

EXTENSIONS OF VEECH GROUPS I: A HYPERBOLIC ACTION

SPENCER DOWDALL, MATTHEW G. DURHAM, CHRISTOPHER J. LEININGER,
AND ALESSANDRO SISTO

ABSTRACT. Given a lattice Veech group in the mapping class group of a closed surface S , this paper investigates the geometry of Γ , the associated $\pi_1 S$ -extension group. We prove that Γ is the fundamental group of a bundle with a singular Euclidean-by-hyperbolic geometry. Our main result is that collapsing “obvious” product regions of the universal cover produces an action of Γ on a hyperbolic space, retaining most of the geometry of Γ . This action is a key ingredient in the sequel where we show that Γ is hierarchically hyperbolic and quasi-isometrically rigid.

1. INTRODUCTION

Let S be a closed, connected, oriented surface of genus at least 2. For any group G , a $\pi_1 S$ -extension of G is a group Γ fitting into a short exact sequence:

$$1 \rightarrow \pi_1 S \rightarrow \Gamma \rightarrow G \rightarrow 1.$$

Such a group Γ is the fundamental group of an S -bundle and is determined up to isomorphism by its monodromy homomorphism $G \rightarrow \text{Mod}^\pm(S) \cong \text{Out}(\pi_1 S)$ to the extended mapping class group; see e.g. [Mor01, ST20]. An important problem is to decide how geometric properties of Γ (or the bundle) are reflected in properties of the monodromy.

A motivating success story for this line of inquiry is Thurston’s geometrization theorem (see e.g. [Ota98]), stating that when $G \cong \mathbb{Z}$, the associated surface bundle over the circle admits a hyperbolic structure if and only if the image of the monodromy is generated by a pseudo-Anosov mapping class. A generalization of this to the world of coarse geometry states that a $\pi_1 S$ -extension of any group G is *word-hyperbolic* if and only if the monodromy has finite kernel and its image in $\text{Mod}(S)$ is *convex cocompact* in the sense of Farb and Mosher (see [FM02, Ham, MS12]).

In that spirit, this paper initiates an analysis of the geometry of extension groups of *lattice Veech subgroups* of the mapping class group. There are two important perspectives motivating this analysis. Firstly, in the classical setting of Kleinian groups, convex cocompactness is a special case of *geometric finiteness*. While there is as yet no precise analogue of this latter notion in the context of mapping class groups—nor any proposal for how such a notion should relate, as in the case of convex cocompactness, to the (coarse) geometry of associated extension groups—it is evident that lattice Veech groups should qualify as geometrically finite with respect to *any* definition (c.f. Mosher [Mos06, §6]). Taking these as prototypes for geometric finiteness suggests that their extension groups should model compelling geometric features. Our main theorem shows this is indeed the case, in that the extension becomes hyperbolic after collapsing naturally occurring *vertex subgroups* that are virtually isomorphic to fundamental groups of Seifert manifolds; see §3.6.

Theorem 1.1. *Suppose $G < \text{Mod}(S)$ is a lattice Veech group with extension group Γ and let $\Upsilon_1, \dots, \Upsilon_k < \Gamma$ be representatives of the conjugacy classes of vertex subgroups. Then Γ admits an isometric action on a hyperbolic space \hat{E} , quasi-isometric to the Cayley graph of Γ coned off along the cosets of $\Upsilon_1, \dots, \Upsilon_k$. Furthermore, every element not conjugate into one of $\Upsilon_1, \dots, \Upsilon_k$ acts loxodromically on \hat{E} .*

The hyperbolic space \hat{E} is constructed from bundle E on which Γ acts nicely as described below (and in particular, it has more structure than the coned-off Cayley graph); see Definition 3.1.

In the sequel [DDLS21], we further analyze how the coned-off cosets interact with each other and prove that Γ admits a hierarchically hyperbolic structure. In fact, we show that this structure is Γ invariant, thus proving the following.

Theorem 1.2 ([DDLS21]). *For any lattice Veech group $G < \text{Mod}(S)$, the extension group Γ of G admits the structure of a hierarchically hyperbolic group for which the maximal hyperbolic space is \hat{E} .*

We note that Theorem 1.1 is a key step in the proof of Theorem 1.2 and refer the reader to our second paper [DDLS21] for a detailed description of the hierarchically hyperbolic group structure, a general discussion of geometric finiteness in mapping class groups, and several applications of Theorems 1.1–1.2. One such application, which we mention here to highlight the significance of \hat{E} , is that the action of the extension group Γ on \hat{E} is a universal acylindrical action. We also remark that while the characterization of loxodromics in Theorem 1.1 can formally be deduced from Theorem 1.2, we also give a direct elementary proof here in §4.9.

A second motivation for studying extensions of lattice Veech groups is that their associated S -bundles naturally carry a 4-dimensional singular Euclidean-by-hyperbolic geometry in the sense of Thurston [Thu97]. This is similar to the fact, which has been used to great effect (e.g., [Bow13, CT07, Koz16, Min94, Mos03]) in the study of surface bundles and 3-manifolds, that hyperbolic plane bundles over Teichmüller geodesics carry singular Sol geometries. Indeed, as explained in §5.4, the 4-dimensional geometry exhibited here can be viewed as a direct generalization of this 3-dimensional Sol phenomenon.

This perspective is central to our entire approach. The proof of Theorem 1.1 involves the construction of a bundle E on which the extension group Γ of a Veech group G acts geometrically by bundle automorphisms, together with a detailed analysis of the geometry of E . The bundle E is derived from the two interpretations of G as simultaneously the affine group of a Euclidean cone metric on S and as the stabilizer of an isometrically embedded hyperbolic plane in Teichmüller space (the associated Teichmüller disk); see §2.9 and §3.2 for definitions. Away from a codimension 2 singular locus, E is modeled on the homogeneous Euclidean-by-hyperbolic geometry mentioned above; in §5 we consider this geometry more closely and prove the following:

Theorem 1.3. *Suppose Γ is the extension group of a Veech group and E/Γ is the associated bundle. Away from the singular locus, E/Γ is an $(\text{SAff}^\pm(\mathbb{R}^2), \mathfrak{X})$ -manifold, where $\mathfrak{X} = \text{SAff}^\pm(\mathbb{R}^2)/K$ and $K \cong \text{O}(2)$ is a maximal compact subgroup.*

The space E fibers over a Teichmüller disk D and the Veech group G acts on D as a finite co-area Fuchsian group. This action is necessarily *not* cocompact, and truncating D by removing a G -invariant family of horoballs gives a space $\bar{D} \subset D$

on which G acts cocompactly. The extension group Γ acts cocompactly on the subbundle $\bar{E} \subset E$ over \bar{D} , and thus Γ and \bar{E} are quasi-isometric. For simplicity, assume Γ is torsion free. Then the subbundle over a boundary component of \bar{D} is naturally a copy of the universal cover of a graph manifold (in fact, this graph manifold is finitely covered by the mapping torus of a Dehn multi-twist). The Seifert manifolds in the graph manifold have universal covers that are products which thus obstructs hyperbolicity of \bar{E} (and hence Γ). Via the quasi-isometry between \bar{E} and Γ , these product spaces correspond to the cosets of the vertex groups, and the space \hat{E} in Theorem 1.1 is alternatively (coarsely) defined by coning off these products. The theorem thus says that after doing so, one obtains a hyperbolic space.

We end this introduction with a final application of Theorem 1.2 and the analysis in §5 of this paper, which serves to illustrate the naturality of the geometry of \bar{E} .

Theorem 1.4 ([DDLS21]). *The homomorphism $\text{Isom}(\bar{E}) \rightarrow \text{QI}(\bar{E}) \cong \text{QI}(\Gamma)$ is an isomorphism and $\text{Isom}(\bar{E})$ contains Γ as a finite index subgroup.*

Remark 1.5. The space \bar{E} depends on \bar{D} whose definition includes choices (specifically, choices of horoballs in D which are removed), and we remark that Theorem 1.4 is in fact only true for a particular choice of \bar{E} .

1.1. Outline and proofs. Let us briefly outline the paper and comment on the main structure of the proofs. See the beginning of each section for more discussion. In §2 we review the necessary background material. Then in §3 we introduce all the objects and notation that will be used in the rest of the paper. In particular, we define the space E and its truncation \bar{E} , which is a quasi-isometric model for the extension group Γ of a Veech group G . This section also describes some of the key geometric features and subspaces of E and \bar{E} , and introduces the (coned-off) space \hat{E} .

In §4 we prove the main theorem, showing that \hat{E} is hyperbolic, and classify the elements of Γ acting loxodromically. Roughly, \hat{E} is something like a bundle over a coned-off Teichmüller disk, where all fibers are objects related to flat geometry on surfaces. More precisely, each fiber is either the universal cover of S with a flat metric, or a tree dual to a certain foliation (combinatorially, these are the Bass-Serre trees dual to the JSJ splittings of the graph manifolds, or more generally, graph orbifolds, mentioned above, explaining why the subgroups in Theorem 1.1 are called vertex groups). To prove hyperbolicity, we exploit features of flat geometry to construct explicit paths that form thin triangles, and then use the “Guessing geodesics” criterion (Proposition 2.2) due to Masur–Schleimer [MS13] and Bowditch [Bow14]. The proofs involve a careful analysis of various types of triangles in the universal cover of a flat surface (which can be more complicated than one might expect). This section ends with §4.9 where we prove that any element not conjugate into a vertex subgroup acts loxodromically on \hat{E} . For this, we quickly reduce to the case that the element projected to the Veech group G is not loxodromic, and hence up to a power, fixes one of the trees mentioned above. Not being in a vertex stabilizer, the element has an axis in the tree, and the proof is completed by constructing an equivariant Lipschitz retraction from \hat{E} to the axis in the tree.

The final section, §5, shifts gears to describe the four-dimensional Thurston-type geometry on which the nonsingular part of E and \bar{E} are modeled. After describing this geometry, we briefly discuss the connection between the geometry of E and \bar{E} , and singular Sol metrics on fibered hyperbolic 3-manifolds. Some of the results

in this section will be important in the sequel [DDLS21]. We also note that it is the similarity between this singular **Sol** geometry and the geometry of E that motivated our “geodesic guesses” in the proof of hyperbolicity of \hat{E} .

1.2. Notation and guides. This paper contains notation that is used throughout. For the reader’s convenience, we have included both a Notation Index §A as well as a detailed diagram illustrating the key objects and maps §3.8.

