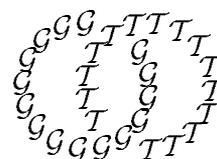


Geometry & Topology
 Volume 6 (2002) 91–152
 Published: 14 March 2002



Convex cocompact subgroups of mapping class groups

BENSON FARB
 LEE MOSHER

Department of Mathematics, University of Chicago
 5734 University Ave, Chicago, IL 60637, USA
 and

Department of Mathematics and Computer Science
 Rutgers University, Newark, NJ 07102, USA

Email: farb@math.uchicago.edu and mosher@andromeda.rutgers.edu

Abstract

We develop a theory of convex cocompact subgroups of the mapping class group MCG of a closed, oriented surface S of genus at least 2, in terms of the action on Teichmüller space. Given a subgroup G of MCG defining an extension $1 \rightarrow \pi_1(S) \rightarrow \Gamma_G \rightarrow G \rightarrow 1$, we prove that if Γ_G is a word hyperbolic group then G is a convex cocompact subgroup of MCG . When G is free and convex cocompact, called a *Schottky subgroup* of MCG , the converse is true as well; a semidirect product of $\pi_1(S)$ by a free group G is therefore word hyperbolic if and only if G is a Schottky subgroup of MCG . The special case when $G = \mathbf{Z}$ follows from Thurston's hyperbolization theorem. Schottky subgroups exist in abundance: sufficiently high powers of any independent set of pseudo-Anosov mapping classes freely generate a Schottky subgroup.

AMS Classification numbers Primary: 20F67, 20F65

Secondary: 57M07, 57S25

Keywords: Mapping class group, Schottky subgroup, cocompact subgroup, convexity, pseudo-Anosov

Proposed: Walter Neumann
 Seconded: Shigeyuki Morita, Robion Kirby

Received: 20 October 2001
 Accepted: 20 February 2002

1 Introduction

1.1 Convex cocompact groups

A *convex cocompact* subgroup of $\text{Isom}(\mathbf{H}^n)$, the isometry group of hyperbolic n -space, is a discrete subgroup $G < \text{Isom}(\mathbf{H}^n)$, with limit set $\Lambda_G \subset \partial\mathbf{H}^n$, such that G acts cocompactly on the convex hull $\text{Hull}_G \subset \mathbf{H}^n$ of its limit set Λ_G . It follows that G is a word hyperbolic group with model geometry Hull_G and Gromov boundary Λ_G . Given any finitely generated, discrete subgroup $G < \text{Isom}(\mathbf{H}^n)$, G is convex cocompact if and only if any orbit of G is a quasiconvex subset of \mathbf{H}^n . Convex cocompact subgroups satisfy several useful properties: every infinite order element of G is loxodromic; Λ_G is the smallest nontrivial G -invariant closed subset of $\overline{\mathbf{H}^n} = \mathbf{H}^n \cup \partial\mathbf{H}^n$; the action of G on $\partial\mathbf{H}^n \setminus \Lambda_G$ is properly discontinuous; assuming $\Lambda_G \neq \partial\mathbf{H}^n$, the stabilizer subgroup of Λ_G is a finite index supergroup of G , and it is the relative commensurator of G in $\text{Isom}(\mathbf{H}^n)$.

A *Schottky group* is a convex cocompact subgroup of $\text{Isom}(\mathbf{H}^n)$ which is free. Schottky subgroups of $\text{Isom}(\mathbf{H}^n)$ exist in abundance and can be constructed using the classical ping-pong argument, attributed to Klein: if ϕ_1, \dots, ϕ_n are loxodromic elements whose axes have pairwise disjoint endpoints at infinity, then sufficiently high powers of ϕ_1, \dots, ϕ_n freely generate a Schottky group.¹

We shall investigate the notions of convex cocompact groups and Schottky groups in the context of Teichmüller space. Given a closed, oriented surface S of genus ≥ 2 , the mapping class group MCG acts as the full isometry group of the Teichmüller space \mathcal{T} [45].² This action extends to the Thurston compactification $\overline{\mathcal{T}} = \mathcal{T} \cup \mathcal{PMF}$ [16]. Teichmüller space is *not* Gromov hyperbolic [34], no matter what finite covolume, equivariant metric one picks [10], and yet it exhibits many aspects of a hyperbolic metric space [38] [32]. A general theory of limit sets of finitely generated subgroups of MCG is developed in [36].

In this paper we develop a theory of convex cocompact subgroups and Schottky subgroups of MCG acting on \mathcal{T} , and we show that Schottky subgroups exist in abundance. We apply this theory to relate convex cocompactness of subgroups of MCG with the large scale geometry of extensions of surface groups by subgroups of MCG .

¹The term ‘‘Schottky group’’ sometimes refers explicitly to a subgroup of $\text{Isom}(\mathbf{H}^n)$ produced by the ping-pong argument, but the broader reference to free, convex cocompact subgroups has become common.

²In this paper, MCG includes orientation reversing mapping classes, and so represents what is sometimes called the ‘‘extended’’ mapping class group.

Our first result establishes the concept of convex cocompactness for subgroups of MCG , by proving the equivalence of several properties:

Theorem 1.1 (Characterizing convex cocompactness) *Given a finitely generated subgroup $G < MCG$, the following statements are equivalent:*

- *Some orbit of G is quasiconvex in \mathcal{T} .*
- *Every orbit of G is quasiconvex in \mathcal{T} .*
- *G is word hyperbolic, and there is a G -equivariant embedding $\partial f: \partial G \rightarrow \mathbf{PMF}$ with image Λ_G such that the following properties hold:*
 - *Any two distinct points $\xi, \eta \in \Lambda_G$ are the ideal endpoints of a unique geodesic $\overleftrightarrow{(\xi, \eta)}$ in \mathcal{T} .*
 - *Let WH_G be the “weak hull” of G , namely the union of the geodesics $\overleftrightarrow{(\xi, \eta)}$, $\xi \neq \eta \in \Lambda_G$. Then the action of G on WH_G is cocompact, and if $f: G \rightarrow \text{WH}_G$ is any G -equivariant map then f is a quasi-isometry and the following map is continuous:*

$$\bar{f} = f \cup \partial f: G \cup \partial G \rightarrow \bar{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$$

Any such subgroup G is said to be *convex cocompact*. This theorem is proved in Section 3.3.

A convex cocompact subgroup $G < MCG$ shares many properties with convex cocompact subgroups of $\text{Isom}(\mathbf{H}^n)$. Every infinite order element of G is pseudo-Anosov (Proposition 3.1). The limit set Λ_G is the smallest nontrivial closed subset of $\bar{\mathcal{T}}$ invariant under the action of G , and the action of G on $\mathbf{PMF} - \Lambda_G$ is properly discontinuous (Proposition 3.2); this depends on work of McCarthy and Papadoupoulos [36]. The stabilizer of Λ_G is a finite index supergroup of G in MCG , and it is the relative commensurator of G in MCG (Corollary 3.3).

A *Schottky subgroup* of $MCG = \text{Isom}(\mathcal{T})$ is defined to be a convex cocompact subgroup which is free of finite rank. In Theorem 1.4 we prove that if ϕ_1, \dots, ϕ_n are pseudo-Anosov elements of MCG whose axes have pairwise disjoint endpoints in \mathbf{PMF} , then for all sufficiently large positive integers a_1, \dots, a_n the mapping classes $\phi_1^{a_1}, \dots, \phi_n^{a_n}$ freely generate a Schottky subgroup of MCG .

Warning Our formulation of convex cocompactness in \mathcal{T} is not as strong as in \mathbf{H}^n . Although there is a general theory of limit sets of finitely generated subgroups of MCG [36], we have no general theory of their convex hulls. Such a theory would be tricky, and unnecessary for our purposes. In particular, when

G is convex cocompact, we do not know whether there is a closed, convex, G -equivariant subset of \mathcal{T} on which G acts cocompactly. One could attempt to construct such a subset by adding to WH_G any geodesics with endpoints in WH_G , then adding geodesics with endpoints in that set, etc, continuing transfinitely by adding geodesics and taking closures until the result stabilizes; however, there is no guarantee that G acts cocompactly on the result.

1.2 Surface group extensions

There is a natural isomorphism of short exact sequences

$$\begin{array}{ccccccc}
 1 & \longrightarrow & \pi_1(S, p) & \xrightarrow{\iota} & MCG(S, p) & \xrightarrow{q} & MCG(S) \longrightarrow 1 \\
 & & \parallel & & \Downarrow & & \Downarrow \\
 1 & \longrightarrow & \pi_1(S, p) & \longrightarrow & \text{Aut}(\pi_1(S, p)) & \longrightarrow & \text{Out}(\pi_1(S, p)) \longrightarrow 1
 \end{array}$$

where $MCG(S, p)$ is the mapping class group of S punctured at the base point p . In the bottom sequence, the inclusion $\pi_1(S, p)$ is obtained by identifying $\pi_1(S, p)$ with its group of inner automorphisms, an injection since $\pi_1(S, p)$ is centerless. For each based loop ℓ in S , $\iota(\ell)$ is the punctured mapping class which “pushes” the base point p around the loop ℓ (see Section 2.2 for the exact definition). The homomorphism q is the map which “forgets” the puncture p . Exactness of the top sequence is proved in [7]. The isomorphism $MCG(S) \approx \text{Out}(\pi_1(S, p))$ follows from work of Dehn–Nielsen [43], Baer [3], and Epstein [13]. As a consequence, either of the above sequences is natural for extensions of $\pi_1(S)$, in the following sense. For any group homomorphism $G \rightarrow MCG(S)$, by applying the fiber product construction to the homomorphisms

$$\begin{array}{ccc}
 MCG(S, p) & & G \\
 & \searrow & \swarrow \\
 & MCG(S) &
 \end{array}$$

we obtain a group Γ_G and a commutative diagram of short exact sequences

$$\begin{array}{ccccccc}
 1 & \longrightarrow & \pi_1(S) & \longrightarrow & \Gamma_G & \longrightarrow & G \longrightarrow 1 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 1 & \longrightarrow & \pi_1(S) & \longrightarrow & MCG(S, p) & \longrightarrow & MCG(S) \longrightarrow 1
 \end{array}$$

Note that we are suppressing the homomorphism $G \rightarrow MCG(S)$ in the notation Γ_G . If G is free then the top sequence splits and we can write $\Gamma_G = \pi_1(S) \rtimes G$.

G , where again our notation suppresses a lift $G \rightarrow \text{Aut}(\pi_1(S))$ of the given homomorphism $G \rightarrow MCG(S) \approx \text{Out}(\pi_1(S))$.

Every group extension $1 \rightarrow \pi_1(S) \rightarrow E \rightarrow G \rightarrow 1$ arises from the above construction, because the given extension determines a homomorphism $G \rightarrow \text{Out}(\pi_1(S)) \approx MCG(S)$ which in turn determines an extension $1 \rightarrow \pi_1(S) \rightarrow \Gamma_G \rightarrow G \rightarrow 1$ isomorphic to the given extension.

When P is a cyclic subgroup of MCG , Thurston's hyperbolization theorem for mapping tori (see, eg, [44]) shows that $\pi_1(S) \rtimes P$ is the fundamental group of a closed, hyperbolic 3-manifold if and only if P is a pseudo-Anosov subgroup. In particular, $\pi_1(S) \rtimes P$ is a word hyperbolic group if and only if P is a convex cocompact subgroup of MCG . Our results about the extension groups Γ_G are aimed towards generalizing this statement as much as possible. The theme of these results is that the geometry of Γ_G is encoded in the geometry of the action of G on \mathcal{T} .

From [39] it follows that if Γ_G is word hyperbolic then G is word hyperbolic. Our next result gives much more precise information:

Theorem 1.2 (Hyperbolic extension has convex cocompact quotient) *If Γ_G is word hyperbolic then the homomorphism $G \rightarrow MCG$ has finite kernel and convex cocompact image.*

This theorem is proved in Section 5. Finiteness of the kernel K of $G \rightarrow MCG$ is easy to prove, using the fact that $\pi_1(S) \times K$ is a subgroup of Γ_G . If K is infinite, then either it is a torsion group, or it has an infinite order element and so Γ_G has a $\mathbf{Z} \oplus \mathbf{Z}$ subgroup; in either case, Γ_G cannot be word hyperbolic. Because one can mod out by a finite kernel without affecting word hyperbolicity of the extension group, this brings into focus the extensions defined by inclusion of subgroups of MCG .

We are particularly interested in free subgroups of MCG . A finite rank, free, convex cocompact subgroup is called a *Schottky subgroup*. For Schottky subgroups we have a converse to Theorem 1.2, giving a complete characterization of word hyperbolic groups Γ_F when $F < MCG$ is free:

Theorem 1.3 (Surface-by-Schottky group has hyperbolic extension) *If F is a finite rank, free subgroup of MCG then the extension group $\Gamma_F = \pi_1(S) \rtimes F$ is word hyperbolic if and only if F is a Schottky group.*

This is proved in Section 6. Some special cases of this theorem are immediate. It is not hard to see that $\pi_1(S) \rtimes F$ has a $\mathbf{Z} \oplus \mathbf{Z}$ subgroup if and only if there

exists a nontrivial element $f \in F$ which is not pseudo-Anosov. Such an element f , being infinite order, must be reducible. Assuming $f \in F$ is nontrivial and reducible, the group $\pi_1(S) \rtimes F$ contains the subgroup $\pi_1(S) \rtimes \langle f \rangle$ which is the fundamental group of a closed 3-manifold that contains an incompressible torus. Conversely, when $\pi_1(S) \rtimes F$ has a $\mathbf{Z} \oplus \mathbf{Z}$ subgroup then that subgroup must map onto an infinite cyclic subgroup $\langle f \rangle \subset F$ whose action on $\pi_1(S)$ preserves a nontrivial conjugacy class, and so f is not pseudo-Anosov. Theorem 1.3 is therefore mainly about free, pseudo-Anosov subgroups of MCG (see Question 1.5 below).

The abundance of word hyperbolic extensions of the form $\pi_1(S) \rtimes F$ was proved in [40]. It was shown by McCarthy [35] and Ivanov [23] that if ϕ_1, \dots, ϕ_n are pseudo-Anosov elements of MCG which are pairwise independent, meaning that their axes have distinct endpoints in the Thurston boundary \mathbf{PMF} , then sufficiently high powers of these elements freely generate a pseudo-Anosov subgroup F . The main result of [40] shows in addition that, after possibly making the powers higher, the group $\pi_1(S) \rtimes F$ is word hyperbolic. The nature of the free subgroups $F < MCG$ produced in [40] was somewhat mysterious, but Theorems 1.2 and 1.3 clear up this mystery by characterizing the subgroups F using an intrinsic property, namely convex cocompactness.

By combining [40] and Theorem 1.3, we immediately have the following result:

Theorem 1.4 (Abundance of Schottky subgroups) *If $\phi_1, \dots, \phi_n \in MCG$ are pairwise independent pseudo-Anosov elements, then for all sufficiently large positive integers a_1, \dots, a_n the mapping classes $\phi_1^{a_1}, \dots, \phi_n^{a_n}$ freely generate a Schottky subgroup F of MCG .*

Finally, we shall show in Section 7 that all of the above results generalize to the setting of closed hyperbolic 2-orbifolds. These generalized results find application in the results of [15], as we now recall.

1.3 An application

In the paper [15] we apply our theory of Schottky subgroups of MCG to investigate the large-scale geometry of word hyperbolic surface-by-free groups:

Theorem [15] *Let $F \subset MCG(S)$ be Schottky. Then the group $\Gamma_F = \pi_1(S) \rtimes F$ is quasi-isometrically rigid in the strongest sense:*

- Γ_F embeds with finite index in its quasi-isometry group $\mathbf{QI}(\Gamma_F)$.

It follows that:

- Let H be any finitely generated group. If H is quasi-isometric to Γ_F , then there exists a finite normal subgroup $N \triangleleft H$ such that H/N and Γ_F are abstractly commensurable.
- The abstract commensurator group $\text{Comm}(\Gamma_F)$ is isomorphic to $\text{QI}(\Gamma_F)$, and can be computed explicitly.

The computation of $\text{Comm}(\Gamma_F) \approx \text{QI}(\Gamma_F)$ goes as follows. Among all orbifold subcovers $S \rightarrow \mathcal{O}$ there exists a unique minimal such subcover such that the subgroup $F < MCG(S)$ descends isomorphically to a subgroup $F' < MCG(\mathcal{O})$. The whole theory of Schottky groups extends to general closed hyperbolic orbifolds, as we show in Section 7 of this paper. In particular, F' is a Schottky subgroup of $MCG(\mathcal{O})$. By Corollary 3.3 it follows that F' has finite index in its relative commensurator $N < MCG(\mathcal{O})$, which can be regarded as a virtual Schottky group. The inclusion $N < MCG(\mathcal{O})$ determines a canonical extension $1 \rightarrow \pi_1(\mathcal{O}) \rightarrow \Gamma_N \rightarrow N \rightarrow 1$, and we show in [15] that the extension group Γ_N is isomorphic to $\text{QI}(\Gamma_F)$.

1.4 Some questions

Our results on convex cocompact and Schottky subgroups of MCG motivate several questions.

Proposition 3.1 implies that if F is a Schottky subgroup of MCG then every nontrivial element of F is pseudo-Anosov.

Question 1.5 *Suppose $F < MCG$ is a finite rank, free subgroup all of whose nontrivial elements are pseudo-Anosov. Is F convex cocompact? In other words, is F a Schottky group?*

A non-Schottky example F would be very interesting for the following reasons. There exist examples of infinite, finitely presented groups which are not word hyperbolic and whose solvable subgroups are all virtually cyclic, but all known examples fail to be of finite type; see for example [9]. If there were a non-Schottky subgroup $F < MCG$ as in Question 1.5, then the group $\pi_1(S) \rtimes F$ would be of finite type (being the fundamental group of a compact aspherical 3-complex), it would not be word hyperbolic (since F is not Schottky), and every nontrivial solvable subgroup $H < \pi_1(S) \rtimes F$ would be infinite cyclic. To see why the latter holds, since $\pi_1(S) \rtimes F$ is a torsion free subgroup of

$MCG(S, p)$ it follows by [8] that the subgroup H is finite rank free abelian. Under the homomorphism $H \rightarrow F$, the groups $\text{image}(H \rightarrow F) < F$ and $\text{kernel}(H \rightarrow F) < \pi_1(S)$ each are free abelian of rank at most 1, and so it suffices to rule out the case where the image and kernel both have rank 1. But in that case we would have a pseudo-Anosov element of $MCG(S)$ which fixes the conjugacy class of some infinite order element of $\pi_1 S$, a contradiction.

Note that Question 1.5 has an analogue in the theory of Kleinian groups: if G is a discrete, cocompact subgroup of $\text{Isom}(\mathbf{H}^3)$, is every free subgroup of G a Schottky subgroup? More generally, if G is a discrete, cofinite volume subgroup of $\text{Isom}(\mathbf{H}^3)$, is every free loxodromic subgroup of G a Schottky group? The first question, at least, would follow from Simon's tame ends conjecture [11].

For a source of free, pseudo-Anosov subgroups on which to test question 1.5, consider Whittlesey's group [47], an infinite rank, free, normal, pseudo-Anosov subgroup of the mapping class group of a closed, oriented surface of genus 2.

Question 1.6 *Is every finitely generated subgroup of Whittlesey's group a Schottky group?*

Concerning non-free subgroups of MCG , note first that Question 1.5 can also be formulated for any finitely generated subgroup of MCG , though we have no examples of non-free pseudo-Anosov subgroups. This invites comparison with the situation in $\text{Isom}(\mathbf{H}^n)$ where it is known for any $n \geq 2$ that there exist convex cocompact subgroups which are not Schottky, indeed are not virtually Schottky.

Question 1.7 *Does there exist a convex cocompact subgroup $G < MCG$ which is not Schottky, nor is virtually Schottky?*

The converse to Theorem 1.2, while proved for free subgroups in Theorem 1.3, remains open in general. This issue becomes particularly interesting if Question 1.7 is answered affirmatively:

Question 1.8 *If $G < MCG$ is convex cocompact, is the extension group Γ_G word hyperbolic?*

Surface subgroups of mapping class groups are interesting. Gonzalez-Diez and Harvey showed that MCG can contain the fundamental group of a closed, oriented surface of genus ≥ 2 [19], but their construction always produces subgroups containing mapping classes that are not pseudo-Anosov.

If questions 1.7 and 1.8 were true, it would raise the stakes on the fascinating question of whether there exist surface-by-surface word hyperbolic groups:

Question 1.9 *Does there exist a convex cocompact subgroup $G < MCG$ isomorphic to the fundamental group of a closed, oriented surface S_g of genus $g \geq 2$? If so, is the surface-by-surface extension group Γ_G word hyperbolic?*

Misha Kapovich shows in [25] that when G is a surface group, the extension group Γ_G cannot be a lattice in $\text{Isom}(\mathbf{CH}^2)$.

1.5 Sketches of proofs

Although Teichmüller space \mathcal{T} is not hyperbolic in any reasonable sense [34], [10], nevertheless it possesses interesting and useful hyperbolicity properties. To formulate these, recall that the action of MCG by isometries on \mathcal{T} is smooth and properly discontinuous, with quotient orbifold $\mathcal{M} = \mathcal{T}/MCG$ called the *moduli space* of S . The action is *not* cocompact, and we define a subset $A \subset \mathcal{T}$ to be *cobounded* if its image under the universal covering map $\mathcal{T} \rightarrow \mathcal{M}$ has compact closure in \mathcal{M} , equivalently there is a compact subset of \mathcal{T} whose translates under $\text{Isom}(\mathcal{T})$ cover A .

In [38], Minsky proves (see Theorem 3.6 below) that if ℓ is a cobounded geodesic in \mathcal{T} then any projection $\mathcal{T} \rightarrow \ell$ that takes each point of \mathcal{T} to a closest point on ℓ satisfies properties similar to a closest point projection from a δ -hyperbolic metric space onto a bi-infinite geodesic. This projection property is a key step in the proof of the Masur–Minsky theorem [32] that Harvey’s curve complex is a δ -hyperbolic metric space. These results say intuitively that \mathcal{T} exhibits hyperbolicity as long as one focusses only on cobounded aspects. Keeping this in mind, the tools of [38] and [32] can be used to prove Theorem 1.1 along the classical lines of the proof for subgroups of $\text{Isom}(\mathbf{H}^n)$.

The proof of Theorem 1.3, that $\pi_1(S) \rtimes F$ is word hyperbolic if F is Schottky, uses the Bestvina–Feighn combination theorem [6]. Consider a tree \mathfrak{t} on which F acts freely and cocompactly, and choose an F -equivariant mapping $\phi: \mathfrak{t} \rightarrow \mathcal{T}$. Let $\mathcal{H} \rightarrow \mathcal{T}$ be the canonical hyperbolic plane bundle over Teichmüller space. Pulling back via ϕ we obtain a hyperbolic plane bundle $\pi: \mathcal{H}_{\mathfrak{t}} \rightarrow \mathfrak{t}$, and $\pi_1(S) \rtimes F$ acts properly discontinuously and cocompactly on $\mathcal{H}_{\mathfrak{t}}$. This shows that $\mathcal{H}_{\mathfrak{t}}$ is a model geometry for the group $\pi_1(S) \rtimes F$, and in particular $\mathcal{H}_{\mathfrak{t}}$ is a δ -hyperbolic metric space if and only if $\pi_1(S) \rtimes F$ is word hyperbolic.

By the Bestvina–Feighn combination theorem [6] and its converse due to Gersten [18], hyperbolicity of $\mathcal{H}_{\mathfrak{t}}$ is equivalent to δ -hyperbolicity of each “hyperplane” $\mathcal{H}_{\ell} = \pi^{-1}(\ell)$, where ℓ ranges over all the bi-infinite lines in \mathfrak{t} and δ is independent of ℓ .

