremum of lengths of arcs of $|\lambda| \cap (S - |\kappa|)$. This is finite since λ is filling (by Theorem 3.10). We claim that $L(\cdot,\cdot)$ is bounded on $\Lambda_G \times \mathcal{K}$.

Suppose to the contrary that there are sequences $[\lambda_i]$ in Λ_G and $[\kappa_i]$ in $\mathcal{K}$ such that $L([\lambda_i],[\kappa_i])$ tends to infinity with i. Since Λ_G and $\mathcal{K}$ are compact, we may assume that the $[\lambda_i]$ tend to a lamination $[\lambda]$ in Λ_G, the $[\kappa_i]$ to a lamination $[\kappa]$ in $\mathcal{K}$. Since the $L([\lambda_i],[\kappa_i])$ are tending to infinity, we have a sequence of geodesic arcs α_i in $|\lambda_i| \cap (S - |\kappa_i|)$ whose lengths are tending to infinity. The Hausdorff limit of the α_i is a diagonal extension of $|\lambda|$; we conclude that $|\kappa|$ does not transversely intersect that extension, and hence κ has zero intersection number with λ. But this means that $[\kappa]$ is an element of $Z\Lambda_G$, contradicting the fact that $\mathcal{K} \cap Z\Lambda_G$ is empty.

So, there is a constant $L = L(G,\mathcal{K})$ such that for any $[\lambda]$ in Λ_G and any $[\kappa]$ in $\mathcal{K}$, the length of any arc in $|\lambda| \cap (S - |\kappa|)$ is bounded above by L.

A similar argument shows that there is a constant $\Theta = \Theta(G,\mathcal{K}) > 0$ such that for all $[\lambda]$ in Λ_G and $[\kappa]$ in $\mathcal{K}$, the angle $\theta([\lambda],[\kappa])$ between leaves of $|\lambda|$ and those of $|\kappa|$ is at least Θ.

Note that if Y is a proper domain that is not an annulus and λ is a geodesic lamination, the projection $\pi_Y(\lambda)$ may be obtained by realizing the boundary components of Y as geodesics and intersecting $|\lambda|$ with $\mathrm{int}(Y)$. Therefore, by the Keen–Halpern Collar Lemma [31, 26], there is a constant $M = M(G,\mathcal{K})$ such that for any pair $[\lambda_-]$ and $[\lambda_+]$ in Λ_G and any proper subdomain Y with $\xi(Y) \geq 4$ and ∂Y in $\mathcal{K}$,

$$i(\pi_Y(\lambda_-), \pi_Y(\lambda_+)) \leq M.$$

This implies the existence of a constant $B = B(G,\mathcal{K})$ such that

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq B$$

whenever Y is not an annulus and ∂Y is an element of $\mathcal{K}$.

When Y is an annulus,

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq 2\left\lceil 2\cosh^{-1}(1/\sin(\Theta))/\ell(\partial Y)\right\rceil + 2$$

whenever ∂Y is in $\mathcal{K}$, by Lemma 2.6. The injectivity radius of our chosen hyperbolic metric bounds $\ell(\partial Y)$ from below, and so there is a constant $C = C(G, \mathcal{K})$ such that

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq C$$

whenever Y is an annulus and ∂Y lies in $\mathcal{K}$.

Letting $D = D(G, \mathcal{K}) = \max\{B, C\}$, we conclude that

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq D$$

whenever Y is a proper domain with ∂Y in $\mathcal{K}$.

If there is a compact fundamental domain $\mathcal{K}$ for the action of G on Δ_G, we have the bound

$$\mathrm{d}_Y(\lambda_-, \lambda_+) \leq D$$

for all Y with ∂Y in Δ_G and all pairs $[\lambda_-]$ and $[\lambda_+]$ in Λ_G, since Λ_G is G–invariant and $[\lambda_-]$ and $[\lambda_+]$ were arbitrary.

As the laminations in Λ_G are filling by Theorem 3.10, given a proper domain Y, ∂Y is an element of Δ_G, and we have the desired bound for all proper subdomains.

By Rafi's Theorem, the geodesics joining distinct points in Λ_G are uniformly cobounded.

Let $[\lambda]$ be an element of Λ_G. We may find a $[\lambda']$ in Λ_G such that λ and λ' bind S. To see this, first note that G is either finite, in which case the conclusions of the theorem are trivial, or G contains a pseudo-Anosov mapping class; for if not, then $Z\Lambda_G$ would contain a simple closed curve, which is prohibited by the filling hypothesis. Now, let $[\lambda_-]$ and $[\lambda_+]$ be the unstable and stable laminations of a pseudo-Anosov mapping class g in G. If $[\lambda]$ is an element of $\{[\lambda_-], [\lambda_+]\}$, we are done. If not, $[\lambda]$ and $[\lambda_-]$ bind S. In any case, there is a Teichmüller geodesic with directions $[\lambda]$ and $[\lambda']$. This geodesic is cobounded by the previous paragraph and Masur's criterion tells us that $[\lambda]$ is uniquely ergodic. We conclude that $Z\Lambda_G = \Lambda_G$.

Since every lamination in Λ_G is uniquely ergodic, every pair of points in Λ_G are joined by a Teichmüller geodesic. So the weak hull of Λ_G is defined and it is cobounded by the above.

Since $\mathrm{Mod}(S)$ possesses a torsion free subgroup of finite index [62], so does G, and so, to complete the proof, it suffices to show that G contains no reducible element. But iterating such an element on Λ_G would produce a non-filling lamination in Λ_G, which is excluded by the above. □

4 Hulls

4.1 Minsky's quasi-projections to Teichmüller geodesics

Following Minsky [51], given a closed set C in $\mathcal{T}(S)$, we define a closest–points projection from $\mathcal{T}(S)$ to the set of subsets of C,

$$\pi_{\mathrm{C}} \colon \mathcal{T}(S) \to \mathcal{P}(\mathrm{C}),$$

by demanding that

$$\pi_{\mathrm{C}}(X) = \{Y \in \mathrm{C} \,|\, \mathrm{d}_{\mathcal{T}}(X,Y) = \mathrm{d}_{\mathcal{T}}(X,\mathrm{C})\}$$

where $\mathrm{d}_{\mathcal{T}}(X,\mathrm{C}) = \inf_{Y\in\mathrm{C}} \mathrm{d}_{\mathcal{T}}(X,Y)$. Given a set $\mathcal{X} \subset \mathcal{T}(S)$, we abuse notation and refer to $\cup_{X\in\mathcal{X}} \pi_{\mathrm{C}}(X) \subset \mathrm{C}$ as $\pi_{\mathrm{C}}(\mathcal{X})$.

Minsky has proven that these projections behave in much the same way as such projections in $\mathbb{H}^3$.

Theorem 4.1 (Contraction Theorem [51]). *For every $\varepsilon > 0$ there is a constant b, depending only on ε and the topological type of S such that for any ε–cobounded geodesic τ and X in $\mathcal{T}(S)$,*

$$\mathrm{diam}(\pi_\tau(\mathcal{N}_{\mathrm{d}_{\mathcal{T}}(X,\tau)}(X))) \leq b.$$

Theorem 4.2 (Corollary 4.1 of [51]). *For every $\varepsilon > 0$ there is a constant b, depending only on ε and the topological type of S such that the following holds. Let τ be an ε–cobounded geodesic in $\mathcal{T}(S)$. If $R > 0$ and points $X,Y \in \mathcal{T}(S)$ are connected by a path of length T that remains outside an R–neighborhood of τ, then*

$$\mathrm{diam}(\pi_\tau(X) \cup \pi_\tau(Y)) \leq \frac{b}{R}T + b.$$

Furthermore, for any $X,Y \in \mathcal{T}(S)$,

$$\mathrm{diam}(\pi_\tau(X) \cup \pi_\tau(Y)) \leq \mathrm{d}_{\mathcal{T}}(X,Y) + 4b.$$

Theorem 4.3 (Theorem 4.2 of [51]). *For every $\varepsilon > 0$, $K \geq 1$ and $C \geq 0$, there exists a constant D depending on ε, K, C and the topological type of S such that the following holds. Let ψ be a (K,C)–quasi-geodesic path in $\mathcal{T}(S)$ whose endpoints in $\mathcal{T}(S)$ are connected by an ε–cobounded geodesic τ. Then ψ remains in the D–neighborhood of τ.*

4.2 Thin triangles and the hull

We need the following general fact about cobounded geodesic triangles in Teichmüller space. A different proof has been discovered by M. Duchin, see [19].

Theorem 4.4 (Thick triangles are thin). *For every $\varepsilon > 0$ there is a $\delta > 0$ such that if $\triangle$ is a geodesic triangle with vertices in $\mathcal{T}(S) \cup \mathbb{P}\mathcal{ML}(S)$ whose sides are ε–cobounded, then $\triangle$ is δ–thin.*

Proof. Let X, Y, and Z be points in $\mathcal{T}(S)$. Let P be a point in the geodesic segment $[X,Y]$ that minimizes the distance between Z and that segment. It is shown in the proof of Lemma 7.2 of [44] that the path $[X,P] \cup [P,Z]$ is a $(3,0)$–quasi-geodesic—in fact, this is true in any geodesic metric space. We include the proof here for the reader's convenience.

If P and X coincide, there is nothing to do. So suppose that $P \neq X$. Let U be a point in $[X,P]$, V a point in $[P,Z]$. By the triangle inequality and choice of P, P also minimizes the distance between V and $[X,Y]$. So $\mathrm{d}_{\mathcal{T}}(U,V) \geq \mathrm{d}_{\mathcal{T}}(P,V)$.