Acknowledgments. The authors would like to thank MSRI and its Fall 2016 program on *Geometric Group Theory*, where this work began. We also gratefully acknowledge NSF grants DMS 1107452, 1107263, 1107367 (the GEAR Network) for supporting travel related to this project. Dowdall was partially supported by NSF grants DMS-1711089 and DMS-2005368. Durham was partially supported by NSF grant DMS-1906487. Leininger was partially supported by NSF grants DMS-1510034, DMS-1811518, and DMS-2106419. Sisto was partially supported by the Swiss National Science Foundation (grant #182186). The authors would also like to thank the anonymous referees for many helpful suggestions.

2. DEFINITIONS AND BACKGROUND

We briefly recall some essential background material on coarse geometry, mapping class groups, Teichmüller spaces, quadratic differentials and Veech groups. As we expect that most readers will already be familiar with this material, we largely defer to other sources for details.

2.1. Basic metric geometry. A *path* in a metric space (X, d) is a continuous function $\gamma: [a, b] \rightarrow X$ from some (possibly degenerate) compact interval in the real line. Such a path *connects* or *joins* its endpoints $\gamma(a), \gamma(b)$ and has *length*

$$\text{length}(\gamma) = \sup \left\{ \sum_{i=1}^n d(\gamma(t_{i-1}), \gamma(t_i)) \mid a = t_0 \leq \dots \leq t_n = b \right\}.$$

A path is *rectifiable* if its length is finite, and is a *geodesic* if its length equals the distance between its endpoints. The metric space is said to be *geodesic* if each pair of points in it are connected by a geodesic. It is customary to denote by $[x, y]$ any choice of a geodesic connecting x to y . We will often conflate geodesics with their images, and same for paths.

A subset $Z \subset X$ is *rectifiably path connected* if any pair of points in Z are connected by a rectifiable path $[a, b] \rightarrow Z$. In this case Z inherits an *induced path metric* where the distance between points $z_1, z_2 \in Z$ is defined as the infimum of the lengths of all paths in Z connecting z_1 and z_2 . We call d a *path metric* if it agrees with the induced path metric on X itself; that is, if for all x, y the distance $d(x, y)$ coincides with the infimum of the lengths of paths connecting x to y .

Given two metrics d and d' on a set X , we say that d is *coarsely bounded* by d' if there exists a monotone function $N: [0, \infty) \rightarrow [0, \infty)$ so that $d(x, y) \leq N(d'(x, y))$, for all $x, y \in X$. If d is coarsely bounded by d' and d' is coarsely bounded by d , we say that d and d' are *coarsely equivalent*.

To any (set-theoretic) quotient $\pi: X \rightarrow Q$ of the metric space (X, d) , there is an associated *quotient pseudo-metric* d_Q on Q defined as

$$d_Q(p, q) = \inf \sum_{i=1}^n d(x_i, y_i),$$

where the infimum is over all chains $x_1, y_1, \dots, x_n, y_n \in X$ satisfying $\pi(x_1) = p$, $\pi(x_n) = q$, and $\pi(y_i) = \pi(x_{i+1})$ for all $i = 1, \dots, n-1$; see [BH99, I.5.19]. Call such a chain *tight* if $y_i \neq x_{i+1}$ for each $1 \leq i < n$. Observe that if a chain is not tight because $y_i = x_{i+1}$, then $d(x_i, y_{i+1}) \leq d(x_i, y_i) + d(x_{i+1}, y_{i+1})$ and we may delete y_i, x_{i+1} to get a new chain with potentially smaller sum. Thus $d_Q(p, q)$ is in fact an infimum over tight chains.

2.2. Graph approximation. The $C \geq 0$ neighborhood of a subset A of a metric space (X, d) will be denoted

$$N_C(A) = \{x \in X \mid d(x, a) \leq C \text{ for some } a \in A\}.$$

We say that A is C -dense if $X = N_C(A)$. A *quasi-isometric embedding* $f: X \rightarrow Y$ is a map of metric spaces for which there exist constants $K \geq 1, C \geq 0$ such that

$$\frac{1}{K}d_Y(f(x), f(y)) - C \leq d_X(x, y) \leq Kd_Y(f(x), f(y)) + C$$

for all $x, y \in X$. The map is furthermore a *quasi-isometry* if it is coarsely surjective, meaning that $f(X)$ is R -dense in Y for some $R \geq 0$. It is well known that every quasi-isometry $f: X \rightarrow Y$ admits a *coarse inverse* quasi-isometry $g: Y \rightarrow X$ for which $g \circ f$ and $f \circ g$ are bounded distance from the identities on X and Y . A quasi-isometric embedding with domain an interval in \mathbb{R} or \mathbb{Z} is called a *quasi-geodesic*.

The following well-known fact will be useful in a few places.

Proposition 2.1. *Suppose Ω is a path metric space and $\Upsilon \subset \Omega$ an R -dense subset, where $R > 0$. Fix $R' > 3R$ and consider a graph \mathcal{G} with vertex set Υ such that:*

- all pairs of elements of Υ within distance $3R$ are joined by an edge in \mathcal{G} ,
- if an edge in \mathcal{G} joins points $w, w' \in \Upsilon$, then $d_\Omega(w, w') \leq R'$.

Then the inclusion of Υ into Ω extends to a quasi-isometry $\mathcal{G} \rightarrow \Omega$.

Proof. We regard Υ as a subset of both Ω and \mathcal{G} , where we denote the respective metrics by $d_\Omega, d_\mathcal{G}$. Let $f: \mathcal{G} \rightarrow \Omega$ be any map that includes the vertex set Υ into Ω and sends each edge to a constant speed path connecting the endpoints of length at most R' . The map f is clearly R' -Lipschitz and the image is R -dense.

To prove the remaining inequality, consider now arbitrary vertices $w, w' \in \Upsilon$ of \mathcal{G} . Let α be a path connecting w to w' in Ω of length $Rk \leq d_\Omega(w, w') + R$ for some integer k . Subdividing α into subpaths of length R and considering points of Ω within distance R of the endpoints of the subpaths, we can find a sequence of points $w = w_0, \dots, w_k = w'$ of Υ so that $d_\Omega(w_i, w_{i+1}) \leq 3R$. In particular,

$$d_\mathcal{G}(w, w') \leq k \leq d_\Omega(w, w')/R + 1,$$

and the proof is complete. \square

2.3. Hyperbolicity. A metric space (X, d) is called Gromov *hyperbolic* if there is a constant $\delta' \geq 0$ such that for all $x, y, z, w \in X$ one has

$$(x|z)_w \geq \min \{(x|y)_w, (y|z)_w\} - \delta',$$

where $(a|b)_c$ denotes $(d(a, c) + d(b, c) - d(a, b))/2$. In the case that X is geodesic, this is equivalent to the existence of a constant $\delta \geq 0$ so that for all $x, y, z \in X$ and all choices of geodesics $[x, y], [y, z], [z, x]$ we have

$$[x, y] \subseteq N_\delta([y, z] \cup [z, x]).$$

When we want to specify the constant δ , we say that the space X is δ -hyperbolic. We call δ a hyperbolicity constant for X . It is a standard result that the property of being hyperbolic is preserved by quasi-isometries of path metric spaces; see e.g., [Väi05]. It is also well-known that the hyperbolic plane \mathbb{H}^2 is, say, 2-hyperbolic.

We will need the following criterion for hyperbolicity, which is an easy modification of one due to Masur–Schleimer [MS13, Theorem 3.11], see also [Bow14, Proposition 3.1]. For the statement, given a metric space Ω and designated subsets $L(x, y) \subseteq \Omega$ for each pair $x, y \in \Omega$, we say that the subsets $L(x, y)$ form δ -slim triangles if for all $x, y, z \in \Omega$ we have $L(x, y) \subseteq N_\delta(L(x, z) \cup L(z, y))$. So, by definition, a geodesic metric space is hyperbolic if the set of geodesics form slim triangles.

Proposition 2.2 (Guessing geodesics). *Suppose Ω is a path metric space, $\Upsilon \subset \Omega$ an R -dense subset for some $R > 0$, and $\delta \geq 0$ a constant such that for all pairs $x, y \in \Upsilon$ there are rectifiably path-connected sets $L(x, y) \subset \Omega$ containing x, y satisfying the properties:*

- (1) *the $L(x, y)$ form δ -slim triangles, and*
- (2) *if $x, y \in \Upsilon$ have $d(x, y) \leq 3R$, then the diameter of $L(x, y)$ is at most δ .*

Then there exists δ' , depending only on R and δ , so that Ω is δ' -hyperbolic. Moreover, any geodesic connecting x, y has Hausdorff distance at most δ' to $L(x, y)$.

Proof. For the special case of a graph \mathcal{G} with Υ the 1-dense set of vertices and the subsets $L(x, y)$ connected subgraphs, this statement is precisely the criterion of Bowditch [Bow14, Proposition 3.1]. To translate this criterion for graphs into the condition we have here, we construct the graph \mathcal{G} from Proposition 2.1 and prove that \mathcal{G} is hyperbolic; this suffices since \mathcal{G} is quasi-isometric to Ω .

Let $g: \Omega \rightarrow \Upsilon$ be a map fixing Υ and assigning to $x \in \Omega$ some point of Υ within distance R of x . Note that we will think of elements of Υ as contained both in \mathcal{G} and in Ω . For $x, y \in \Upsilon$, let $\mathcal{L}(x, y)$ be the span of $g(L(x, y))$; that is, the largest subgraph whose vertex set is precisely $g(L(x, y))$. Since $L(x, y)$ is rectifiably path connected, given any two points $w, w' \in L(x, y)$, there is a rectifiable path connecting them. Partition this path into subpaths of length at most R , at points $w = w_0, w_1, \dots, w_n = w'$. Thus $d_\Omega(w_i, w_{i+1}) \leq R$, and so $d_\Omega(g(w_i), g(w_{i+1})) \leq 3R$, ensuring that $g(w_i)$ and $g(w_{i+1})$ are adjacent in \mathcal{G} . Consequently, $w = w_0$ is connected by a path of length n to $w' = w_n$ in \mathcal{G} with all vertices in $g(L(x, y))$, and hence $\mathcal{L}(x, y)$ is connected.