Recall that for each Teichmüller geodesic g , the canonical marked Riemann surface bundle \mathcal{S}_g over g carries a natural *singular SOLV metric*; the bundle \mathcal{S}_g equipped with this metric is denoted $\mathcal{S}_g^{\text{SOLV}}$. Lifting the metric to the universal cover \mathcal{H}_g we obtain a singular SOLV space denoted $\mathcal{H}_g^{\text{SOLV}}$.

When F is a Schottky group, convex cocompactness tells us that for each bi-infinite geodesic ℓ in \mathfrak{t} , the map $\ell \xrightarrow{\phi} \mathcal{T}$ is a quasigeodesic and there is a unique Teichmüller geodesic g within finite Hausdorff distance from $\phi(\ell)$. This feeds into Proposition 4.2, a basic construction principle for quasi-isometries which will be used several times in the paper. The conclusion is:

Fact 1.10 *The hyperplane \mathcal{H}_ℓ is uniformly quasi-isometric to the singular SOLV-space $\mathcal{H}_g^{\text{SOLV}}$, by a quasi-isometry which is a lift of a closest point map $\ell \rightarrow g$.*

Uniform hyperbolicity of singular SOLV-spaces $\mathcal{H}_g^{\text{SOLV}}$, where g is a uniformly cobounded geodesic in \mathcal{T} , is then easily checked by another application of the Bestvina–Feighn combination theorem, and Theorem 1.3 follows.

For Theorem 1.2, we first outline the proof in the special case of a free subgroup of MCG . As noted above, using Gersten’s converse to the Bestvina–Feighn combination theorem, word hyperbolicity of $\pi_1(S) \rtimes F$ implies uniform hyperbolicity of the hyperplanes \mathcal{H}_ℓ . Now we use a result of Mosher [41], which shows that from uniform hyperbolicity of the hyperplanes \mathcal{H}_ℓ it follows that the lines ℓ are uniform quasigeodesics in \mathcal{T} , and each ℓ has uniformly finite Hausdorff distance from some Teichmüller geodesic g . Piecing together the geodesics g in \mathcal{T} , one for each geodesic ℓ in \mathfrak{t} , we obtain the data we need to prove that F is Schottky.

The general proof of Theorem 1.2 follows the same outline, except that we cannot apply Gersten’s converse to the Bestvina–Feighn combination theorem. That result applies only to the setting of groups acting on trees, not to the setting of Theorem 1.2 where Γ_G acts on the Cayley graph of G . To handle this problem we need a new idea: a generalization of Gersten’s converse to the Bestvina–Feighn combination theorem, which holds in a much broader setting. This generalization is contained in Lemma 5.2. The basis of this result is an analogy between the “flaring property” of Bestvina–Feighn and the divergence of geodesics in a word hyperbolic group [12].

Acknowledgements We are grateful to the referee for a thorough reading of the paper, and for making numerous useful comments.

Both authors are supported in part by the National Science Foundation.

2 Background

2.1 Coarse language

Quasi-isometries and uniformly proper maps Given a metric space X and two subsets $A, B \subset X$, the *Hausdorff distance* $d_{\text{Haus}}(A, B)$ is the infimum of all real numbers r such that each point of A is within distance r of a point of B , and vice versa.

A *quasi-isometric embedding* between two metric spaces X, Y is a map $f: X \rightarrow Y$ such that for some $K \geq 1$, $C \geq 0$, we have

$$\frac{1}{K}d(x, y) - C \leq d(fx, fy) \leq Kd(x, y) + C$$

for each $x, y \in X$. To refer to the constants we say that f is a K, C -quasi-isometric embedding.

For example, a quasigeodesic embedding $\mathbf{R} \rightarrow X$ is called a *quasigeodesic line* in X . We also speak of *quasigeodesic rays or segments* with the domain is a half-line or a finite segment, respectively. Since every map of a segment is a quasi-isometry, it usually behooves one to include the constants and speak about a (K, C) -quasi-isometric segment.

A *quasi-isometry* between two metric spaces X, Y is a map $f: X \rightarrow Y$ which, for some $K \geq 1$, $C \geq 0$ is a K, C quasi-isometry and has the property that $\text{image}(f)$ has Hausdorff distance $\leq C$ from Y . Every quasi-isometry $f: X \rightarrow Y$ has a *coarse inverse*, which is a quasi-isometry $\bar{f}: Y \rightarrow X$ such that $\bar{f} \circ f: X \rightarrow X$ is a bounded distance in the sup norm from Id_X , and similarly for $f \circ \bar{f}: Y \rightarrow Y$; the sup norm bounds and the quasi-isometry constants of \bar{f} depend only on the quasi-isometry constants of f .

More general than a quasi-isometric embedding is a *uniformly proper embedding* $f: X \rightarrow Y$, which means that there exists $K \geq 1$, $C \geq 0$, and a function $r: [0, \infty) \rightarrow [0, \infty)$ satisfying $r(t) \rightarrow \infty$ as $t \rightarrow \infty$, such that

$$r(d(x, y)) \leq d(fx, fy) \leq Kd(x, y) + C$$

for each $x, y \in X$.

Geodesic and quasigeodesic metric spaces A metric space is *proper* if closed balls are compact. A metric d on a space X is called a *path metric* if for any $x, y \in X$ the distance $d(x, y)$ is the infimum of the path lengths of rectifiable paths between x and y , and d is called a *geodesic metric* if $d(x, y)$

equals the length of some rectifiable path between x and y . The following fact is an immediate consequence of the Ascoli–Arzela theorem:

Fact 2.1 *A compact path metric space is a geodesic metric space. More generally, a proper path metric is a geodesic metric.* \square

The Ascoli–Arzela theorem also shows that for any proper geodesic metric space X , every path homotopy class contains a shortest path. This implies that the metric on X lifts to a geodesic metric on any covering space of X .

A metric space X is called a *quasigeodesic metric space* if there exists constants λ, ϵ such that for any $x, y \in X$ there exists an interval $[a, b] \subset \mathbf{R}$ and a λ, ϵ quasigeodesic embedding $\sigma: [a, b] \rightarrow X$ such that $\sigma(a) = x$ and $\sigma(b) = y$. For example, if Y is a geodesic metric space and X is a subset of Y such that $d_{\text{Haus}}(X, Y) < \infty$ then X is a quasigeodesic metric space.

The fundamental theorem of geometric group theory, first known to Efremovich, to Schwarz, and to Milnor, can be given a general formulation as follows. Let X be a proper, quasigeodesic metric space, and let the group G act on X properly discontinuously and cocompactly, by an action denoted $(g, x) \mapsto g \cdot x$. Then G is finitely generated, and for any base point $x_0 \in X$ the map $G \rightarrow X$ defined by $g \mapsto g \cdot x_0$ is a quasi-isometry between the word metric on G and the metric space X .

Uniform families of quasi-isometries The next lemma says a family of geodesic metrics which is “compact” in a suitable sense has the property that any two metrics in the family are uniformly quasi-isometric, with respect to the identity map.

Given a compact space X , let $M(X)$ denote the space of metrics generating the topology of X , regarded as a subspace of $[0, \infty)^{X \times X}$ with the topology of uniform convergence.

Lemma 2.2 *Let X be a compact, path connected space with universal cover \tilde{X} . Let $D \subset M(X)$ be a compact family of geodesic metrics. Let \tilde{D} be the set of lifted metrics on \tilde{X} . Then there exist $K \geq 1, C \geq 0$ such that for any $\tilde{d}, \tilde{d}' \in \tilde{D}$ the identity map on \tilde{X} is a K, C quasi-isometry between (\tilde{X}, \tilde{d}) and (\tilde{X}, \tilde{d}') .*

Proof By compactness of D , the metric spaces X_d have a uniform injectivity radius—that is, there exists $\epsilon > 0$ such that for each $d \in D$ every homotopically

nontrivial closed curve in X_d has length $> 4\epsilon$, and it follows that every closed ϵ ball in X_d lifts isometrically to \tilde{X}_d . Let $P \subset \tilde{X} \times \tilde{X}$ be the set of pairs $(x, y) \in \tilde{X} \times \tilde{X}$ such that for some $d \in \tilde{D}$ we have $d(x, y) \leq \epsilon$. Evidently $\pi_1(X)$ acts cocompactly on P , and so we have a finite supremum

$$A = \sup\{\tilde{d}(x, y) \mid \tilde{d} \in \tilde{D} \text{ and } (x, y) \in P\}$$

Given $\tilde{d} \in \tilde{D}$ and $x, y \in \tilde{X}$, choose a \tilde{d} -geodesic γ from x to y and let $x = x_0, x_1, \dots, x_{n-1}, x_n = y$ be a monotonic sequence along γ such that $d(x_{i-1}, x_i) = \epsilon$ for $i = 1, \dots, n-1$ and $d(x_{n-1}, x_n) \leq \epsilon$. For any $\tilde{d}' \in \tilde{D}$ it follows that:

$$\tilde{d}'(x, y) \leq An = A \left\lceil \frac{\tilde{d}(x, y)}{\epsilon} \right\rceil \leq \frac{A}{\epsilon} \tilde{d}(x, y) + A$$

Setting $K = \frac{A}{\epsilon}$ and $C = A$ the lemma follows. \square

Hyperbolic metric spaces A geodesic metric space X is *hyperbolic* if there exists $\delta \geq 0$ such that for any $x, y, z \in X$ and any geodesics \overline{xy} , \overline{yz} , \overline{zx} , any point on \overline{xy} has distance $\leq \delta$ from some point on $\overline{yz} \cup \overline{zx}$. A finitely generated group is *word hyperbolic* if the Cayley graph of some (any) finite generating set, equipped with the geodesic metric making each edge of length 1, is a hyperbolic metric space.

If X is δ -hyperbolic, then for any $\lambda \geq 1$, $\epsilon \geq 0$ there exists A , depending only on $\delta, \lambda, \epsilon$, such that the following hold: for any $x, y \in X$, any λ, ϵ quasigeodesic segment between x and y has Hausdorff distance $\leq A$ from any geodesic segment between x and y ; for any $x \in X$, any λ, ϵ quasigeodesic ray starting at x has Hausdorff distance $\leq A$ from some geodesic ray starting at x ; and any λ, ϵ quasigeodesic line in X has Hausdorff distance $\leq A$ from some geodesic line in X .

The *boundary* of X , denoted ∂X , is the set of coarse equivalence classes of geodesic rays in X , where two rays are coarsely equivalent if they have finite Hausdorff distance. For any $\xi \in \partial X$ and $x_0 \in X$, there is a ray based at x_0 representing ξ ; we denote such a ray $\overrightarrow{[x_0, \xi]}$. For any $\xi \neq \eta \in \partial X$ there is a geodesic line ℓ in X such that any point on ℓ divides it into two rays, one representing ξ and the other representing η .

Assuming X is proper, there is a compact topology on $X \cup \partial X$ in which X is dense, which is characterized by the following property: a sequence $\xi_i \in X \cup \partial X$ converges to $\xi \in \partial X$ if and only if, for any base point $p \in X$, if $\overrightarrow{[p, \xi_i]}$ denotes

either a segment from p to $\xi_i \in X$, or a ray from p with ideal endpoint $\xi_i \in \partial X$, then any subsequential limit of the sequence $\overrightarrow{[p, \xi_i]}$ is a ray with ideal endpoint ξ . It follows that any quasi-isometric embedding between δ -hyperbolic geodesic metric spaces extends to a continuous embedding of boundaries. In particular, if X is hyperbolic then the action of $\text{Isom}(X)$ on X extends continuously to an action on $X \cup \partial X$.

The following fundamental fact is easily proved by considering what happens to geodesics in a δ -hyperbolic metric space under a quasi-isometry.

Lemma 2.3 *For all $\delta \geq 0$, $K \geq 1$, $C \geq 0$ there exists $A \geq 0$ such that the following holds. If X, Y are two δ -hyperbolic metric spaces and if $f, g: X \rightarrow Y$ are two K, C quasi-isometries such that $\partial f = \partial g: \partial X \rightarrow \partial Y$, then:*

$$d_{\text{sup}}(f, g) = \sup_{x \in X} d(f(x), g(x)) \leq A \quad \square$$

2.2 Teichmüller space and the Thurston boundary

Fix once and for all a closed, oriented surface S of genus $g \geq 2$. Let \mathcal{C} be the set of isotopy classes of nontrivial simple closed curves on S .

The fundamental notation for the paper is as follows. Let \mathcal{T} be the Teichmüller space of S . Let \mathcal{MF} be the space of measured foliations on S , and let \mathbf{PMF} be the space of projective measured foliations on S , with projectivization map $\mathbf{P}: \mathcal{MF} \rightarrow \mathbf{PMF}$. The Thurston compactification of Teichmüller space is $\overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$. Let MCG be the mapping class group of S , and let $\mathcal{M} = \mathcal{T}/MCG$ be the moduli space of S . Definitions of these objects are all recalled below.

The Teichmüller space \mathcal{T} is the set of hyperbolic structures on S modulo isotopy, with the structure of a smooth manifold diffeomorphic to \mathbf{R}^{6g-6} given by Fenchel–Nielsen coordinates. The Riemann mapping theorem associates to each conformal structure on S a unique hyperbolic structure in that conformal class, and hence we may naturally identify \mathcal{T} with the set of conformal structures on S modulo isotopy. Given a conformal structure or a hyperbolic structure σ , we will often confuse σ with its isotopy class by writing $\sigma \in \mathcal{T}$.

There is a length pairing $\mathcal{T} \times \mathcal{C} \rightarrow \mathbf{R}_+$ which associates to each $\sigma \in \mathcal{T}$, $C \in \mathcal{C}$ the length of the unique simple closed geodesic on the hyperbolic surface σ in the isotopy class C . We obtain a map $\mathcal{T} \rightarrow [0, \infty)^{\mathcal{C}}$ which is an embedding with image homeomorphic to an open ball of dimension $6g - 6$. Moreover, under projectivization $[0, \infty)^{\mathcal{C}} \rightarrow \mathbf{P}[0, \infty)^{\mathcal{C}}$, \mathcal{T} embeds in $\mathbf{P}[0, \infty)^{\mathcal{C}}$ with precompact image.

Thurston's boundary A *measured foliation* \mathcal{F} on S is a foliation with finitely many singularities equipped with a positive transverse Borel measure, with the property that for each singularity s there exists $n \geq 3$ such that in a neighborhood of s the foliation \mathcal{F} is modelled on the horizontal measured foliation of the quadratic differential $z^{n-2}dz^2$ in the complex plane. A *saddle connection* of \mathcal{F} is a leaf segment connecting two distinct singularities; collapsing a saddle connection to a point yields another measured foliation on S . The set of measured foliations on S modulo the equivalence relation generated by isotopy and saddle collapse is denoted \mathcal{MF} . Given a measured foliation \mathcal{F} , its equivalence class is denoted $[\mathcal{F}] \in \mathcal{MF}$; elements of \mathcal{MF} will often be represented by the letters X, Y, Z .

For each measured foliation \mathcal{F} , there is a function $\ell_{\mathcal{F}}: \mathcal{C} \rightarrow [0, \infty)$ defined as follows. Given a simple closed curve c , we may pull back the transverse measure on \mathcal{F} to obtain a measure on c , and then integrate over c to obtain a number $\int_c \mathcal{F}$. Define $\ell_{\mathcal{F}}(c) = i(\mathcal{F}, c)$ to be the infimum of $\int_{c'} \mathcal{F}$ as c' ranges over the isotopy class of c . The function $\ell_{\mathcal{F}}$ is well-defined up to equivalence, thereby defining an embedding $\mathcal{MF} \rightarrow [0, \infty)^{\mathcal{C}}$ whose image is homeomorphic to $\mathbf{R}^{6g-6} - \{0\}$.

Given a measured foliation \mathcal{F} , multiplying the transverse measure by a positive scalar r defines a measured foliation denoted $r \cdot \mathcal{F}$, yielding a positive scalar multiplication operation $\mathbf{R} \times \mathcal{MF} \rightarrow \mathcal{MF}$. With respect to the equivalence relation $\mathcal{F} \sim r \cdot \mathcal{F}$, $r > 0$, the set of equivalence classes is denoted \mathbf{PMF} and the projection is denoted $\mathbf{P}: \mathcal{MF} \rightarrow \mathbf{PMF}$. We obtain an embedding $\mathbf{PMF} \rightarrow \mathbf{P}[0, \infty)^{\mathcal{C}}$ whose image is homeomorphic to a sphere of dimension $6g - 7$. We often use the letters ξ, η, ζ to represent points of \mathbf{PMF} .

Thurston's compactification theorem [16] says, by embedding into $\mathbf{P}[0, \infty)^{\mathcal{C}}$, that there is a homeomorphism of triples:

$$(\overline{\mathcal{T}}, \mathcal{T}, \mathbf{PMF}) \approx (B^{6g-6}, \text{int}(B^{6g-6}), S^{6g-7})$$

We will also need the standard embedding $\mathcal{C} \rightarrow \mathcal{MF}$, defined on $[c]$ as follows. Take an embedded annulus $A \subset S$ foliated by circles in the isotopy class $[c]$, and assign total transverse measure 1 to the annulus. Choose a deformation retraction of each component of the closure of $S - A$ onto a finite 1-complex, and extend to a map $f: S \rightarrow S$ homotopic to the identity and which is an embedding on $\text{int}(A)$. The measured foliation on A pushes forward under f to the desired measured foliation on S , giving a well-defined point in \mathcal{MF} depending only on $[c]$.

The intersection number $\mathcal{MF} \times \mathcal{C} \xrightarrow{i(\cdot, \cdot)} [0, \infty)$ extends continuously to $\mathcal{MF} \times$

$\mathcal{MF} \xrightarrow{i(\cdot, \cdot)} [0, \infty)$. This intersection number is most efficaciously defined in terms of measured geodesic laminations.

Marked surfaces Having fixed once and for all the surface S , a *marked surface* is a pair (F, ϕ) where F is a surface and $\phi: S \rightarrow F$ is a homeomorphism. Thus we may speak about a marked hyperbolic surface, a marked Riemann surface, a marked measured foliation on a surface, etc.

Given a marked hyperbolic surface $\phi: S \rightarrow F$, pulling back via ϕ determines a hyperbolic structure on S and a point of \mathfrak{t} . Two marked hyperbolic surfaces $\phi: S \rightarrow F$ and $\phi': S \rightarrow F'$ give the same element of \mathcal{T} if and only if they are equivalent in the following sense: there exists an isometry $h: F \rightarrow F'$ such that $\phi'^{-1} \circ h \circ \phi: S \rightarrow S$ is isotopic to the identity. In this manner, we can identify the collection of marked hyperbolic surfaces up to equivalence with the Teichmüller space \mathcal{T} of S . This allows us the freedom of representing a point of \mathcal{T} by a hyperbolic structure on some other surface F , assuming implicitly that we have a marking $\phi: S \rightarrow F$. The same discussion holds for marked Riemann surfaces, marked measured foliations on surfaces, etc.

Given two marked surfaces $\phi: S \rightarrow F$, $\phi': S \rightarrow F'$, a *marked map* is a homeomorphism $\psi: F \rightarrow F'$ such that $\psi \circ \phi$ is isotopic to ϕ' .

Mapping class groups and moduli space Let $\text{Homeo}(S)$ be the group of homeomorphisms of S , let $\text{Homeo}_0(S)$ be the normal subgroup consisting of homeomorphisms isotopic to the identity, and let $MCG = MCG(S) = \text{Homeo}(S)/\text{Homeo}_0(S)$ be the *mapping class group* of S . Pushing a hyperbolic structure on S forward via an element of $\text{Homeo}(S)$ gives a well-defined action of MCG on \mathcal{T} . This action is smooth and properly discontinuous but not cocompact. It follows that the *moduli space* $\mathcal{M} = \mathcal{T}/MCG$ is a smooth, noncompact orbifold with fundamental group MCG and universal covering space \mathcal{T} .

Let $\text{Homeo}(S, p)$ be the group of homeomorphisms of S preserving a base point p , let $\text{Homeo}_0(S, p)$ be the normal subgroup consisting of those homeomorphisms which are isotopic to the identity leaving p stationary, and let $MCG(S, p) = \text{Homeo}(S, p)/\text{Homeo}_0(S, p)$. Recall the short exact sequence:

$$1 \rightarrow \pi_1(S, p) \xrightarrow{\iota} MCG(S, p) \xrightarrow{q} MCG(S) \rightarrow 1$$

The map q is the map which “forgets” the puncture p . To define the map ι , for each closed loop $\ell: [0, 1] \rightarrow S$ based at p , choose numbers $0 = x_0 < x_1 < \dots <$

$x_n = 1$ and embedded open balls $B_1, \dots, B_n \subset S$ so that $\ell[x_{i-1}, x_i] \subset B_i$ for $i = 1, \dots, n$, and let $\pi_i: S \rightarrow S$ be a homeomorphism which is the identity on $S - B_i$ and such that $\pi_i(\ell(x_{i-1})) = \ell(x_i)$. Then $\iota(\ell)$ is defined to be the isotopy class rel p of the homeomorphism $\pi_n \circ \pi_{n-1} \circ \dots \circ \pi_1: (S, p) \rightarrow (S, p)$, which we say is obtained by “pushing” the point p around the loop ℓ . The mapping class $\iota(\ell)$ is well-defined independent of the choices made, and independent of the choice of ℓ in its path homotopy class. When ℓ is simple, $\iota(\ell)$ may also be described as the composition of opposite Dehn twists on the two boundary components of a regular neighborhood of ℓ . For details see [7].

As noted in the introduction, by the Dehn–Nielsen–Baer–Epstein theorem, the above sequence is naturally isomorphic to the sequence

$$1 \rightarrow \pi_1(S, p) \rightarrow \text{Aut}(\pi_1(S, p)) \rightarrow \text{Out}(\pi_1(S, p)) \rightarrow 1$$

Canonical bundles Over the Teichmüller space \mathcal{T} of S there is a *canonical marked hyperbolic surface bundle* $\mathcal{S} \rightarrow \mathcal{T}$, defined as follows. Topologically $\mathcal{S} = S \times \mathcal{T}$, with the obvious marking $S \xrightarrow{\cong} S \times \sigma = \mathcal{S}_\sigma$ for each $\sigma \in \mathcal{T}$. As σ varies over \mathcal{T} , one can assign a hyperbolic structure on S in the class of σ , varying continuously in the C^∞ topology on Riemannian metrics; this follows from the description of Fenchel–Nielsen coordinates. It follows that on each fiber \mathcal{S}_σ of \mathcal{S} there is a hyperbolic structure which varies continuously in σ . Note that by the Riemann mapping theorem we can also think of \mathcal{S} as the canonical marked Riemann surface bundle over \mathcal{T} .

The action of MCG on \mathcal{T} lifts uniquely to an action on \mathcal{S} , such that for each fiber \mathcal{S}_σ and each $[h] \in MCG$ the map

$$\mathcal{S}_\sigma \xrightarrow{[h]} \mathcal{S}_{[h](\sigma)}$$

is an isometry, and the map

$$S \xrightarrow{\cong} \mathcal{S}_\sigma \xrightarrow{[h]} \mathcal{S}_{[h](\sigma)} \xrightarrow{\cong} S$$

is in the mapping class $[h]$.

The universal cover of \mathcal{S} is called the canonical \mathbf{H}^2 –bundle over \mathcal{T} , denoted $\mathcal{H} \rightarrow \mathcal{T}$. There is a fibration preserving, isometric action of the once-punctured mapping class group $MCG(S, p)$ on the total space \mathcal{H} such that the quotient action of $MCG(S, p)$ on \mathcal{S} has kernel $\pi_1(S, p)$, and corresponds to the given action of $MCG(S) = MCG(S, p)/\pi_1(S, p)$ on \mathcal{S} . Also, the action of $\pi_1(S, p)$ on any fiber of \mathcal{H} is conjugate to the action on the universal cover \tilde{S} by deck transformations. Bers proved in [4] that \mathcal{H} is a Teichmüller space in its own

right: there is an $MCG(S, p)$ equivariant homeomorphism between \mathcal{H} and the Teichmüller space of the once-punctured surface $S - p$.