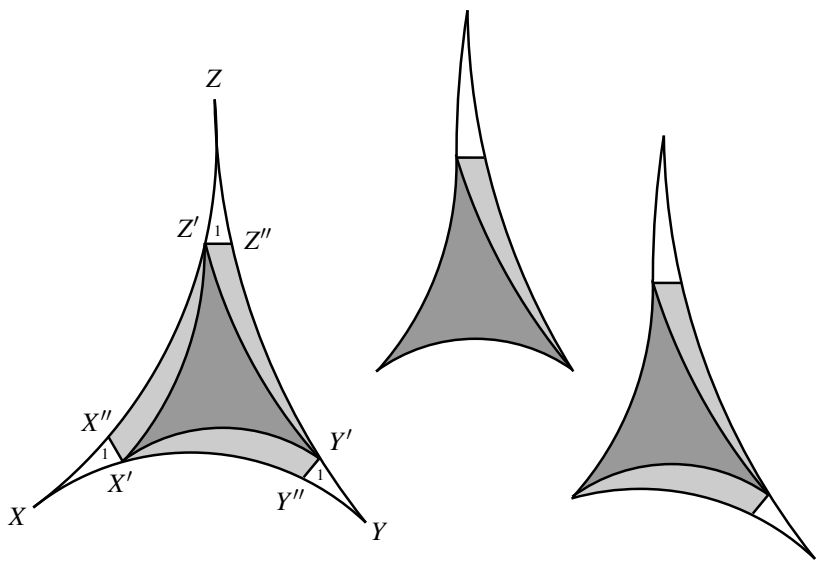

Figure 3: Three potential ideal triangles, their polygons, and inscribed triangles.

By the triangle inequality, $\mathrm{d}_{\mathcal{T}}(U,V) \geq \mathrm{d}_{\mathcal{T}}(U,P) - \mathrm{d}_{\mathcal{T}}(P,V)$. Together with the previous inequality, we have $3\mathrm{d}_{\mathcal{T}}(U,V) \geq \mathrm{d}_{\mathcal{T}}(U,P) + \mathrm{d}_{\mathcal{T}}(P,V)$. So, $[X,P] \cup [P,Z]$ is indeed a $(3,0)$–quasi-geodesic.

Now suppose that the geodesic segments $[X,Y]$, $[X,Z]$, and $[Y,Z]$ are ε–cobounded.

By Theorem 4.3, the path $[X,P] \cup [P,Z]$ remains in a $D(3,0,\varepsilon)$–neighborhood of $[X,Z]$. By symmetry, the path $[Y,P] \cup [P,Z]$ lies in the $D(3,0,\varepsilon)$–neighborhood of $[Y,Z]$. In particular, the segment $[X,Y]$ lies in the $D(3,0,\varepsilon)$–neighborhood of the union $[X,Z] \cup [Y,Z]$. Symmetry guarantees that the triangle $\triangle[X,Y,Z]$ is $D(3,0,\varepsilon)$–thin.

We continue to assume that the sides of the triangle are cobounded and turn to the case where at least one of X, Y and Z lie in $\mathbb{PML}(S)$. Suppose that $W \in \{X,Y,Z\}$ is such that $W \in \mathbb{PML}(S)$ and let W' and W'' be points lying in the interiors of the sides incident to W as pictured in Figure 3. Since the ray $[W',W)$ is cobounded, W is uniquely ergodic, and by Masur's Asymptotic Rays Theorem, the rays $[W',W)$ and $[W'',W)$ are asymptotic. We re-choose W' and W'' so that $\mathrm{d}_{\mathcal{T}}(W',W'') \leq 1$ and the rays $[W',W)$ and $[W'',W)$ are contained in the 1–neighborhoods of each other.

We truncate the triangle $\triangle[X,Y,Z]$ at any such W to obtain a geodesic polygon inscribed with the triangle $\triangle[X',Y',Z']$, where $W' = W$ if $W \in \mathcal{T}(S)$—the possibilities are depicted in Figure 3. The result is composed of the triangle $\triangle[X',Y',Z']$ and at most three geodesic triangles of a special type: each has an ε–cobounded side, a side of length at most 1, and a side of $\triangle[X',Y',Z']$. The union of the latter two sides is a $(1,1)$–quasi-geodesic—as is the union of the former two—and since the remaining side is ε–cobounded, this union is contained in the $D(1,1,\varepsilon)$–neighborhood of that side, by Theorem 4.3. Moreover, this implies that the sides of $\triangle[X',Y',Z']$ are ε'–cobounded, for some ε' depending only on ε and the topological type of S. By

our previous argument, the triangle $\triangle[X',Y',Z']$ is $D(3,0,\varepsilon')$–thin. Yet another application of Theorem 4.3 shows that the special triangles above are M–thin, where $M = \max\{D(1,1,\varepsilon'), D(1,1,\varepsilon)\}$.

In any case, it is easily seen that our triangle $\triangle[X,Y,Z]$ is δ–thin for any δ larger than $2M + D(3,0,\varepsilon') + 1$. □

Theorem 4.5. *For every $\varepsilon > 0$, there is an $A \geq 0$ such that if a weak hull $\mathfrak{H} = \mathfrak{H}_{\mathfrak{A}}$ is ε–cobounded, then $\mathfrak{H} \cup \mathfrak{A}$ is A–quasi-convex. Moreover, any two points in $\mathfrak{H}$ are within 2δ of a bi-infinite geodesic in $\mathfrak{H}$ (where δ is given by Theorem 4.4).*

Proof. Note that since $\mathfrak{H}$ is cobounded, every lamination in $\mathfrak{A}$ is uniquely ergodic by Masur's criterion and so the end of any geodesic ray in $\mathfrak{H}$ converges in $\mathcal{T}(S) \cup \mathbb{PML}(S)$ to its direction, by the Two Boundaries Theorem. In particular, every pair of distinct points in $\mathfrak{A}$ binds S.

Let δ be the constant given by Theorem 4.4 and let X and Y be two points in $\mathfrak{H}$.

We begin by finding a bi-infinite geodesic γ in $\mathfrak{H}$ such that X and Y are both within 2δ of γ. If X and Y lie in a geodesic contained in $\mathfrak{H}$, we are done. Otherwise, there are two bi-infinite geodesics σ and τ in $\mathfrak{H}$ containing X and Y, respectively.

There are two cases to consider, when σ and τ have an endpoint in common, and when they do not.

In the first case, σ and τ are two sides of a geodesic triangle contained in $\mathfrak{H}$. By Theorem 4.4, this triangle is δ–thin, and the desired geodesic is easily found.

In the second case, σ and τ determine four points in $\mathbb{PML}(S)$. Join every pair of these points by a Teichmüller geodesic. The resulting union of geodesics in $\mathcal{T}(S) \cup \mathbb{PML}(S)$ is the 1–skeleton of a 3–simplex, and we refer to the four geodesic triangles in this configuration as the **faces of the simplex**, the geodesics themselves as **edges**. All of the edges are contained in $\mathfrak{H}$ and so all of the faces of the simplex are δ–thin.

Since σ and τ are each incident to two faces of the simplex, for $W \in \{X,Y\}$ there are at least three edges at a distance at most 2δ from W. If for some W there are four edges at such a distance, we know that X and Y are 2δ away from a common edge, by the pigeon–hole principle. Now, it is easy to see that if for each $W \in \{X,Y\}$ there are exactly three edges a distance at most 2δ from W, then these three edges share a vertex. But two such triples of edges in a 3-simplex must share a common edge.

In any case, the common edge is the desired geodesic γ, and the second half of the theorem follows.

Joining the geodesic segment $[X,Y]$ to γ by geodesic segments yields a $(1,4\delta)$–quasi-geodesic, which must lie in the $D(1,4\delta,\varepsilon)$–neighborhood of γ, where $D(1,4\delta,\varepsilon)$ is the constant given by Theorem 4.3. In particular, the segment $[X,Y]$ lies in the $D(1,4\delta,\varepsilon)$–neighborhood of $\mathfrak{H}$.

If both X and Y lie in $\mathbb{PML}(S)$, they are the negative and positive directions of a geodesic contained in $\mathfrak{H}$.

If Y, say, lies in $\mathbb{PML}(S)$ and X lies in $\mathcal{T}(S)$, let σ be a bi-infinite geodesic in $\mathfrak{H}$ containing X. Joining the endpoints of σ to Y by geodesics in $\mathfrak{H}$ yields a triangle that is δ–thin. So, one of the geodesics containing Y is within δ of X. Call this geodesic γ and let $Z \in \gamma$ be within δ of X. By the Asymptotic Rays Theorem, $[X,Y)$ and $[Z,Y)$ are

asymptotic and so we may choose points X' and Z' on these rays, respectively, so that $\mathrm{d}_{\mathcal{T}}(X',Z') \leq 1$ and the rays $[X',Y)$ and $[Z',Y)$ are contained in the 1–neighborhoods of each other. The path $[Z,X] \cup [X,X'] \cup [X',Z']$ is a $(1,\delta+1)$–quasi-geodesic and so, by Theorem 4.3, it lies in the $D(1,\delta+1,\varepsilon)$–neighborhood of γ. We conclude that the ray $[X,Y)$ lies in the $(D(1,\delta+1,\varepsilon)+1)$–neighborhood of γ.

Setting $A = D(1,4\delta,\varepsilon) + D(1,\delta+1,\varepsilon) + 1$ completes the proof. □

Corollary 4.6. *If a subgroup G of* $\mathrm{Mod}(S)$ *has a limit set whose weak hull $\mathfrak{H}_G$ is defined and G acts cocompactly on $\mathfrak{H}_G$, then G is convex cocompact.*

Proof. If G acts cocompactly on $\mathfrak{H}_G$, then a G–orbit in $\mathfrak{H}_G$ is B–dense for some positive number B. Since $\mathfrak{H}_G$ is A–quasi-convex for some A, we conclude that the G–orbit is $(A+B)$–quasi-convex. □

We also have

Corollary 4.7. *Let $\mathfrak{H}$ be an ε–cobounded weak hull and let A be the constant given by Theorem 4.5. Let $\mathfrak{H}^A = \mathcal{N}_A(\mathfrak{H})$ equipped with the induced path metric. Then $\mathfrak{H}^A$ is a proper δ–hyperbolic metric space for some δ.*

Proof. By the choice of A, the restriction of the metric on $\mathfrak{H}^A$ to $\mathfrak{H}$ agrees with the restriction of the Teichmüller metric and every geodesic triangle with vertices in $\mathfrak{H}$ has Teichmüller geodesic edges. Let $\triangle$ be such a triangle. Since $\mathfrak{H}$ is ε–cobounded and $\triangle$ is contained in its A–neighborhood, $\triangle$ is ε'–cobounded for some ε'. By Theorem 4.4, there is a δ' depending only on ε' and S such that $\triangle$ is δ'–thin. As $\mathfrak{H}$ is A–dense in $\mathfrak{H}^A$, we conclude that $\mathfrak{H}^A$ is δ–hyperbolic for some δ. □

5 Kleinian manifolds

The following is part of Proposition 5.1 of [33].