Notice that g is a quasi-isometry (since it is a coarse inverse of the quasi-isometry induced by inclusion from Proposition 2.1). Hence, since the $L(w, w')$ form δ -slim triangles, we have that $\mathcal{L}(w, w')$ also form slim triangles (for a possibly larger constant). Similarly, the bound on the diameter of the sets $L(w, w')$ implies a bound on the diameter of the sets $\mathcal{L}(w, w')$. By the hyperbolicity criterion for graphs, it follows that \mathcal{G} , and hence Ω , is hyperbolic. The uniformity of the hyperbolicity constant and the final sentence of the proposition follow from the corresponding statements for \mathcal{G} . \square

2.4. Surfaces and mapping class groups. We next describe a variety of geometric, topological, and analytic objects associated to a closed surface S of genus at least two. Firstly, we write $\text{Mod}(S) = \pi_0(\text{Homeo}^+(S))$ for its *mapping class group*, the group of components of its orientation preserving homeomorphisms. Let \dot{S} denote the surface S with a marked point and $\text{Mod}(\dot{S})$ its mapping class group (in this case, all homeomorphisms must preserve the marked point). Birman proved that

by forgetting the marked point, these mapping class groups fit into the following *Birman exact sequence*,

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

Given a subgroup $G < \text{Mod}(S)$, we let $\Gamma_G < \text{Mod}(\dot{S})$ denote the preimage in $\text{Mod}(\dot{S})$, which thus also fits into a short exact sequence

$$(1) \quad 1 \rightarrow \pi_1 S \rightarrow \Gamma_G \rightarrow G \rightarrow 1$$

mapping into the one above by inclusion. See e.g. [Bir74, Bir69, FM10].

2.5. Teichmüller spaces. Two complex structures X and Y on S are equivalent if there is a biholomorphic map $(S, X) \rightarrow (S, Y)$ isotopic to the identity on S . The *Teichmüller space* of S , denoted $\mathcal{T}(S)$, is the space of equivalence classes of complex structures on S . Given a complex structure X on S , we also write $X \in \mathcal{T}(S)$ for its equivalence class. We equip $\mathcal{T}(S)$ with the Teichmüller metric, for which the distance from X to Y measures the maximal quasiconformal dilatation between the two complex structures at every point, minimized over all representatives of the equivalence classes. See e.g. [Gar87, IT92, Ber60].

Any complex structure X on S also determines a point of $\mathcal{T}(\dot{S})$, where equivalence is defined via isotopies that fix the marked point. Forgetting the marked point defines a fibration of Teichmüller spaces called the *Bers fibration*,

$$(2) \quad \tilde{\mathcal{T}} \rightarrow \mathcal{T}(\dot{S}) \rightarrow \mathcal{T}(S).$$

Changing the complex structure on \dot{S} by an isotopy that does *not* fix the base point gives different points in the fiber over $X \in \mathcal{T}(S)$. The isotopy to the identity lifts to an isotopy to the identity of the universal covering \tilde{S} , and tracking the location of a fixed lift of the marked point gives the identification of the fiber with \tilde{S} . The fibration is equivariant with respect to the homomorphisms in the Birman exact sequence. See [Ber73, FM02].

2.6. Quadratic differentials. Given a complex structure X on S , by a *quadratic differential for X* we mean a nonzero holomorphic section of the square of the canonical line bundle over (S, X) . In a local coordinate z for X , this is represented as $q(z)dz^2$, where q is a holomorphic function. Integrating the square root of a quadratic differential q for X (away from the zeros) gives a local coordinate ζ for X called a *preferred coordinate for q* in which q is represented as $d\zeta^2$. The transition functions for overlapping preferred coordinates are locally given by $z \mapsto \pm z + c$ for some $c \in \mathbb{C}$. The Euclidean metric, being invariant by such transformations, pulls back to a metric on the S minus the zeros of q . The metric completion with respect to this metric is a nonpositively curved Euclidean cone metric, called a *flat metric*, obtained by filling the zeros back in, so that there is a cone point of cone angle $k\pi$ at each zero of q of order $k - 2$. We will write q for both the quadratic differential as well as the associated flat metric. See e.g. [Gar87, Str80].


Remark 2.3. Using the notation q for both the quadratic differential as well as the flat metric is imprecise because the latter only determines the former up to multiplication by a unit modulus complex number. When passing back and forth between the two, a choice of specific quadratic differential will either be irrelevant or the choice will be clear from the context.

2.7. Flat geodesics, foliations, and directions. Consider a complex structure X and flat metric q (from a quadratic differential of the same name). The geodesics for q are Euclidean straight lines away from the cone points. When a geodesic passes through a cone point, it subtends two angles, one on “each side” of the cone point, each of which is at least π . In fact, this characterizes q -geodesics: a path which is a Euclidean straight line away from the cone points, and makes angle at least π on both sides of any cone point it meets is necessarily a geodesic. Geodesic segments between pairs of cone points of q that contain no cone points in their interior are called *saddle connections*.

Because the transition functions for the preferred coordinates defining the flat metric have the form $z \mapsto \pm z + c$, any line in the tangent space of a nonsingular point can be parallel translated around the surface, in the complement of the cone points, to produce a parallel line field. Such line fields are in a one-to-one correspondence with the projective tangent space at any non-cone point, and we denote this space of directions as $\mathbb{P}^1(q)$.

Integrating a parallel line field gives a foliation by geodesics, which extends to a singular foliation over the entire surface. Thus for every $\alpha \in \mathbb{P}^1(q)$, we have a singular foliation $\mathcal{F}(\alpha)$ in direction α . An important special case occurs when $\mathcal{F}(\alpha)$ has all nonsingular leaves being closed geodesics. In this case, S is a union of Euclidean cylinders, the interior of each foliated by parallel geodesic core curves in direction α , and with boundary curves a geodesic concatenation of saddle connections (also in direction α). In this case, we say that $\mathcal{F}(\alpha)$ defines a *cylinder decomposition*.

Example 2.4. Consider the genus 2 surface S obtained from the polygon shown below with the sides identified by translations according to the numbering indicated. This surface is built from three squares and has an obvious 3-fold branched cover map to the square torus.



The Euclidean metric on the polygon descends to a flat metric on S determined by a complex structure X and quadratic differential q . The vertices shown project to a single cone point of cone angle 6π . The horizontal foliation illustrated defines a cylinder decomposition with core curves isotopic to the dotted and dashed curves and with boundaries made of the three horizontal saddle connections. In fact, any direction of rational slope defines a cylinder decomposition, as can be seen from the projection to the square torus.

The complex structure and quadratic differential can be pulled back to the universal covering \tilde{S} , and will be denoted by the same names. The metric on \tilde{S} is

CAT(0) (since all cone angles are greater than 2π), and in particular, any two points are connected by a unique geodesic segment. The characterization of geodesics in \tilde{S} still holds, and saddle connections are defined similarly (or equivalently, as lifts of saddle connections). The preimage of a cylinder on S is an *infinite strip* in \tilde{S} : this is an isometric embedding of a strip $[0, w] \times \mathbb{R}$ of some width $w > 0$. The universal covering determines a canonical identification of the directions for q on \tilde{S} with those on S , and we refer to either one as $\mathbb{P}^1(q)$.

2.8. Teichmüller disks. Suppose q is any quadratic differential for X with preferred coordinates $\{\zeta_j\}$ covering the complement of the zeros of q . For any matrix $A \in \mathrm{SL}_2(\mathbb{R})$, viewed as a real linear transformation of \mathbb{C} , we obtain a new atlas $\{A \circ \zeta_j\}$ defining a new complex structure and quadratic differential,

$$A \cdot (X, q) = (A \cdot X, A \cdot q),$$

with the same set of zeros of the same orders. (A caveat: the notation $A \cdot X$ only makes sense in the presence of a quadratic differential q for X). In this setting, the identity map $(S, X, q) \rightarrow (S, A \cdot X, A \cdot q)$ is *affine* in the respective preferred coordinates for q and $A \cdot q$, and Teichmüller's theorem states that the distance between X and $A \cdot X$ in $\mathcal{T}(S)$ is precisely the logarithm $\log(\|A\|)$ of the operator norm

$$\|A\| = \max\{\|A(v)\| \mid v \in \mathbb{C}, \|v\| = 1\}.$$

We call q and $A \cdot q$ the *initial* and *terminal flat metrics*, respectively. We note that one usually refers to the specific quadratic differentials as initial and terminal, but this is unimportant for us; compare Remark 2.3. As the identity $(S, X, q) \rightarrow (S, A \cdot X, A \cdot q)$ is affine, it sends every q -geodesic to an $(A \cdot q)$ -geodesic.

If $B \in \mathrm{SO}(2)$ and $A \in \mathrm{SL}_2(\mathbb{R})$, then $A \cdot X$ and $BA \cdot X$ are equivalent (since a rotation is holomorphic), and so the map $\mathrm{SO}(2)A \mapsto A \cdot X$ defines a homeomorphism from $\mathrm{SO}(2) \backslash \mathrm{SL}_2(\mathbb{R})$ to the orbit $D_q \subset \mathcal{T}(S)$. On the other hand, $\mathrm{SO}(2) \backslash \mathrm{SL}_2(\mathbb{R})$ may be identified with the (upper half-plane model of the) hyperbolic plane, \mathbb{H}^2 , via the homeomorphism $\mathrm{SO}(2)A \mapsto A^{-1}(i)$, and we do so. By Teichmüller's theorem, the push-forward of the Poincaré metric is the restriction of the Teichmüller metric to D_q ; in particular, D_q is totally geodesic in $\mathcal{T}(S)$. We call D_q the *Teichmüller disk* of the quadratic differential q . For any other point $(X', q') = A \cdot (X, q)$, we have $D_q = D_{q'}$ and we have a canonical identification $\mathbb{P}^1(q) \cong \mathbb{P}^1(q')$ from the affine identity map $id_S: (S, X, q) \rightarrow (S, X', q')$. See [GL00] for details.

Any geodesic through X in D_q is given by $t \mapsto A_t \cdot (X, q)$, where $\{A_t\}_{t \in \mathbb{R}} < \mathrm{SL}_2(\mathbb{R})$ is a symmetric, 1-parameter, hyperbolic subgroup. The (orthogonal) eigenlines for the common eigenvalues of the nontrivial elements give two direction in $\mathbb{P}^1(q)$, and we use this to identify the circle at infinity of D_q with $\mathbb{P}^1(q)$. Specifically, we identify the direction α of the contracting eigenline of A_t , $t > 0$, with the endpoint of the positive ray. In particular, along this ray the flat-lengths of any saddle connection or closed geodesic in direction α contracts exponentially. Furthermore, the horocycle through X based at the point at infinity associated to α is given by $t \mapsto B_t \cdot (X, q)$, where $\{B_t\}_{t \in \mathbb{R}} < \mathrm{SL}_2(\mathbb{R})$ is a 1-parameter parabolic subgroup with all nontrivial elements having common eigenline α . Consequently, a saddle connection or closed geodesic in direction α has constant length along this horocycle, as does the transverse measure to the foliation $\mathcal{F}(\alpha)$. Combining these properties, we have the following useful fact.