The tangent bundle $T\mathcal{S}$ has a smooth 2-dimensional *vertical sub-bundle* $T_v\mathcal{S}$ consisting of the tangent planes to fibers of the fibration $\mathcal{S} \rightarrow \mathcal{T}$. A *connection* on the bundle $\mathcal{S} \rightarrow \mathcal{T}$ is a smooth codimension-2 sub-bundle of $T\mathcal{S}$ complementary to $T_v\mathcal{S}$. The existence of an MCG -equivariant connection on \mathcal{S} can be derived following standard methods, as follows. Choose a locally finite, equivariant open cover of \mathcal{T} , and an equivariant partition of unity dominated by this cover. For each MCG -orbit of this cover, choose a representative $U \subset \mathcal{T}$ of this orbit, and choose a linear retraction $T\mathcal{S}_U \rightarrow T_v\mathcal{S}_U$. Pushing these retractions around by the action of MCG and taking a linear combination using the partition of unity, we obtain an equivariant linear retraction $T\mathcal{S} \rightarrow T_v\mathcal{S}$, whose kernel is the desired connection.

By lifting to \mathcal{H} we obtain a connection on the bundle $\mathcal{H} \rightarrow \mathcal{T}$, equivariant with respect to the action of the group $MCG(S, p)$.

Notation Given any subset $A \subset \mathcal{T}$, or more generally any continuous map $A \rightarrow \mathcal{T}$, by pulling back the bundle $\mathcal{S} \rightarrow \mathcal{T}$ we obtain a bundle $\mathcal{S}_A \rightarrow A$, as shown in the following diagram:

$$\begin{array}{ccc} \mathcal{S}_A & \xrightarrow{\dots\dots\dots} & \mathcal{S} \\ \downarrow \dots\dots\dots & & \downarrow \\ A & \xrightarrow{\quad\quad\quad} & \mathcal{T} \end{array}$$

Similarly, the pullback of the bundle $\mathcal{H} \rightarrow \mathcal{T}$ is denoted $\mathcal{H}_A \rightarrow A$.

Quadratic differentials Given a conformal structure σ on S , a *quadratic differential* q on \mathcal{S}_σ assigns to each conformal coordinate z an expression of the form $q(z)dz^2$ where $q(z)$ is a complex valued function on the domain of the coordinate system, and

$$q(z) \left(\frac{dz}{dw} \right)^2 = q(w), \quad \text{for overlapping coordinates } z, w.$$

We shall always assume that the functions $q(z)$ are holomorphic, in other words, our quadratic differentials will always be “holomorphic” quadratic differentials. A quadratic differential q is *trivial* if $q(z)$ is always the zero function.

Given a nontrivial quadratic differential q on \mathcal{S}_σ , a point $p \in \mathcal{S}_\sigma$ is a zero of q in one coordinate if and only if it is a zero in any coordinate; also, the order of

the zero is well-defined. If p is not a zero then there is a coordinate z near p , unique up to multiplication by ± 1 , such that p corresponds to the origin and such that $q(z) \equiv 1$; this is called a *regular canonical coordinate*. If p is a zero of order $n \geq 1$ then up to multiplication by the $(n+2)^{\text{nd}}$ roots of unity there exists a unique coordinate z in which p corresponds to the origin and such that $q(z) = z^n$; this is called a *singular canonical coordinate*. There is a well-defined *singular Euclidean metric* $|q(z)||dz|^2$ on S , which in any regular canonical coordinate $z = x + iy$ takes the form $dx^2 + dy^2$. In any singular canonical coordinate this metric has finite area, and so the total area of S in this singular Euclidean metric is finite, denoted $\|q\|$. We say that q is *normalized* if $\|q\| = 1$.

By the Riemann–Roch theorem, the quadratic differentials on \mathcal{S}_σ form a complex vector space QD_σ of complex dimension $3g - 3$, and these vector spaces fit together, one for each $\sigma \in \mathcal{T}$, to form a complex vector bundle over \mathcal{T} denoted $\text{QD} \rightarrow \mathcal{T}$. Teichmüller space has a complex structure whose cotangent bundle is canonically isomorphic to the bundle QD . The Teichmüller metric on \mathcal{T} induces a Finsler metric on the (real) tangent bundle of \mathcal{T} , and the norm $\|q\|$ is dual to this metric. The normalized quadratic differentials form a sphere bundle $\text{QD}^1 \rightarrow \mathcal{T}$ of real dimension $6g - 7$ embedded in QD .

Corresponding to each quadratic differential q on \mathcal{S}_σ there is a pair of measured foliations, the *horizontal foliation* $\mathcal{F}_x(q)$ and the *vertical foliation* $\mathcal{F}_y(q)$. In a regular canonical coordinate $z = x + iy$, the leaves of $\mathcal{F}_x(q)$ are parallel to the x -axis and have transverse measure $|dy|$, and the leaves of $\mathcal{F}_y(q)$ are parallel to the y -axis and have transverse measure $|dx|$. The foliations $\mathcal{F}_x(q)$, $\mathcal{F}_y(q)$ have the zero set of q as their common singularity set, and at each zero of order n both have an $(n+2)$ -pronged singularity, locally modelled on the singularity at the origin of the horizontal and vertical measured foliations of $z^n dz^2$.

Conversely, consider a *transverse pair of measured foliations* $(\mathcal{F}_x, \mathcal{F}_y)$ on S which means that $\mathcal{F}_x, \mathcal{F}_y$ have the same singular set, are transverse at all regular points, and at each singularity s there is a number $n \geq 3$ such that \mathcal{F}_x and \mathcal{F}_y are locally modelled on the horizontal and vertical measured foliations of $z^{n-2} dz^2$. Associated to the pair $\mathcal{F}_x, \mathcal{F}_y$ there are a conformal structure and a quadratic differential defined as follows. Near each regular point, there is an oriented coordinate $z = x + iy$ in which \mathcal{F}_x is the horizontal foliation with transverse measure $|dy|$, and \mathcal{F}_y is the vertical foliation with transverse measure $|dx|$. These regular coordinates have conformal overlap. Near any singularity s , at which $\mathcal{F}_x, \mathcal{F}_y$ are locally modelled on the the horizontal and vertical foliations of $z^n dz^2$, the coordinate z has conformal overlap with any regular coordinate. We therefore obtain a conformal structure $\sigma(\mathcal{F}_x, \mathcal{F}_y)$ on S , on

which we have a quadratic differential $q(\mathcal{F}_x, \mathcal{F}_y)$ defined in regular coordinates by dz^2 .

A pair of measured foliations $(X, Y) \in \mathcal{MF}(F) \times \mathcal{MF}(F)$ is said to *jointly fill* the surface F if, for every $Z \in \mathcal{MF}(F)$, either $i(X, Z) \neq 0$ or $i(Y, Z) \neq 0$. This condition is invariant under positive scalar multiplication on $\mathcal{MF}(F)$, and so joint filling is well-defined for a pair of points in $\mathbf{PMF}(F)$. A basic fact is that a pair $X, Y \in \mathcal{MF}(F)$ jointly fills F if and only if there exist a transverse pair of measured foliations $\mathcal{F}_x, \mathcal{F}_y$ representing X, Y ; moreover, such a pair $\mathcal{F}_x, \mathcal{F}_y$ is unique up to *joint isotopy*, meaning that for any other transverse pair $\mathcal{F}'_x, \mathcal{F}'_y$ representing X, Y respectively, there exists $h \in \text{Homeo}_0(S)$ such that $\mathcal{F}'_x = h(\mathcal{F}_x)$, $\mathcal{F}'_y = h(\mathcal{F}_y)$. These facts may be proved by passing back and forth between measured geodesic laminations and measured foliations.

By uniqueness up to joint isotopy as just described, it follows that for each jointly filling pair $(X, Y) \in \mathcal{MF}(F) \times \mathcal{MF}(F)$ there is a conformal structure $\sigma(\mathcal{F}_x, \mathcal{F}_y)$ and quadratic differential $q(\mathcal{F}_x, \mathcal{F}_y)$ on $\sigma(X, Y)$, well-defined up to isotopy independent of the choice of a transverse pair $\mathcal{F}_x, \mathcal{F}_y$ representing X, Y . We thus have a well-defined point $\sigma(X, Y) \in \mathcal{T}(F)$ and a well-defined element $q(X, Y) \in \text{QD}_{\sigma(X, Y)} \mathcal{T}(F)$.

Geodesics and a metric on \mathcal{T} We shall describe geodesic lines in \mathcal{T} following [17] and [21]; of course everything depends on Teichmüller's theorem (see eg, [1] or [22]).

Let $\mathcal{FP} \subset \mathcal{MF} \times \mathcal{MF}$ denote the set of jointly filling pairs, and let \mathbf{PFP} be the image of \mathcal{FP} under the product of projection maps $\mathbf{P} \times \mathbf{P}: \mathcal{MF} \times \mathcal{MF} \rightarrow \mathbf{PMF} \times \mathbf{PMF}$.

Associated to each jointly filling pair $(\xi, \eta) \in \mathbf{PFP}$ we associate a *Teichmüller line* $\overleftrightarrow{(\xi, \eta)}$, following [17]. Choosing a transverse pair of measured foliations $\mathcal{F}_x, \mathcal{F}_y$ representing ξ, η respectively, we obtain a *parameterized Teichmüller geodesic* given by the map $t \mapsto \sigma(e^{-t}\mathcal{F}_x, e^t\mathcal{F}_y)$; it follows from Teichmüller's theorem that this map is an embedding $\mathbf{R} \rightarrow \mathcal{T}$. Uniqueness of $\mathcal{F}_x, \mathcal{F}_y$ up to joint isotopy and positive scalar multiplication imply that the map $t \mapsto \sigma(e^{-t}\mathcal{F}_x, e^t\mathcal{F}_y)$ is well-defined up to translation of the t -parameter, as is easily checked. Thus, the image of this map is well defined and is denoted $\overleftrightarrow{(\xi, \eta)}$; in addition, parameter difference between points on the line is well-defined, and there is a well-defined orientation. The *positive direction* of the geodesic is defined to be the point $\eta = \mathbf{P}\mathcal{F}_y \in \mathbf{PMF}$, the projective class of the vertical measured foliation; the negative direction is the point $\xi = \mathbf{P}\mathcal{F}_x \in \mathbf{PMF}$. Note that as

$t \rightarrow +\infty$ the vertical measured foliation becomes “exponentially thicker” and so dominates over the horizontal foliation which becomes “exponentially thinner”, a useful mnemonic for remembering which direction is which.

Teichmüller’s theorem says that any two distinct points of \mathcal{T} lie on a unique Teichmüller line: for any $\sigma \neq \tau \in \mathcal{T}$ there exists a unique pair $(\xi, \eta) \in \mathbf{P}\mathcal{F}\mathcal{P}$ such that $\sigma, \tau \in \overrightarrow{(\xi, \eta)}$. Moreover, if $d(\sigma, \tau)$ is the parameter difference between σ and τ along this geodesic, then d is a metric on \mathcal{T} , called the *Teichmüller metric*. In particular, each line $\overrightarrow{(\xi, \eta)}$ is, indeed, a geodesic for the Teichmüller metric. It is also true that the segment $[\sigma, \tau] \subset \overrightarrow{(\xi, \eta)}$ is the unique geodesic segment connecting σ to τ , and hence geodesic segments are uniquely extensible. Thus we obtain a 1–1 correspondence between oriented geodesic segments and the set $\mathcal{T} \times \mathcal{T}$. Also, every bi-infinite geodesic line in \mathcal{T} is uniquely expressible in the form $\overrightarrow{(\xi, \eta)}$, and so we obtain a 1–1 correspondence between oriented geodesic lines and the set $\mathbf{P}\mathcal{F}\mathcal{P} \subset \mathbf{P}\mathcal{M}\mathcal{F} \times \mathbf{P}\mathcal{M}\mathcal{F}$.

There is also a 1–1 correspondence between geodesic rays in \mathcal{T} and the set $\mathcal{T} \times \mathbf{P}\mathcal{M}\mathcal{F}$: for any $\sigma \in \mathcal{T}$ and $\eta \in \mathbf{P}\mathcal{M}\mathcal{F}$ there is a unique geodesic ray, denoted $\overrightarrow{[\sigma, \eta]}$, whose endpoint is σ and whose direction is $\eta \in \mathbf{P}\mathcal{M}\mathcal{F}$, and every geodesic ray has this form. This is an immediate consequence of the Hubbard–Masur theorem [21], which says that for each $\sigma \in \mathcal{T}$ the map $\text{QD}_\sigma \rightarrow \mathcal{M}\mathcal{F}$ taking $q \neq 0 \in \text{QD}_\sigma$ to $[\mathcal{F}_y(q)]$ is a homeomorphism.

Throughout the paper, the term “geodesic” will refer to any geodesic segment, ray, or line in \mathcal{T} . Geodesics in \mathcal{T} are *uniquely extendable*: any geodesic segment or ray is contained in a unique geodesic line. Since \mathcal{T} is a complete metric space, an argument using the Ascoli–Arzela theorem shows that any sequence of geodesics, each element of which intersects a given bounded subset of \mathcal{T} , has a subsequence converging pointwise to a geodesic.

By unique extendability of geodesics it follows that \mathcal{T} is a proper, geodesic metric space. From the definitions it follows that the action of MCG on \mathcal{T} is isometric, and so the metric on \mathcal{T} descends to a proper, geodesic metric on $\mathcal{M} = \mathcal{T}/MCG$.

The reader is cautioned that a geodesic ray $\overrightarrow{[\sigma, \eta]}$ is *not known* to converge in $\overline{\mathcal{T}}$ to its direction $\eta \in \mathbf{P}\mathcal{M}\mathcal{F}$. However, consider the case where η is *uniquely ergodic*, which means that for any measured foliation \mathcal{F} representing η , every transverse measure on the underlying singular foliation of \mathcal{F} is a scalar multiple of the given measure on \mathcal{F} . In this case the ray $\overrightarrow{[\sigma, \eta]}$ does converge to η , as is proved by Masur [30], and so in this situation the direction η is also called the *end* or *endpoint* of the ray.

Cobounded geodesics in \mathcal{T} A subset $A \subset \mathcal{T}$ is *cobounded* if the image of A under the projection $\mathcal{T} \rightarrow \mathcal{M}$ is a bounded subset of \mathcal{M} ; equivalently, there is a bounded subset of \mathcal{T} whose translates by the action of MCG cover A . If the bounded set $\mathcal{B} \subset \mathcal{M}$ contains the projected image of A then we also say that A is \mathcal{B} -cobounded. Since \mathcal{M} is a proper metric space it follows that A is cobounded in \mathcal{T} if and only if A is “co-precompact”, meaning that the projection of A to \mathcal{M} has compact closure.

One common gauge for coboundedness, as noted by Mumford [42], is the injectivity radius of a hyperbolic structure, or to put it another way, the length $\ell(\sigma)$ of the shortest closed geodesic in a hyperbolic structure σ .³ For each $\epsilon > 0$ the “ ϵ -thick subset” of \mathcal{T} , namely the set $\mathcal{T}_\epsilon = \{\sigma \in \mathcal{T} \mid \ell(\sigma) \geq \epsilon\}$, is an MCG equivariant subset of \mathcal{T} projecting to a compact subset of \mathcal{M} , and as $\epsilon \rightarrow 0$ this gives an exhaustion of \mathcal{M} by compact sets. A subset of \mathcal{T} is therefore cobounded if and only if it is contained in the ϵ -thick subset of \mathcal{T} for some $\epsilon > 0$.

Extremal length, rather than hyperbolic length, is used to obtain another common gauge of coboundedness, and is comparable to the length of the shortest geodesic by Maskit’s work [27].

We rarely use any particular gauge for coboundedness. Instead, the primary way in which we use coboundedness is in carrying out compactness arguments over closed, bounded subsets. For this reason we rarely refer to any gauge, instead sticking with coboundedness as the more primitive mathematical concept.

One important fact we need is that if $\rho = \overrightarrow{[\sigma, \eta]}$ is a cobounded geodesic ray in Teichmüller space then ρ converges to η in Thurston’s compactification $\overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$. This follows from two theorems of Masur. First, since ρ is cobounded, the direction $\eta \in \mathbf{PMF}$ is uniquely ergodic; this result, proved in [29], was later sharpened in [31] to show that if η is not uniquely ergodic then the projection of $\overrightarrow{[\sigma, \eta]}$ to moduli space leaves every compact subset. Second, when η is uniquely ergodic, any ray with direction η converges to η in Thurston’s compactification. This is a small part of a Masur’s Two Boundaries Theorem [30], concerning relations between the Teichmüller boundary and the Thurston boundary of \mathcal{T} (we will use the full power of this theorem in the proof of Theorem 1.1).

The following result is essentially a consequence of [38]:

³Also called the “systole” in the differential geometry literature.

Lemma 2.4 (End Uniqueness) *If $\overrightarrow{[\sigma, \xi]}$, $\overrightarrow{[\tau, \eta]}$ are two cobounded rays in \mathcal{T} which have finite Hausdorff distance in \mathcal{T} then $\xi = \eta$. If $\overleftarrow{(\xi, \xi')}$, $\overleftarrow{(\eta, \eta')}$ are two cobounded lines in \mathcal{T} which have finite Hausdorff distance then, up to relabelling the ends of one of the lines, we have $\xi = \eta$ and $\xi' = \eta'$, and so $\overleftarrow{(\xi, \xi')} = \overleftarrow{(\eta, \eta')}$.*

Proof For the proof we review briefly notions of extremal length, in the classical setting of simple closed curves, as well as Kerckhoff’s extension to the setting of measured foliations [26].

Recall that for any conformal structure on an open annulus A there is a unique Euclidean annulus of the form $S^1 \times (0, M)$ conformally equivalent to A , with $M \in \mathbf{R}_+ \cup \{\infty\}$; the modulus of A , denoted $M(A)$, is defined to be the number M . For any Riemann surface \mathcal{S}_σ and any isotopy class of simple closed curves $[c] \in \mathcal{C}$, the *extremal length* $\ell_{\text{ext}}(\sigma, [c])$ is the infimum of $1/M(A)$ taken over all annuli $A \subset F$ whose core is in the isotopy class $[c]$.

Kerckhoff proved [26] that the function $\ell_{\text{ext}} : \mathcal{T} \times (\mathbf{R}_+ \cdot \mathcal{C}) \rightarrow (0, \infty)$ defined by $\ell_{\text{ext}}(\sigma, r[c]) \mapsto r\ell_{\text{ext}}(\sigma, [c])$ extends continuously to a function $\ell_{\text{ext}} : \mathcal{T} \times \mathcal{MF} \rightarrow [0, \infty)$. Moreover, for any transverse pair of measured foliations $\mathcal{F}_x, \mathcal{F}_y$ with associated conformal structure $\sigma = \sigma(\mathcal{F}_x, \mathcal{F}_y)$ and quadratic differential $q = q(\mathcal{F}_x, \mathcal{F}_y)$, we have

$$\ell_{\text{ext}}(\sigma, \mathcal{F}_y) = \sqrt{\|q\|}$$

Given $X \in \mathcal{MF}$, the *extremal length horoball* based at X is defined to be $H(X) = \{\sigma \in \mathcal{T} \mid \ell_{\text{ext}}(\sigma, X) \leq 1\}$. Note for example that, setting $\xi = \mathbf{P}X$, for every $\eta \in \mathbf{P}\mathcal{MF}$ the extremal length of X at points of $\overleftarrow{(\eta, \xi)}$ decreases strictly monotonically to zero as the point moves towards ξ , and so every Teichmüller geodesic with positive direction $\mathbf{P}X$ eventually enters $H(X)$ in the positive direction and, once in, never leaves. Given $\xi \in \mathbf{P}\mathcal{MF}$, there is a one parameter family of extremal length horoballs based at ξ , namely $H(X)$ for all $X \in \mathcal{MF}$ such that $\mathbf{P}X = \xi$.

For the first sentence of the theorem, consider two geodesic rays $\overrightarrow{[\sigma, \xi]}$, $\overrightarrow{[\tau, \eta]}$ such that $\xi \neq \eta \in \mathbf{P}\mathcal{MF}$. Pick any extremal length horoball H based at η . The proof of Theorem 4.3 of [38] shows that $H \cap \overrightarrow{[\sigma, \xi]}$ is bounded. However, $H \cap \overrightarrow{[\tau, \eta]}$ is an infinite subray of $\overrightarrow{[\tau, \eta]}$, and moreover as a point $p \in \overrightarrow{[\tau, \eta]}$ travels to infinity in $\overrightarrow{[\tau, \eta]}$ the horoball H contains a larger and larger ball in \mathcal{T} centered on p . It follows that $\overrightarrow{[\sigma, \xi]}$ and $\overrightarrow{[\tau, \eta]}$ have infinite Hausdorff distance in \mathcal{T} .

The second sentence follows from the first, by dividing each line into two rays. \square

Remark Combining results of Masur mentioned above, one can show that even more is true: two cobounded geodesic rays which have finite Hausdorff distance are asymptotic, meaning that as they go to ∞ , the distance between the rays approaches zero. To see why, as mentioned earlier Masur proves that if $\overrightarrow{[\sigma, \eta]}$ is cobounded then η is uniquely ergodic. Furthermore, two rays $\overrightarrow{[\sigma, \eta]}$, $\overrightarrow{[\tau, \eta]}$ with uniquely ergodic endpoint η are asymptotic, according to [28].

2.3 Singular SOLV spaces

Consider a geodesic $g = \overleftarrow{(\xi, \eta)}$ in \mathcal{T} , and let $\mathcal{S}_g \rightarrow g$ be the canonical marked Riemann surface bundle over g , obtained by pulling back the canonical marked Riemann surface bundle $\mathcal{S} \rightarrow \mathcal{T}$. Topologically we identify $\mathcal{S}_g = S \times g$. Choosing a transverse pair of measured foliations $\mathcal{F}_x, \mathcal{F}_y$ representing ξ, η respectively, we have $g(t) = \sigma(e^{-t}\mathcal{F}_x, e^t\mathcal{F}_y)$. Let $|dy|$ be the transverse measure on the horizontal measured foliation \mathcal{F}_x and let $|dx|$ be the transverse measure on the vertical measured foliation \mathcal{F}_y . We may assume that the pair $\mathcal{F}_x, \mathcal{F}_y$ is normalized, meaning that the Euclidean area equals 1:

$$\|q(\mathcal{F}_x, \mathcal{F}_y)\| = \int_S |dx| \times |dy| = 1$$

and hence for all $t \in \mathbf{R}$ the pair $e^{-t}\mathcal{F}_x, e^t\mathcal{F}_y$ is normalized:

$$\|q(e^{-t}\mathcal{F}_x, e^t\mathcal{F}_y)\| = \int_S |e^t dx| \times |e^{-t} dy| = 1$$

Note that the singular Euclidean metric on each fiber $\mathcal{S}_{g(t)}$, may be expressed as

$$ds_\sigma^2 = e^{2t} |dx|^2 + e^{-2t} |dy|^2$$

Define the *singular SOLV metric* on \mathcal{S}_g to be the singular Riemannian metric given by the formula:

$$ds_g^2 = e^{2t} |dx|^2 + e^{-2t} |dy|^2 + dt^2$$

We use the notation $\mathcal{S}_g^{\text{SOLV}}$ to denote \mathcal{S}_g equipped with this metric. The universal cover of \mathcal{S}_g is the canonical Poincaré disc bundle \mathcal{H}_g over g , and lifting the singular SOLV metric from $\mathcal{S}_g^{\text{SOLV}}$ to \mathcal{H}_g we obtain a singular SOLV space denoted $\mathcal{H}_g^{\text{SOLV}}$. The singular locus of $\mathcal{S}_g^{\text{SOLV}} = S \times g$ is the union of the singular lines

$s \times g$, one for each singularity s of the pair $\mathcal{F}_x, \mathcal{F}_y$. Away from the singular lines, $\mathcal{S}_g^{\text{SOLV}}$ and $\mathcal{H}_g^{\text{SOLV}}$ are locally modelled on 3-dimensional SOLV-geometry. On each singular line the metric is locally modelled by gluing together several copies of the half-plane $y \geq 0$ in SOLV-geometry.