Proposition 5.1 (Klarreich). *Let X_n and Y_n be sequences in $\mathcal{T}(S)$ that converge in Teichmüller's compactification to filling laminations $[\mu]$ and $[\lambda]$. Let τ_n be the sequence of Teichmüller geodesic segments joining X_n to Y_n. Then for every accumulation point $[\nu]$ of $\{\tau_n\}$ in Teichmüller's boundary one has $i(\nu,\mu) = 0$ or $i(\nu,\lambda) = 0$.* □

We need

Proposition 5.2. *Let $\mathfrak{A} \subset \mathbb{PML}(S)$ be a closed set consisting entirely of uniquely ergodic laminations. Let $[\mu_n]$ and $[\lambda_n]$ be sequences in $\mathfrak{A}$ converging to $[\mu]$ and $[\lambda]$, and, for each n, let τ_n be the Teichmüller geodesic with negative and positive directions $[\mu_n]$ and $[\lambda_n]$. Then the set of accumulation points of $\{\tau_n\}$ in Thurston's boundary is contained in $\{[\mu],[\lambda]\}$.*

Proof. By Masur's Two Boundaries Theorem, the ends of the geodesics τ_n converge to their directions $[\mu_n]$ and $[\lambda_n]$. The proposition now follows from Proposition 5.1 by a diagonal argument and the unique ergodicity of $[\mu]$ and $[\lambda]$. □

Lemma 5.3. *Let G be a subgroup of* $\mathrm{Mod}(S)$ *such that every lamination in Λ_G is uniquely ergodic. Then the weak hull $\mathfrak{H}_G$ of Λ_G is closed in $\mathcal{T}(S)\cup\Delta_G$. In particular, $\mathfrak{H}_G\cup\Lambda_G$ is closed in $\mathcal{T}(S)\cup\mathbb{PML}(S)$.*

Proof. Let X_n be a sequence in $\mathfrak{H}_G$ and let τ_n be a sequence of bi-infinite geodesics in $\mathfrak{H}_G$ containing the X_n.

Suppose that the X_n accumulate at a point X in $\mathcal{T}(S)$. We may pass to a subsequence so that the X_n converge to X. The Arzelà–Ascoli Theorem allows us to pass to a further subsequence so that the τ_n converge to a geodesic through X. Since Λ_G is closed, the limiting geodesic lies in the weak hull $\mathfrak{H}_G$.

If the X_n accumulate at a point $[\nu]$ in $\mathbb{PML}(S)$, pass to a subsequence so that the X_n converge to $[\nu]$ and so that the ends of the geodesics τ_n converge to projective measured laminations $[\lambda]$ and $[\mu]$. By Proposition 5.2, $[\nu]\in\{[\mu],[\lambda]\}\subset\Lambda_G$. □

Theorem 5.4. *If G is a subgroup of* $\mathrm{Mod}(S)$ *that acts cocompactly on $\mathcal{T}(S)\cup\Delta_G$, then G is convex cocompact.*

Proof. Suppose that $\dot{M}_G=(\mathcal{T}(S)\cup\Delta_G)/G$ is compact. Then Δ_G/G is compact. By Theorem 3.9, every lamination in Λ_G is uniquely ergodic, $Z\Lambda_G=\Lambda_G$, and the weak hull $\mathfrak{H}_G$ of the limit set Λ_G is defined. By Lemma 5.3, $\mathfrak{H}_G$ is closed in $\mathcal{T}(S)\cup\Delta_G$.

As G acts cocompactly and properly discontinuously on $\mathcal{T}(S)\cup\Delta_G$ and $\mathfrak{H}_G$ is closed therein, G acts cocompactly on $\mathfrak{H}_G$. The theorem now follows from Corollary 4.6. □

6 Hulls revisited: quasi-projections

With suitably adjusted constants, Minsky's quasi-projection theorems hold for cobounded weak hulls.

Theorem 6.1 (Hull contraction). *Given $\varepsilon>0$ there is a constant c depending only on ε and the topological type of S such that for any ε–cobounded weak hull $\mathfrak{H}$ and any point X in $\mathcal{T}(S)$,*

$$\mathrm{diam}(\pi_{\mathfrak{H}}(\mathcal{N}_L(X)))\le c,$$

where $L=\mathrm{d}_{\mathcal{T}}(X,\mathfrak{H})$.

Proof. Fix an ε–cobounded weak hull $\mathfrak{H}$ and let X be a point in $\mathcal{T}(S)$. By Theorem 4.5, $\mathfrak{H}$ is A–quasi-convex for some A.

If $L=\mathrm{d}_{\mathcal{T}}(X,\mathfrak{H})<A$, $\mathrm{diam}(\pi_{\mathfrak{H}}(\mathcal{N}_L(X)))$ is at most $2A$, and so we suppose that $L\ge A$.

We begin by bounding the distance between two points in $\pi_{\mathfrak{H}}(X)$. Strictly speaking, this follows from the proof of the theorem given below. As it is a basic ingredient in the proof, we include the argument in the interest of clarity. We refer the reader to Figure 4 for a diagram of the following.

Let Y and Z be two points in $\pi_{\mathfrak{H}}(X)$. The geodesic segment $[Y,Z]$ is contained in the A–neighborhood of $\mathfrak{H}$. Let P be a point on $[Y,Z]$ minimizing the distance between X and that segment, and let $s=\mathrm{d}_{\mathcal{T}}(X,P)$. Note that $s\ge L-A$, for if not, $\mathrm{d}_{\mathcal{T}}(X,\mathfrak{H})<L$ as $[Y,Z]\subset\mathcal{N}_A(\mathfrak{H})$ and this contradicts the fact that $\mathrm{d}_{\mathcal{T}}(X,\mathfrak{H})=L$.

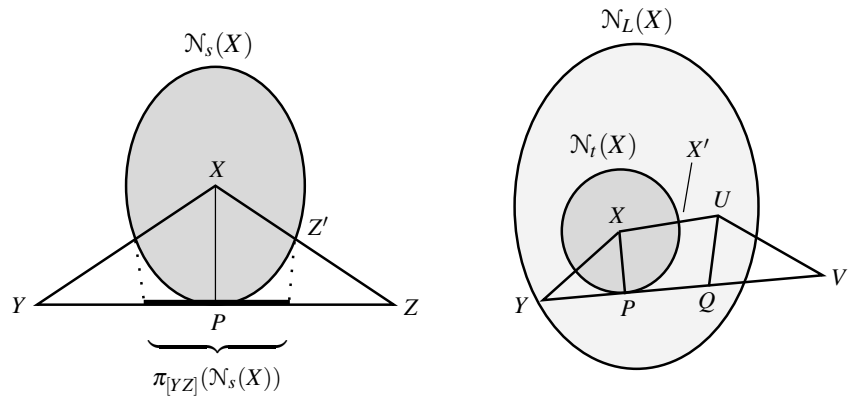

Figure 4: Bounding the diameter of the projection.

By Minsky's Contraction Theorem, there is a constant b depending only on ε and the topological type of S such that

$$\operatorname{diam}(\pi_{[Y,Z]}(\mathcal{N}_s(X))) \leq b.$$

Consider the geodesic segment $[X,Z]$. This is composed of a segment $[X,Z']$ of length s and a segment $[Z',Z]$ of length at most A, as the whole segment has length L. By the above, the diameter of $\pi_{[Y,Z]}([X,Z'])$ is at most b. By Theorem 4.2, the diameter of $\pi_{[Y,Z]}([Z',Z])$ is at most $A+4b$. So the diameter of $\pi_{[Y,Z]}([X,Z])$ is at most $A+5b$.

Since $P \in \pi_{[Y,Z]}(X)$ and $\pi_{[Y,Z]}(Z) = \{Z\}$, we conclude that $\mathrm{d}_{\mathcal{T}}(P,Z) \leq A+5b$. Symmetry yields the same bound for the distance between P and Y, and we conclude that $\mathrm{d}_{\mathcal{T}}(Y,Z) \leq 2A+10b$.

We now turn to the proof of the theorem, continuing to let Y denote a point in $\pi_{\mathfrak{H}}(X)$, letting U be a point different from X in $\mathcal{N}_L(X)$, and letting V be an element of $\pi_{\mathfrak{H}}(U)$.

Let Q be a point of $[Y,V]$ minimizing the distance between U and that segment. An argument similar to the above shows that $\mathrm{d}_{\mathcal{T}}(Q,V) \leq A+5b$. Let P be a point on $[Y,Q]$ minimizing the distance between X and $[Y,Q]$. As $[Y,V] \subset \mathcal{N}_A(\mathfrak{H})$, so is $[Y,Q]$, and we have $t = \mathrm{d}_{\mathcal{T}}(X,P) \geq L-A$. Again, we have that $\mathrm{d}_{\mathcal{T}}(Y,P) \leq A+5b$.

Now, since $U \in \mathcal{N}_L(X)$, $\mathrm{d}_{\mathcal{T}}(X,U) \leq L$ and so $[X,U]$ is composed of two segments $[X,X']$ and $[X',U]$, where $[X,X']$ lies in the $(L-A)$–neighborhood of X and $[X',U]$ has length at most A. We conclude that

$$\operatorname{diam}(\pi_{[P,Q]}([X,X'])) \leq b,$$

by Minsky's Contraction Theorem, and that

$$\operatorname{diam}(\pi_{[P,Q]}([X',U])) \leq A+4b,$$

by Theorem 4.2. So, the distance between Y and V is at most $3A+15b$, and we conclude that

$$\operatorname{diam}(\pi_{\mathfrak{H}}(\mathcal{N}_L(X))) \leq 6A+30b.$$

Letting $c = 6A + 30b$ completes the proof. $\square$

As in [51], this easily yields analogs of Theorems 4.2 and 4.3 for weak hulls, which we state for completeness.