Proposition 2.5. *Given a quadratic differential q for some complex structure X on S , and any saddle connection or simple closed geodesic σ with respect to q in direction $\alpha \in \mathbb{P}^1(q)$, the horoballs (respectively, horocycles) based at α in D_q are sublevel sets (respectively, level sets) of the length of σ . If there is a maximal cylinder in direction α , then its width is constant on horocycles based at α . \square*

2.9. Veech groups. Given a quadratic differential q for X and associated Teichmüller disk $D_q \subset \mathcal{T}(S)$, the stabilizer of D_q , denoted $G_q < \text{Mod}(S)$ is called the *Veech group* of q (or D_q). This is equivalently the subgroup of $\text{Mod}(S)$ consisting of all mapping classes represented by homeomorphisms that are affine in preferred coordinates for q . The derivative of any such affine homeomorphism in preferred coordinates is well defined, up to sign, and defines a homomorphism $G_q \rightarrow \text{PSL}_2(\mathbb{R})$. The action of G_q on D_q gives a homomorphism from $G_q \rightarrow \text{Isom}^+(D_q) \cong \text{PSL}_2(\mathbb{R})$, and up to conjugation, this is precisely the derivative homomorphism. Since the action of $\text{Mod}(S)$ on $\mathcal{T}(S)$ is properly discontinuous, so is the action of G_q on D_q , and hence the image in $\text{PSL}_2(\mathbb{R})$ is discrete. We say that (S, X, q) is a *lattice surface* if G_q is a lattice (i.e. the quotient D_q/G_q has finite area).

An element of G_q is called *elliptic*, *parabolic*, or *hyperbolic*, respectively, if its image under the derivative map is of that type. Elliptic elements fix a point of D_q , while parabolic and hyperbolic elements $g \in G_q$ fix one or two points, respectively, in $\mathbb{P}^1(q)$. If g is parabolic with fixed point $\alpha \in \mathbb{P}^1(q)$, then the foliation $\mathcal{F}(\alpha)$ defines a g -invariant cylinder decomposition, and a power of g is a composition τ_α of Dehn twists in the core curves of the cylinders. This power of g is the identity on the saddle connections in direction α , which form a union of spines for the complementary components of the core curves of the cylinders. The set of *parabolic directions*, denoted $\mathcal{P}(q) \subset \mathbb{P}^1(q)$, is the fixed points of parabolic elements in G_q .

Given $\alpha \in \mathcal{P}(q)$, the boundary of every cylinder in direction α consists of saddle connections. The key result for us in the next theorem is that all saddle connections arise in this way in the case of interest to us. See [MT02, Thu88] for details.

Theorem 2.6 (Veech dichotomy). *Suppose (S, X, q) is a lattice surface. Then every saddle connection is contained in the boundary of a cylinder in direction $\alpha \in \mathcal{P}(q)$. Moreover, for every $\alpha \notin \mathcal{P}(q)$, every half leaf of $\mathcal{F}(\alpha)$ is dense in S .*

Example 2.7. Consider the Veech group G_q for the surface (S, X, q) of Example 2.4. There are two parabolic elements $g, h \in G_q$ with derivatives

$$dg = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad dh = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}.$$

The element g preserves the horizontal foliation and is isotopic to a composition of Dehn twists about the core curves of the corresponding cylinder decomposition. The element h similarly preserves the vertical foliation and is given by a composition of Dehn twists. The derivatives of the elements $g, h \in G_q$ generate a lattice in $\text{PSL}_2(\mathbb{R})$. It follows that $\langle g, h \rangle < G_q$ is a finite index subgroup, and (S, X, q) is a lattice surface. Since the rational slopes are precisely the directions of cylinder decompositions, we see from the Veech dichotomy that $\mathcal{P}(q) = \mathbb{P}^1(\mathbb{Q})$ in this case.

Remark 2.8. On the one hand, it might seem that lattice Veech groups are fairly special; only countably many Teichmüller disks can define a lattice Veech group. On the other hand, they actually exist in abundance. For instance, the example described above is a “square tiled” flat metric. The Teichmüller disks defined by

square tiled flat metrics are dense in Teichmüller space and the maximal Veech group for each one is a lattice (see [Zor02]).

3. SETUP AND NOTATION

We now fix a flat metric (i.e. quadratic differential) q for a complex structure X_0 in the Teichmüller space $\mathcal{T}(S)$ of a closed surface S of genus at least 2, and assume (S, X_0, q) is a lattice surface. Let $D = D_q \subset \mathcal{T}(S)$ be the Teichmüller disk of q and let ρ be the Poincaré metric on D . For any $X \in D$, let q_X be the terminal flat metric from the Teichmüller mapping $(S, X_0, q) \rightarrow (S, X, q_X)$. Let $G = G_q \leq \text{Mod}(S)$ be the lattice Veech group of q and $\mathcal{P} = \mathcal{P}(q) \subset \mathbb{P}^1(q)$ be the set of parabolic directions on q , which are precisely the directions of the saddle connections by Theorem 2.6.

3.1. Horoballs. For each direction $\alpha \in \mathcal{P}$, we fix a closed horoball $B_\alpha \subset D$ that is invariant by the corresponding parabolic subgroup of G . We choose these so that the family $\{B_\alpha\}_{\alpha \in \mathcal{P}}$ is G -invariant and 1-separated, meaning that the pairwise distance between any two horoballs is at least 1. For each $\alpha \in \mathcal{P}$, we let $c_\alpha: D \rightarrow B_\alpha$ be the ρ -closest-point projection map. This map is 1-Lipschitz, as can be seen by observing that the derivative of c_α is norm non-increasing at each point $X \in D$.

Let $p: D \rightarrow \hat{D}$ denote the quotient space obtained by collapsing each B_α to a point, and give \hat{D} the quotient pseudo-metric $\hat{\rho}$. It is straightforward to see that this is indeed a metric and in fact, according to [BH99, Lemma I.5.20], a path metric. The 1-separated assumption implies that \hat{D} is quasi-isometric to the electric space obtained by coning each B_α to a point, and that $p(B_\alpha)$ and $p(B_{\alpha'})$ are at least $\hat{\rho}$ -distance 1 apart if $\alpha \neq \alpha'$.

We also consider the associated truncated Teichmüller disk

$$\bar{D} = D \setminus \bigcup_{\alpha \in \mathcal{P}} B_\alpha^\circ,$$

where B_α° is the interior of B_α . The induced path metric $\bar{\rho}$ on \bar{D} is CAT(0) by [BH99, Theorem II.11.27], and the group G acts cocompactly on \bar{D} . The projections c_α above now restrict to maps $c_\alpha: \bar{D} \rightarrow \partial B_\alpha \subset \bar{D}$ that remain 1-Lipschitz.

3.2. Bundles and the total space. Let $\pi: E \rightarrow D$ be the pull-back bundle of the Bers fibration (2) via the inclusion $D \subset \mathcal{T}(S)$. We identify $E \subset \mathcal{T}(\hat{S})$. Let $\Gamma = \Gamma_G < \text{Mod}(\hat{S})$ be the extension of G as in (1). The group Γ acts on E and the quotient E/Γ is a noncompact (orbifold) S -bundle over D/G . To rectify this noncompactness, we also consider the pull-back bundle over \bar{D}

$$\bar{E} = \pi^{-1}(\bar{D}) = E \setminus \bigcup_{\alpha \in \mathcal{P}} B_\alpha^\circ,$$

where $B_\alpha^\circ = \pi^{-1}(B_\alpha^\circ)$ is the interior of the horoball preimage $\mathcal{B}_\alpha = \pi^{-1}(B_\alpha)$ for $\alpha \in \mathcal{P}$. The space \bar{E} is a 4-manifold with boundary

$$\partial \bar{E} = \pi^{-1}\left(\bigcup_{\alpha \in \mathcal{P}} \partial B_\alpha\right) = \bigcup_{\alpha \in \mathcal{P}} \partial \mathcal{B}_\alpha,$$

and the group Γ acts on \bar{E} with quotient \bar{E}/Γ a compact S -bundle over \bar{D}/G :

$$S \rightarrow \bar{E}/\Gamma \rightarrow \bar{D}/G.$$

3.3. Moving between fibers. For any $X \in D$ let $E_X = \pi^{-1}(X) \subset E$ be the fiber over X , with $E_0 = E_{X_0}$ denoting the fiber over X_0 . The fiber E_X is canonically identified with the universal cover \tilde{S} of S equipped with the (pulled back) complex structure X and flat metric q_X .

Given $X, Y \in D$, let $f_{X,Y}: E_Y \rightarrow E_X$ be the lift of the Teichmüller mapping to the universal covering, which is affine with respect to the flat metrics q_X and q_Y on the domain and range, respectively. We note that $f_{X,Y}f_{Y,Z} = f_{X,Z}$ and that $f_{X,Y}$ is $e^{\rho(X,Y)}$ -biLipschitz (in preferred Euclidean coordinates, $f_{X,Y}$ is given by $(u, v) \mapsto (e^{\rho(X,Y)}u, e^{-\rho(X,Y)}v)$). Varying over all $Y \in D$, these determine a map

$$f_X: E \rightarrow E_X, \quad \text{where } f_X|_{E_Y} = f_{X,Y} \quad \text{for any } Y \in D.$$

In the case of X_0 , we simply write $f = f_{X_0}: E \rightarrow E_0$. For any $X \in D$ and $x \in E_X$, the fiber $D_x = f_X^{-1}(x)$ is the unique Teichmüller disk in $\mathcal{T}(\tilde{S})$ through x that covers D via the projection π . We also consider these objects in \bar{E} over \bar{D} , as well, writing $\bar{D}_x = D_x \cap \bar{E}$, which is a truncated Teichmüller disk mapping bijectively to \bar{D} .

We note that E is product, being a bundle over the contractible space D . The product structure $E \cong D \times \tilde{S}$ is quite natural; indeed, after identifying E_X with \tilde{S} (for any $X \in D$), the maps π and f_X determine the projections onto the two factors.

For any $\alpha \in \mathcal{P}$, let $f_\alpha: \bar{E} \rightarrow \partial\mathcal{B}_\alpha \subset \bar{E}$ be defined by

$$f_\alpha(x) = f_{c_\alpha(\pi(x))}(x),$$

for any $x \in \bar{E}$. In words, f_α restricted to any fiber E_X is the canonical map $f_{Y,X}$ to the fiber E_Y , where $Y = c_\alpha(X)$ is the ρ -closest point on $\partial\mathcal{B}_\alpha$ to X .

3.4. Trees. For every $\alpha \in \mathcal{P}$ there is a simplicial tree T_α dual to foliation of E_0 in direction α . The action of $\pi_1 S$ on E_0 determines an action on T_α , identifying it as the Bass-Serre tree dual to the (core curves of the) cylinder decomposition of S in direction α .