2.4 Comparing hyperbolic and singular Euclidean structures

Given $\sigma \in \mathcal{T}$, the Riemann surface \mathcal{S}_σ has several important metrics in its conformal class: a unique hyperbolic metric; and one singular Euclidean metric for each $q \in \text{QD}_\sigma$. These lift to the universal cover \mathcal{H}_σ . Given $\sigma, \tau \in \mathcal{T}$, if each Riemann surface $\mathcal{S}_\sigma, \mathcal{S}_\tau$ is given either its unique hyperbolic metric or one of its singular Euclidean metrics, then for any marked map $\phi: \mathcal{S}_\sigma \rightarrow \mathcal{S}_\tau$, each lift $\tilde{\phi}: \mathcal{H}_\sigma \rightarrow \mathcal{H}_\tau$ is a quasi-isometry. We are interested in how the quasi-isometry constants of $\tilde{\phi}$ compare to the Teichmüller distance $d(\sigma, \tau)$, although we need only the crudest estimates. Proposition 2.5 shows how to bound the quasi-isometry constants in terms of $d(\sigma, \tau)$. Part 1 of this proposition was first proved by Minsky in [37], Lemma 3.3; we give a quicker proof using Lemma 2.2.

Proposition 2.5 *For each bounded subset $\mathcal{B} \subset \mathcal{M}$ and each $r > 0$ there exists $K \geq 1, C \geq 0, A \geq 0$ such that the following hold:*

- (1) *Suppose that $\sigma, \tau \in \mathcal{T}$ are each \mathcal{B} -cobounded and $d(\sigma, \tau) \leq r$. Let $f_{\sigma\tau}: \mathcal{S}_\sigma \rightarrow \mathcal{S}_\tau$ be the canonical marked map $\mathcal{S}_\sigma = S \times \sigma \rightarrow S \times \tau = \mathcal{S}_\tau$. If we impose on \mathcal{S}_σ and \mathcal{S}_τ either the hyperbolic metric or the singular Euclidean metric associated to some normalized quadratic differential, then any lift $\tilde{f}_{\sigma\tau}: \mathcal{H}_\sigma \rightarrow \mathcal{H}_\tau$ of $f_{\sigma\tau}$ is a K, C quasi-isometry.*
- (2) *Let $\sigma_i \in \mathcal{T}$, $i = 1, 2, 3$, be \mathcal{B} -cobounded and have pairwise distances $\leq r$, let metrics be imposed on \mathcal{S}_{σ_i} as above, and let $f_{ij}: \mathcal{S}_{\sigma_i} \rightarrow \mathcal{S}_{\sigma_j}$, etc. be the marked maps as above, with K, C -quasi-isometric lifts $\tilde{f}_{ij}: \mathcal{H}_{\sigma_i} \rightarrow \mathcal{H}_{\sigma_j}$. If \tilde{f}_{13} is the unique lift of f_{13} such that*

$$\partial \tilde{f}_{23} \circ \partial \tilde{f}_{12} = \partial \tilde{f}_{13},$$

then

$$d_{\text{sup}}(\tilde{f}_{23} \circ \tilde{f}_{12}, \tilde{f}_{13}) \leq A.$$

Proof Part (1) is an easy consequence of Lemma 2.2, as follows. Choose a compact subset $\mathcal{A} \subset \mathcal{T}$ whose image in \mathcal{M} covers \mathcal{B} and such that over any point of \mathcal{B} there exists a point $\sigma \in \mathcal{A}$ such that $\mathcal{B}_\mathcal{T}(\sigma, r) \subset \mathcal{A}$. It follows that the points σ, τ in (1) may be translated to lie in \mathcal{A} . Identifying $\mathcal{S}_\mathcal{A}$ diffeomorphically with

$S \times \mathcal{A}$, compactness of \mathcal{A} produces a compact family of hyperbolic metrics on S , and compactness of the restriction of QD^1 to \mathcal{A} produces a compact family of singular Euclidean metrics. Now apply Lemma 2.2.

For part (2), note that by compactness of \mathcal{A} and of the compactness of the restriction of QD^1 to \mathcal{A} , there exists a uniform δ such that any hyperbolic metric and any normalized singular Euclidean structure determined by an element $\sigma \in \mathcal{A}$ has a δ -hyperbolic universal cover. Part (2) is now a direct consequence of Lemma 2.3. \square

3 Convex cocompact subgroups of $\text{Isom}(\mathcal{T})$

3.1 Variations of convex cocompactness

Given a proper, geodesic metric space X , a subset $L \subset X$ is *quasiconvex* if there exists $A \geq 0$ such that every geodesic segment in X with endpoints in L is contained in the A -neighborhood of L .

When G is a finitely generated, discrete subgroup of the isometry group of \mathbf{H}^n , it is well known that the following properties of G are all equivalent to each other:

Orbit Quasiconvexity Any orbit of G is a quasiconvex subset of \mathbf{H}^n .

Single orbit quasiconvexity There exists an orbit of G which is quasiconvex in \mathbf{H}^n .

Convex cocompact G acts cocompactly on the convex hull of its limit set Λ .

Moreover, these properties imply that G is word hyperbolic, and there is a continuous G -equivariant embedding of the Gromov boundary ∂G into $\partial \mathbf{H}^n$ whose image is the limit set Λ . Similar facts hold for finitely generated groups acting discretely on any Gromov hyperbolic space, for example finitely generated subgroups of Gromov hyperbolic groups.

In this section we prove Theorem 1.1, which is a list of similar equivalences for finitely generated subgroups of the isometry group of the Teichmüller space \mathcal{T} of S . In this case the entire isometry group $\text{Isom}(\mathcal{T})$ acts discretely on \mathcal{T} , and in fact by Royden's Theorem [45], [24] the canonical homomorphism $MCG \rightarrow \text{Isom}(\mathcal{T})$ is an isomorphism, except in genus 2 where the kernel is cyclic of order 2.

Although \mathcal{T} fails to be negatively curved in any reasonable sense, nevertheless one can say that it behaves in a negatively curved manner as long as one focusses only on cobounded aspects. This, at least, is one way to interpret the projection properties introduced by Minsky in [38] and further developed by Masur and Minsky in [32]. Given a \mathcal{B} -cobounded geodesic g in \mathcal{T} , Minsky's projection property says that a closest point projection map of \mathcal{T} onto g behaves in a negatively curved manner, such that the quality of the negative curvature depends only on \mathcal{B} . See Theorem 3.6 for the precise statement.

For a finitely generated subgroup $G \subset \text{Isom}(\mathcal{T})$ we can obtain equivalences as above, as long as we tack on an appropriate uniform coboundedness property; in some cases the desired property comes for free by uniform coboundedness of the action of G on any of its orbits.

First we have some properties of G which are variations on orbit quasiconvexity:

Orbit quasiconvexity Any orbit of G is quasiconvex in \mathcal{T} .

Single orbit quasiconvexity There exists an orbit of G that is quasiconvex in \mathcal{T} .

Weak orbit quasiconvexity There exists a constant A and an orbit \mathcal{O} of G , and for each $x, y \in \mathcal{O}$ there exists a geodesic segment $[x', y']$ in \mathcal{T} , such that $d(x, x') \leq A$, $d(y, y') \leq A$, and $[x', y']$ is in the A -neighborhood of \mathcal{O} .

The latter is a more technical version of orbit quasiconvexity which is quite useful in several settings.

Another property of G is a version of convex cocompactness, into which we incorporate the hyperbolicity properties mentioned above:

Convex cocompact The group G is word hyperbolic, and there exists a continuous G -equivariant embedding $f_\infty: \partial G \rightarrow \mathbf{PMF}$ with image Λ_G , such that $\Lambda_G \times \Lambda_G - \Delta \subset \mathbf{PF}\mathcal{P}$, and the following holds. Letting

$$\text{WH}_G = \cup \{ \overleftrightarrow{(\zeta, \zeta')} \mid \zeta \neq \zeta' \in \Lambda_G \}$$

be the *weak hull* of Λ_G , if $f: G \rightarrow \text{WH}_G$ is any G -equivariant map, then f is a quasi-isometry and the map $\bar{f} = f \cup f_\infty: G \cup \partial G \rightarrow \text{WH}_G \cup \Lambda_G$ is continuous.

In this definition, WH_G is metrized by restricting the Teichmüller metric on \mathcal{T} , which *a posteriori* has the effect of making WH_G into a quasigeodesic metric space. The definition implies that G acts cocompactly on WH_G : since $\Lambda_G \times \Lambda_G - \Delta$ is a closed subset of $\mathbf{PF}\mathcal{P}$ it follows that WH_G is a closed subset of \mathcal{T} ;

and since G acts coboundedly on itself it follows that G acts coboundedly on WH_G ; thus, the image of WH_G in moduli space is closed and bounded, hence compact.

3.2 Properties of convex cocompact subgroups

In this section we prove several properties of convex cocompact subgroups of $\text{Isom}(\mathcal{T})$ which are analogues of well known properties in $\text{Isom}(\mathbf{H}^n)$.

Proposition 3.1 *Every infinite order element g of a convex cocompact subgroup $G < \text{Isom}(\mathcal{T}) \approx MCG$ is a pseudo-Anosov mapping class.*

Proof Any infinite order element of a word hyperbolic group has source–sink dynamics on its Gromov boundary, and so g has source–sink dynamics on $\partial G \approx \Lambda_G$. It follows that g has an axis in WH_G . But the elements of $\text{Isom}(\mathcal{T}) \approx MCG$ having an axis in \mathcal{T} are precisely the pseudo-Anosovs [5]. \square

The following is a consequence of work of McCarthy and Papadopoulos [36].

Proposition 3.2 *If G is a convex cocompact subgroup of $\text{Isom}(\mathcal{T})$ then:*

- (1) Λ_G is the smallest nontrivial closed subset of $\overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$ invariant under G .
- (2) The action of G on $\mathbf{PMF} \setminus \Lambda_G$ is properly discontinuous.

Proof The Gromov boundary of a word hyperbolic group is the closure of the fixed points of infinite order elements in the group, and so by Proposition 3.1 the set Λ_G is the closure of the fixed points of the pseudo-Anosov elements of G . Item (1) now follows from Theorem 4.1 of [36].

To prove (2), let

$$Z(\Lambda) = \{\zeta \in \mathbf{PMF} \mid \text{there exists } \zeta' \in \Lambda \text{ such that } i(\zeta, \zeta') = 0\}$$

Theorem 6.16 of [36] says that G acts properly discontinuously on $\mathbf{PMF} - Z(\Lambda)$, and so it suffices to prove that $\Lambda = Z(\Lambda)$. Each point $\zeta' \in \Lambda$ is the ideal endpoint of a cobounded geodesic ray, which implies that ζ' is uniquely ergodic and fills the surface [29], and so if $i(\zeta, \zeta') = 0$ then $\zeta = \zeta'$. \square

Remark One theme of [36] is that for a general finitely generated subgroup $G < MCG$, there are several different types of “limit sets” for the action of G on \mathbf{PMF} . Assuming that G contains a pseudo-Anosov element, the two sets mentioned in the proof above play key roles in [36]: $\Lambda(G)$ which is the closure of the fixed points of pseudo-Anosov elements of the subgroup, and is also the smallest nontrivial closed G -invariant subset; and the set $Z(\Lambda(G))$. What we have proved is that for a convex cocompact subgroup G , these two sets are identical. Henceforth we refer to Λ_G as *the limit set* for the action of G on \mathbf{PMF} .

The analogue of the following result is true for convex cocompact discrete subgroups of \mathbf{H}^n , as well as for word hyperbolic groups [2]; the proof here is similar.

Proposition 3.3 *Let G be a convex cocompact subgroup of $\text{Isom}(\mathcal{T})$, and let N_G and Comm_G be the normalizer and the relative commensurator of G in $\text{Isom}(\mathcal{T})$. Then each of the inclusions $G < N_G < \text{Comm}_G$ is of finite index, and we have $\text{Comm}_G = \text{Stab}(\Lambda_G) = \text{Stab}(\text{WH}_G)$.*

Proof Let Λ_G be the limit set of G , with weak hull WH_G , and note that we trivially have $\text{Stab}(\text{WH}_G) = \text{Stab}(\Lambda_G)$.

Note that $\text{Stab}(\text{WH}_G)$ acts properly on WH_G . Indeed, $\text{Isom}(\mathcal{T})$ acts properly on \mathcal{T} , and so any subgroup of $\text{Isom}(\mathcal{T})$ acts properly on any subset of \mathcal{T} which is invariant under that subgroup. Since $G \subset \text{Stab}(\text{WH}_G)$, and since G acts cocompactly on WH_G , it follows that G is contained with finite index in $\text{Stab}(\text{WH}_G)$. This implies that $\text{Stab}(\text{WH}_G) \subset \text{Comm}_G$. To complete the proof we only have to prove the reverse inclusion $\text{Comm}_G \subset \text{Stab}(\text{WH}_G)$.

Given $g \in \text{Isom}(\mathcal{T})$, suppose that $g \in \text{Comm}_G$, and choose finite index subgroups $H, K < G$ such that $g^{-1}Hg = K$. By the definition of convex cocompactness it follows that $\text{WH}_H = \text{WH}_G = \text{WH}_K$. Since $g(\text{WH}_K) = \text{WH}_H$ it follows that $g \in \text{Stab}(\text{WH}_G)$. \square

Remark Another natural property for subgroups $G < MCG$ is quasiconvexity with respect to the word metric on MCG . It seems possible to us that this is not equivalent to orbit quasiconvexity of G in $\text{Isom}(\mathcal{T})$. Masur and Minsky [33] give an example of an infinite cyclic subgroup of $\text{Isom}(\mathcal{T})$ which is not orbit quasiconvex, and yet this subgroup is quasi-isometrically embedded in MCG [14]; it may also be quasiconvex in MCG , but we have not investigated this.

3.3 Equivalence of definitions: Proof of Theorem 1.1

Here is our main result equating the various quasiconvexity properties with convex cocompactness:

Theorem 1.1 *If G is a finitely generated subgroup of $\text{Isom}(\mathcal{T})$, the following are equivalent:*

- (1) *Orbit quasiconvexity*
- (2) *Single orbit quasiconvexity*
- (3) *Weak orbit quasiconvexity*
- (4) *Convex cocompactness*

Because of this theorem we are free to refer to “quasiconvexity” or “convex cocompactness” of G without any ambiguity.

Proof of Theorem 1.1 The key ingredients in the proof are results of Minsky from [38] concerning projections from balls and horoballs in \mathcal{T} to geodesics in \mathcal{T} , and results of Masur–Minsky [32] characterizing δ –hyperbolicity of proper geodesic metric spaces in terms of projections properties to paths.

To begin with, note that the implications (1) \Rightarrow (2) \Rightarrow (3) are obvious. We now prove that (3) \Rightarrow (1).

Suppose we have an orbit \mathcal{O} of G and a constant A , and for each $x, y \in \mathcal{O}$ we have two points $x', y' \in \mathcal{T}$, endpoints of a unique geodesic segment $[x', y']$ in \mathcal{T} , such that $d(x, x') \leq A$, $d(y, y') \leq A$, and $[x', y'] \subset N_A(\mathcal{O})$. The set \mathcal{O} maps to a single point in \mathcal{M} and so the projection of $N_A(\mathcal{O})$ to \mathcal{M} is a bounded set \mathcal{B} . It follows that each $[x', y']$ is \mathcal{B} –cobounded. Now consider an arbitrary orbit \mathcal{O}_1 of G ; we must prove that \mathcal{O}_1 is quasiconvex in \mathcal{T} . The orbits $\mathcal{O}, \mathcal{O}_1$ have finite Hausdorff distance C in \mathcal{T} . Given $x_1, y_1 \in \mathcal{O}_1$, choose $x, y \in \mathcal{O}$ within distance C of x_1, y_1 , respectively, and consider the geodesic segment $[x', y']$ and the piecewise geodesic path

$$\gamma = [x', x] * [x, x_1] * [x_1, y_1] * [y_1, y] * [y, y']$$

Of the five subsegments of γ , all but the middle subsegment have length $\leq \text{Max}\{A, C\}$, and it follows that γ is a $(1, D)$ –quasigeodesic in \mathcal{T} , with D depending only on A, C . Since the geodesic $[x', y']$ is \mathcal{B} –cobounded we can apply the following result of Minsky [38] to obtain δ , depending only on \mathcal{B} and D , such that $\gamma \subset N_\delta[x', y']$.

Theorem 3.4 (Stability of cobounded geodesics) *For any bounded subset \mathcal{B} of \mathcal{M} and any $K \geq 1, C \geq 0$ there exists $\delta \geq 0$ such that if γ is a K, C quasigeodesic in \mathcal{T} with endpoints x, y , and if $[x, y]$ is \mathcal{B} -cobounded, then $\gamma \subset N_\delta[x, y]$. \square*

It follows that $[x_1, y_1] \subset \gamma \subset N_{\delta+A}\mathcal{O} \subset N_{\delta+A+C}\mathcal{O}_1$, proving quasiconvexity of \mathcal{O}_1 in \mathcal{T} .

Weak orbit quasiconvexity implies convex cocompactness Fix an orbit \mathcal{O} of G , and so \mathcal{O} is quasiconvex in \mathcal{T} . Let \mathcal{G} be the set of all geodesic segments, rays, and lines that are obtained as pointwise limits of sequences of geodesics with endpoints in \mathcal{O} . Let $\cup\mathcal{G} \subset \mathcal{T}$ be the union of the elements of \mathcal{G} . The left action of G on \mathcal{O} is evidently cobounded. By quasiconvexity of \mathcal{O} it follows that the action of G on the union of geodesic segments with endpoints in \mathcal{O} is cobounded, which implies in turn that the action of G on $\cup\mathcal{G}$ is cobounded. Since $\cup\mathcal{G}$ is closed and \mathcal{T} is locally compact, it follows that the G action on $\cup\mathcal{G}$ is cocompact. The set $\cup\mathcal{G}$ therefore projects to a compact subset of \mathcal{M} which we denote \mathcal{B} . All geodesics in \mathcal{G} are therefore \mathcal{B} -cobounded.

Let $\cup\mathcal{G}$ be equipped with the restriction of the Teichmüller metric. Note that while $\cup\mathcal{G}$ is not a geodesic metric space, it is a quasigeodesic metric space: there exists $A \geq 0$ such that any $x, y \in \cup\mathcal{G}$ are within distance A of points $x', y' \in \mathcal{O} \subset \cup\mathcal{G}$, and the geodesic $[x', y']$ is contained in $\cup\mathcal{G}$.

To prepare for the proof that G is word hyperbolic, fix a finite generating set for G with Cayley graph Γ , and fix a G -equivariant map $f: \Gamma \rightarrow \cup\mathcal{G}$ taking the vertices of Γ to \mathcal{O} and taking each edge of Γ to an element of \mathcal{G} . Since G acts properly and coboundedly on both Γ and $\cup\mathcal{G}$, and since both are quasigeodesic metric spaces, it follows that the equivariant map f is a quasi-isometry between Γ and $\cup\mathcal{G}$; pick a coarse inverse $F: \cup\mathcal{G} \rightarrow \Gamma$.

By definition the group G is word hyperbolic if and only if the Cayley graph Γ is δ -hyperbolic for some $\delta \geq 0$. Our proof that G is word hyperbolic will use a result of Masur and Minsky, Theorem 2.3 of [32]:

Theorem 3.5 *Let X be a geodesic metric space and suppose that there is a set of paths \mathbf{P} in X with the following properties:*

Coarse transitivity *There exists $C \geq 0$ such that for any $x, y \in X$ with $d(x, y) \geq C$ there is a path in \mathbf{P} joining x and y .*

Contracting projections: *There exist $a, b, c > 0$, and for each path $\gamma: I \rightarrow X$ in \mathbf{P} there exists a map $\pi: X \rightarrow I$ such that:*

Coarse projection For each $t \in I$ we have $\text{diam}(\gamma[t, \pi(\gamma t)]) \leq c$.

Coarse lipschitz If $d(x, y) \leq 1$ then $\text{diam}(\gamma[\pi x, \pi y]) \leq c$.

Contraction If $d(x, \gamma(\pi x)) \geq a$ and $d(x, y) \leq b \cdot d(x, \gamma(\pi x))$ then

$$\text{diam}(\gamma[\pi x, \pi y]) \leq c$$

Then X is δ -hyperbolic for some $\delta \geq 0$.

To prove that G is δ -hyperbolic we take \mathbf{P} to be the set of geodesic segments in G , and we look at the set of paths $f \circ \mathbf{P} = \{f \circ \gamma \mid \gamma \in \mathbf{P}\}$ in $\cup \mathcal{G}$. Using some results of Minsky [38], we will show that $f \circ \mathbf{P}$ satisfies the hypotheses of Theorem 3.5. Then we shall pull the hypotheses back to \mathbf{P} and apply Theorem 3.5.

The first result of Minsky that we need is the main theorem of [38]:

Theorem 3.6 (Contraction Theorem) For every bounded subset \mathcal{B} of \mathcal{M} there exists $c > 0$ such that if γ is any \mathcal{B} -cobounded geodesic in \mathcal{T} then the closest point projection $\mathcal{T} \rightarrow \gamma$ satisfies the (a, b, c) contracting projection property with $(a, b) = (0, 1)$.

In our context, where we have a uniform \mathcal{B} such that each geodesic in \mathcal{G} is \mathcal{B} -cobounded, it follows that there is a uniform c such that each geodesic in \mathcal{G} satisfies the $(0, 1, c)$ contracting projection property.

Now consider $\gamma = [x_0, x_1, \dots, x_n]$ a geodesic in the Cayley graph Γ , mapping via f to a piecewise geodesic $f\gamma = [fx_0, fx_1] \cup \dots \cup [fx_{n-1}, fx_n]$ in $\cup \mathcal{G}$, with each subsegment $[fx_i, fx_{i+1}]$ an element of \mathcal{G} . It follows that $f\gamma$ is a K, C quasigeodesic in \mathcal{T} , for $K \geq 1, C \geq 0$ independent of the given geodesic in Γ . The \mathcal{T} -geodesic $[fx_0, fx_n]$ is \mathcal{B} -cobounded. Applying Theorem 3.4 it follows that $f\gamma \subset N_D[fx_0, fx_n]$, where D depends only on \mathcal{B}, K, C . As noted above, closest point projection from \mathcal{T} onto $[fx_0, fx_n]$ satisfies the $(0, 1, c)$ contracting projection property. From this it follows that closest point projection $\pi: \mathcal{T} \rightarrow f\gamma$ satisfies the (a', b', c') contraction property where (a', b', c') depend only on \mathcal{B}, K, C . Now define the projection $\Gamma \rightarrow \gamma$ to be the composition $\Gamma \xrightarrow{f} \cup \mathcal{G} \xrightarrow{\pi} f\gamma \xrightarrow{F} \Gamma \rightarrow \gamma$ where the last map is closest point projection in Γ . This composition clearly satisfies the (a'', b'', c'') projection property where (a'', b'', c'') depend only on (a', b', c') and the quasi-isometry constants and coarse inverse constants for f, F .

Geodesics in Γ are clearly coarsely transitive, and applying Theorem 3.5 it follows that G is word hyperbolic. This means that geodesic triangles in Γ are

uniformly thin, and it implies that for each K, C there is a δ such that K, C quasigeodesic triangles in Γ are δ -thin. Applying the quasi-isometry between Γ and $\cup \mathcal{G}$, it follows that there is a uniform δ such that for each $x, y, z \in \mathcal{O}$ the geodesic triangle $\Delta[x, y, z]$ in $\cup \mathcal{G}$ is δ -thin; we fix this δ for the arguments below.