Theorem 6.2. *For every $\varepsilon > 0$ there is a constant c, depending only on ε and the topological type of S such that the following holds. Let $\mathfrak{H}$ be an ε–cobounded weak hull in $\mathcal{T}(S)$. If $R > 0$ and points $X, Y \in \mathcal{T}(S)$ are connected by a path of length T that remains outside an R–neighborhood of $\mathfrak{H}$, then*

$$\operatorname{diam}(\pi_{\mathfrak{H}}(X) \cup \pi_{\mathfrak{H}}(Y)) \leq \frac{c}{R}T + c.$$

Furthermore, for any $X, Y \in \mathcal{T}(S)$,

$$\operatorname{diam}(\pi_{\mathfrak{H}}(X) \cup \pi_{\mathfrak{H}}(Y)) \leq \mathrm{d}_{\mathcal{T}}(X, Y) + 4c.$$

Theorem 6.3. *Let ψ be a (K, C)–quasi-geodesic path in $\mathcal{T}(S)$ whose endpoints in $\mathcal{T}(S)$ are contained in an ε–cobounded weak hull $\mathfrak{H}$. Then ψ remains in a $E(K, C, \varepsilon)$ neighborhood of $\mathfrak{H}$.*

Remark. Theorem 6.3 also follows directly from Theorem 4.3 and Theorem 4.5.

7 Quasi-isometric embedding in $\mathcal{C}(S)$

Let G be a subgroup of $\mathrm{Mod}(S)$ with finite generating set $\mathcal{U}$ and word metric $\mathrm{d}_{\mathcal{U}}$. For any $v \in \mathcal{C}(S)$, the G-orbit Gv of v defines a map $\Phi_v \colon G \to \mathcal{C}(S)$. We have the following

Theorem 7.1. *For any $v \in \mathcal{C}(S)$, if Φ_v is a quasi-isometric embedding, then G is convex cocompact.*

Remark. See also U. Hamenstädt [27].

If Φ_v is a (K, C)–quasi-isometric embedding, then for any u in $\mathcal{C}(S)$, the map Φ_u is a (K, C')–quasi-isometric embedding, where $C' = C + 2\mathrm{d}_{\mathcal{C}}(u, v)$—in particular, we may assume that v is any point of $\mathcal{C}(S)$, when a choice of v is convenient.

Given a point X in $\mathcal{T}(S)$, the G–orbit GX of G defines a map $\Psi_X \colon G \to \mathcal{T}(S)$. It so happens that Φ_v being a quasi-isometric embedding implies that Ψ_X is as well. We record this in the following

Lemma 7.2. *If Φ_v is a quasi-isometric embedding for some $v \in \mathcal{C}(S)$, then for any point X in $\mathcal{T}(S)$, the map $\Psi_X \colon G \to \mathcal{T}(S)$ is a quasi-isometric embedding.*

Proof. Since Ψ_X is defined by taking an element h to hX, the desired upper bound is an immediate consequence of the finite generation of G.

The Teichmüller space sits naturally in the electric Teichmüller space $\mathcal{T}_{\mathrm{el}}(S)$, see the proof of Theorem 7.6 for a definition. By Lemma 3.1 of [44], and its proof, the electric space and $\mathcal{C}(S)$ are $\mathrm{Mod}(S)$–equivariantly quasi-isometric. We may assume that $v \in \mathcal{C}(S)$ is the image of X under such a quasi-isometry. Since Φ_v is a quasi-isometric embedding and the inclusion $\mathcal{T}(S) \to \mathcal{T}_{\mathrm{el}}(S)$ is Lipschitz, we obtain the desired lower bound. $\square$

7.1 The boundary and ending laminations

By Theorem 1.1 of [44], see also [11], $\mathcal{C}(S)$ is δ'–hyperbolic for some δ'. If $\Phi_v\colon G \to \mathcal{C}(S)$ is a quasi-isometric embedding then G is δ–hyperbolic for some δ, the map Φ_v has a continuous extension

$$\Phi_v\colon\ G \cup \partial G \to \mathcal{C}(S) \cup \partial\mathcal{C}(S),$$

and the restriction

$$\partial\Phi_v\colon\ \partial G \to \partial\mathcal{C}(S)$$

is a topological embedding, see Théorème 2.2 of [18].

By the stability of quasi-geodesics in δ–hyperbolic metric spaces, see Théorème 1.2 of [18], there exists an $A > 0$ such that for any geodesic $\mathcal{G}$ in G, the quasi-geodesic $\Phi_v(\mathcal{G})$ and any geodesic joining its endpoints have Hausdorff distance at most A. In particular, for any distinct pair of points $x, y \in \Phi_v(G \cup \partial G)$, any geodesic between x and y is contained in $\mathcal{N}_A(\Phi_v(G))$—thus $\Phi_v(G)$ is A-quasi-convex.

In the next section, we find estimates required to cobound the weak hull (see Corollary 7.5). To do this, we must recall the geometric description of $\partial\mathcal{C}(S)$.

Let $\mathcal{L}_{\text{fill}}(S)$ be the set of filling laminations in $\mathbb{P}\mathcal{ML}(S)$ and let $\mathrm{F}\colon\ \mathcal{L}_{\text{fill}}(S) \to \mathcal{UML}(S)$ be the map that forgets transverse measures. The image of F is the space of potential ending laminations for hyperbolic 3–manifolds homeomorphic to $S \times \mathbb{R}$ and is denoted here by $\mathcal{EL}(S)$. It is a theorem of E. Klarreich [33] that $\partial\mathcal{C}(S)$ is naturally homeomorphic to $\mathcal{EL}(S)$ so that if a quasi-geodesic limits to $|\mu| \in \partial\mathcal{C}(S)$, then every accumulation point in $\mathbb{P}\mathcal{ML}(S)$ of its vertices—viewed as elements in $\mathbb{P}\mathcal{ML}(S)$—projects to $|\mu|$ under F. In particular, for any $m \in \partial G$, $\Phi_v(m)$ is naturally identified with a lamination in $\mathcal{EL}(S)$.

7.2 Bounding the subsurface projection coefficients

We make repeated use of the following theorem of Masur and Minsky [45].

Theorem 7.3 (Masur–Minsky). *There exists a constant $M = M(S)$ with the following property. Let Y be a proper domain of S with $\xi(Y) \neq 3$ and let γ be a geodesic segment, ray, or bi-infinite line in $\mathcal{C}(S)$, such that $\pi_Y(v) \neq \emptyset$ for every vertex v of γ. Then*

$$\mathrm{diam}_Y(\gamma) \leq M.$$

The main theorem allowing us to cobound the hull is

Theorem 7.4 (Quasi-isometric projection bound). *If Φ_v is a quasi-isometric embedding for some $v \in \mathcal{C}(S)$, then there exists a constant $D > 0$ such that for any two distinct points $m_-, m_+ \in \partial G$ and any proper domain $Y \subset S$ with $\xi(Y) \neq 3$, we have*

$$\mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+)) \leq D.$$

As the proof is technical, we pause to sketch the argument.

Given distinct points m_-, m_+ in ∂G, there is a geodesic in G joining them. This geodesic is carried to a quasi-geodesic in $\mathcal{C}(S)$, which is uniformly close to a geodesic γ joining $\Phi_v(m_-)$ and $\Phi_v(m_+)$. To bound a coefficient $\mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+))$, it suffices to bound $\mathrm{diam}_Y(\gamma)$.

If Y is a proper domain whose boundary is far from γ, Theorem 7.3 provides a bound on $\mathrm{diam}_Y(\gamma)$. If ∂Y is close to γ, it is close to $\Phi_v(G)$. In fact, we may assume that ∂Y is close to $\Phi_v(\mathbf{1})$ by translating. Since the two ends of γ diverge, γ may be decomposed into three parts: a finite segment γ_0 near ∂Y and two infinite rays γ_- and γ_+ far from ∂Y. Theorem 7.3 again bounds $\mathrm{diam}_Y(\gamma_\pm)$. The segment γ_0 fellow travels the image of a geodesic segment in G lying in a fixed neighborhood of $\mathbf{1}$. Finiteness of this neighborhood allows us to bound $\mathrm{diam}_Y(\gamma_0)$. The triangle inequality provides the bound on $\mathrm{diam}_Y(\gamma)$.

Proof. We let Ω denote the set of pairs of distinct points in ∂G:

$$\Omega = \{(m_-, m_+) \,|\, m_-, m_+ \in \partial G \text{ and } m_- \neq m_+\}.$$

We assume Φ_v is a (K,C)–quasi-isometry and as noted above, $\Phi_v(G)$ is A–quasi-convex. It is convenient to assume further that we have chosen A sufficiently large so that for any geodesic $\mathcal{G}$ in $G \cup \partial G$ and any geodesic γ connecting the endpoints of $\Phi_v(\mathcal{G})$, any closest point projection map from $\Phi_v(\mathcal{G})$ to γ is A**–coarsely order preserving**: if $u_0, u_1, u_2 \in \gamma$ are pairwise separated by a distance at least A and $u_0 < u_1 < u_2$, then for every triple $h_0, h_1, h_2 \in \mathcal{G}$ for which $\Phi_v(h_i)$ is a point closest to u_i, $i = 0,1,2$, we have $h_0 < h_1 < h_2$.