There is a natural $\pi_1 S$ -equivariant map $E_0 \rightarrow T_\alpha$ that pushes forward the transverse measure of the foliation to a metric on T_α which on each edge is a multiple of the Euclidean metric from the simplicial structure. For any other point $Y \in D$, $f_{X_0,Y}: E_Y \rightarrow E_0$ maps the foliation of E_Y in direction α to the foliation of E_0 in direction α . Composing the map $E_0 \rightarrow T_\alpha$ with $f_{X_0,Y}$ gives another “natural map”, and the induced metrics differ only by dilation (according to the change in transverse measures). We assemble these maps all together into

$$t_\alpha: E \rightarrow T_\alpha,$$

given as the composition $f: E \rightarrow E_0$ with the map $E_0 \rightarrow T_\alpha$. From the last sentence of Proposition 2.5, it follows that any two points X, Y on the same horocycle based at α define the same metric on T_α .

3.5. Cone points. For $X \in D$, let $\Sigma_X \subset E_X$ denote the set of cone points of the flat structure q_X on E_X . We write $\Sigma_0 = \Sigma_{X_0}$, as a special case, and define

$$\Sigma = \bigcup_{X \in D} \Sigma_X \quad \text{and} \quad \bar{\Sigma} = \Sigma \cap \bar{E}.$$

Note that Σ is the union of all Teichmüller disks through all cone points and $\bar{\Sigma}$ the union of truncated Teichmüller disks through cone points.

3.6. Bundle metrics. We give E a metric d defined as follows. In $E - \Sigma$, the metric is Riemannian given by the orthogonal direct sum of the flat metric in the fiber and Poincaré metric on Teichmüller disks, and then E can be identified as the metric completion. As explained in §5, away from the singular locus, the metric is locally homogeneous; see Proposition 5.3. This metric is invariant by Γ . The induced *path* metric in each fiber E_X is its flat metric, but the fibers are distorted. The induced metric in each Teichmüller disk D_x is a Poincaré metric, and each such D_x is isometrically embedded with π restricting to an isometry $D_x \rightarrow D$.

The subspace $\bar{E} \subset E$ is given the induced path metric which we denote \bar{d} . Since Γ acts isometrically and cocompactly on (\bar{E}, \bar{d}) , this metric space will serve as our quasi-isometric model for the extension group Γ .

We note that $\bar{\Sigma}$ is r -dense for some $r > 0$, which follows from the compactness of the quotient metric space \bar{E}/Γ . Using the Arzelà–Ascoli Theorem and the compactness of \bar{E}/Γ , it is easy to see that any sequence of paths joining x and y in \bar{E} whose lengths converge to $\bar{d}(x, y)$ has a subsequential limit that is, necessarily, a geodesic. Thus (\bar{E}, \bar{d}) is a geodesic metric space.

From the G -invariant quotient $p: D \rightarrow \hat{D}$ we construct a Γ -invariant quotient $P: E \rightarrow \hat{E}$ as follows. Let T_α be the tree and $t_\alpha: E \rightarrow T_\alpha$ the map described in §3.4. We endow T_α with the metric d_α coming from any point on the horocycle ∂B_α , making T_α into an \mathbb{R} -tree. For concreteness (and for later use) we also fix a point $X_\alpha \in \partial B_\alpha$ so that d_α is obtained from the push forward of the transverse measure on the foliation of E_{X_α} in direction α .

Definition 3.1 (The hyperbolic space). Let $P: E \rightarrow \hat{E}$ denote the quotient obtained by collapsing each set B_α onto T_α via the map $t_\alpha|_{B_\alpha}$. Equip \hat{E} with the quotient pseudo-metric \hat{d} obtained from \bar{d} under the (surjective) restriction

$$\bar{P} = P|_{\bar{E}}: \bar{E} \rightarrow \hat{E}.$$

We note that there is a uniform lower and upper bound of the length of every edge of every tree T_α since there are only finitely-many Γ -orbits of edges.

Lemma 3.2 below shows that \hat{d} is in fact a path metric. The main part of Theorem 1.1 is that the metric space \hat{E} is hyperbolic (see Theorem 4.1). Let $\mathcal{V} \subset \hat{E}$ be the set of all vertices of all trees T_α over all $\alpha \in \mathcal{P}$. We write $\alpha: \mathcal{V} \rightarrow \mathcal{P}$ to denote the map that associates to v the direction $\alpha(v) \in \mathcal{P}$ so that $v \in T_{\alpha(v)}$. Given $v \in \mathcal{V}$, we will write $B_v = B_{\alpha(v)}$, $\partial B_v = \partial B_{\alpha(v)}$, etc.

The map $\bar{\pi} = \pi|_{\bar{E}}: \bar{E} \rightarrow \bar{D}$ descends to a quotient map $\hat{\pi}: \hat{E} \rightarrow \hat{D}$ that is 1-Lipschitz by construction (as it is the descent of a 1-Lipschitz map). For every $x \in E$, the image \hat{D}_x of D_x in \hat{E} is obtained by collapsing $B_\alpha \cap D_x$ to a point, for each $\alpha \in \mathcal{P}$, and hence $\hat{\pi}|_{\hat{D}_x}: \hat{D}_x \rightarrow \hat{D}$ is a bijection. In particular, each \hat{D}_x with its path metric is isometric to \hat{D} , and isometrically embedded in \hat{E} . Objects in E , \bar{E} , and \hat{E} are called *vertical* if they are contained in a fiber of π , $\bar{\pi}$, or $\hat{\pi}$, respectively, and *horizontal* if they are contained in D_x , \bar{D}_x , or \hat{D}_x , for some $x \in E, \bar{E}$. The

following commutative diagram summarizes the situation:

$$(3) \quad \begin{array}{ccccc} & & E & & \\ & \nearrow & \downarrow \pi & \searrow P & \\ \bar{E} & \xrightarrow{\bar{P}} & \hat{E} & & \\ & \searrow \bar{\pi} & D & \nearrow p & \\ & & \hat{D} & & \end{array}$$

Some key features of the various metrics that we will use are highlighted in the following two lemmas. In particular, the next lemma states that the trees T_α inside of \hat{E} behave exactly as expected.

Lemma 3.2. *The quotient pseudo-metric \hat{d} on \hat{E} is a path metric. The map $\bar{P}: \bar{E} \rightarrow \hat{E}$ is 1-Lipschitz and is a local isometry at every point $x \in \bar{E} - \partial\bar{E}$. Furthermore, for every $\alpha \in \mathcal{P}$,*

- *The induced path metric on $P(\partial\mathcal{B}_\alpha) = T_\alpha$ is d_α , the \mathbb{R} -tree metric determined by the horocycle $\partial\mathcal{B}_\alpha$.*
- *The subspace topology on $T_\alpha \subset \hat{E}$ agrees with the \mathbb{R} -tree topology on T_α .*

Proof. The map \bar{P} is 1-Lipschitz by definition of \hat{d} . According to [BH99, Lemma I.5.20], since \bar{d} is a path metric, \hat{d} is necessarily a path metric provided it is a metric. Recall that for $\hat{x}, \hat{y} \in \hat{E}$, the quotient pseudo-metric $\hat{d}(\hat{x}, \hat{y})$ is the infimum of $\sum_i \bar{d}(x_i, y_i)$ over all tight chains $x_1, y_1, \dots, x_n, y_n$ from \hat{x} to \hat{y} , meaning that $\bar{P}(x_1) = \hat{x}$, $\bar{P}(y_n) = \hat{y}$, and $\bar{P}(y_i) = \bar{P}(x_{i+1})$ for all $1 \leq i < n$ with $y_i \neq x_{i+1}$.

First suppose $x \in \bar{E} \setminus \partial\bar{E}$ and set $3\epsilon = \bar{d}(x, \partial\bar{E}) > 0$. Let $U \subset \bar{E}$ be the ball of radius ϵ about x and let $z \in U$ be arbitrary. For any $y \in \bar{E}$, consider any tight chain $z_1, y_1, \dots, z_n, y_n$ from $\hat{z} = \bar{P}(z)$ to $\hat{y} = \bar{P}(y)$. If $y_1 \neq y$, then tightness implies $y_1 \in \partial\bar{E}$ (since \bar{P} is injective on $\bar{E} \setminus \partial\bar{E}$) and hence $\sum_i \bar{d}(z_i, y_i) \geq \bar{d}(z_1, y_1) \geq 2\epsilon$. Otherwise we have the trivial chain $z = z_1, y_1 = y$ with sum $\bar{d}(z, y)$. This proves

$$\bar{d}(z, y) \geq \hat{d}(\hat{z}, \hat{y}) \geq \min\{2\epsilon, \bar{d}(z, y)\} \quad \text{for all } z \in U \text{ and } y \in \bar{E}.$$

In particular, $\hat{d}(\hat{z}, \hat{y}) = \bar{d}(z, y)$ for all $z, y \in U$, proving that \bar{P} is a local isometry on $\bar{E} \setminus \partial\bar{E}$. We also see that $\hat{d}(\hat{z}, \hat{y}) = 0$ forces $z = y$ and hence $\hat{z} = \hat{y}$, establishing the positive definiteness of \hat{d} in the case that one point lies in $\bar{P}(\bar{E} \setminus \partial\bar{E})$.