Now we turn to a description of the “limit set” $\Lambda \subset \mathbf{PMF}$ of G , with the ultimate goal of identifying it with the Gromov boundary ∂G .

Each geodesic ray in \mathcal{G} has the form $\overrightarrow{[x, \eta]}$, for some $x \in \mathcal{O}$, $\eta \in \mathbf{PMF}$; define $\Lambda \subset \mathbf{PMF}$ be the set of all such points η , over all geodesic rays in \mathcal{G} . The set Λ is evidently G -equivariant.

Fact 1 For any $x \in \mathcal{O}$, $\eta \in \Lambda$, the ray $\overrightarrow{[x, \eta]}$ in \mathcal{T} is an element of \mathcal{G} .

To prove this, by definition of Λ there exists a ray $\overrightarrow{[y, \eta]}$ in \mathcal{G} for some $y \in \mathcal{O}$. Choose a sequence $y_1, y_2, \dots \in \mathcal{O}$ staying uniformly close to $\overrightarrow{[y, \eta]}$ and going to infinity. Pass to a subsequence so that the sequence of segments $[x, y_n]$ converges to some ray $\overrightarrow{[x, \eta']} \in \mathcal{G}$; it suffices to show that $\eta' = \eta$. Since x is fixed and the points y_n stay uniformly close to $\overrightarrow{[y, \eta]}$, it follows by Theorem 3.4 that the segments $[x, y_n]$ stay uniformly close to $\overrightarrow{[y, \eta]}$, and so $\overrightarrow{[x, \eta']}$ is in a finite neighborhood of $\overrightarrow{[y, \eta]}$. The reverse inclusion, that $\overrightarrow{[y, \eta]}$ is in a finite neighborhood of $\overrightarrow{[x, \eta']}$, is a standard argument: as points move to infinity in $\overrightarrow{[x, \eta']}$ taking bounded steps, uniformly nearby points move to infinity in $\overrightarrow{[y, \eta]}$ also taking bounded steps, and thus must come uniformly close to an arbitrary point of $\overrightarrow{[y, \eta]}$. This shows that the rays $\overrightarrow{[x, \eta']}$, $\overrightarrow{[y, \eta]}$ have finite Hausdorff distance, and applying Lemma 2.4 (End Uniqueness) shows that $\eta = \eta'$.

Note that in the proof of Fact 1 we have established a little more, namely that for any $x, y \in \mathcal{O}$ and $\eta \in \Lambda$ the rays $\overrightarrow{[x, \eta]}$ and $\overrightarrow{[y, \eta]}$ have finite Hausdorff distance. This will be useful below.

Fact 2 For any $\eta \neq \zeta \in \Lambda$ there exists a line $\overleftrightarrow{(\eta, \zeta)}$ contained in \mathcal{G} .

From Fact 2 it immediately follows that $\Lambda \times \Lambda - \Delta \subset \mathbf{PF}\mathcal{P}$, that the weak hull WH_G of Λ is defined, and that G acts coboundedly on WH_G , since G acts coboundedly on $\cup \mathcal{G}$.

To prove Fact 2, pick a point $x \in \mathcal{O}$, and note that by Fact 1 we have two rays $\overrightarrow{[x, \eta]}$, $\overrightarrow{[x, \zeta]}$ in \mathcal{G} . Pick a sequence $y_n \in \mathcal{O}$ staying uniformly close to

$\overrightarrow{[x, \eta]}$ and going to infinity, and a sequence $z_n \in \mathcal{O}$ staying uniformly close to $\overrightarrow{[x, \zeta]}$ and going to infinity. We have a sequence of triangles $[x, y_n, z_n]$ in \mathcal{G} , all δ -thin. Applying Theorem 3.4 there is a D such that the sides $[x, y_n]$ are contained in the D -neighborhood of $\overrightarrow{[x, \eta]}$, and the sides $[x, z_n]$ are contained in the D -neighborhood of $\overrightarrow{[x, \zeta]}$. Each side $[y_n, z_n]$, being contained in the δ -neighborhood of $\overrightarrow{[x, y_n]} \cup \overrightarrow{[x, z_n]}$, is therefore contained in the $D + \delta$ -neighborhood of $\overrightarrow{[x, \eta]} \cup \overrightarrow{[x, \zeta]}$.

We claim that the point x is uniformly close to the segments $[y_n, z_n]$. If not, then from uniform thinness of the triangles $[x, y_n, z_n]$ it follows that there are points $y'_n \in [x, y_n]$ and $z'_n \in [x, z_n]$ such that the segments $[x, y'_n]$ and $[x, z'_n]$ get arbitrarily long while the Hausdorff distance between them stays uniformly bounded. This implies that there are sequences $y''_n \in \overrightarrow{[x, \eta]}$ going to infinity and $z''_n \in \overrightarrow{[x, \zeta]}$ going to infinity such that the Hausdorff distance between the segments $[x, y''_n]$ and $[x, z''_n]$ stays uniformly bounded, which implies in turn that the rays $\overrightarrow{[x, \eta]}$ and $\overrightarrow{[x, \zeta]}$ have finite Hausdorff distance. Applying End Uniqueness 2.4, it follows that $\eta = \zeta$, contradicting the hypothesis of Fact 2, and the claim follows.

Passing to a subsequence and applying Ascoli–Arzela it follows that $[y_n, z_n]$ converges to a line in \mathcal{G} . One ray of this line is Hausdorff close to $\overrightarrow{[x, \eta]}$ and so has endpoint η , and the other ray is Hausdorff close to $\overrightarrow{[x, \zeta]}$ and so has endpoint ζ , by End Uniqueness. We therefore have $\lim[y_n, z_n] = \overleftarrow{(\eta, \zeta)}$, completing the proof of Fact 2.

Now we define a map $f_\infty: \partial G \rightarrow \Lambda$. Recall that the relation of finite Hausdorff distance is an equivalence relation on geodesic rays in the Cayley graph Γ of G , and ∂G is the set of equivalence classes. Consider then a point $\xi \in \partial G$ represented by two geodesic rays $[x_0, x_1, \dots)$ and $[y_0, y_1, \dots)$ with finite Hausdorff distance in Γ . These map to piecewise geodesic, quasigeodesic rays $\rho = [fx_0, fx_1] \cup [fx_1, fx_2] \cup \dots$ and $\sigma = [fy_0, fy_1] \cup [fy_1, fy_2] \cup \dots$ with finite Hausdorff distance in $\cup \mathcal{G}$. The sequence of geodesic segments $[fx_0, fx_n]$ in \mathcal{G} has a subsequence converging to some ray $\overrightarrow{[fx_0, \zeta]}$ in \mathcal{G} , and $[fy_0, fy_n]$ has a subsequence converging to some ray $\overrightarrow{[fy_0, \zeta']}$ in \mathcal{G} . To obtain a well defined map $\partial G \rightarrow \Lambda$ it suffices to prove that $\zeta = \zeta'$, and then we can set $f_\infty(\xi) = \zeta$.

To prove that $\zeta = \zeta'$ it suffices, by End Uniqueness 2.4, to prove that the rays $\overrightarrow{[fx_0, \zeta]}$ and $\overrightarrow{[fy_0, \zeta']}$ have finite Hausdorff distance in \mathcal{T} . Since the piecewise geodesic rays ρ, σ have finite Hausdorff distance in \mathcal{T} , it suffices to prove that ρ

has finite Hausdorff distance from $\overrightarrow{[fx_0, \zeta]}$, and similarly σ has finite Hausdorff distance from $\overrightarrow{[fy_0, \zeta']}$. Consider a point $p \in \rho$. For sufficiently large n we have $p \in \rho_n = [fx_0, fx_1] \cup \dots \cup [fx_{n-1}, fx_n]$. Applying Theorem 3.4 there is a uniform constant D such that $\rho_n \subset N_D([fx_0, fx_n])$, and so p is within distance D of some point in $[fx_0, fx_n]$. Since $\overrightarrow{[fx_0, \zeta]}$ is the pointwise limit of $\overrightarrow{[fx_0, fx_n]}$ as $n \rightarrow \infty$ it follows that p is within a uniformly bounded distance of $\overrightarrow{[fx_0, \zeta]}$. This shows that ρ is within a finite neighborhood of $\overrightarrow{[fx_0, \zeta]}$. The reverse inclusion is a standard argument: as points move along ρ towards the end taking bounded steps, uniformly nearby points move along $\overrightarrow{[fx_0, \zeta]}$ towards the end also taking bounded steps, and thus must come uniformly close to some point of $\overrightarrow{[fx_0, \zeta]}$.

Hence $f_\infty: \partial G \rightarrow \Lambda$ is well defined. Observe that a similar argument proves a little more: if $x_i \in G$ converges to $\xi \in \partial G$ then the segments $[fx_0, fx_i]$ converge in the compact-open topology to the ray $\overrightarrow{[fx_0, f\xi]}$; details are left to the reader.

We now turn to verifying required properties of f_∞ .

To see that f_∞ is surjective, consider a point $\eta \in \Lambda$ and pick a ray $\overrightarrow{[x, \eta]}$ in \mathcal{G} . It follows that $\rho = F(\overrightarrow{[x, \eta]})$ is a quasigeodesic ray in Γ . Since Γ is δ -hyperbolic it follows that ρ has finite Hausdorff distance from some geodesic ray ρ' in Γ , with endpoint $\zeta' \in \partial G$. As shown above, $f(\rho')$ has finite Hausdorff distance from some geodesic ray $\overrightarrow{[x', f_\infty \zeta']}$. Since f, F are coarse inverses it follows that $\overrightarrow{[x, \eta]}$ has finite Hausdorff distance from $\overrightarrow{[x', f_\infty \zeta]}$, and so by End Uniqueness it follows that $\eta = f_\infty \zeta'$.

To see that f_∞ is injective, consider two points $\eta, \zeta \in \partial G$ and suppose that $f_\infty(\eta) = f_\infty(\zeta)$; let $\xi \in \Lambda$ be this point. Pick rays ρ, σ in Γ representing η, ζ respectively. As we have just seen, the images $f(\rho), f(\sigma)$ have finite Hausdorff distance in \mathcal{T} to rays $\overrightarrow{[y, \xi]}, \overrightarrow{[z, \xi]}$ in \mathcal{G} , respectively. As noted at the end of the proof of Fact 1, the rays $\overrightarrow{[y, \xi]}$ and $\overrightarrow{[z, \xi]}$ have finite Hausdorff distance in \mathcal{T} ; applying the coarse inverse F it follows that ρ, σ have finite Hausdorff distance in Γ and therefore $\eta = \zeta$.

We have shown that f_∞ is a bijection between ∂G and Λ . We want to prove that f_∞ is a homeomorphism, and that the extension $\overline{f} = f \cup f_\infty: G \cup \partial G \rightarrow \overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$ is continuous. For this purpose first we establish:

Fact 3 Λ is a closed subset of \mathbf{PMF} , and therefore compact.

To prove this, choose a sequence $\zeta_n \in \Lambda$ so that $\lim \zeta_n = \zeta_\infty$ in \mathbf{PMF} ; we must prove that $\zeta_\infty \in \Lambda$. Choose a point $x \in \mathcal{O}$, and apply Fact 1 to obtain rays $\overrightarrow{[x, \zeta_n]}$. Passing to a subsequence these converge to a limiting ray $\lim \overrightarrow{[x, \zeta_n]} = \overrightarrow{[x, \zeta'_\infty]}$ in \mathcal{G} , and so $\zeta'_\infty \in \Lambda$. Looking in the unit tangent bundle of \mathcal{T} at the point x it follows that $\lim \zeta_n = \zeta'_\infty$, and so $\zeta_\infty = \zeta'_\infty \in \Lambda$.

Fact 4 $f_\infty: \partial G \rightarrow \Lambda$ is a homeomorphism.

Since both the domain and range are compact Hausdorff spaces it suffices to prove continuity in one direction. Continuity of f_∞^{-1} follows by simply noting that for fixed $x \in \mathcal{O}$ and for a convergent sequence $\xi_n \rightarrow \xi$ in $\Lambda \subset \mathbf{PMF}$, the sequence of rays $\overrightarrow{[x, \xi_n]}$ converges in the compact open topology to the ray $\overrightarrow{[x, \xi]}$.

Fact 5 The map $\overline{f} = f \cup f_\infty: G \cup \partial G \rightarrow \overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$ is continuous.

To be precise, this map is continuous using the Thurston compactification $\overline{\mathcal{T}}$ of \mathcal{T} . We prove this by showing first that the map is continuous using the Teichmüller compactification, and then we apply Masur's Two Boundaries Theorem [30] which says that the map from the Teichmüller compactification to the Thurston compactification is continuous at uniquely ergodic points of \mathbf{PMF} .

First we recall the Teichmüller compactification in a form convenient for our current purposes. There are actually many different Teichmüller compactifications, one for each choice of a base point in \mathcal{T} ; we shall fix a base point $z = f(x) \in \mathcal{O}$ for some $x \in G$. As we have seen, there is a unique geodesic segment $[z, z']$ for each $z' \in \mathcal{T}$, and a unique geodesic ray $\overrightarrow{[z, \zeta]}$ for each $\zeta \in \mathbf{PMF}$. The Teichmüller topology on $\overline{\mathcal{T}} = \mathcal{T} \cup \mathbf{PMF}$ restricts to the standard topologies on \mathcal{T} and on \mathbf{PMF} , it has \mathcal{T} as a dense open subset, and a sequence $z_i \in \mathcal{T}$ converges to $\zeta \in \mathbf{PMF}$ if and only if the sequence of segments $[z, z_i]$ converges to the ray $\overrightarrow{[z, \zeta]}$ in the compact open topology; equivalently, letting B denote the unit ball in \mathcal{T} centered on z , the distance $d(z, z_i)$ goes to infinity and the set $[z, z_i] \cap B$ converges to the set $\overrightarrow{[z, \zeta]} \cap B$ in the Hausdorff topology.

We already proved in Fact 4 that f_∞ is continuous; for this we implicitly used the fact that the Thurston topology on \mathbf{PMF} is identical to the Teichmüller topology, defined by identifying \mathbf{PMF} with the unit tangent bundle at x . We also observed earlier, after the proof that f_∞ is well-defined, that if $x_i \in G$ converges to $\xi \in \partial G$, then $f(x_i) \in \mathcal{T}$ converges to $f_\infty(\xi) \in \mathbf{PMF}$ in

the Teichmüller topology on $\overline{\mathcal{T}}$. Putting these together it follows that \overline{f} is continuous using the Teichmüller topology on $\overline{\mathcal{T}}$. Since $\Lambda = f_\infty(\partial G)$ consists entirely of uniquely ergodic points in \mathbf{PMF} , Masur's Two Boundaries Theorem [30] implies that the identity map on $\overline{\mathcal{T}}$ is continuous from the Teichmüller topology to the Thurston topology at each point of Λ , and so \overline{f} is continuous using the Thurston topology on $\overline{\mathcal{T}}$.

We now put the pieces together to complete the proof of convex cocompactness. Let $f': G \rightarrow \text{WH}_G$ be an arbitrary G -equivariant map, and define $f'_\infty: \partial G \rightarrow \mathbf{PMF}$ to be equal to f_∞ . We must prove that f' is a quasi-isometry and that the extension $\overline{f'} = f' \cup f'_\infty: G \cup \partial G \rightarrow \text{WH}_G \cup \Lambda_G$ is continuous. From Facts 1–5 above, it follows that the quasi-isometry $f: G \rightarrow \cup \mathcal{G}$ has continuous extension $\overline{f}: G \cup \partial G \rightarrow \cup \mathcal{G} \cup \Lambda$, and so $\cup \mathcal{G}$ is a Gromov hyperbolic metric space with Gromov compactification $\cup \mathcal{G} \cup \Lambda$. Since $\text{WH}_G \subset \cup \mathcal{G}$ is a G -invariant subset, it follows that WH_G is Gromov hyperbolic with Gromov compactification $\text{WH}_G \cup \Lambda$. The map f' is a G -equivariant map between quasigeodesic metric spaces on which G acts properly and coboundedly by isometries, and hence f' is a quasi-isometry. Since $d(f'(x), f(x))$ is uniformly bounded for $x \in G$, then from the fact that $f'_\infty = f_\infty$ it follows that $\overline{f'}$ is continuous.

This completes the proof that weak orbit quasiconvexity implies convex cocompactness.

Convex cocompact implies weak orbit quasiconvexity Assuming G is convex cocompact, pick a finite generating set for G with Cayley graph Γ and G -equivariant, coarsely inverse quasi-isometries $f: \Gamma \rightarrow \text{WH}_G$, $\overline{f}: \text{WH}_G \rightarrow \Gamma$.

Let \mathcal{O} be an orbit of G in \mathcal{T} . Since G acts coboundedly on WH_G it follows that \mathcal{O} has finite Hausdorff distance from WH_G in \mathcal{T} . It suffices to show that for any two points $x, y \in \mathcal{O}$ there is a geodesic line whose infinite ends are in Λ such that x, y come within a uniformly finite distance of that line.

Pick a G -equivariant map $g: \Gamma \rightarrow \mathcal{T}$ taking the vertices of Γ bijectively to \mathcal{O} and each edge of Γ to a geodesic segment, so f and g differ by a bounded amount. Since Γ is δ -hyperbolic it follows that there is a constant A such that any two vertices of Γ lie within distance A of some bi-infinite geodesic. Pick $x, y \in \mathcal{O}$, and pick a bi-infinite geodesic γ in Γ such that $g^{-1}(x), g^{-1}(y)$ are within distance A of γ . Let $\xi, \eta \in \partial G$ be the two ends of γ . By the statement of convex cocompactness, there is a K, C quasigeodesic line in Γ of the form $\overline{f} \left(\overleftarrow{(f_\infty \xi, f_\infty \eta)} \right)$ whose two infinite ends are ξ, η , where K, C are independent

of ξ, η . It follows that γ and $\overline{f}(\overleftarrow{(f_\infty\xi, f_\infty\eta)})$ are uniformly close, and so $f(\gamma)$ and $\overleftarrow{(f_\infty\xi, f_\infty\eta)}$ are uniformly close, and so the points x, y are uniformly close to $\overleftarrow{(f_\infty\xi, f_\infty\eta)}$. \square

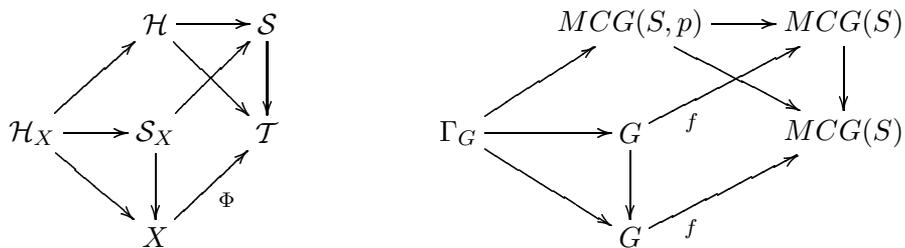
4 Hyperbolic surface bundles over graphs

In this section our goal is to give an explicit construction of model geometries for surface group extensions, and to study regularity properties of these geometries. Here is a brief outline; detailed constructions follow.

Consider a finitely generated group G and a homomorphism $f: G \rightarrow \text{Isom}(\mathcal{T}) \approx MCG$. Let X be a Cayley graph for G . Choose a map $\Phi: X \rightarrow \mathcal{T}$ which is equivariant with respect to the homomorphism f , that is, $\Phi(g \cdot x) = f(g) \cdot \Phi(x)$, $x \in X, g \in G$, where we use the \cdot notation to denote an action. By pulling back the canonical marked hyperbolic surface bundle $\mathcal{S} \rightarrow \mathcal{T}$ via the map Φ we obtain a marked hyperbolic surface bundle $\mathcal{S}_X \rightarrow X$. By pulling back the canonical hyperbolic plane bundle $\mathcal{H} \rightarrow \mathcal{T}$ we obtain a hyperbolic plane bundle $\mathcal{H}_X \rightarrow X$, and a covering map $\mathcal{H}_X \rightarrow \mathcal{S}_X$ with deck transformation group $\pi_1(\mathcal{S})$. There is an action of the extension group Γ_G on \mathcal{H}_X such that the covering map $\mathcal{H}_X \rightarrow \mathcal{S}_X$ is equivariant with respect to the homomorphism $\Gamma_G \rightarrow G$.

By imposing a G -equivariant, proper, geodesic metric on \mathcal{S}_X and lifting to \mathcal{H}_X , we can then use \mathcal{H}_X as a model geometry for the extension group Γ_G .

We may summarize all this in the following commutative diagrams:



Each group in the right hand diagram acts on the corresponding space in the left hand diagram, and each map in the left hand diagram is equivariant with respect to the corresponding group homomorphism in the right hand diagram.

We will impose several Γ_G -equivariant structures on the space \mathcal{H}_X , by finding appropriate G -equivariant structures on \mathcal{S}_X and lifting.

For example, we put an equivariant, proper, geodesic metric on \mathcal{H}_X by lifting an equivariant, proper, geodesic metric on \mathcal{S}_X . These metrics will have the property that the topological fibrations $\mathcal{S}_X \rightarrow X$, $\mathcal{H}_X \rightarrow X$ are also “metric fibrations” in the following sense. In a metric space Z , given subsets $A, B \subset Z$, denote the *min distance* by $d_{\min}(A, B) = \inf\{d(a, b) \mid a \in A, b \in B\}$, and the *Hausdorff distance* by $d_{\text{Haus}}(A, B) = \inf\{r \mid A \subset N_r(B), B \subset N_r(A)\}$.

Metric fibration property A map of metric spaces $f: Z \rightarrow Y$ satisfies the *metric fibration property* if Y is covered by neighborhoods U such that if $y, z \in U$ then

$$d_{\min}(f^{-1}(y), f^{-1}(z)) = d_{\text{Haus}}(f^{-1}(y), f^{-1}(z)) = d_Y(y, z)$$

4.1 Metrics and connections on surface bundles over paths

The marked hyperbolic surface bundle over a path in \mathcal{T} Consider first a smooth path $\alpha: I \rightarrow \mathcal{T}$, defined on a closed connected subset $I \subset \mathbf{R}$, that is, a closed interval, a closed ray, or the whole line. Pulling back the canonical marked hyperbolic surface bundle $\mathcal{S} \rightarrow \mathcal{T}$ via the map α we obtain a marked hyperbolic surface bundle $\mathcal{S}_\alpha \rightarrow I$. We impose a Riemannian metric on \mathcal{S}_α as follows.

Recall that we have chosen a connection on the bundle $\mathcal{S} \rightarrow \mathcal{T}$. By pulling back the connection on the bundle $\mathcal{S} \rightarrow \mathcal{T}$ we obtain a connection on the bundle $\mathcal{S}_\alpha \rightarrow I$, that is, a 1-dimensional sub-bundle of $T\mathcal{S}_\alpha$ which is complementary to the vertical sub-bundle $T_v\mathcal{S}_\alpha$. There is a unique vector field V on \mathcal{S}_α parallel to the connection such that the projection map $\mathcal{S}_\alpha \rightarrow I$ takes each vector of V to a positive unit vector in the tangent bundle of $I \subset \mathbf{R}$. There is now a unique Riemannian metric on \mathcal{S} whose restriction to $T_v\mathcal{S}_\alpha$ is the given hyperbolic metric along leaves of \mathcal{S}_α , and such that V is a unit vector field orthogonal to $T_v\mathcal{S}_\alpha$. Since I is closed subset of \mathbf{R} , the path metric on \mathcal{S}_α induced from this Riemannian metric is proper, and so by Fact 2.1 we may regard \mathcal{S}_α as a geodesic metric space.