We partition the proper domains of S into two classes

$$\mathrm{Dom}(\mathrm{far}) = \{Y \subset S \,|\, \mathrm{d}_{\mathcal{C}}(\partial Y, \Phi_v(G)) \geq A+2\},$$
$$\mathrm{Dom}(\mathrm{near}) = \{Y \subset S \,|\, \mathrm{d}_{\mathcal{C}}(\partial Y, \Phi_v(G)) < A+2\},$$

and define

$$\mathrm{Dom}(0) = \{Y \subset S \,|\, \mathrm{d}_{\mathcal{C}}(\partial Y, v) < A+2\} \subset \mathrm{Dom}(\mathrm{near}).$$

Let Y be an element of $\mathrm{Dom}(\mathrm{far})$. By Lemma 5.14 of [50], for any pair (m_-, m_+) in Ω, there exists a geodesic γ between $\Phi_v(m_-)$ and $\Phi_v(m_+)$. By our choice of A, this lies in $\mathcal{N}_A(\Phi_v(G))$ and so $\mathrm{d}_{\mathcal{C}}(\partial Y, \gamma) \geq 2$. In particular, $\pi_Y(u) \neq \emptyset$ for every $u \in \gamma$, and Theorem 7.3 implies that

$$\mathrm{diam}_Y(\gamma) \leq M.$$

Let $\{v_i\}$ be the vertices of γ. Since $\Phi_v(m_+)$ is a sub-lamination of any accumulation point of $\{v_i\}_{i=0}^{\infty}$ in the Hausdorff topology, it follows that for any subsurface Y, we have

$$\pi_Y(\Phi_v(m_+)) \subset \pi_Y(v_i)$$

for sufficiently large i. Along with a similar argument for $\Phi_v(m_-)$, this yields

$$\mathrm{diam}_Y(\gamma) = \mathrm{diam}_Y(\{v_i\} \cup \{\Phi_v(m_\pm)\}),$$

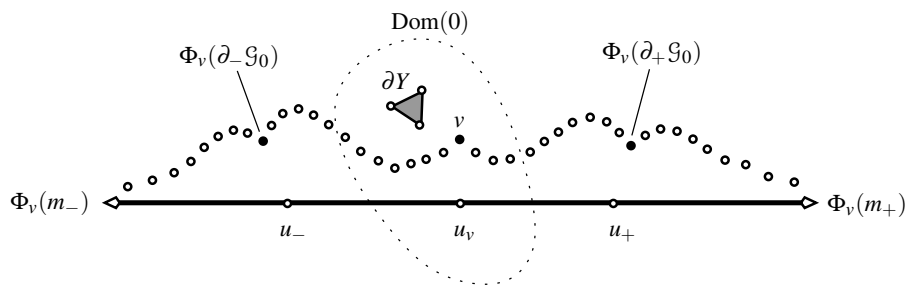

Figure 5: The quasi-geodesic $\Phi_v(\mathcal{G})$, its fellow–traveling geodesic γ, and some points of interest.

and we conclude that

$$\mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+)) \leq M.$$

Now let Y be a domain in Dom(near), (m_-, m_+) a pair in Ω, and $\mathcal{G}$ a geodesic in G joining m_- to m_+. Let $h \in \mathcal{G}$ be a point for which $\mathrm{d}_{\mathcal{C}}(\partial Y, \Phi_v(h)) < A + 2$. Then $h^{-1}(\partial Y) = \partial(h^{-1}(Y))$, $\mathbf{1} \in h^{-1}(\mathcal{G})$, and $h^{-1}(Y) \in \mathrm{Dom}(0)$. Furthermore, $h^{-1}(\mathcal{G})$ has endpoints $h^{-1}(m_-)$ and $h^{-1}(m_+)$ and

$$\mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+)) = \mathrm{d}_{h^{-1}(Y)}(\Phi_v(h^{-1}(m_-)), \Phi_v(h^{-1}(m_+))).$$

So it suffices to find a constant D' such that

$$\mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+)) \leq D' \tag{3}$$

whenever $Y \in \mathrm{Dom}(0)$ and (m_-, m_+) is a pair joined by a geodesic through $\mathbf{1}$. Setting $D = \max\{D', M\}$ will complete the proof.

Finding D'. We fix a constant R satisfying

$$R \geq K(4A + 5 + C) \tag{4}$$

and refer the reader to Figure 5 for a schematic of what follows.

We fix a pair (m_-, m_+) in Ω and a geodesic $\mathcal{G}$ through $\mathbf{1}$ joining m_- and m_+. Let $\mathcal{G}_0 \subset \mathcal{G}$ denote the intersection of $\mathcal{G}$ with the ball of radius R about $\mathbf{1}$.

Next, let $\partial_- \mathcal{G}_0$ and $\partial_+ \mathcal{G}_0$ denote the initial and terminal points of $\mathcal{G}_0$, respectively. Since $\Phi_v(\mathbf{1}) = v$ we have

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_\pm \mathcal{G}_0), v) &\geq \frac{1}{K} \mathrm{d}_U(\partial_\pm \mathcal{G}_0, \mathbf{1}) - C \\ &= \frac{R}{K} - C \\ &\geq 4A + 5. \end{aligned} \tag{5}$$

Similarly, we observe

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_-\mathcal{G}_0), \Phi_v(\partial_+\mathcal{G}_0)) &\geq 8A + 10 + C \\ &\geq 8A + 10. \end{aligned} \tag{6}$$

Again by [50], there is a geodesic γ with endpoints $\Phi_v(m_-)$ and $\Phi_v(m_+)$. This has Hausdorff distance at most A from $\Phi_v(\mathcal{G})$, and we let $u_- = u_-(\gamma)$ and $u_+ = u_+(\gamma)$ denote a pair of closest points on γ to $\Phi_v(\partial_-\mathcal{G}_0)$ and $\Phi_v(\partial_+\mathcal{G}_0)$, respectively. Note that

$$\mathrm{d}_{\mathcal{C}}(u_\pm, \Phi_v(\partial_\pm\mathcal{G}_0)) \leq A. \tag{7}$$

By (5) and (7) (and the triangle inequality) we have

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(u_\pm, v) &\geq \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_\pm\mathcal{G}_0), v) - \mathrm{d}_{\mathcal{C}}(u_\pm, \Phi_v(\partial_\pm\mathcal{G}_0)) \\ &\geq 4A + 5 - A \\ &= 3A + 5. \end{aligned} \tag{8}$$

Then, by (6) and (7),

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(u_-, u_+) &\geq \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_-\mathcal{G}_0), \Phi_v(\partial_+\mathcal{G}_0)) \\ &\quad - \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_-\mathcal{G}_0), u_-) \\ &\quad - \mathrm{d}_{\mathcal{C}}(\Phi_v(\partial_+\mathcal{G}_0), u_+) \\ &\geq 8A + 10 - 2A \\ &= 6A + 10. \end{aligned} \tag{9}$$

In particular, u_-, u_+ decomposes γ into a pair of geodesic rays $\gamma_\pm$ and a geodesic segment γ_0. The endpoints of γ_-, γ_0, and γ_+ are $\{\Phi_v(m_-), u_-\}$, $\{u_-, u_+\}$, and $\{u_+, \Phi_v(m_+)\}$, respectively.

Let $u_v = u_v(\gamma)$ denote a closest point on γ to v, which is a distance at most A from v. By (8) we have

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(u_\pm, u_v) &\geq \mathrm{d}_{\mathcal{C}}(u_\pm, v) - \mathrm{d}_{\mathcal{C}}(u_v, v) \\ &\geq 3A + 5 - A \\ &= 2A + 5. \end{aligned} \tag{10}$$

Thus, because any closest point projection to γ is A–coarsely order preserving, and since $\partial_-\mathcal{G}_0 < \mathbf{1} < \partial_+\mathcal{G}_0$, it must be that $u_- < u_v < u_+$, and $u_v \in \gamma_0$.

Moreover, by (10) and because γ is a geodesic, we have, for every $u \in \gamma_\pm$,

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(u, v) &\geq \mathrm{d}_{\mathcal{C}}(u, u_v) - \mathrm{d}_{\mathcal{C}}(v, u_v) \\ &\geq \mathrm{d}_{\mathcal{C}}(u_\pm, u_v) - \mathrm{d}_{\mathcal{C}}(v, u_v) \\ &\geq 2A + 5 - A \\ &\geq A + 5. \end{aligned} \tag{11}$$

Now suppose that $Y \in \mathrm{Dom}(0)$ and $u \in \gamma_\pm$. By (11) we have

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(\partial Y, u) &\geq \mathrm{d}_{\mathcal{C}}(u, v) - (\mathrm{d}_{\mathcal{C}}(\partial Y, v) + \mathrm{diam}(\partial Y)) \\ &\geq A + 5 - (A + 2 + 1) \\ &= 2 \end{aligned} \tag{12}$$

and therefore $\pi_Y(u) \neq \emptyset$ for every $u \in \gamma_\pm$ and hence

$$\operatorname{diam}_Y(\gamma_\pm) \leq M. \tag{13}$$

As before, we have

$$\mathrm{d}_Y(\Phi_v(m_\pm), u_\pm) \leq \operatorname{diam}_Y(\gamma_\pm).$$

Next, suppose that $\zeta_\pm$ are geodesics connecting $\Phi_v(\partial_\pm \mathcal{G}_0)$ to $u_\pm$. These geodesics have length at most A and therefore for every $u \in \zeta_\pm$, by (5), we have

$$\begin{aligned} \mathrm{d}_{\mathcal{C}}(\partial Y, u) &\geq \mathrm{d}_{\mathcal{C}}(v, \Phi_v(\partial_\pm \mathcal{G}_0)) - \mathrm{d}_{\mathcal{C}}(u, \Phi_v(\partial_\pm \mathcal{G}_0)) \\ &\quad - \mathrm{d}_{\mathcal{C}}(v, \partial Y) - \operatorname{diam}(\partial Y) \\ &\geq 4A + 5 - A - (A+2) - 1 \\ &= 2A + 2. \end{aligned} \tag{14}$$

In particular, we see that $\pi_Y(\zeta_\pm) \neq \emptyset$ and so, by Theorem 7.3,

$$\operatorname{diam}_Y(\zeta_\pm) \leq M. \tag{15}$$

Thus, by (13) and (15) we obtain

$$\begin{aligned} \mathrm{d}_Y(\Phi_v(m_-), \Phi_v(m_+)) &\leq \operatorname{diam}_Y(\gamma_-) + \operatorname{diam}_Y(\zeta_-) \\ &\quad + \mathrm{d}_Y(\Phi_v(\partial_- \mathcal{G}_0), \Phi_v(\partial_+ \mathcal{G}_0)) \\ &\quad + \operatorname{diam}_Y(\zeta_+) + \operatorname{diam}_Y(\gamma_+) \\ &\leq 4M + \mathrm{d}_Y(\Phi_v(\partial_- \mathcal{G}_0), \Phi_v(\partial_+ \mathcal{G}_0)). \end{aligned}$$

Note that this last expression depends only on Y, v, and $\mathcal{G}_0 \subset \mathcal{N}_R(\mathbf{1})$.