For $\alpha \in \mathcal{P}$, let $\bar{\ell}_\alpha$ be the path metric on $\partial\mathcal{B}_\alpha \subset \bar{E}$, and ℓ_α the path pseudo metric on $T_\alpha \subset \hat{E}$. Note that $\ell_\alpha \leq d_\alpha$ by construction. For any $X \in \partial\mathcal{B}_\alpha$, the restriction of $\bar{P}|_{\partial\mathcal{B}_\alpha} = t_\alpha|_{\partial\mathcal{B}_\alpha}$ to E_X is just the projection $E_X \rightarrow T_\alpha$ onto the \mathbb{R} -tree dual to the foliation of E_X in direction α . Since the path metric induced by \bar{d} on $\partial\mathcal{B}_\alpha$ is the path metric coming from the Riemannian orthogonal sum of flat metrics in fibers and horocycle length in Teichmüller disks, we see that every path γ in $\partial\mathcal{B}_\alpha$ satisfies

$$\bar{\ell}_\alpha\text{-length}(\gamma) \geq d_\alpha\text{-length}(\bar{P}(\gamma)).$$

Next suppose $x \in \partial\bar{E}$, and say $x \in \partial\mathcal{B}_\alpha$. Since $\hat{\pi}: \hat{E} \rightarrow \hat{D}$ is 1-Lipschitz, we have $\hat{d}(\bar{P}(x), \bar{P}(y)) \geq \hat{\rho}(\hat{\pi}(\bar{P}(x)), \hat{\pi}(\bar{P}(y))) > 0$ for any $y \notin \partial\mathcal{B}_\alpha$. For the purposes of proving \hat{d} is a metric, it therefore suffices to fix some $0 < \epsilon < \frac{1}{2}$ and consider a point $y \in \partial\mathcal{B}_\alpha$ such that $\hat{d}(\bar{P}(x), \bar{P}(y)) < \epsilon$. Take any tight chain x_1, \dots, y_n from $\hat{x} = \bar{P}(x)$ to $\hat{y} = \bar{P}(y)$ with $\sum_i \bar{d}(x_i, y_i) < \hat{d}(\hat{x}, \hat{y}) + \epsilon < 2\epsilon$. By tightness, $x_i, y_i \in \partial\bar{E}$ for each $1 \leq i \leq n$. If x_i and y_i lie in distinct components of $\partial\bar{E}$,

then $\bar{d}(x_i, y_i) \geq \bar{\rho}(\bar{\pi}(x), \bar{\pi}(y)) \geq 1$ by the fact that our horoballs $\{B_\alpha\}_{\alpha \in \mathcal{P}}$ are 1-separated. As this contradicts the assumption on $\sum_i \bar{d}(x_i, y_i)$, we must have $x_i, y_i \in \partial \mathcal{B}_\alpha$ for all $1 \leq i \leq n$. Now let γ_i be a geodesic joining x_i to y_i in \bar{E} of length $\bar{d}(x_i, y_i) < 2\epsilon$. Hence γ_i lies in the ϵ -neighborhood of $\partial \mathcal{B}_\alpha$. The map $f_\alpha: \bar{E} \rightarrow \partial \mathcal{B}_\alpha$ is e^ϵ -Lipschitz on this neighborhood, as can be seen by noting that this holds in the fiber direction (since $f_{X,Y}$ is $e^{\rho(X,Y)}$ -Lipschitz) and c_α is 1-Lipschitz. Therefore $f_\alpha(\gamma_i)$ is a path in $\partial \mathcal{B}_\alpha$ of length at most $e^\epsilon \bar{d}(x_i, y_i)$. Since $\bar{P}(f_\alpha(y_i)) = \bar{P}(f_\alpha(x_{i+1}))$, the images $\bar{P}(f_\alpha(\gamma_i))$ concatenate to give a path γ in T_α from \hat{x} to \hat{y} satisfying

$$\begin{aligned} \hat{d}(\hat{x}, \hat{y}) &\leq \ell_\alpha(\hat{x}, \hat{y}) \leq d_\alpha(\hat{x}, \hat{y}) \leq d_\alpha\text{-length}(\gamma) \\ &= \sum_{i=1}^n d_\alpha\text{-length}(\bar{P}(f_\alpha(\gamma_i))) \leq e^\epsilon \sum_{i=1}^n \bar{d}(x_i, y_i). \end{aligned}$$

Since this holds for all tight chains, we conclude that $e^{-\epsilon} d_\alpha(\hat{x}, \hat{y})$ is a lower bound on $\hat{d}(\hat{x}, \hat{y})$ whenever $\hat{x}, \hat{y} \in T_\alpha$ satisfy $\hat{d}(\hat{x}, \hat{y}) < \epsilon$. This proves \hat{d} is a metric, since we now see that $\hat{d}(\hat{x}, \hat{y}) = 0$ implies $d_\alpha(\hat{x}, \hat{y}) = 0$ and hence $\hat{x} = \hat{y}$.

Since any path in T_α may be subdivided into pieces whose successive endpoints satisfy $\hat{d}(\hat{x}_i, \hat{x}_{i+1}) < \epsilon$, this also proves that $\ell_\alpha \geq e^{-\epsilon} d_\alpha$ for all small $\epsilon > 0$. Therefore $\ell_\alpha = d_\alpha$ as claimed by the lemma. Finally, this argument establishes Lipschitz inequalities $e^{-\epsilon} d_\alpha \leq \hat{d} \leq d_\alpha$ for nearby points in T_α and, specifically, proves that for each $\hat{x} \in T_\alpha$ and all small $\epsilon > 0$ we have

$$\{\hat{y} \in T_\alpha \mid d_\alpha(\hat{x}, \hat{y}) < \epsilon\} \subset \{\hat{y} \in T_\alpha \mid \hat{d}(\hat{x}, \hat{y}) < \epsilon\} \subset \{\hat{y} \in T_\alpha \mid d_\alpha(\hat{x}, \hat{y}) < e^\epsilon \epsilon\}.$$

Thus the \mathbb{R} -tree and subspace topologies on T_α agree, and the lemma holds. \square

Although the metric on \hat{E} is defined from the metric on \bar{E} , the next lemma shows that the map from E to \hat{E} is equally well-behaved.

Lemma 3.3. *The map $P: E \rightarrow \hat{E}$ is 1-Lipschitz.*

Proof. Suppose $x, y \in E$ are any two points. Given any $\epsilon > 0$, there is a path γ from x to y with length $(1 + \epsilon)d(x, y)$ that decomposes as a concatenation $\gamma = \gamma_1 \cdots \gamma_k$, with each γ_i contained in \bar{E} or in \mathcal{B}_α , for some $\alpha \in \mathcal{P}$. Since ϵ is arbitrary, it suffices to prove that the length of $P(\gamma)$ is no greater than the length of γ . If γ_i is a path in \bar{E} , then the d -length is equal to the \bar{d} -length (since \bar{d} is the path metric induced by d), and since \bar{P} is 1-Lipschitz, the length of $P(\gamma_i)$ is no greater than the length of γ_i . On the other hand, if γ_i is contained in \mathcal{B}_α , then it maps by P to the tree T_α . The restriction of P to each fiber E_X , for $X \in \mathcal{B}_\alpha$ is 1-Lipschitz, and the horizontal directions collapse completely, and so it is easy to see that the length of $P(\gamma_i)$ is no more than that of γ_i , thus completing the proof. \square

3.7. Spines. For any $\alpha \in \mathcal{P}$ and $X \in D$, we consider the union of the set of all saddle connections in direction α in E_X . The components of this space are precisely the preimages of vertices of $v \in T_\alpha$ under the map $t_\alpha|_{E_X} = P \circ f_{X_0, X}: E_X \rightarrow T_\alpha$. For a vertex $v \in T_\alpha^{(0)} \subset \mathcal{V}$ and $X \in D$, we write $\theta_X^v = (t_\alpha|_{E_X})^{-1}(v) \subset E_X$ for this component, which we call the v -spine in E_X . The closure of each component of

$$E_X - \bigcup_{v \in T_\alpha^{(0)}} \theta_X^v$$

is an infinite strip covering a Euclidean cylinder in direction α . For each v , let Θ_X^v be the union of θ_X^v together with all these infinite strips that meet θ_X^v , which we

call the *thickened v -spine* in E_X . Note that $f_{X,Y}$ maps θ_Y^v and Θ_Y^v , respectively, to θ_X^v and Θ_X^v , respectively.

The union of the v -spines and thickened v -spines over ∂B_v are denoted

$$\theta^v = \bigcup_{X \in \partial B_v} \theta_X^v \quad \Theta^v = \bigcup_{X \in \partial B_v} \Theta_X^v.$$

These are bundles over ∂B_v which we call the *v -spine bundle* and the *thickened v -spine bundle*. These bundles are only used for organizational purposes in this paper, but will play a more fundamental role in the sequel [DDLS21].

Lemma 3.4. *There exists a constant $M > 0$ such that*

- (1) *For every $X \in \bar{D}$, every saddle connection in E_X has length at least $\frac{1}{M}$.*
- (2) *For each $v \in \mathcal{V}$ and $X \in \partial B_v$, every saddle connection in θ_X^v has length at most M and every strip in Θ_X^v has width at most M and at least $\frac{1}{M}$. In particular, for points $X \in \partial B_\alpha$, the saddle connections and strips of E_X in direction $\alpha \in \mathcal{P}$ have, respectively, uniformly bounded lengths and widths.*
- (3) *For every $X \in \bar{D}$, there is a triangulation of E_X by saddle connections so that the triangles have diameters at most M and interior angles at least $\frac{1}{M}$.*

Proof. For each $X \in D$, the fiber E_X is canonically identified with the universal cover of the closed surface S equipped with the flat metric for the quadratic differential q_X . The existence of a (finite) triangulation of (S, X, q_X) by saddle connections follows from [MS91], and we lift this to a triangulation of E_X .

Since (S, X, q_X) has finitely many cone points, there is a lower bound on the flat distance between any two cone points and hence a lower bound on the length of any saddle connection on E_X . For each $\alpha \in \mathcal{P}$, the flat surface (S, X, q_X) decomposes into finitely many cylinders in direction α whose boundary curves are geodesic concatenations of saddle connections in direction α . In particular, (S, X, q_X) has only finitely many saddle connections in the α direction. Since the strips and saddle connections in the α direction on E_X are precisely the preimages of these finitely many cylinders and saddle connections on (S, q_X) , there is a maximal width/length of any strip/saddle connection on E_X in the α direction.

Items (1) and (3) follow from compactness of \bar{D}/G . Item (2) follows from the facts that widths/lengths of strips/saddle connections in direction α over a ∂B_α are constant (Proposition 2.5) and that there are only finitely many G -orbits in \mathcal{P} . \square

3.8. A useful diagram. Figure 1 collects many of the main pieces of the setup for the paper. The truncated Teichmüller disk \bar{D} is the base of the \mathbb{H}^2 -bundle \bar{E} . The rest of the pieces include:

- (a) Three 1-separated horoballs.
- (b) A horoball in the direction β with the chosen basepoint X_β . Twisting via the multitwist $\tau_\beta \in G$ moves X_β along the horocycle.
- (c) Two fibers in the universal cover, $\partial \mathcal{B}_\beta$, of the graph manifold which is the bundle over the horocycle in direction β . The fibers E_{X_β} and $E_{\tau_\beta(X_\beta)}$ are above X_β and $\tau_\beta(X_\beta)$, respectively. The multitwist τ_β induces shearing in the fibers.
- (d) A horoball in the direction α , with the chosen basepoint X_α and its horocycle ∂B_α .
- (e) The fiber E_0 over the basepoint X_0 . A flat geodesic γ between cone points in the fiber E_0 is a concatenation of saddle connections. The fiber E_0 maps