Here is another description of the Riemannian metric on \mathcal{S}_α . Integration of the connection sub-bundle defines a 1-dimensional foliation on \mathcal{S}_α transverse to the surface fibration, whose leaves are called *connection paths*. Choosing a base leaf of the fibration $\mathcal{S}_\alpha \rightarrow I$, and identifying this base leaf with S , we may project along connection paths to define a fibration $\mathcal{S}_\alpha \rightarrow S$. Combining this with the fibration $\mathcal{S}_\alpha \rightarrow I$ we obtain a diffeomorphism $\mathcal{S}_\alpha \approx S \times I$. Letting g_t

be the given Riemannian metric of curvature -1 on the leaf $\mathcal{S}_t \approx S \times t$, $t \in I$, we obtain the Riemannian metric on \mathcal{S}_α via the formula

$$ds^2 = g_t^2 + dt^2.$$

Remark The metric on \mathcal{S}_α depends on the choice of a connection on the bundle $\mathcal{S} \rightarrow \mathcal{T}$. However, when α is cobounded, two different connections on $\mathcal{S} \rightarrow \mathcal{T}$ will induce metrics on \mathcal{S}_α which are bilipschitz equivalent, with bilipschitz constant depending only on the pair of connections and on the coboundedness of α , not on α itself.

For each $s, t \in I$ we have a *connection map* $h_{st}: \mathcal{S}_s \rightarrow \mathcal{S}_t$, defined by moving each point of \mathcal{S}_s along a connection path until it hits \mathcal{S}_t . Clearly we have $h_{st} \circ h_{rs} = h_{rt}$, ($r, s, t \in I$). Notice that the map h_{st} takes each point of \mathcal{S}_s to the unique closest point on \mathcal{S}_t , and that point is at distance $|s - t|$. In fact, starting from an arbitrary point on \mathcal{S}_s , all paths to \mathcal{S}_t have length $\geq |s - t|$, and the connection path is the unique one with length $= |s - t|$. It follows that the map $\mathcal{S}_\alpha \rightarrow I$ satisfies the metric fibration property.

Consider more generally a piecewise smooth path $\alpha: I \rightarrow \mathcal{T}$. On each subinterval $I' \subset I$ over which α is smooth, there is a Riemannian metric as constructed above. At a point $t \in I$ where two such subintervals meet, the Riemannian metrics on the two sides agree when restricted to \mathcal{S}_t . We therefore have a piecewise Riemannian metric on \mathcal{S}_α , inducing a proper geodesic metric. The connection paths which are defined over each smooth subinterval $I' \subset I$ piece together to give connection paths on all of \mathcal{S}_α , and we obtain connection maps $h_{st}: \mathcal{S}_s \rightarrow \mathcal{S}_t$ for all $s, t \in I$.

Note that since the connection on $\mathcal{S} \rightarrow \mathcal{T}$ is equivariant with respect to the action of MCG , the piecewise Riemannian metric on each \mathcal{S}_α is *natural*, meaning that for any $h \in MCG$, the induced map $\mathcal{S}_\alpha \rightarrow \mathcal{S}_{h\circ\alpha}$ is an isometry. Similarly, the connection paths and connection maps are also natural.

Each connection map $h_{st}: \mathcal{S}_s \rightarrow \mathcal{S}_t$ is clearly a diffeomorphism, and since its domain is compact it follows that h_{st} is bilipschitz. The next proposition exhibits some regularity, bounding the bilipschitz constant of h_{st} by a function of $|s - t|$ that depends only on the coboundedness of the path $\alpha: I \rightarrow \mathcal{T}$, and a Lipschitz constant for α . For technical reasons we state the lemma only for paths $\alpha: I \rightarrow \mathcal{T}$ which are *piecewise affine*, meaning that I is a concatenation of subintervals I' such that $\alpha|_{I'}$ is an *affine* path, a constant speed reparameterization of a Teichmüller geodesic. Piecewise affine paths are sufficient for all of what follows.

Lemma 4.1 *For each bounded subset $\mathcal{B} \subset \mathcal{M}$ and each $\rho \geq 1$ there exists $K \geq 1$ such that the following happens. If $\alpha: I \rightarrow \mathcal{T}$ is a \mathcal{B} -cobounded, ρ -lipschitz, piecewise affine path, then for each $s, t \in I$ the connection map $h_{st}: \mathcal{S}_s \rightarrow \mathcal{S}_t$ is $K^{|s-t|}$ -bilipschitz.*

In what follows we shall describe the conclusion of this proposition by saying that K is a *bilipschitz constant* for the connection maps on \mathcal{S}_α .

Proof A standard lemma found in most O.D.E. textbooks shows that if Φ is a smooth flow on a compact manifold then there is a constant $K \geq 1$ such that $\|\Phi_t(v)\| \leq K^{|t|} \|v\|$. We can plug into this argument as follows.

The conclusion of the lemma is local, and so it suffices to prove it under the assumption that $I = [0, 1]$ and that α is affine. There exists a compact subset $\mathcal{A} \subset \mathcal{T}$ such that any \mathcal{B} -cobounded, ρ -lipschitz path $\alpha: [0, 1] \rightarrow \mathcal{T}$, can be translated by the action of MCG to lie in the set \mathcal{A} . Let $C(\mathcal{A}, \rho)$ be the set of all ρ -lipschitz affine paths $[0, 1] \mapsto \mathcal{A}$, a compact space in the compact open topology. By naturality of the metric on \mathcal{S}_α , it suffices to prove the lemma for $\alpha \in C(\mathcal{A}, \rho)$. For each $\alpha \in C(\mathcal{A}, \rho)$ and each vector \vec{w} tangent to a fiber \mathcal{S}_s , $s \in [0, 1]$, define:

$$l(\vec{w}) = \lim_{t \rightarrow 0} \frac{1}{t} \log \left(\frac{\|Dh_{s,s+t}(\vec{w})\|}{\|\vec{w}\|} \right) = \left. \frac{d}{dt} \right|_{t=0} \log \left(\frac{\|Dh_{s,s+t}(\vec{w})\|}{\|\vec{w}\|} \right)$$

Since $l(c\vec{w}) = l(\vec{w})$ for $c \neq 0$, we may regard $l(\vec{w})$ as a function defined on the projective tangent bundle of S crossed with I , a compact space. As \vec{w} varies, and as α varies over the compact space $C(\mathcal{A}, \rho)$, the function $l(\vec{w})$ varies continuously, and so by compactness $l(\vec{w})$ has a finite upper bound l . Setting $K = e^l$, it now follows by standard methods that $\|h_{s,s+t}(\vec{w})\| \leq K^{|t|} \|\vec{w}\|$ when \vec{w} is tangent to \mathcal{S}_s , and so $h_{s,s+t}$ is $K^{|t|}$ bilipschitz. \square

The hyperbolic plane bundle over a path in \mathcal{T} Letting $\alpha: I \rightarrow \mathcal{T}$ be a piecewise affine path as above, by pulling back the canonical hyperbolic plane bundle $\mathcal{H} \rightarrow \mathcal{T}$ we obtain a bundle $\mathcal{H}_\alpha \rightarrow I$. Note that there is a universal covering map $\mathcal{H}_\alpha \rightarrow \mathcal{S}_\alpha$ with deck transformation group $\pi_1(S)$ such that the composition $\mathcal{H}_\alpha \rightarrow \mathcal{S}_\alpha \rightarrow \mathcal{S}$ equals the composition $\mathcal{H}_\alpha \rightarrow \mathcal{H} \rightarrow \mathcal{S}$, and also the composition $\mathcal{H}_\alpha \rightarrow \mathcal{S}_\alpha \rightarrow I$ equals the fibration map $\mathcal{H}_\alpha \rightarrow I$. By lifting the piecewise Riemannian metric from \mathcal{S}_α we obtain a piecewise Riemannian metric on \mathcal{H}_α , inducing a proper, geodesic metric. The map $\mathcal{H}_\alpha \rightarrow I$ satisfies the metric fibration property. The connection paths on \mathcal{S}_α lift to connection

paths on \mathcal{H}_α , and we obtain connection maps $h_{st}: \mathcal{H}_s \rightarrow \mathcal{H}_t$. By applying Lemma 4.1 it follows that if α is \mathcal{B} -cobounded and ρ -lipschitz then the same constant $K = K(\mathcal{B}, \rho)$ is a bilipschitz constant for the connection maps on \mathcal{H}_α .

4.2 Metrics and connections on surface bundles over graphs

Let $f: G \rightarrow MCG$ be a homomorphism defined on a finitely generated group G . We have a canonical extension $1 \rightarrow \pi_1(S) \rightarrow \Gamma_G \rightarrow G \rightarrow 1$.

Fix once and for all a Cayley graph X for G , on which G acts cocompactly with quotient a rose. Fix a geodesic metric on X with each edge having length 1. Choose a G -equivariant map $\Phi: X \rightarrow \mathcal{T}$ taking each edge of X to an affine path in \mathcal{T} . Letting $\|\Phi\|$ be the maximum speed of the map Φ , ie, the maximal length of the image of an edge of X under Φ , it follows that Φ is a $\|\Phi\|$ -lipschitz map. Evidently the image of Φ is a cobounded subset of \mathcal{T} , because the vertices of X map to a single orbit and each edge of X maps to a geodesic of length $\leq \|\Phi\|$. Choose a compact set $\mathcal{B} \subset \mathcal{M}$ so that $\text{image}(\Phi)$ is \mathcal{B} -cobounded.

Using the method of Section 4.1, for each edge e of X we have a bundle $\mathcal{S}_e \rightarrow e$ equipped with a Riemannian metric. Given any vertex v of X , for any two edges e, e' incident to v the Riemannian metrics on \mathcal{S}_e and $\mathcal{S}_{e'}$ fit together isometrically at \mathcal{S}_v . We may therefore paste together the Riemannian metrics on \mathcal{S}_e for all edges e to obtain a marked hyperbolic surface bundle $\mathcal{S}_X \rightarrow X$ equipped with a piecewise Riemannian metric. The induced path metric on \mathcal{S}_X is a proper, geodesic metric. By naturality of the metrics on the bundles \mathcal{S}_e , the action of G on X lifts to an isometric action on \mathcal{S}_X .

By lifting the metric from \mathcal{S}_X to its universal cover \mathcal{H}_X we obtain a hyperbolic plane bundle $\mathcal{H}_X \rightarrow X$ on which the extension group Γ_G acts cocompactly, equipped with a Γ_G equivariant, piecewise Riemannian metric, inducing a proper, geodesic metric on \mathcal{H}_X . Note in particular that Γ_G is thus quasi-isometric to \mathcal{H}_X .

Note that this construction produces bundles $\mathcal{S}_X \rightarrow X$ and $\mathcal{H}_X \rightarrow X$ isomorphic to the pullback bundles described at the beginning of Section 4. Since each map $\mathcal{S}_e \rightarrow e$, $\mathcal{H}_e \rightarrow e$ satisfies the metric fibration property, it follows that the maps $\mathcal{S}_X \rightarrow X$, $\mathcal{H}_X \rightarrow X$ also satisfy that property.

The connections on the spaces \mathcal{S}_e , for edges e of X , piece together to define a G -equivariant connection on \mathcal{S}_X . To make sense out of this, we consider only the connection map defined for a piecewise path $\gamma: [a, b] \rightarrow X$, as follows. The bundle $\mathcal{S}_X \rightarrow X$ pulls back to give a bundle $\mathcal{S}_\gamma \rightarrow [a, b]$, and the connection

paths over each edge of X piece together to give connection paths on \mathcal{S}_γ , with an induced connection map $h_\gamma: \mathcal{S}_{\gamma(a)} \rightarrow \mathcal{S}_{\gamma(b)}$. It follows immediately from Lemma 4.1 that h_γ is $K^{\text{len}(\gamma)}$ -bilipschitz, where $K = K(\mathcal{B}, \|\Phi\|)$.

By lifting to \mathcal{H}_X , for each piecewise geodesic path $\gamma: [a, b] \rightarrow X$ we similarly obtain a $K^{\text{len}(\gamma)}$ bilipschitz connection map $\tilde{h}_\gamma: \mathcal{H}_{\gamma(a)} \rightarrow \mathcal{H}_{\gamma(b)}$.

4.3 Large scale geometry of surface bundles over paths

Our goal now is to compare metrics on \mathcal{H}_γ and \mathcal{H}_β for paths γ, β in \mathcal{T} which are closely related.

Given a metric space Z , two paths $\gamma, \beta: I \rightarrow Z$, and a constant $A \geq 0$, we say that γ, β are A -fellow travellers if $d(\gamma(t), \beta(t)) \leq A$ for all $t \in I$. More generally, given paths $\gamma: I \rightarrow Z$, $\beta: J \rightarrow Z$, a constant $A \geq 0$, and constants $\lambda \geq 1, \epsilon \geq 0$, we say that γ, β are asynchronous A -fellow travellers with respect to a λ, ϵ quasi-isometry $\phi: I \rightarrow J$ if the paths γ and $\beta \circ \phi$ are A -fellow travellers. It is a well known and simple fact that given a quasigeodesic $\gamma: I \rightarrow Z$ and another path $\beta: J \rightarrow Z$, the following are equivalent:

- (1) β is a quasigeodesic and β, γ have finite Hausdorff distance;
- (2) β is an asynchronous fellow traveller of γ .

Moreover, the constants are uniformly related: in $1 \implies 2$, there exist asynchronous fellow traveller constants A, λ, ϵ depending only on the quasigeodesic constants for β and the Hausdorff distance of β, γ ; in $2 \implies 1$, there exist quasigeodesic constants for β and a bound on the Hausdorff distance between β and γ depending only on the asynchronous fellow traveller constants.

The following proposition says that if $\gamma: I \rightarrow \mathcal{T}$, $\beta: J \rightarrow \mathcal{T}$ are asynchronous fellow travellers in \mathcal{T} , then there is a fiber preserving quasi-isometry $\mathcal{H}_\gamma \rightarrow \mathcal{H}_\beta$. Moreover, if γ is a geodesic, and if instead of \mathcal{H}_γ we use the singular SOLV space $\mathcal{H}_\gamma^{\text{SOLV}}$, then there is a fiber preserving quasi-isometry $\mathcal{H}_\gamma^{\text{SOLV}} \rightarrow \mathcal{H}_\beta$.

Proposition 4.2 *For each bounded subset $\mathcal{B} \subset \mathcal{M}$, and each $\rho \geq 1$, $\lambda \geq 1$, $\epsilon \geq 0$, $A \geq 0$, $K \geq 1$, there exists $K' \geq 1$, $C' \geq 0$ such that the following hold. Suppose that $\gamma: I \rightarrow \mathcal{T}$, $\beta: J \rightarrow \mathcal{T}$ are \mathcal{B} -cobounded ρ -Lipschitz, piecewise affine paths in \mathcal{T} . Suppose also that γ, β are asynchronous A -fellow travellers, with respect to a λ, ϵ quasi-isometry $\phi: I \rightarrow J$. Then:*

(1) There exists a commutative diagram

$$\begin{array}{ccc} \mathcal{S}_\gamma & \xrightarrow{\Phi} & \mathcal{S}_\beta \\ \downarrow & & \downarrow \\ I & \xrightarrow{\phi} & J \end{array}$$

such that the top row preserves markings, and such that any lifted map $\tilde{\Phi}: \mathcal{H}_\gamma \rightarrow \mathcal{H}_\beta$ is a K', C' quasi-isometry.

(2) If γ is a geodesic, then there exists a commutative diagram

$$\begin{array}{ccc} \mathcal{S}_\gamma^{\text{SOLV}} & \xrightarrow{\Phi} & \mathcal{S}_\beta \\ \downarrow & & \downarrow \\ I & \xrightarrow{\phi} & J \end{array}$$

such that the top row preserves markings, and such that any lifted map $\tilde{\Phi}: \mathcal{H}_\gamma^{\text{SOLV}} \rightarrow \mathcal{H}_\beta$ is a K', C' quasi-isometry.

One way to interpret item (1) of this proposition is that a cobounded, lipschitz path in Teichmüller space has a well-defined geometry associated to it: approximate the given path by a piecewise affine path and take the associated hyperbolic plane bundle; the metric on that bundle is well-defined up to quasi-isometry, independent of the approximation. A further argument shows that the geometry is independent of the choice of an equivariant connection on the bundle $\mathcal{S} \rightarrow \mathcal{T}$: any two equivariant connections are related in a uniformly bilipschitz manner over any cobounded subset of \mathcal{T} .

Proof Both (1) and (2) are proved in the same manner using Proposition 2.5; we prove only (1).

To smooth the notation in the proof we denote $t' = \phi(t)$, we let \mathcal{S}_t denote the fiber $\mathcal{S}_{\gamma(t)}$ of \mathcal{S}_γ , we let $\mathcal{S}'_{t'}$ denote the corresponding fiber $\mathcal{S}_{\beta(\phi(t'))}$ of \mathcal{S}_β , etc.

To prove (1), by applying Proposition 2.5(1) we choose for each $t \in \mathbf{R}$ a marked map $\Phi_t: \mathcal{S}_t \rightarrow \mathcal{S}'_{t'}$ for which any lift $\tilde{\Phi}_t: \mathcal{H}_t \rightarrow \mathcal{H}'_{t'}$ is a K_1, C_1 quasi-isometry, where the constants K_1, C_1 depend only on \mathcal{B}, A . Since each Φ_t preserves markings we may choose the lifts $\tilde{\Phi}_t$ so that for any s, t we have a commutative diagram of induced boundary maps:

$$\begin{array}{ccc} \partial\mathcal{H}_s & \xrightarrow{\partial\tilde{\Phi}_s} & \partial\mathcal{H}'_{s'} \\ \partial\tilde{h}_{st} \downarrow & & \downarrow \partial\tilde{h}'_{s't'} \\ \partial\mathcal{H}_t & \xrightarrow{\partial\tilde{\Phi}_t} & \partial\mathcal{H}'_{t'} \end{array}$$

Applying Proposition 2.5(2) it follows that if we strip off the ∂ symbols from the above diagram, and if we choose s, t so that $|s - t| \leq 1$, then we obtain the following diagram, a coarsely commutative diagram in the sense that the two paths around the diagram differ in the sup norm by a constant C_2 depending only on $\mathcal{B}, \rho, \lambda, \epsilon, A, K$:

$$\begin{array}{ccc}
 \mathcal{H}_s & \xrightarrow{\tilde{\Phi}_s} & \mathcal{H}'_{s'} \\
 \tilde{h}_{s,t} \downarrow & & \downarrow \tilde{h}_{s't'} \\
 \mathcal{H}_t & \xrightarrow{\tilde{\Phi}_t} & \mathcal{H}'_{t'}
 \end{array}$$

Define $\tilde{\Phi}: \mathcal{H}_\gamma \rightarrow \mathcal{H}_\beta$ so that $\tilde{\Phi}|_{\mathcal{H}_s} = \tilde{\Phi}_s$. To prove that $\tilde{\Phi}$ is a quasi-isometry we need only show that if $x, y \in \mathcal{H}_\gamma$ satisfy $d(x, y) \leq 1$ then $d(\tilde{\Phi}(x), \tilde{\Phi}(y))$ is bounded by a constant depending only on $\mathcal{B}, \rho, \lambda, \epsilon, A, K$, and then carry out the similar argument with inverses.

Given $x, y \in \mathcal{H}_\gamma$ with $d(x, y) \leq 1$, choose s, t so that $x \in \mathcal{H}_s, y \in \mathcal{H}_t$. By the metric fibration property we have $|s - t| \leq 1$. Changing notation if necessary we may assume that $s \leq t$. Let α be the geodesic in \mathcal{H}_γ connecting x and y , and by the metric fibration property note that $\alpha \subset \mathcal{H}_{[s-1, t+1]}$. Consider the map $p: \mathcal{H}_{[s-1, t+1]} \rightarrow \mathcal{H}_t$ whose restriction to \mathcal{H}_r is the connection map \tilde{h}_{rt} ; it follows that p is bilipschitz with constant $K^{t-s+2} \leq K^3$. The distance in \mathcal{H}_t between the point $p(x) = h_{st}(x)$ and the point y is therefore at most K^3 . Mapping over to \mathcal{H}_β we have

$$\begin{aligned}
 d(\tilde{\Phi}(x), \tilde{\Phi}(y)) &\leq d(\tilde{\Phi}(x), h_{s't'}(\tilde{\Phi}(x))) + d(h_{s't'}(\tilde{\Phi}(x)), \tilde{\Phi}(h_{st}(x))) \\
 &\quad + d(\tilde{\Phi}(h_{st}(x)), \tilde{\Phi}(y)) \\
 &\leq |s' - t'| + C_2 + (K_1 K^3 + C_1)
 \end{aligned}$$

and since $|s' - t'| \leq \lambda |s - t| + \epsilon \leq \lambda + \epsilon$, the proof is done. □

5 Hyperbolic extension implies convex cocompact quotient

In this section we prove Theorem 1.2.

Fix a homomorphism $f: G \rightarrow MCG$ defined on a finitely generated group G , and suppose that the extension group Γ_G is word hyperbolic. We must prove that f has finite kernel and that $f(G)$ is a convex cocompact subgroup of MCG .

Fix a Cayley graph X for G and an f -equivariant map $\Phi: X \rightarrow G$ which is affine on edges of X . Choose a bounded subset $\mathcal{B} \subset \mathcal{M}$ and a number $\rho \geq 1$ such that Φ is \mathcal{B} -cobounded and ρ -lipschitz. We have a hyperbolic plane bundle $\mathcal{H}_X \rightarrow X$, and an action of Γ_G on \mathcal{H}_X , such that the fibration $\mathcal{H}_X \rightarrow X$ is equivariant with respect to the homomorphism $\Gamma_G \rightarrow G$. We also have a piecewise Riemannian metric for which $\mathcal{H}_X \rightarrow X$ satisfies the metric fibration property. We also have a connection on \mathcal{H}_X , in the form of a connection map $h_\gamma: \mathcal{H}_{\gamma(a)} \rightarrow \mathcal{H}_{\gamma(b)}$ for any geodesic path $\gamma: [a, b] \rightarrow X$. The connection and metric are each equivariant with respect to Γ_G . Since \mathcal{H}_X is a proper geodesic metric space, it follows that \mathcal{H}_X is a model geometry for Γ_G . Since Γ_G is word hyperbolic, it follows that \mathcal{H}_X is δ -hyperbolic for some $\delta \geq 0$.

Fact 5.1 *For each point $x \in X$, the inclusion map $\mathcal{H}_x \hookrightarrow \mathcal{H}_X$ is uniformly proper, with uniform properness data independent of x .*

Proof This follows because the subgroup of Γ_G stabilizing \mathcal{H}_x is the normal subgroup $\pi_1(S)$, and the inclusion map $\pi_1(S) \hookrightarrow \Gamma_G$ is uniformly proper with respect to word metrics, a fact that holds for any finitely generated subgroup of a finitely generated group. \square

For each geodesic path $\gamma: I \rightarrow X$, I a closed, connected subset of \mathbf{R} , we obtain a piecewise affine path $\Phi \circ \gamma: I \rightarrow \mathcal{T}$ and a hyperbolic plane bundle $\mathcal{H}_\gamma \rightarrow I$, which can be regarded either as the pullback of the bundle $\mathcal{H} \rightarrow \mathcal{T}$ via $\Phi \circ \gamma$, or as the restriction of the bundle $\mathcal{H}_X \rightarrow X$ to γ . In either case, we obtain a piecewise Riemannian metric and connection on \mathcal{H}_γ , natural with respect to the action of $\pi_1(S)$. The connection on \mathcal{H}_γ has bilipschitz constant K depending only on \mathcal{B} and ρ , meaning that for any $s, t \in \mathbf{R}$, the connection map $h_{st}: \mathcal{H}_s \rightarrow \mathcal{H}_t$ is $K^{|s-t|}$ -bilipschitz.

Here is an outline of the proof of Theorem 1.2.