Since $\Phi_v(\mathcal{N}_R(\mathbf{1}))$ is finite, there is a constant D'' such that for each pair u and w in $\Phi_v(\mathcal{N}_R(\mathbf{1}))$, the intersection number $i(u,w)$ is at most D''. As a consequence, there is a constant D''' such that $\mathrm{d}_Y(u,w) \leq D'''$ for all proper domains Y.

Setting $D' = 4M + D'''$, condition (3) is satisfied, and the proof is complete. □

7.3 Proof of Theorem 7.1

In [33], Klarreich shows that $\mathrm{F}\colon \mathcal{L}_{\mathrm{fill}}(S) \to \mathcal{EL}(S)$ is a closed map. Combining this with Theorem 7.4, we obtain

Corollary 7.5. *If Φ_v is a quasi-isometric embedding, then for any $m \in \partial G$, the lamination $\Phi_v(m)$ is uniquely ergodic. Moreover, we can lift $\partial\Phi_v$ to a continuous G–equivariant map*

$$\partial\Psi\colon\ \partial G \to \Lambda_G \subset \mathcal{L}_{\mathrm{fill}}(S) \subset \mathbb{P}\mathcal{ML}(S)$$

parameterizing the limit set, and the weak hull $\mathfrak{H}_G$ of Λ_G is defined and cobounded.

Proof. Let m_- and m_+ be any pair of points in ∂G. If $[\mu_\pm] \in \mathrm{F}^{-1}(\Phi_v(m_\pm))$, then μ_- and μ_+ bind S and so correspond to a Teichmüller geodesic $\tau(\mu_-, \mu_+)$. Since $[\mu_\pm]$ have supporting laminations $\Phi_v(m_\pm)$, Theorem 7.4 produces a constant D such that for any proper domain $Y \subset S$ with $\xi(Y) \neq 3$, the subsurface projection coefficient $\mathrm{d}_Y(\mu_-, \mu_+)$ is at most D. By Rafi's Theorem, Theorem 2.7 here, there is an $\varepsilon > 0$ depending only on S and D such that the Teichmüller geodesic $\tau(\mu_-, \mu_+)$ is ε–cobounded. By Masur's criterion, the laminations $[\mu_\pm]$ (or equivalently, $\Phi_v(m_\pm)$) are uniquely ergodic. In particular note that $\mathrm{F}^{-1}(\Phi_v(m_\pm))$ are in fact singletons $\{[\mu_\pm]\}$. Therefore, for any $m \in \partial G$, the lift $\partial\Psi(m)$ given by $\mathrm{F}^{-1}(\Phi_v(m))$ is defined.

Since $\partial\Psi$ is the unique lift of $\partial\Phi_v$, it follows that if V is any closed set in $\mathbb{PML}(S)$, $\partial\Phi_v^{-1}(\mathrm{F}(V)) = \partial\Psi^{-1}(V)$. Because F is a closed map and $\partial\Phi_v$ is continuous, we see that $\partial\Psi$ is continuous. Furthermore, as fixed points of hyperbolic elements of G are dense in ∂G, fixed points of pseudo-Anosov elements of G are dense in $\partial\Psi(\partial G)$. Therefore, being the image of a compact set, $\partial\Psi(\partial G)$ must agree with Λ_G. In particular, $\mathfrak{H}_G$ is defined and ε–cobounded. □

We now give the

Proof of Theorem 7.1. According to Corollary 7.5, $\mathfrak{H}_G$ is defined and ε–cobounded for some ε. Let A be the constant given by Theorem 4.5.

By Corollary 4.7, $\mathcal{N}_A(\mathfrak{H}_G)$ (with the induced path metric) is a δ-hyperbolic metric space for some δ. Moreover, the inclusion $\mathfrak{H}_G \to \mathcal{T}(S)$ is an isometric embedding and every geodesic in $\mathcal{T}(S)$ connecting a pair of points in $\mathfrak{H}_G$ is contained in $\mathcal{N}_A(\mathfrak{H}_G)$, by Theorem 4.5.

Let X be a point in $\mathfrak{H}_G$. By Lemma 7.2, $\Psi_X\colon G \to \mathcal{T}(S)$ is a quasi-isometric embedding and so $\Psi_X\colon G \to \mathcal{N}_A(\mathfrak{H}_G)$ is a quasi-isometric embedding by the above. By the stability of quasi-geodesics in a δ–hyperbolic metric space, $\Psi_x(G) \subset \mathfrak{H}_G$ is quasi-convex in $\mathcal{N}_A(\mathfrak{H}_G)$, and so in $\mathcal{T}(S)$. □

7.4 Electricity and the converse

Theorem 7.6. *If $G < \mathrm{Mod}(S)$ is convex cocompact, then for any $v \in \mathcal{C}(S)$, Φ_v is a quasi-isometric embedding.*

Remark. See also Hamenstädt, [27].

Proof. Let $\varepsilon > 0$. For $\alpha \in \mathcal{C}_0(S)$, let

$$\mathrm{thin}(\alpha, \varepsilon) = \{X \in \mathcal{T}(S) \,|\, \mathrm{ext}_X(\alpha) \leq \varepsilon\}$$

and let

$$\mathrm{thin}(\varepsilon) = \bigcup_{\alpha \in \mathcal{C}_0} \mathrm{thin}(\alpha, \varepsilon).$$

Let G be a convex cocompact subgroup of $\mathrm{Mod}(S)$. Let ε_0 be a number small enough so that the nerve of the family $\{\mathrm{thin}(\alpha, \varepsilon_0)\}$ is the complex of curves and so that $\mathfrak{H}_G$ is ε_0–cobounded.

Recall that $\mathcal{T}_{\mathrm{el}}(S)$ is obtained from $\mathcal{T}(S)$ in the following way. For each $\mathrm{thin}(\alpha,\varepsilon_0)$, we create a new point u_α and adjoin an interval of length $\frac{1}{2}$ between each point in $\mathrm{thin}(\alpha,\varepsilon_0)$ and u_α. Taking the induced path metric yields the space $\mathcal{T}_{\mathrm{el}}(S)$.

Since $\mathcal{C}(S)$ is quasi-isometric to $\mathcal{T}_{\mathrm{el}}(S)$, it suffices to show that an orbit in $\mathcal{T}_{\mathrm{el}}(S)$ defines a quasi-isometric embedding $G \to \mathcal{T}_{\mathrm{el}}(S)$. We claim that there are constants $K \geq 1$ and $C \geq 0$ such that the bi-infinite Teichmüller geodesics in $\mathfrak{H}_G$ are all (K,C)–quasi-geodesics in $\mathcal{T}_{\mathrm{el}}(S)$. As any two points in $\mathfrak{H}_G$ are within $2\delta = 2\delta(\varepsilon_0)$ of a bi-infinite geodesic in $\mathfrak{H}_G$ (by Theorem 4.5) and the orbit map from G to $\mathfrak{H}_G$ is a quasi-isometry, this will complete the proof.

Let $\mathbb{R}$ be the set of real numbers equipped with its usual metric. We electrify $\mathbb{R}$ as follows. Let $\mathcal{J}(r)$ be the collection of all closed intervals $\mathbb{J}$ of length $r \geq 1$. For each $\mathbb{J}$ in $\mathcal{J}(r)$, we create a new point $w_{\mathbb{J}}$ and adjoin an interval of length $\frac{1}{2}$ between each point in $\mathbb{J}$ and $w_{\mathbb{J}}$. We equip this space with the induced path metric and call the result $\mathbb{R}_{\mathrm{el}}(r)$. It is not difficult to see that the inclusion $\mathbb{R} \to \mathbb{R}_{\mathrm{el}}(r)$ is an $(r,0)$–quasi-isometry.

Let τ be a bi-infinite geodesic in $\mathfrak{H}_G$. Identifying τ with $\mathbb{R}$, we consider the projection $\pi_\tau\colon \mathcal{T}(S) \to \mathcal{P}(\tau)$ as a map $\pi_\tau\colon \mathcal{T}(S) \to \mathcal{P}(\mathbb{R})$ and, using the Axiom of Choice, we replace this with a map

$$\Pi_\tau\colon \mathcal{T}(S) \to \mathbb{R}$$

by demanding that $\Pi_\tau(X)$ be some element of $\pi_\tau(X)$, noting that for any two points X and Y in $\mathcal{T}(S)$,

$$\mathrm{d}_{\mathbb{R}}(\Pi_\tau(X),\Pi_\tau(Y)) \leq \mathrm{diam}(\pi_\tau(X) \cup \pi_\tau(Y)).$$

By Theorem 4.3 of [51], there is a constant $B \geq 1$, depending only on ε_0, such that

$$\mathrm{diam}(\Pi_\tau(\mathrm{thin}(\alpha,\varepsilon_0))) \leq B.$$

We extend the projection Π_τ to a projection

$$\widehat{\Pi}_\tau\colon \mathcal{T}_{\mathrm{el}}(S) \to \mathbb{R}_{\mathrm{el}}(B)$$

by demanding that the restriction $\widehat{\Pi}_\tau|_{\mathcal{T}(S)} = \Pi_\tau$, that for each α in $\mathcal{C}_0(S)$, $\widehat{\Pi}_\tau(u_\alpha) = w_{\mathbb{J}}$ for some interval $\mathbb{J}$ in $\mathcal{J}(B)$ containing $\Pi_\tau(\mathrm{thin}(\alpha,\varepsilon_0))$—again using the Axiom of Choice—and that an electric edge between u_α and a point X in $\mathrm{thin}(\alpha,\varepsilon_0)$ be sent isometrically to an electric edge joining $w_{\mathbb{J}}$ and $\Pi_\tau(X)$.