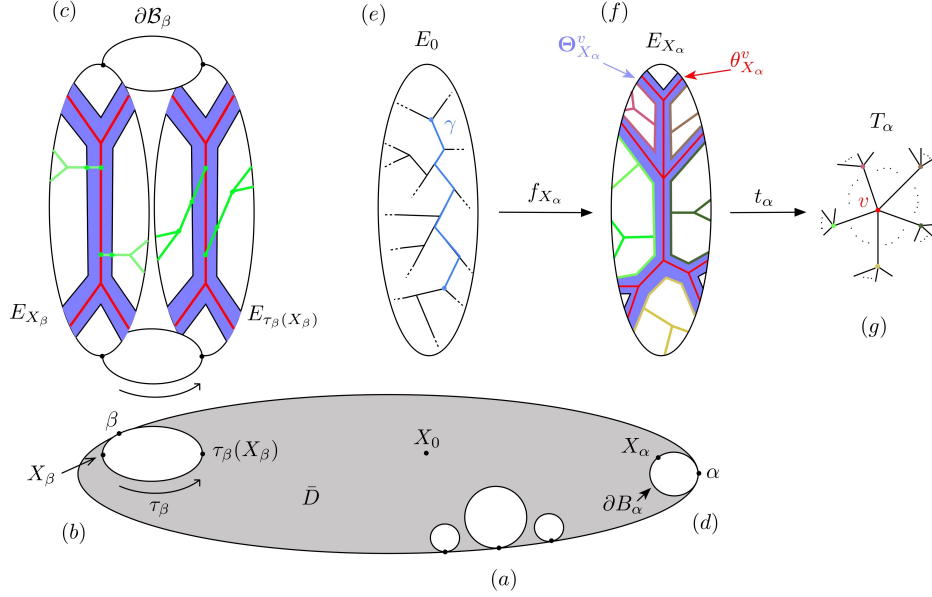


FIGURE 1. A cartoon of key aspects of E and \bar{E} over D and \bar{D} , respectively.

to the fiber E_{X_α} by the lift f_{X_α} of the Teichmüller mapping between X_0 and X_α .

- (f) The fiber E_{X_α} over the horocyclic point X_α . The spine $\theta_{X_\alpha}^v$ in direction α is in red, with the thickened spine neighborhood $\Theta_{X_\alpha}^v$ indicated in lavender. Other spines stick off from the boundary of $\Theta_{X_\alpha}^v$.
- (g) The tree T_α for the direction α with the map $t_\alpha|_{E_{X_\alpha}} : E_{X_\alpha} \rightarrow T_\alpha$. The various spines are collapsed to the vertices and two spines are connected by an edge in T_α if their neighborhoods meet. The length of an edge of T_α is the width of the strip in E_{X_α} mapping to it.

3.9. Combinatorial paths and distances. In studying \bar{E} and \hat{E} , it will be helpful to utilize certain well-behaved paths that, in particular, allow us to understand when pairs of points are bounded distance apart. The following lemma provides coarse geometric information about distances in combinatorial terms. It will be used primarily in the sequel [DDLS21].

Lemma 3.5. *There exists $R > 0$ so that any two points $x, y \in \bar{\Sigma}$ are connected in \bar{E} by a path of length at most $Rd(x, y)$ that is a concatenation of at most $Rd(x, y) + 1$ pieces, each of which is either a saddle connection of length at most R in a vertical vertical fiber, or a horizontal geodesic segment in \bar{E} .*

Proof. Take a geodesic γ' in \bar{E} joining $x, y \in \bar{\Sigma}$. Let $L \leq \bar{d}(x, y)$ be the length of the projected path $\bar{\pi}(\gamma')$ in \bar{D} , and choose an integer n so that $n - 1 \leq L < n$. Divide γ' into n subpaths $\gamma' = \gamma'_1 \cdots \gamma'_n$ such that $\bar{\pi}(\gamma'_i)$ has length $L/n \leq 1$. Let x'_{i-1}, x'_i be the endpoints of γ'_i and set $X_i = \pi(x'_i) \in \bar{D}$. For each $0 < i < n$, the diameter bound for a triangulation as in Lemma 3.4 allows us to connect $x'_i \in E_{X_i}$ to some cone point $x_i \in \Sigma_{X_i}$ by a vertical path of length at most M . Append these

to the endpoints of γ'_i to obtain a path γ_i from x_{i-1} to x_i with $\bar{\pi}(\gamma_i) = \bar{\pi}(\gamma'_i)$. These concatenate to give a path $\gamma = \gamma_1 \cdots \gamma_n$ in \bar{E} joining x to y of total length

$$\text{length}(\gamma) \leq \bar{d}(x, y) + 2M(n-1) \leq \bar{d}(x, y) + 2ML \leq (2M+1)\bar{d}(x, y).$$

For each $0 < i \leq n$, set $y_i = f_{X_i, X_{i-1}}(x_{i-1})$ and connect x_{i-1} to y_i by a horizontal geodesic h_i . Note that h_i has length at most that of $\bar{\pi}(\gamma'_i)$. As noted in §3.3, $f_{X, Y}$ is $e^{\rho(X, Y)}$ -biLipschitz, and hence on $\pi^{-1}(B_1(X_i))$, f_{X_i} is e -Lipschitz. Therefore $\mu'_i = f_{X_i}(\gamma_i)$ a path in E_{X_i} whose length is at most e times the length of γ_i . Pushing μ'_i into the 1-skeleton of the triangulation of E_{X_i} produces a new path μ_i , whose length grows by another fixed factor, that connects y_i to x_i via a sequence of vertical saddle connections of length at most M .

Putting it all together, we have a path $h_1\mu_1 \cdots h_n\mu_n$ from x to y consisting of horizontal geodesics h_i and vertical saddle connections of length at most M . The horizontal pieces contribute total length

$$\sum_{i=1}^n \bar{d}\text{-length}(h_i) \leq \sum_{i=1}^n \bar{\rho}\text{-length}(\bar{\pi}(\gamma'_i)) \leq \bar{d}\text{-length}(\gamma') = \bar{d}(x, y).$$

Similarly, the vertical pieces contribute total length at most a fixed multiple of $\bar{d}\text{-length}(\gamma) \leq (2M+1)\bar{d}(x, y)$. Hence the length of $h_1\mu_1 \cdots h_n\mu_n$ is bounded as desired. Since each saddle connection has length bounded below (Lemma 3.4), the number of saddle connections in this concatenation is linearly bounded by $\bar{d}(x, y)$, and the number of horizontal pieces is $n \leq L+1 \leq \bar{d}(x, y) + 1$. \square

The following lemma allows us to approximate arbitrary points in \hat{E} by the “nicer” subset $\mathcal{V} \subset \hat{E}$.

Lemma 3.6. *There exists $R_0 > 0$ so that \mathcal{V} is R_0 -dense in \hat{E} .*

Proof. This follows from the fact that \mathcal{V} is Γ -invariant and that \hat{E}/Γ is compact, being the continuous image of \bar{E}/Γ under the descent of $\bar{P}: \bar{E} \rightarrow \hat{E}$. \square

We will control distances in \hat{E} with the following type of nicely behaved paths.

Definition 3.7. A *horizontal jump* in \hat{E} is the image under \bar{P} of a geodesic in \bar{D}_z , for some $z \in \bar{\Sigma}$, that connects two components of $\partial\bar{D}_z$ and whose interior is disjoint from $\partial\bar{D}_z$. A *combinatorial path* in \hat{E} is a concatenation of horizontal jumps.

Note that every horizontal jump in \hat{E} connects points of \mathcal{V} , by construction, and has length at least 1. Indeed, for distinct $\alpha, \alpha' \in \mathcal{P}$, the points $\hat{\pi}(P(\mathcal{B}_\alpha))$ and $\hat{\pi}(P(\mathcal{B}_{\alpha'}))$ have distance at least 1. Since $\hat{\pi}$ is 1-Lipschitz, any path joining $\partial\mathcal{B}_\alpha$ to $\partial\mathcal{B}_{\alpha'}$ in \bar{E} thus projects to a path of length at least 1 in \hat{E} . In particular, the number of jumps in a combinatorial path is bounded by the path’s total length.

Lemma 3.8. *There is constant $C > 0$ such that any pair of points $x, y \in \mathcal{V}$ may be connected by a combinatorial path of length at most $C\hat{d}(x, y)$ that, in particular, consists of at most $C\hat{d}(x, y)$ horizontal jumps.*

Proof. For $\alpha \in \mathcal{P}$, let $\bar{\ell}_\alpha$ be the path metric on $\partial\mathcal{B}_\alpha \subset \bar{E}$. Also let ℓ_α be the path metric on $T_\alpha \subset \hat{E}$ which, recall from Lemma 3.2, is the \mathbb{R} -tree metric dual to the foliation of E_{X_α} . Finally, let $\bar{\Sigma}_\alpha = \bar{\Sigma} \cap \partial\mathcal{B}_\alpha$ and note that there exists some $K \geq 30$, independent of $\alpha \in \mathcal{P}$, such that each $z \in \partial\mathcal{B}_\alpha$ satisfies $\bar{d}(z, w) \leq K/30$ for some $w \in \bar{\Sigma}_\alpha$; indeed by Lemma 3.4(3) we may take $K = 30M$. As a first step towards the lemma, we show how to jump between points of $\bar{P}(\bar{\Sigma}_\alpha)$:

Claim 3.9. *There exists $R' > 0$ such that any pair of points $v_1, v_2 \in \bar{P}(\bar{\Sigma}_\alpha) \subset \mathcal{V}$ may be connected by a combinatorial path of length at most $R'\ell_\alpha(v_1, v_2)$.*

Proof of Claim 3.9. By decomposing a T_α -geodesic from v_1 to v_2 into its edges, it suffices to suppose v_1 and v_2 are adjacent vertices of T_α . Pick any point $X \in \partial B_\alpha$ and consider the restriction $\bar{P}|_{E_X}: E_X \rightarrow T_\alpha$. The preimage of v_1, v_2 under this map are adjacent spines $\theta_X^{v_1}, \theta_X^{v_2}$ which are separated by a strip whose width is bounded by Lemma 3.4. Therefore we may pick a bounded-length saddle connection $\sigma \subset E_X$ joining cone points $y_1 \in \theta_X^{v_1}$ to $y_2 \in \theta_X^{v_2}$.