Our main task will be to prove that for each geodesic path $\gamma: I \rightarrow X$, the space \mathcal{H}_γ is a δ' -hyperbolic metric space, for some constant δ' depending only on \mathcal{B} , ρ , and δ . Of course, when I is a finite segment the space \mathcal{H}_γ is quasi-isometric to the hyperbolic plane and so \mathcal{H}_γ is a hyperbolic metric space, but uniformity of the hyperbolicity constant δ' is crucial. This is obtained using the concept of *flaring*, introduced by Bestvina and Feighn for their combination theorem [6], and further developed by Gersten in [18]. The combination theorem says, in an appropriate context, that flaring implies hyperbolicity. Gersten's converse, proved in the same context, says that hyperbolicity implies flaring. We shall

give a new technique for proving the converse, which applies in a much broader, “higher-dimensional” context, and using this technique we show that since \mathcal{H}_X is δ -hyperbolic it follows that each \mathcal{H}_γ satisfies flaring, with uniformity of constants. Then we shall apply the Bestvina–Feighn combination theorem in its original context to conclude that \mathcal{H}_γ is δ' -hyperbolic.

Next we will apply a result of Mosher [41] which says that since \mathcal{H}_γ is hyperbolic, the path $\Phi \circ \gamma: I \rightarrow \mathcal{T}$ is a quasigeodesic which is Hausdorff close to a Teichmüller geodesic, again with uniformity of constants. This will quickly imply finiteness of the kernel of f . The collection of these Teichmüller geodesics, one for each geodesic γ in X , will be used to verify the orbit quasiconvexity property for the group $f(G)$.

In what follows, a path $I \xrightarrow{\gamma} X$ will often be confused with the composed path $I \xrightarrow{\gamma} X \xrightarrow{\Phi} \mathcal{T}$; the context should make the meaning clear.

Remark The context of the Bestvina–Feighn combination theorem, and Gersten’s converse, is the following. Consider a finite graph of groups Γ , with word hyperbolic vertex and edge groups, such that each edge-to-vertex group injection is a quasi-isometric embedding. Associated to this is the Bass–Serre tree T , and a graph of spaces $X \rightarrow T$ on which $\pi_1\Gamma$ acts properly discontinuously and cocompactly. For each path in the tree T , Bestvina–Feighn define a flaring condition on the portion of X lying over that path. The combination theorem combined with Gersten’s converse says that flaring is satisfied uniformly over all paths in the Bass–Serre tree if and only if $\pi_1\Gamma$ is word hyperbolic. When G is a free group mapped to MCG then the extension $1 \rightarrow \pi_1 S \rightarrow \Gamma_G \rightarrow G \rightarrow 1$ fits into this context, because Γ_G is the fundamental group of a graph of groups with edge and vertex groups isomorphic to $\pi_1 S$, and with isomorphic edge-to-vertex injections, where the underlying graph is a rose with fundamental group G . This was the technique used in [40] to construct examples where Γ_G is word hyperbolic. When G is not free then this doesn’t work, motivating our “higher-dimensional” version of Gersten’s result.

5.1 Flaring

Motivated by the statement of the Bestvina–Feighn combination theorem, we make the following definitions.

Consider a sequence of positive real numbers $(r_j)_{j \in J}$, indexed by a subinterval J of \mathbf{Z} .

The L -lipschitz condition says that $r_i/r_j < L^{|i-j|}$ for all i, j , or equivalently $r_i/r_j < L$ whenever $|i - j| = 1$.

Given $\kappa > 1$, an integer $n \geq 1$, and $A \geq 0$, we say that (r_j) satisfies the (κ, n, A) -flaring property if, whenever the three integers $j - n, j, j + n$ are all in J , we have:

$$r_j > A \implies \text{Max}\{r_{j-n}, r_{j+n}\} \geq \kappa \cdot r_j$$

The number A is called the *flaring threshold*. Having a positive flaring threshold A allows the sequence to stay bounded by A on arbitrarily long intervals. However, at any place where the sequence has a value larger than A , exponential growth kicks in inexorably, in either the positive or the negative direction.

Consider a piecewise affine, cobounded, lipschitz path $\gamma: I \rightarrow \mathcal{T}$ and the corresponding hyperbolic plane bundle $\mathcal{H}_\gamma \rightarrow I$. A λ -quasivertical path in \mathcal{H}_γ is a λ -lipschitz path $\alpha: I' \rightarrow \mathcal{H}_\gamma$, defined on a subinterval $I' \subset I$, which is a section of the projection map $\mathcal{H}_\gamma \rightarrow I$. For example, a λ -quasivertical path is a connection path if and only if it is 1-quasivertical. Note that each λ -quasivertical path is a $(\lambda, 0)$ -quasigeodesic.

The *vertical flaring* property for the fibration $\mathcal{H}_\gamma \rightarrow \gamma$ says that there exists $\kappa > 1$, an integer $n \geq 1$, and a function $A(\lambda): [1, \infty) \rightarrow (0, \infty)$, such that if $\alpha, \beta: I \rightarrow \mathcal{H}_\gamma$ are two λ -quasivertical paths with the same domain I' , then setting $J = I' \cap \mathbf{Z}$ the sequence

$$d_j(\alpha(j), \beta(j)), \quad j \in J$$

satisfies the $\kappa, n, A(\lambda)$ flaring property, where d_j is the distance function on \mathcal{H}_j , $j \in J$. One can check that if the vertical flaring property holds for some function $A(\lambda)$ then it holds for a function which grows linearly.

Lemma 5.2 (Hyperbolicity of \mathcal{H}_X implies vertical flaring of \mathcal{H}_γ)

With notation as above, for every δ there exists $\kappa, n, A(\lambda)$ such that if \mathcal{H}_X is δ -hyperbolic then for each bi-infinite geodesic γ in X the fibration $\mathcal{H}_\gamma \rightarrow I$ satisfies $\kappa, n, A(\lambda)$ vertical flaring.

The intuition behind the proof is that the flaring property is exactly analogous to the geodesic divergence property in hyperbolic groups, described by Cannon in [12]. The geodesic divergence property says that in a δ -hyperbolic metric space, if p is a base point and if α, β are a pair of geodesic rays based at p , and if d_i is the shortest length of a path between $\alpha(i)$ and $\beta(i)$ that stays outside of the ball of radius i centered on p , then the sequence d_i satisfies a flaring property with constants independent of α, β . In our context, α and β will no

longer have one endpoint in common. But the quasivertical property together with the metric fibration property give us just what we need to adapt Cannon's proof of geodesic divergence given in [12], substituting the geodesic triangles in Cannon's proof with geodesic rectangles.

Proof We use d for the metric on \mathcal{H}_X .

First observe that any λ -quasivertical path α in \mathcal{H}_γ is a $(\lambda, 0)$ -quasigeodesic in \mathcal{H}_X , in fact

$$|s - t| \leq d(\alpha(s), \alpha(t)) \leq \lambda |s - t|$$

The upper bound is just the fact that α is λ -lipschitz, and the lower bound follows from the metric fibration property for $\mathcal{H}_X \rightarrow X$, together with the fact that γ is a geodesic in X .

Consider then a pair of λ -quasivertical paths $\alpha, \beta: I' \rightarrow \mathcal{H}_\gamma$ defined on a subinterval $I' \subset I$, and let $J = I' \cap Z = \{j_-, \dots, j_+\}$. We assume that $j_+ - j_-$ is even and let $j_0 = \frac{j_+ + j_-}{2} \in J$. For each $j \in J$ we have a fiber \mathcal{H}_j isometric to \mathbf{H}^2 , with metric denoted d_j . We must prove that the sequence $D_j = d_j(\alpha(j), \beta(j))$ satisfies κ, n, A flaring, with κ, n independent of λ and with κ, n, A independent of α, β , and γ .

For $j, k \in J$ let $h_{jk}: \mathcal{H}_j \rightarrow \mathcal{H}_k$ be the connection map, a $K^{|j-k|}$ bilipschitz map.

For each $j \in J$ we have an \mathcal{H}_j geodesic $\rho_j: [0, D_j] \rightarrow \mathcal{H}_j$ with endpoints $\alpha(j), \beta(j)$.

Claim 5.3 *There is a family of quasivertical paths v described as follows:*

- For each $j \in J$ and each $t \in [0, D_j]$ the family contains a unique quasivertical path $v_{jt}: [j_-, j_+] \rightarrow \mathcal{H}_\gamma$ that passes through the point $\rho_j(t)$. If we fix $j \in J$, we thus obtain a parameterization of the family v_{jt} by points $t \in [0, D_j]$.
- The ordering of the family v_{jt} induced by the order on $t \in [0, D_j]$ is independent of j . The first path v_{j0} in the family is identified with α , and the last path v_{jD_j} is identified with β .
- Each v_{jt} is λ' -quasivertical, where λ' depends only on λ and K .

When j is assumed fixed, we write v_t for the path v_{jt} .

Proof of claim Given $j-1, j \in J$, consider the following $(K, 0)$ -quasigeodesic in \mathcal{H}_j :

$$\rho'_j = h_{j-1,j} \circ \rho_{j-1}: [0, D_{j-1}] \rightarrow \mathcal{H}_j$$

Since connection paths are geodesics, and since α, β are λ -quasivertical, it follows that the endpoint $\rho'_j(0) = h_{j-1,j}(\alpha(j-1))$ and the corresponding endpoint $\rho_j(0) = \alpha(j)$ have distance in \mathcal{H}_X at most $\lambda + 1$, and similarly for the opposite endpoints $\rho'_j(D_{j-1}) = h_{j-1,j}(\beta(j-1))$ and $\rho_j(D_j) = \beta(j)$. Each endpoint of ρ'_j and the corresponding endpoint of ρ_j therefore have distance in \mathcal{H}_j bounded by a constant depending only on λ ; this follows from Fact 5.1. Since the spaces \mathcal{H}_j are all isometric to \mathbf{H}^2 , it follows that the Hausdorff distance between ρ_j and ρ'_j in \mathcal{H}_j is bounded by a constant depending only on K, λ , which implies in turn that there is a quasi-isometric reparameterization $r_j: [0, D_{j-1}] \rightarrow [0, D_j]$ such that

$$d_j(\rho'_j(t), \rho_j(r_j(t))) \leq D$$

where the constant D and the quasi-isometry constants for r_j depend only on K, λ . By possibly increasing the quasi-isometry constants we may assume furthermore that r_j is an orientation preserving homeomorphism. It follows that we may connect the point $\rho_{j-1}(t)$ to the point $\rho_j(r_j(t))$ by a λ' -quasivertical path defined over the interval $[j-1, j] \subset \mathbf{R}$, where λ' depends only on K, λ ; when $t = 0$ we may choose the path to be $\alpha \upharpoonright [j-1, j]$, and when $t = D_{j-1}$ we may choose the path $\beta \upharpoonright [j-1, j]$. By piecing together these paths as j varies over J , we obtain the required family of paths v . □

We use δ -hyperbolicity of \mathcal{H}_X in the following manner. First, for any geodesic rectangle $a * b * c * d$ in \mathcal{H}_X it follows that any point on a is within distance 2δ of $b \cup c \cup d$. Second, for any $(\lambda', 0)$ quasigeodesic in \mathcal{H}_X , the Hausdorff distance to any geodesic with the same endpoints is bounded by a constant δ_1 depending only on δ, λ' . For any rectangle of the form $v * \sigma * w * \sigma'$ where σ, σ' are geodesics and v, w are $(\lambda', 0)$ quasigeodesics, it follows that any point on v is within distance $\delta_2 = 2\delta + 2\delta_1$ of $\sigma \cup w \cup \sigma'$.

By Fact 5.1 there exists a constant δ_3 such that:

$$\text{for all } j \in J, x, y \in \mathcal{H}_j, \text{ if } d(x, y) \leq (1 + \lambda')\delta_2 \text{ then } d_j(x, y) \leq \delta_3$$

We are now ready to define the flaring parameters κ, n, A . Let

$$\begin{aligned} \kappa &= \frac{3}{2} \\ n &= \lfloor \delta_2 + 3\delta_3 \rfloor + 1 \\ A &= \delta_3 \end{aligned}$$

where $\lfloor x \rfloor$ is the greatest integer $\leq x$. Assuming as we may that $j_{\pm} = j_0 \pm n$ (and so the Hausdorff distance between \mathcal{H}_{j_0} and $\mathcal{H}_{j_{\pm}}$ in \mathcal{H}_X equals n), we must prove:

- if $D_{j_0} > A$ then $\max\{D_{j_-}, D_{j_+}\} \geq \kappa D_{j_0}$.

Case 1 $\max\{D_{j_-}, D_{j_+}\} \leq 6\delta_3$ It follows that there is a rectangle in \mathcal{H}_X of the form $\alpha * \sigma_- * \beta * \sigma_+$ where σ_{\pm} is a geodesic in \mathcal{H}_X with the same endpoints as $\rho_{j_{\pm}}$, and where σ_{\pm} has length $\leq 6\delta_3$. Consider now the point $\alpha(j_0)$, whose distance from some point $z \in \sigma_- \cup \beta \cup \sigma_+$ is at most δ_2 . If $z \in \sigma_-$ then it follows that

$$d(\alpha(j_0), \mathcal{H}_{j_+}) \leq \delta_2 + \frac{6\delta_3}{2} < n,$$

a contradiction. We reach a similar contradiction if $z \in \sigma_+$. Therefore $z \in \beta$. It follows that $z = \beta(s) \in \mathcal{H}_s$ for some s such that $|s - j_0| \leq \delta_2$, and so by following along β a length at most $\lambda'\delta_2$ we reach the point $\beta(j_0)$. This shows that $d(\alpha(j_0), \beta(j_0)) \leq (1 + \lambda')\delta_2$, and so $D_{j_0} \leq \delta_3$, that is, $D_{j_0} \leq A$.

Case 2 $\max\{D_{j_-}, D_{j_+}\} \geq 3\delta_3$ In the family v , we claim that there is a discrete subfamily $\alpha = v_{t_0}, v_{t_1}, \dots, v_{t_K} = \beta$, with $t_0 < t_1 < \dots < t_K$, such that the following property is satisfied: for each $k = 1, \dots, K$, letting

$$\Delta_{k\pm} = d_{j_{\pm}}(v_{t_{k-1}}(j_{\pm}), v_{t_k}(j_{\pm}))$$

then we have

$$\max\{\Delta_{k-}, \Delta_{k+}\} \in [3\delta_3, 6\delta_3].$$

By assumption of Case 2, the subfamily $\{\alpha = v_{t_0}, \beta = v_{t_1}\}$ has the property $\max\{\Delta_{k-}, \Delta_{k+}\} = \max\{D_{j_-}, D_{j_+}\} \geq 3\delta_3$ (for $k = 1$). Suppose by induction that we have a subfamily $\alpha = v_{t_0}, v_{t_1}, \dots, v_{t_K} = \beta$, with $t_0 < t_1 < \dots < t_K$, such that $\max\{\Delta_{k-}, \Delta_{k+}\} \geq 3\delta_3$ for all k , but suppose that $\max\{\Delta_{k-}, \Delta_{k+}\} > 6\delta_3$ for some k . If, say, $\Delta_{k+} > 6\delta_3$, then we subdivide the geodesic segment $\rho_{j_+}[v_{t_{k-1}}(j_+), v_{t_k}(j_+)]$ in half at a point $t \in \rho_{j_+}$, yielding two subsegments of length $> 3\delta_3$, and we add the path $v_{j_+,t}$ to our subfamily; similarly, if $\Delta_{k-} > 6\delta_3$ then we subdivide the interval $\rho_{j_-}[v_{t_{k-1}}(j_-), v_{t_k}(j_-)]$ in half. This process must eventually stop, because

$$K \leq \frac{1}{3\delta_3} (D_{j_-} + D_{j_+})$$

thereby proving the claim.

From the exact same argument as in Case 1, using the fact that

$$\max\{\Delta_{k-}, \Delta_{k+}\} \leq 6\delta_3,$$

it now follows that

$$\Delta_{k0} = d_{j_0}(v_{t_{k-1}}(j_0), v_{t_k}(j_0)) \leq \delta_3$$

for all $k = 1, \dots, K$.

We therefore have:

$$\begin{aligned} D_{j_0} &= \sum_{k=1}^K \Delta_{k0} \leq K\delta_3 \\ D_{j_-} + D_{j_+} &= \sum_{k=1}^K \Delta_{k-} + \Delta_{k+} \geq \sum_{k=1}^K \max\{\Delta_{k-}, \Delta_{k+}\} \\ &\geq K \cdot 3\delta_3 \\ \max\{D_{j_-}, D_{j_+}\} &\geq \frac{3}{2}K\delta_3 \\ &\geq \frac{3}{2}D_{j_0} \end{aligned}$$

This completes the proof of Lemma 5.2. \square

Remark The argument given in Lemma 5.2, while stated explicitly only for groups of the form Γ_G , generalizes to a much broader context. Graphs of groups, the context for the Bestvina–Feighn combination theorem [6] and Gersten’s converse [18], have been generalized to triangles of groups by Gersten and Stallings [46], and to general complexes of groups by Haefliger [20]. The arguments of Lemma 5.2 will also apply to show that a developable complex of groups with word hyperbolic fundamental group satisfies a flaring property over any geodesic in the universal covering complex. A converse would also be nice, giving a higher dimensional generalization of the Bestvina–Feighn combination theorem, but we do not know how to prove such a converse, nor do we have any examples to which it might apply (see Question 1.7 in the introduction).

Next we have:

Lemma 5.4 (Flaring implies hyperbolic) *For each bounded subset $\mathcal{B} \subset \mathcal{M}$, each $\rho \geq 1$, and each set of flaring data $\kappa > 1$, $n \geq 1$, $A(\lambda)$, there exists $\delta \geq 0$ such that the following holds. If $\gamma: I \rightarrow \mathcal{T}$ is a \mathcal{B} -cobounded, ρ -lipschitz, piecewise affine path defined on a subinterval $I \subset \mathbf{R}$, and if the metric fibration $\mathcal{H}_\gamma \rightarrow I$ satisfies $\kappa, n, A(\lambda)$ vertical flaring, then \mathcal{H}_γ is δ -hyperbolic.*

Proof This is basically an immediate application of the Bestvina–Feighn combination theorem [6]. To be formally correct, some remarks are needed to translate from our present geometric setting, of a hyperbolic plane bundle $\mathcal{H}_\gamma \rightarrow I$, to the combinatorial setting of [6], and to justify that our vertical flaring property for \mathcal{H}_γ corresponds to the “hallways flare condition” of [6].

We may assume that the endpoints of the interval I , if any, are integers.

The first observation is that there is a $\pi_1(S)$ –equivariant triangulation $\tilde{\tau}$ of \mathcal{H}_γ with the following properties:

Graph of spaces

- For each $n \in J = I \cap \mathbf{Z}$ there is a 2-dimensional subcomplex $\tilde{\tau}_n$ which is a triangulation of the hyperbolic plane \mathcal{H}_n .
- Each 1-cell of $\tilde{\tau}$ is either *horizontal* (a 1-cell of some τ_n), or *vertical* (connecting a vertex of some τ_n to a vertex of some τ_{n+1});
- each 2-cell of $\tilde{\tau}$ is either horizontal (a 2-cell of some $\tilde{\tau}_n$), or vertical (meaning that the boundary contains exactly two vertical 1-cells).

Bounded combinatorics There is an upper bound depending only on \mathcal{B} , ρ for the valence of each 0-cell and the number of sides of each 2-cell.

Quasi-isometry The inclusion of the 1-skeleton of $\tilde{\tau}$ into \mathcal{H}_γ is a quasi-isometry with constants depending only on \mathcal{B} and ρ .

To see why $\tilde{\tau}$ exists as described, consider the marked hyperbolic surface bundle $\mathcal{S}_\gamma \rightarrow I$. For each hyperbolic surface \mathcal{S}_n , $n \in J$, there is a geodesic triangulation τ_n of \mathcal{S}_n with one vertex, whose edges have length bounded only in terms of \mathcal{B} . It follows that there are constants K' , C' depending only on \mathcal{B} , such that if $\tilde{\tau}_n$ is the lifted triangulation in \mathcal{H}_n , then the inclusion of the 1-skeleton of $\tilde{\tau}_n$ into \mathcal{H}_n is a (K', C') quasi-isometry. Then, regarding $\bigcup_{n \in J} \tau_n$ as a triangulation of $\bigcup_{n \in J} \mathcal{S}_n$, we can extend to a cell-decomposition τ of \mathcal{S}_γ which is a graph of spaces of bounded combinatorics. The existence of τ uses the fact that each connection map $h_{n,n+1}: \mathcal{S}_n \rightarrow \mathcal{S}_{n+1}$ is K –bilipschitz, so by moving each vertex of τ_n along a connection path into \mathcal{S}_{n+1} and then moving a finite distance to a vertex of τ_{n+1} we obtain a (K'', C'') –quasi-isometry $h'_{n,n+1}: \tilde{\tau}_n \rightarrow \tilde{\tau}_{n+1}$, with (K'', C'') depending only on K , and from this we easily construct τ so that its lift $\tilde{\tau}$ has the desired properties.

The second observation is that vertical flaring in \mathcal{H}_γ is equivalent to the “hallway flare condition” of [6] for $\tilde{\tau}$, and this equivalence is uniform with respect to the parameters in each property. To see why, note that quasivertical paths in \mathcal{H}_γ correspond to *thin paths* in $\tilde{\tau}$ as defined implicitly in [6] Section 2: an edge

path $\alpha: I' = [m, n] \rightarrow \tilde{\tau}$ is ρ -thin if the restriction of α to each subinterval $[i, i+1]$ lies in $\tilde{\tau}_{[i, i+1]}$ and is a concatenation of at most ρ edges. Under the quasi-isometry $\tilde{\tau} \rightarrow \mathcal{H}_\gamma$ and its coarse inverse $\mathcal{H}_\gamma \rightarrow \tilde{\tau}$, λ -quasivertical paths in \mathcal{H}_γ correspond to ρ -thin paths with a uniform relation between λ and ρ .

In order to complete the translation from the geometric setting to the combinatorial setting, while the results of [6] are stated only when $\tilde{\tau}$ is the universal cover of a finite graph of spaces, nevertheless, the proofs hold as stated for any graph of spaces with uniformly bounded combinatorics: *all* the steps in the proof extend to such graphs of spaces, regardless of the presence of a deck transformation group with compact quotient. The conclusion of the combination theorem is the δ' -hyperbolicity of the 1-skeleton of $\tilde{\tau}$, with δ' depending only on the flaring constants for $\tilde{\tau}$, which depend in turn only on \mathcal{B} , ρ , and the flaring constants for \mathcal{H}_γ . It follows that \mathcal{H}_γ is δ hyperbolic with the correct dependency for the constant δ . \square

5.2 Proof of Theorem 1.2

We adopt the notation from the beginning of Section 5: a homomorphism $f: G \rightarrow MCG$ determining the group Γ_G , a Cayley graph X for G , and a piecewise affine f -equivariant map $\Phi: X \rightarrow \mathcal{T}$ which is \mathcal{B} -cobounded and ρ -lipschitz. We have already proved, in Section 1.2, that word hyperbolicity of Γ_G implies finiteness of the kernel of f .

Letting X^0 be the 0-skeleton, on which G acts transitively, it follows that $\Phi(X^0)$ is an orbit of $f(G)$ in \mathcal{T} . We prove that $f(G)$ is convex cocompact by proving that $\Phi(X^0)$ satisfies orbit quasiconvexity.

Choose two points $x, y \in X^0$. Let $\gamma: I \rightarrow X$ be a geodesic segment connecting x to y . Consider the composed path $I \xrightarrow{\gamma} X \xrightarrow{\Phi} \mathcal{T}$, which by abuse of notation we shall also denote γ . There is a corresponding hyperbolic plane bundle $\mathcal{H}_\gamma \rightarrow I$. Recall that γ is \mathcal{B} -cobounded and ρ -lipschitz in \mathcal{T} , with \mathcal{B} , ρ independent of γ . Now apply Lemmas 5.2 and 5.4, to conclude that \mathcal{H}_γ is δ -hyperbolic, with δ independent of γ .