We claim that this projection is coarsely Lipschitz. To see this, let X and Y be elements of $\mathcal{T}_{\mathrm{el}}(S)$ with $\mathrm{d}_{\mathcal{T}_{\mathrm{el}}}(X,Y) \leq 1$. An electric geodesic γ joining X and Y is a concatenation of Teichmüller geodesic segments in $\mathcal{T}(S)$ and paths in unions of electric edges. Since $\mathrm{d}_{\mathcal{T}_{\mathrm{el}}}(X,Y) \leq 1$, we may write γ as a concatenation of paths $[X,X'] \cup [X',Y'] \cup [Y',Y]$, where $[X,X']$ and $[Y',Y]$ are paths in unions of electric edges, and $[X',Y']$ is a Teichmüller geodesic segment—any of which may be a constant path.

Since $\mathfrak{H}_G$ is ε_0–cobounded, so is τ. By Theorem 4.2, for any two points X and Y in $\mathcal{T}(S)$,

$$\mathrm{d}_{\mathbb{R}}(\Pi_\tau(X),\Pi_\tau(Y)) \leq \mathrm{d}_{\mathcal{T}}(X,Y) + 4b,$$

where $b = b(\varepsilon_0)$, and since the inclusion $\mathbb{R} \to \mathbb{R}_{\mathrm{el}}(B)$ is 1–Lipschitz, we have

$$\mathrm{d}_{\mathbb{R}_{\mathrm{el}}}(\widehat{\Pi}_\tau(X'), \widehat{\Pi}_\tau(Y')) \leq 1 + 4b.$$

Since $\widehat{\Pi}_\tau$ conducts electricity, we have

$$\mathrm{d}_{\mathbb{R}_{\mathrm{el}}}(\widehat{\Pi}_\tau(Z), \widehat{\Pi}_\tau(Z')) \leq 1$$

when $Z \in \{X, Y\}$. So,

$$\mathrm{d}_{\mathbb{R}_{\mathrm{el}}}(\widehat{\Pi}_\tau(X), \widehat{\Pi}_\tau(Y)) \leq 3 + 4b$$

and we see that $\widehat{\Pi}_\tau$ is (C', C')–Lipschitz, for $C' = 3 + 4b$.

Since $\mathbb{R} \to \mathbb{R}_{\mathrm{el}}(B)$ is a $(B, 0)$–quasi-isometry, for all X and Y in τ

$$\begin{aligned} \mathrm{d}_{\mathcal{T}_{\mathrm{el}}}(X, Y) &\geq \frac{1}{C'} \mathrm{d}_{\mathbb{R}_{\mathrm{el}}}(X, Y) - 1 \\ &\geq \frac{1}{BC'} \mathrm{d}_{\mathbb{R}}(X, Y) - 1 \\ &= \frac{1}{BC'} \mathrm{d}_{\mathcal{T}}(X, Y) - 1. \end{aligned}$$

Letting $K = BC'$ and $C = 1$ completes the proof. □

8 Questions

8.1 The analogy

For a Kleinian group, acting cocompactly on the domain of discontinuity is insufficient to guarantee convex cocompactness. For example, L. Bers first established the existence of **singly degenerate** Kleinian groups isomorphic to the fundamental group of a hyperbolic surface [4]: geometrically infinite groups whose domains of discontinuity are topological disks on which the groups act cocompactly.

When drawing an analogy between $\mathcal{T}(S)$ and $\mathbb{H}^3$ it is in many respects prudent to compare $\mathrm{Mod}(S)$ with a Kleinian group Γ of finite covolume. In this picture, the moduli space $\mathcal{M}(S)$ plays the role of the orbifold $M_\Gamma = \mathbb{H}^3/\Gamma$ and as $\mathcal{M}(S)$ is non-compact—and has finite volume in a certain sense (compare Masur [40], Section 5)—the analogy suggests that M_Γ also be non-compact.

The resolution of Marden's Tameness Conjecture by I. Agol [2] and (independently) D. Calegari and D. Gabai [14], combined with R. Canary's Covering Theorem [17], implies that a finitely generated subgroup of Γ is either geometrically finite or virtually the fiber subgroup of a hyperbolic 3–manifold fibering over the circle—see [16]. In particular, no groups like the ones constructed by Bers can occur in Γ.

A cocompact action of a subgroup of Γ on its domain of discontinuity is still however insufficient to guarantee convex cocompactness. One can construct examples of subgroups $\Gamma_0 < \Gamma$ which are geometrically finite, but for which all cusps are rank 2 (and hence "internal" to the convex core). In this situation, the convex hull of the limit

set of Γ_0 is not cobounded with respect to the covering $\mathbb{H}^3 \to M_\Gamma$. As we have seen, this behavior does not present itself in $\mathrm{Mod}(S)$, in light of Theorem 3.9. So, for subgroups of $\mathrm{Mod}(S)$ it is feasible that a cocompact action on the domain of discontinuity is equivalent to convex cocompactness.

Question 1. If a finitely generated subgroup G of $\mathrm{Mod}(S)$ acts cocompactly on $\Delta_G \neq \emptyset$, is it convex cocompact?

By Theorem 3.9, an affirmative answer to this question would follow from an affirmative answer to the following

Question 2. If G is a finitely generated subgroup of $\mathrm{Mod}(S)$ and $\mathfrak{H}_G$ is cobounded, is G convex cocompact?

When G is convex cocompact, Δ_G is the largest open set in $\mathbb{PML}(S)$ on which G acts properly discontinuously. This is also true for Veech groups.

Question 3. Let G be a finitely generated subgroup of $\mathrm{Mod}(S)$. Is Δ_G the largest open set in $\mathbb{PML}(S)$ on which G acts properly discontinuously?

We note that the action of G on the preimage of Δ_G in $\mathcal{ML}(S)$ is also properly discontinuous, and it has been shown by C. Lecuire [34] that the handlebody group G (in genus at least 3) acts properly discontinuously on a strictly larger set in $\mathcal{ML}(S)$ than the preimage of Δ_G.

Question 4. If one takes $\check{\Delta}_G$ to be the largest open set on which G acts properly discontinuously, is it true that G is convex cocompact if and only if $(\mathcal{T}(S) \cup \check{\Delta}_G)/G$ is compact?

We note that the answer to this question is affirmative if the answer to the previous one is as well.

8.2 Examples

At present, the only known examples of convex cocompact subgroups of $\mathrm{Mod}(S)$ are virtually free. To the authors' knowledge, the only known examples are: groups obtained by taking powers of independent pseudo-Anosov mapping classes; certain free products of finite subgroups of $\mathrm{Mod}(S)$, constructed by Honglin Min [49]; and purely pseudo-Anosov subgroups of graphs of Veech groups, due to the second author [35].

In [42], Masur studies the group of mapping classes of S that extend over a handlebody, called the **handlebody group**.

Question 5. Is every finitely generated purely pseudo-Anosov subgroup of the handlebody group convex cocompact?

Let $\dot{S}$ denote the surface S minus a point. There is a short exact sequence

$$1 \to \pi_1(S) \to \mathrm{Mod}(\dot{S}) \to \mathrm{Mod}(S) \to 1$$

where an element of $\pi_1(S)$ is sent to the mapping class that "spins" the puncture about the corresponding loop in S and $\mathrm{Mod}(\dot{S}) \to \mathrm{Mod}(S)$ is the map forgetting the puncture—see [7].

Question 6. Is every finitely generated purely pseudo-Anosov subgroup of $\pi_1(S)$ a convex cocompact subgroup of $\mathrm{Mod}(\dot{S})$?

An affirmative answer to this question would show that K. Whittlesey's group [66] is locally convex cocompact—this is a normal purely pseudo-Anosov subgroup of the mapping class group of a surface of genus two and is isomorphic to a free group of infinite rank.

Question 7 (Farb–Mosher [20]). Is every finitely generated subgroup of Whittlesey's group convex cocompact?

A more delicate question is

Question 8 (Farb–Mosher [20]). Is there a convex cocompact subgroup G of $\mathrm{Mod}(S)$ that is not virtually free?

And a more daring question is

Question 9 (Reid [60]). Let $m \geq 3$ be less than the virtual cohomological dimension of $\mathrm{Mod}(S)$ and let Γ be a torsion free uniform lattice in $\mathrm{SO}(m,1)$. Is there an injection $\Gamma \hookrightarrow \mathrm{Mod}(S)$ whose image is purely pseudo-Anosov?

Note that if Γ is the fundamental group of a closed fibered hyperbolic 3–manifold with fiber subgroup Σ and Γ injects into $\mathrm{Mod}(S)$ with convex cocompact image, then Σ could not act cocompactly on its weak hull $\mathfrak{H}_\Sigma$, as $\mathfrak{H}_\Sigma$ would equal $\mathfrak{H}_\Gamma$. Such a Σ would resolve Question 2 in the negative and it follows from work in [20] that the associated $\pi_1(S)$–extension of Σ would be a non-hyperbolic group with a finite Eilenberg–Mac Lane space and no Baumslag–Solitar subgroups—see [29] and Question 1.1 of [5].

8.3 The sociology of ending laminations

Theorem 1.3 implies that the Gromov boundary of a convex cocompact G embeds in the boundary of $\mathcal{C}(S)$, the space $\mathcal{EL}(S)$ of potential ending laminations for hyperbolic 3–manifolds homeomorphic to $S \times \mathbb{R}$. So, if $\mathcal{EL}(S)$ is totally disconnected, then every convex cocompact subgroup of $\mathrm{Mod}(S)$ is virtually free. To see this, note that ∂G is compact and so, provided G is not virtually cyclic, total disconnectedness of $\mathcal{EL}(S)$ along with the above embedding implies that ∂G is a Cantor set. Such a group is virtually free [63, 25, 24].

With this in mind, we close with a question of Peter Storm and a related question.

Question 10 (Storm). Is $\mathcal{EL}(S)$ connected? Is it path connected?

This is closely related to connectivity outside large balls in $\mathcal{C}(S)$. Specifically, the following is unknown in general.

Question 11. Does there exists an $A > 0$ such that given any $R > 0$ and any two points $u, v \in \mathcal{C}(S)$ outside a ball of radius R, there is a path connecting u to v that remains outside the ball (with the same center) of radius $R - A$?

The answer to this question has been resolved by S. Schleimer [61] when S is a once-punctured surface with genus at least 2. Indeed, Schleimer shows in this case that the complement of any R-ball is path connected.

References

[1] William Abikoff. *The real analytic theory of Teichmüller space*, volume 820 of *Lecture Notes in Mathematics*. Springer, Berlin, 1980.

[2] Ian Agol. Tameness of hyperbolic 3-manifolds, Preprint. `arXiv:math.GT/0405568`,.

[3] Lars V. Ahlfors. *Lectures on quasiconformal mappings*. Manuscript prepared with the assistance of Clifford J. Earle, Jr. Van Nostrand Mathematical Studies, No. 10. D. Van Nostrand Co., Inc., Toronto, Ont.-New York-London, 1966.

[4] Lipman Bers. On boundaries of Teichmüller spaces and on Kleinian groups. I. *Ann. of Math. (2)*, 91:570–600, 1970.

[5] Mladen Bestvina. Questions in geometric group theory. `http://www.math.utah.edu/~bestvina`.

[6] Mladen Bestvina and Mark Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

[7] Joan S. Birman. Mapping class groups and their relationship to braid groups. *Comm. Pure Appl. Math.*, 22:213–238, 1969.

[8] Francis Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.

[9] Francis Bonahon. Geodesic laminations on surfaces. In *Laminations and foliations in dynamics, geometry and topology (Stony Brook, NY, 1998)*, volume 269 of *Contemp. Math.*, pages 1–37. Amer. Math. Soc., Providence, RI, 2001.

[10] Brian H. Bowditch. Geometrical finiteness for hyperbolic groups. *J. Funct. Anal.*, 113(2):245–317, 1993.

[11] Brian H. Bowditch. Intersection numbers and the hyperbolicity of the curve complex. *J. Reine Angew. Math.*, 598:105–129, 2006.

[12] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[13] Jeffrey Brock and Benson Farb. Curvature and rank of Teichmüller space. Preprint, `arXiv:math.GT/0109045`.

[14] Danny Calegari and David Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. Preprint, `arXiv:math.GT/0407161`.

[15] R. D. Canary, D. B. A. Epstein, and P. Green. Notes on notes of Thurston. In *Analytical and geometric aspects of hyperbolic space (Coventry/Durham, 1984)*, volume 111 of *London Math. Soc. Lecture Note Ser.*, pages 3–92. Cambridge Univ. Press, Cambridge, 1987.

[16] Richard D. Canary. Covering theorems for hyperbolic 3-manifolds. In *Low-dimensional topology (Knoxville, TN, 1992)*, Conf. Proc. Lecture Notes Geom. Topology, III, pages 21–30. Internat. Press, Cambridge, MA, 1994.

[17] Richard D. Canary. A covering theorem for hyperbolic 3-manifolds and its applications. *Topology*, 35(3):751–778, 1996.

[18] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes*, volume 1441 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1990.

[19] Moon Duchin. Thin triangles and a multiplicative ergodic theorem for Teichmüller geometry. Preprint, `arXiv:math.GT/0508046`.

[20] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152 (electronic), 2002.

[21] Benson Farb and Lee Mosher. The geometry of surface-by-free groups. *Geom. Funct. Anal.*, 12(5):915–963, 2002.

[22] A. Fathi, F. Laudenbach, and V. Poénaru. *Travaux de Thurston sur les surfaces*. Société Mathématique de France, Paris, 1991. Séminaire Orsay, Reprint of *Travaux de Thurston sur les surfaces*, Soc. Math. France, Paris, 1979 Astérisque No. 66-67 (1991).

[23] Frederick P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York, 1987. A Wiley-Interscience Publication.

[24] Étienne Ghys and Pierre de la Harpe, editors. *Sur les groupes hyperboliques d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1990. Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.

[25] Mikhaïl Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[26] Noemi Halpern. A proof of the collar lemma. *Bull. London Math. Soc.*, 13(2):141–144, 1981.

[27] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, `arXiv:math.GT/0505244`.

[28] John Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Math.*, 142(3-4):221–274, 1979.

[29] Richard P. Kent IV and Christopher J. Leininger. Subgroups of the mapping class group from the geometrical viewpoint . To appear in *Proceedings of the 2005 Ahlfors–Bers Colloquium*, `arXiv:math.GT/0702034`.

[30] Nikolai V. Ivanov. *Subgroups of Teichmüller modular groups*, volume 115 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992. Translated from the Russian by E. J. F. Primrose and revised by the author.

[31] Linda Keen. Collars on Riemann surfaces. In *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pages 263–268. Ann. of Math. Studies, No. 79. Princeton Univ. Press, Princeton, N.J., 1974.

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

[33] Erica Klarreich. The boundary at infinity of the curve complex and the relative Teichmüller space. Preprint, `http://nasw.org/users/klarreich/publications.htm`,.

[34] Cyril Lecuire. Structure hyperboliques convexes sur les variétés de dimension 3, *thèse de doctorat de l'ENS Lyon*. `http://www.maths.warwick.ac.uk/~clecuire`.

[35] Christopher J. Leininger. Graphs of veech groups. Work in progress.

[36] Gilbert Levitt. Foliations and laminations on hyperbolic surfaces. *Topology*, 22(2):119–135, 1983.

[37] Albert Marden. The geometry of finitely generated kleinian groups. *Ann. of Math. (2)*, 99:383–462, 1974.

[38] Howard Masur. On a class of geodesics in Teichmüller space. *Ann. of Math. (2)*, 102(2):205–221, 1975.

[39] Howard Masur. Uniquely ergodic quadratic differentials. *Comment. Math. Helv.*, 55(2):255–266, 1980.

[40] Howard Masur. Interval exchange transformations and measured foliations. *Ann. of Math. (2)*, 115(1):169–200, 1982.

[41] Howard Masur. Two boundaries of Teichmüller space. *Duke Math. J.*, 49(1):183–190, 1982.

[42] Howard Masur. Measured foliations and handlebodies. *Ergodic Theory Dynam. Systems*, 6(1):99–116, 1986.

[43] Howard Masur. Hausdorff dimension of the set of nonergodic foliations of a quadratic differential. *Duke Math. J.*, 66(3):387–442, 1992.

[44] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.

[45] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[46] Howard A. Masur and Yair N. Minsky. Unstable quasi-geodesics in Teichmüller space. In *In the tradition of Ahlfors and Bers (Stony Brook, NY, 1998)*, volume 256 of *Contemp. Math.*, pages 239–241. Amer. Math. Soc., Providence, RI, 2000.

[47] Howard A. Masur and Michael Wolf. Teichmüller space is not Gromov hyperbolic. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 20(2):259–267, 1995.

[48] John McCarthy and Athanase Papadopoulos. Dynamics on Thurston's sphere of projective measured foliations. *Comment. Math. Helv.*, 64(1):133–166, 1989.

[49] Honglin Min. Work in progress.

[50] Yair N. Minsky. The classification of Kleinian surface groups, I: Models and bounds. to appear *Ann. of Math.*, `arXiv:math.GT/0302208`.

[51] Yair N. Minsky. Quasi-projections in Teichmüller space. *J. Reine Angew. Math.*, 473:121–136, 1996.

[52] Yair N. Minsky. Kleinian groups and the complex of curves. *Geom. Topol.*, 4:117–148, 2000.

[53] Yair N. Minsky. Bounded geometry for Kleinian groups. *Invent. Math.*, 146(1):143–192, 2001.

[54] Mahan Mitra. Ending laminations for hyperbolic group extensions. *Geom. Funct. Anal.*, 7(2):379–402, 1997.

[55] Mahan Mitra. Cannon-Thurston maps for hyperbolic group extensions. *Topology*, 37(3):527–538, 1998.

[56] Mahan Mitra. Cannon-Thurston maps for trees of hyperbolic metric spaces. *J. Differential Geom.*, 48(1):135–164, 1998.

[57] Lee Mosher. A hyperbolic-by-hyperbolic hyperbolic group. *Proc. Amer. Math. Soc.*, 125(12):3447–3455, 1997.

[58] David Mumford. A remark on Mahler's compactness theorem. *Proc. Amer. Math. Soc.*, 28:289–294, 1971.

[59] Kasra Rafi. A characterization of short curves of a Teichmüller geodesic. *Geom. Topol.*, 9:179–202, 2005.

[60] Alan W. Reid. Surface subgroups of mapping class groups. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 257–268. Amer. Math. Soc., Providence, RI, 2006.

[61] Saul Schleimer. The end of the curve complex. Preprint, `arXiv:math.GT/0608505`,.

[62] J.-P. Serre. Rigidité de foncteur d'Jacobi d'échelon $n \geq 3$. In *Séminaire Henri Cartan, 1960/61, Exposé 17 (Appendice)*. Secrétariat mathématique, Paris, 1960/61.

[63] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.

[64] Dennis Sullivan. Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups. *Acta Math.*, 155(3-4):243–260, 1985.

[65] Eric L. Swenson. Quasi-convex groups of isometries of negatively curved spaces. *Topology Appl.*, 110(1):119–129, 2001. Geometric topology and geometric group theory (Milwaukee, WI, 1997).

[66] Kim Whittlesey. Normal all pseudo-Anosov subgroups of mapping class groups. *Geom. Topol.*, 4:293–307 (electronic), 2000.

Department of Mathematics, Brown University, Providence, RI 02912
`rkent@math.brown.edu`

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