Now we quote the following result to obtain a Teichmüller geodesic:

Theorem 5.5 [41] *For every bounded set $\mathcal{B} \subset \mathcal{M}$, $\rho \geq 1$, and $\delta \geq 0$, there exists $\lambda \geq 1$, $\epsilon > 0$, and A such that the following hold. If $\gamma: I \rightarrow \mathcal{T}$ is \mathcal{B} -cobounded and ρ -lipschitz, and if \mathcal{H}_γ is δ -hyperbolic, then γ is a (λ, ϵ) -quasigeodesic, and there exists a Teichmüller geodesic g , sharing any endpoints of γ , such that γ and g have Hausdorff distance at most A . \square*

Letting g be the Teichmüller geodesic connecting x to y provided by the theorem, it follows that g is contained in the $A + \rho$ neighborhood of $\Phi(X^0)$. Since $x, y \in \Phi(X^0)$ are arbitrary, this proves orbit quasiconvexity, and so $f(G)$ is convex cocompact.

6 Schottky groups

Definition A *Schottky subgroup* of MCG is a free, convex cocompact subgroup.

The limit set $\Lambda \subset \mathbf{PMF}$ of a Schottky subgroup is therefore a Cantor set, and every nontrivial element is pseudo-Anosov.

In this section we prove Theorem 1.3, that a surface-by-free group is word hyperbolic if and only if the free group is Schottky. One direction is already proved by Theorem 1.2, and so we need only prove that when $F \subset MCG$ is a Schottky subgroup then $\Gamma_F \approx \pi_1(S) \rtimes F$ is word hyperbolic.

Continuing with earlier notation, let $\Lambda \subset \mathbf{PMF}$ be the limit set of F with weak hull WH_Λ . Let \mathfrak{t} be a Cayley graph for the group F , a tree on which F acts properly discontinuously with quotient a rose. Let $\Phi: \mathfrak{t} \rightarrow \mathcal{T}$ be an F -equivariant map, affine on each edge, and ρ -lipschitz for some $\rho \geq 1$. There is a bounded subset $\mathcal{B} \subset \mathcal{M}$ so that both WH_Λ and $\Phi(\mathfrak{t})$ are \mathcal{B} -cobounded. We have a hyperbolic plane bundle $\mathcal{H}_\mathfrak{t} \rightarrow \mathfrak{t}$, on which $\pi_1(S) \rtimes F$ acts properly discontinuously and cocompactly, and we have a piecewise Riemannian metric on $\mathcal{H}_\mathfrak{t}$ on which $\pi_1(S) \rtimes F$ acts by isometries.

We must prove that $\mathcal{H}_\mathfrak{t}$ is δ -hyperbolic. By the Bestvina–Feighn combination theorem [6], it is enough to show that for each bi-infinite geodesic γ in \mathfrak{t} , the bundle $\mathcal{H}_\gamma \rightarrow \mathbf{R}$ satisfies vertical flaring, with flaring data $\kappa, n, A(\lambda)$ independent of the choice of γ (see the proof of Lemma 5.4 for translating the combinatorial setting of [6] to our present geometric setting).

Since F is convex cocompact, there is a geodesic line g in WH_Λ which has finite Hausdorff distance from $\Phi(\gamma)$. Let $\mathcal{H}_g^{\text{SOLV}}$ be the singular SOLV-space thereby obtained. By Proposition 4.2, the closest point map $\gamma \rightarrow g$ lifts to a quasi-isometry $\mathcal{H}_\gamma \rightarrow \mathcal{H}_g^{\text{SOLV}}$, with quasi-isometry constants independent of γ , depending only on \mathcal{B} and ρ . It therefore suffices to check the flaring condition in $\mathcal{H}_g^{\text{SOLV}}$, with flaring data independent of anything.

Take any κ with $1 < \kappa < \frac{\epsilon^2}{2\sqrt{2}}$, say $\kappa = 2.6$. Let $n = 2$. We show that for any λ there is an A such that any two λ quasivertical lines in $\mathcal{H}_g^{\text{SOLV}}$ satisfy

the $(\kappa, 2, A)$ -flaring condition. For this argument we do not need that g is cobounded (although in that case $\mathcal{H}_g^{\text{SOLV}}$ may not have bounded geometry).

Let $\alpha, \alpha': [-2, 2] \rightarrow \mathcal{H}_g^{\text{SOLV}}$ be two λ quasivertical lines, lying over a length 4 subsegment $[r - 2, r + 2]$ of $g \approx \mathbf{R}$. Let x_i, y_i be the points where α, α' respectively intersect \mathcal{H}_{r+i} . Let $\xi_0 = x_0$ and let ξ_i be obtained by flowing x_0 vertically into \mathcal{H}_{r+i} ; define $\eta_0 = y_0$ and η_i similarly. Note that for $i \in [-2, 2]$ the points ξ_i and x_i are connected in $\mathcal{H}_g^{\text{SOLV}}$ by a path which goes along α from ξ_i to ξ_0 travelling a distance at most 2λ , and then vertically from $\xi_0 = x_0$ to x_i ; the vertical projection of this path into \mathcal{H}_i has length at most $2e^2\lambda$, and so $d_i(x_i, \xi_i) \leq 2e^2\lambda$. Similarly, $d_i(y_i, \eta_i) \leq 2e^2\lambda$.

We turn for the moment to showing that the sequence

$$d_{r+i}(\xi_i, \eta_i), \quad i = -2, -1, 0, 1, 2$$

satisfies the $(\frac{e^2}{2\sqrt{2}}, 2, 0)$ -flaring condition. In the singular Euclidean surface \mathcal{H}_{r+i} , let ℓ_i be the geodesic from ξ_i to η_i , so the above sequence becomes:

$$\text{len}(\ell_i), \quad i = -2, -1, 0, 1, 2$$

The singular Euclidean geodesic ℓ_0 is a concatenation of subsegments of constant slope, two consecutive subsegments meeting at a singularity. If at least half of ℓ_0 has slope of absolute value ≥ 1 then:

$$\frac{1}{2} \text{len}(\ell_0) \cdot \frac{1}{\sqrt{2}} \cdot e^2 \leq \text{len}(\ell_2)$$

If at least half of ℓ_0 has slope of absolute value ≤ 1 , we get a similar inequality but with $\text{len}(\ell_{-2})$ on the right hand side. We have therefore shown:

$$\max\{d_{r+2}(\xi_2, \eta_2), d_{r-2}(\xi_{-2}, \eta_{-2})\} \geq \frac{e^2}{2\sqrt{2}} d_0(\xi_0, \eta_0)$$

It follows that

$$\begin{aligned} \max\{d_{r+2}(x_2, y_2), d_{r-2}(x_{-2}, y_{-2})\} &\geq \frac{e^2}{2\sqrt{2}} d_0(x_0, y_0) - 2e^2\lambda \\ &\geq \kappa d_0(x_0, y_0) \end{aligned}$$

where the last inequality holds as long as:

$$d_0(x_0, y_0) \geq A = \frac{2e^2\lambda}{\frac{e^2}{2\sqrt{2}} - \kappa}$$

This ends the proof that $\pi_1(S) \rtimes F$ is word hyperbolic when F is Schottky.

7 Extending the theory to orbifolds

In this section we sketch how the theory can be extended to 2-dimensional orbifolds. We shall consider only those compact orbifolds whose underlying 2-manifold is closed, and whose orbifold locus therefore consists only of cone points, what we shall call a *cone orbifold*. The reason for this restriction is that if the underlying 2-manifold has nonempty boundary then the orbifold does not support any pseudo-Anosov homeomorphisms, since the isotopy classes of the boundary curves must be permuted.⁴

As it turns out, the mapping class group and Teichmüller space of a cone orbifold depend not on the actual orders of the different cone points, but only on the partition of the set of cone points into subsets of constant order. For example, a spherical orbifold with one $\mathbf{Z}/2$ cone point and three $\mathbf{Z}/4$ cone points has the same mapping class group and Teichmüller space as a spherical orbifold with three $\mathbf{Z}/42$ cone points and one $\mathbf{Z}/1000$ cone point. The relevant structures can therefore be described more directly and economically in the following manner.

Let S be a closed surface, not necessarily orientable. Let $\mathbf{P} = \{P_i\}_{i \in I}$ be a finite, pairwise disjoint collection of finite, nonempty subsets of S . Let $\text{Homeo}(S, \mathbf{P})$ be the group of homeomorphisms of S which leave invariant each of the sets P_i , $i \in I$. Let $\text{Homeo}_0(S, \mathbf{P})$ be the component of the identity of $\text{Homeo}(S, \mathbf{P})$ with respect to the compact open topology; equivalently, $\text{Homeo}_0(S, \mathbf{P})$ consists of all elements of $\text{Homeo}(S, \mathbf{P})$ which are isotopic to the identity through elements of $\text{Homeo}(S, \mathbf{P})$. The mapping class group is $MCG(S, \mathbf{P}) = \text{Homeo}(S, \mathbf{P}) / \text{Homeo}_0(S, \mathbf{P})$.

To define the Teichmüller space, first we must widen the concept of a conformal structure so that it applies to non-orientable surfaces, and we do this by allowing overlap maps which are anticonformal as well as conformal. The Teichmüller space $\mathcal{T}(S, \mathbf{P})$ is then defined to be the set of conformal structures on S modulo the action of $\text{Homeo}_0(S, \mathbf{P})$. Quadratic differentials and measured foliations on (S, \mathbf{P}) are defined using the usual local models at points of $S - \cup \mathbf{P}$, but at a point of \mathbf{P} a quadratic differential can have the local model $z^{n-2}dz^2$ for any $n \geq 1$; the horizontal measured foliation of $z^{n-2}dz^2$ is the local model for an n -pronged singularity of a measured foliation. Thus, at a point of $\cup \mathbf{P}$ a measured foliation can have any number of prongs ≥ 1 , whereas a singularity in $S - \cup \mathbf{P}$ must have ≥ 3 prongs as usual. With these definitions, Teichmüller maps are defined as usual, making $\mathcal{T}(S, \mathbf{P})$ into a proper geodesic metric space

⁴While the monograph [16] develops a kind of pseudo-Anosov theory on a bounded surface, it is *not* appropriate for our present purposes.

on which $MCG(S, \mathbf{P})$ acts properly discontinuously, but not cocompactly; also, pseudo-Anosov homeomorphisms of (S, \mathbf{P}) are defined as usual.

We shall assume that (S, \mathbf{P}) actually supports a pseudo-Anosov homeomorphism which has an n -pronged singularity with $n \neq 2$. This rules out a small number of special cases, as follows. When S is a sphere, $\cup \mathbf{P}$ must have at least four points. When S is a projective plane, $\cup \mathbf{P}$ must have at least two points. When S is a torus or Klein bottle, $\cup \mathbf{P}$ must have at least one point. When S is the surface of Euler characteristic -1 , namely the connected sum of a torus and a projective plane, the curve along which the torus and the projective plane are glued is actually a characteristic curve for S , meaning that it is preserved up to isotopy by any mapping class; therefore, in order for (S, \mathbf{P}) to support a pseudo-Anosov homeomorphism, $\cup \mathbf{P}$ must have at least one point.

Now we apply these concepts to 2-dimensional cone orbifolds. Suppose \mathcal{O} is a cone orbifold with underlying surface S . Let P_n be the set of \mathbf{Z}/n cone points, and let $\mathbf{P} = \{P_n\}_{n \geq 2}$. Then we may define the mapping class group $MCG(\mathcal{O})$ to be $MCG(S, \mathbf{P})$, and the Teichmüller space $\mathcal{T}(\mathcal{O})$ to be $\mathcal{T}(S, \mathbf{P})$. Note that with the restrictions above on the type of (S, \mathbf{P}) , the orbifold \mathcal{O} has negative Euler characteristic. It follows that if $\tilde{\mathcal{O}} \rightarrow \mathcal{O}$ is the orbifold universal covering map, then for any conformal structure on \mathcal{O} the lifted conformal structure is isomorphic to the Riemann disc. It follows that any conformal structure on \mathcal{O} can be uniquely uniformized to produce a hyperbolic structure, with a cone angle of $2\pi/n$ at each \mathbf{Z}/n cone point.

At this stage we must confront the fact that the universal extension for surface groups, as formulated in Section 1.2, must be reformulated before it can be applied to orbifolds. The Dehn–Nielsen–Baer–Epstein theorem is still true, as long as one uses orbifold fundamental groups: if p is a generic point of the cone orbifold \mathcal{O} , and if $\pi_1(\mathcal{O}, p)$ is the orbifold fundamental group, then we have $MCG(\mathcal{O}) \approx \text{Out}(\pi_1(\mathcal{O}, p))$. However, the “once-punctured” mapping class group $MCG(\mathcal{O}, p)$ is *not* isomorphic to $\text{Aut}(\pi_1(\mathcal{O}, p))$. For example, take a based simple loop ℓ which bounds a disc whose interior contains a single \mathbf{Z}/n cone point. In the group $\pi_1(\mathcal{O}, p)$, the loop ℓ represents an element of order n , and under the usual injection $\pi_1(\mathcal{O}, p) \hookrightarrow \text{Aut}(\pi_1(\mathcal{O}, p))$ we obtain an element of order n . However, the element of $MCG(\mathcal{O}, p)$ obtained by pushing p around ℓ has infinite order in $MCG(\mathcal{O}, p)$.

To repair this we need another group to take over the role of $MCG(\mathcal{O}, p)$. Let $\widetilde{\text{Homeo}}(\mathcal{O})$ denote the group of homeomorphisms of $\tilde{\mathcal{O}}$ which are lifts of homeomorphisms of \mathcal{O} , that is, a homeomorphism $\tilde{f}: \tilde{\mathcal{O}} \rightarrow \tilde{\mathcal{O}}$ is in the group $\widetilde{\text{Homeo}}(\mathcal{O})$ if and only if there exists a homeomorphism $f: \mathcal{O} \rightarrow \mathcal{O}$ such that

the following diagram commutes:

$$\begin{array}{ccc}
 \tilde{\mathcal{O}} & \xrightarrow{\tilde{f}} & \tilde{\mathcal{O}} \\
 \downarrow & & \downarrow \\
 \mathcal{O} & \xrightarrow{f} & \mathcal{O}
 \end{array}$$

With respect to the compact open topology, $\widetilde{\text{Homeo}}(\mathcal{O})$ becomes a topological group. Let $\widetilde{\text{Homeo}}_0(\mathcal{O})$ be the component of the identity $\widetilde{\text{Homeo}}(\mathcal{O})$. Equivalently, $\widetilde{\text{Homeo}}_0(\mathcal{O})$ is the subgroup of elements of $\widetilde{\text{Homeo}}(\mathcal{O})$ isotopic to the identity through elements of $\widetilde{\text{Homeo}}(\mathcal{O})$; alternatively it is the subgroup of $\widetilde{\text{Homeo}}(\mathcal{O})$ acting trivially on the circle at infinity of $\tilde{\mathcal{O}} \approx \mathbf{H}^2$. Define

$$\widetilde{MCG}(\mathcal{O}) = \widetilde{\text{Homeo}}(\mathcal{O}) / \widetilde{\text{Homeo}}_0(\mathcal{O}).$$

Note that universal covering map $\tilde{\mathcal{O}} \rightarrow \mathcal{O}$ induces a surjective homomorphism $\widetilde{MCG}(\mathcal{O}) \rightarrow MCG(\mathcal{O})$, and the kernel is the group of deck transformations, isomorphic to $\pi_1(\mathcal{O})$. We now have a natural isomorphism of short exact sequences

$$\begin{array}{ccccccc}
 1 & \longrightarrow & \pi_1(\mathcal{O}) & \longrightarrow & \widetilde{MCG}(\mathcal{O}) & \longrightarrow & MCG(\mathcal{O}) \longrightarrow 1 \\
 & & \parallel & & \Downarrow & & \Downarrow \\
 1 & \longrightarrow & \pi_1(\mathcal{O}) & \longrightarrow & \text{Aut}(\pi_1(\mathcal{O})) & \longrightarrow & \text{Out}(\pi_1(\mathcal{O})) \longrightarrow 1
 \end{array}$$

where we have suppressed the generic base point needed to define $\pi_1(\mathcal{O})$.

We are now in a position to state that our main results, Theorem 1.1, 1.2, 1.3, and 1.4, are true with the orbifold \mathcal{O} in place of the surface S , and the proofs are unchanged. Although the references that we quote are stated solely in terms of surfaces, namely [38] and [32] for Theorem 1.1, [39] for Theorem 1.2, and [40] for Theorem 1.4, nevertheless all the proofs in those references work just as well for orbifolds instead of surfaces.

References

- [1] **W Abikoff**, *The real analytic theory of Teichmüller space*, volume 820 of *Lecture Notes in Mathematics*, Springer (1980)
- [2] **G N Arzhantseva**, *On quasiconvex subgroups of word hyperbolic groups*, *Geometriae Dedicata* 98 (2001) 191–208

- [3] **R Baer**, *Isotopien von Kurven auf orientierbaren, geschlossenen Flächen und ihr Zusammenhang mit der topologischen Deformation der Flächen*, J. Reine Angew. Math. 159 (1928) 101–116
- [4] **L Bers**, *Fiber spaces over Teichmüller spaces*, Acta Math. 130 (1973) 89–126
- [5] **L Bers**, *An extremal problem for quasiconformal mappings and a theorem by Thurston*, Acta Math. 141 (1978) 73–98
- [6] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, J. Diff. Geom. 35 (1992) 85–101
- [7] **J Birman**, *Braids, links, and mapping class groups*, volume 82 of *Annals of Math. Studies*, Princeton University Press (1974)
- [8] **J Birman**, **A Lubotzky**, **J McCarthy**, *Abelian and solvable subgroups of the mapping class groups*, Duke Math. J. 50 (1983) 1107–1120
- [9] **N Brady**, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, J. London Math. Soc. (2) 60 (1999) 461–480
- [10] **J Brock**, **B Farb**, *Curvature and rank of Teichmüller space* (2001), preprint, [arXiv:math.GT/0109045](https://arxiv.org/abs/math.GT/0109045)
- [11] **R D Canary**, *Covering theorems for hyperbolic 3-manifolds*, preprint
- [12] **J Cannon**, *The theory of negatively curved spaces and groups*, from: “Ergodic theory, symbolic dynamics, and hyperbolic spaces”, (C Series T Bedford, M Keane, editor), Oxford Univ. Press (1991)
- [13] **D B A Epstein**, *Curves on 2-manifolds and isotopies*, Acta Math. 115 (1966) 83–107
- [14] **B Farb**, **A Lubotzky**, **Y Minsky**, *Rank one phenomena for mapping class groups*, Duke Math. J. 106 (2001) 581–597
- [15] **B Farb**, **L Mosher**, *The geometry of surface-by-free groups*, Geom. Funct. Anal. (2002), to appear, preprint, [arXiv:math.GR/0008215](https://arxiv.org/abs/math.GR/0008215)
- [16] **A Fathi**, **F Laudenbach**, **V Poenaru**, et al., *Travaux de Thurston sur les surfaces*, volume 66–67 of *Astérisque*, Société Mathématique de France (1979)
- [17] **F Gardiner**, **H Masur**, *Extremal length geometry of Teichmüller space*, Complex Variables Theory Appl. 16 (1991) 209–237
- [18] **S Gersten**, *Cohomological lower bounds for isoperimetric functions on groups*, Topology 37 (1998) 1031–1072
- [19] **G Gonzalez-Diez**, **W J Harvey**, *Surface subgroups inside mapping class groups*, Topology 38 (1999) 57–69
- [20] **A Haefliger**, *Complexes of groups and orbihedra*, from: “Group theory from a geometrical viewpoint (Trieste, 1990)”, World Sci. Publishing, River Edge, NJ (1991) 504–540
- [21] **J Hubbard**, **H Masur**, *Quadratic differentials and foliations*, Acta Math. 142 (1979) 221–274

- [22] **Y Imayoshi, M Taniguchi**, *An introduction to Teichmüller spaces*, Springer (1992)
- [23] **N V Ivanov**, *Subgroups of Teichmüller modular groups*, volume 115 of *Translations of Mathematical Monographs*, Amer. Math. Soc. (1992)
- [24] **N V Ivanov**, *Automorphisms of complexes of curves and Teichmüller spaces*, from: “Progress in knot theory and related topics”, *Travaux en Cours* 56, Hermann, Paris (1997) 113–120
- [25] **M Kapovich**, *On normal subgroups in the fundamental groups of complex surfaces* (1998), preprint, [arXiv:math.GT/9808085](https://arxiv.org/abs/math/9808085)
- [26] **S Kerckhoff**, *The asymptotic geometry of Teichmüller space*, *Topology* 19 (1980) 23–41
- [27] **B Maskit**, *Comparison of hyperbolic and extremal lengths*, *Ann. Acad. Sci. Fenn. Series A I Math.* 10 (1985) 381–386
- [28] **H Masur**, *Uniquely ergodic quadratic differentials*, *Comment. Math. Helv.* 55 (1980) 255–266
- [29] **H Masur**, *Interval exchange transformations and measured foliations*, *Ann. of Math.* 115 (1982) 169–200
- [30] **H Masur**, *Two boundaries of Teichmüller space*, *Duke Math. J.* 49 (1982) 183–190
- [31] **H Masur**, *Hausdorff dimension of the set of nonergodic foliations of a quadratic differential*, *Duke Math. J.* 66 (1992) 387–442
- [32] **H Masur, Y Minsky**, *Geometry of the complex of curves, I. Hyperbolicity*, *Invent. Math.* 138 (1999) 103–149
- [33] **H Masur, Y Minsky**, *Unstable quasigeodesics in Teichmüller space*, from: “In the tradition of Ahlfors and Bers (Stony Brook, NY, 1998)”, *Contemp. Math.* 256, Amer. Math. Soc. (2000) 239–241
- [34] **H Masur, M Wolf**, *Teichmüller space is not Gromov hyperbolic*, *Ann. Acad. Sci. Fenn. Ser. A I Math.* 20 (1995) 259–267
- [35] **J McCarthy**, *A “Tits-alternative” for subgroups of surface mapping class groups*, *Trans. AMS* 291 (1985) 582–612
- [36] **J McCarthy, A Papadopoulos**, *Dynamics on Thurston’s sphere of projective measured foliations*, *Comment. Math. Helv.* 64 (1989) 133–166
- [37] **Y Minsky**, *On rigidity, limit sets, and end invariants of hyperbolic 3-manifolds*, *Jour. AMS* 7 (1994) 539—588
- [38] **Y Minsky**, *Quasi-projections in Teichmüller space*, *J. Reine Angew. Math.* 473 (1996) 121–136
- [39] **L Mosher**, *Hyperbolic extensions of groups*, *J. Pure and Appl. Alg.* 110 (1996) 305–314

- [40] **L Mosher**, *A hyperbolic-by-hyperbolic hyperbolic group*, Proc. AMS 125 (1997) 3447–3455
- [41] **L Mosher**, *Stable quasigeodesics in Teichmüller space and ending laminations* (2001), preprint, [arXiv:math.GR/0107035](https://arxiv.org/abs/math/0107035)
- [42] **D Mumford**, *A remark on Mahler’s compactness theorem*, Proc. Amer. Math. Soc. 28 (1971) 289–294
- [43] **J Nielsen**, *Untersuchungen zur Topologie der geschlossenen zweiseitigen Flächen*, Acta Math. 50 (1927) 189–358
- [44] **J-P Otal**, *Le théorème d’hyperbolisation pour les variétés fibrées de dimension 3*, Astérisque 235, Société Mathématique de France (1996)
- [45] **H Royden**, *Report on the Teichmüller metric*, Proc. Nat. Acad. Sci. U.S.A. 65 (1970) 497–499
- [46] **J Stallings**, *Non-positively curved triangles of groups*, from: “Group theory from a geometrical viewpoint (Trieste, 1990)”, World Sci. Publishing (1991) 491–503
- [47] **K Whittlesey**, *Normal all pseudo-Anosov subgroups of mapping class groups*, Geometry and Topology 4 (2000) 293–307