Let $\beta \in \mathcal{P}$ be the direction of σ . Since σ has bounded length, X is bounded distance in D from some point Y in the horocycle ∂B_β . Let $z_i = f_{Y,X}(y_i)$ and let h_i be the horizontal geodesic in D_{y_i} from y_i to z_i . By definition, the \bar{P} -image of each component of $h_i \cap \bar{D}_{y_i}$ is a horizontal jump in \hat{E} . Each of these jumps has length at most $\rho(X, Y)$ which itself is uniformly bounded. The saddle connection $f_{Y,X}(\sigma)$ in E_Y is collapsed to a point by \bar{P} . Hence taking the jumps from h_1 and then coming back along the jumps from h_2 gives a bounded length combinatorial path in \hat{E} from v_1 to v_2 . Since $\ell_\alpha(v_1, v_2)$ is uniformly bounded below (by the minimal length of an edge in T_α) the claim follows. \square

Claim 3.10. *If $x, y \in \mathcal{V}$ and $r = \hat{d}(x, y) > 0$, then x and y can be connected by a combinatorial path of length $\leq 4R'e^{Kr}$.*

Proof. Let $x_1, y_1, \dots, x_n, y_n \in \bar{E}$ be any tight chain from $\bar{P}(x)$ to $\bar{P}(y)$, as in Lemma 3.2, with $\sum_{i=1}^n \bar{d}(x_i, y_i) \leq 2r$. By tightness and the fact $\bar{P}(x_1), \bar{P}(y_n) \in \mathcal{V}$, we have $x_i, y_i \in \partial \bar{E}$ for all $1 \leq i \leq n$. We note, moreover, that y_i, x_{i+1} lie in a common boundary component $\partial \mathcal{B}_\alpha$ for each $1 \leq i < n$.

We call a pair x_i, y_i along this chain a skip (in contrast to a jump) if $x_i \in \partial \mathcal{B}_\alpha$ and $y_i \in \partial \mathcal{B}_\beta$ for distinct directions $\alpha, \beta \in \mathcal{P}$. Observe that in this case $\bar{d}(x_i, y_i) \geq 1$ since the horoballs B_α, B_β are 1-separated. Choose nearby points $w_i \in \bar{\Sigma}_\alpha$ and $z_i \in \bar{\Sigma}_\beta$ so that $\bar{d}(x_i, w_i), \bar{d}(z_i, w_i) \leq K/30$. Therefore

$$\bar{d}(x_i, w_i) + \bar{d}(w_i, z_i) + \bar{d}(z_i, y_i) \leq \frac{4K}{30} + \bar{d}(x_i, y_i) \leq \frac{K}{6} \bar{d}(x_i, y_i).$$

Now insert copies of w_i, z_i to get a new chain $\dots, x_i, w_i, w_i, z_i, z_i, y_i, \dots$ (of $2(n+2)$ elements) where the skip w_i, z_i is between points of $\bar{\Sigma} \cap \partial \bar{E}$. Making an insertion for each skip and relabeling, as necessary, we henceforth assume our chain x_1, \dots, y_n has $\sum_{i=1}^n \bar{d}(x_i, y_i) \leq Kr/3$, lies in $\partial \bar{E}$, and only skips between points of $\bar{\Sigma}$.

Consider again a skip x_i, y_i of our improved chain, say with $x_i \in \bar{\Sigma}_\alpha$. Let $Z_i \in \bar{D}$ be the closest point on ∂B_α to $\bar{\pi}(y_i)$, and let $z_i \in \partial \mathcal{B}_\alpha \cap \bar{D}_{y_i} \subset \bar{\Sigma}_\alpha$ be its lift to the Teichmüller disk \bar{D}_{y_i} . By construction of \bar{d} , the $\bar{\rho}$ -geodesic from $\bar{\pi}(y_i)$ to $Z_i = \bar{\pi}(z_i)$ lifts to a path from y_i to z_i in \bar{E} with the same length. Since any path γ' from y_i to $\partial \mathcal{B}_\alpha$ has

$$\text{length}(\gamma') \geq \text{length}(\bar{\pi}(\gamma')) \geq \bar{\rho}(\bar{\pi}(z_i), \bar{\pi}(y_i)) = \bar{d}(y_i, z_i),$$

we see that $\bar{d}(z_i, y_i) \leq \bar{d}(x_i, y_i)$. Therefore, by the triangle inequality, we have

$$\bar{d}(x_i, z_i) + \bar{d}(z_i, y_i) \leq \bar{d}(x_i, y_i) + 2\bar{d}(y_i, z_i) \leq 3\bar{d}(x_i, y_i).$$

This means we may insert z_i to get a new chain $\dots, x_i, z_i, z_i, y_i, \dots$ whose distance sum is at most Kr . Making such an adjustment for each skip and relabeling if necessary, we may now assume that our chain x_1, \dots, y_n lies in $\partial \bar{E}$, that $\sum_{i=1}^n \bar{d}(x_i, y_i) \leq Kr$, and that each skip is horizontal and between points of $\bar{\Sigma}$,

meaning that $x_i, y_i \in \bar{\Sigma}$ lie in the same Teichmüller disk $\bar{D}_{x_i} = \bar{D}_{y_i}$ whenever they lie in distinct components of $\partial\bar{E}$.

Now suppose that $\alpha \in \mathcal{P}$ and $1 \leq j \leq k \leq n$ are such that the points $x_j, y_j, \dots, x_k, y_k$ all lie in $\partial\mathcal{B}_\alpha$. For each $j \leq i \leq k$, the pair x_i, y_i may be joined by a path γ_i in \bar{E} of length at most $2\bar{d}(x_i, y_i) \leq 2Kr$. Hence γ_i is contained in the Kr -neighborhood of $\partial\mathcal{B}_\alpha$ in \bar{E} . Since the map $f_\alpha: \bar{E} \rightarrow \partial\mathcal{B}_\alpha$ is e^{Kr} -Lipschitz on this set and restricts to the identity of $\partial\mathcal{B}_\alpha$, we conclude $f_\alpha(\gamma_i)$ has length at most $2e^{Kr}\bar{d}(x_i, y_i)$. Since \bar{P} is 1-Lipschitz, the images of the paths $f_\alpha(\gamma_i)$ under \bar{P} therefore concatenate to yield a path in T_α from $\bar{P}(x_j)$ to $\bar{P}(x_k)$ of length at most

$$\sum_{i=j}^k \text{length}(f_\alpha(\gamma_i)) \leq 2e^{Kr} \sum_{i=j}^k \bar{d}(x_i, y_i).$$

Since $\bar{P}(x_k), \bar{P}(x_j) \in \bar{P}(\bar{\Sigma}_\alpha)$, we may now use Claim 3.9 to construct a combinatorial path from $\bar{P}(x_j)$ to $\bar{P}(x_k)$ of length at most $2R'e^{Kr} \sum_{i=j}^k \bar{d}(x_i, y_i)$.

On the other hand, for each skip x_i, y_i in our chain the horizontal geodesic in $\bar{D}_{x_i} = \bar{D}_{y_i}$ from x_i to y_i projects to a combinatorial path from $\bar{P}(x_i)$ to $\bar{P}(y_i)$ of length at most $\bar{d}(x_i, y_i)$. Concatenating these with the paths produced above for each maximal subchain $x_j, y_j, \dots, x_k, y_k$ in some common component $\partial\bar{\mathcal{B}}_\alpha$, we finally produce a combinatorial path from $x = \bar{P}(x_1)$ to $y = \bar{P}(y_n)$ of length at most $2R'e^{Kr} \sum_{i=1}^n \bar{d}(x_i, y_i) \leq 4R'e^{Kr}r$. \square

We now show lemma holds for $C = 36R'e^{3KR_0}$, where R_0 is the density constant from Lemma 3.6. If $r = \hat{d}(x, y) \leq 3R_0$, then the above produces a combinatorial path from x to y of length at most $\frac{C}{9}\hat{d}(x, y)$, as desired. If $r > 3R_0$, then we may join x to y by a path γ of length at most $2r$. Subdivide γ into $n = \lceil \text{length}(\gamma)/R_0 \rceil$ subsegments of equal length at most R_0 . By Lemma 3.6, each subdivision point is within distance R_0 of \mathcal{V} . In this way, we obtain a sequence $x = x_0, \dots, x_n = y$ in \mathcal{V} with $\hat{d}(x_i, x_{i+1}) \leq 3R_0$ for each i . Connecting each x_i to x_{i+1} by a combinatorial path of length at most $\frac{C}{3}R_0$, we obtain a combinatorial path from x to y of length at most $\frac{C}{3}R_0n$. Since $R_0n < \text{length}(\gamma) + R_0 \leq 3\hat{d}(x, y)$, we are done. \square

4. HYPERBOLICITY OF \hat{E}

The goal of this section is to prove the following.

Theorem 4.1. *The space \hat{E} is hyperbolic.*

This is achieved by applying the hyperbolicity criterion given by Proposition 2.2 (Guessing geodesics) to the collection of (collapsed) *preferred paths*—a special type of combinatorial path as in Definition 3.7—which we now describe.

4.1. Preferred paths. For any two points $x, y \in \Sigma$, our next goal is to construct a particular path $\varsigma(x, y)$ from x to y . Recall that the map $f = f_{X_0}: E \rightarrow E_0 = E_{X_0}$ was defined in §3.3. Further, recall from §3.6 that for each $\alpha \in \mathcal{P}$ we have fixed a point $X_\alpha \in B_\alpha$. To begin, we connect the points $f(x), f(y) \in E_0$ by a geodesic segment $[f(x), f(y)]$ in the flat metric q of E_0 . This geodesic is a concatenation of saddle connections in E_0 ,

$$[f(x), f(y)] = \sigma_1\sigma_2 \cdots \sigma_k,$$

where the saddle connection σ_i has direction $\alpha_i \in \mathcal{P}$. For each $1 \leq i \leq k$, let z_{i-1}, z_i be the initial and terminal endpoints, respectively, of σ_i .

Roughly speaking, $\zeta(x, y)$ is the path constructed as follows. First, start at x and follow a horizontal geodesic in E to the fiber $E_{X_{\alpha_1}} \subset \partial\mathcal{B}_{\alpha_1}$ where saddle connections in direction α_1 are short. The path then traverses the saddle connection σ_1 in $E_{X_{\alpha_1}}$, then continues on along a horizontal geodesic to the fiber $E_{X_{\alpha_2}} \subset \partial\mathcal{B}_{\alpha_2}$ and traverses the next saddle connection σ_2 . The path continues in this way traversing horizontal geodesic segments followed by short saddle connections, as dictated by $[f(x), f(y)]$, until all of the saddle connections σ_i have been traversed and the path can go to y along a horizontal geodesic.

More formally, we let $\zeta(x, y)$ be the concatenation of segments

$$(4) \quad \zeta(x, y) = h_0\gamma_1h_1\gamma_2h_2 \cdots \gamma_k h_k$$

defined as follows. For each $i = 1, \dots, k$, γ_i is the saddle connection

$$\gamma_i = f_{X_{\alpha_i}}(\sigma_i) \subset E_{X_{\alpha_i}}.$$

The paths h_i are horizontal geodesic segments making $\zeta(x, y)$ into a path. More precisely, let γ_i^-, γ_i^+ denote the initial and terminal endpoints of γ_i , respectively, and observe that for $i = 1, \dots, k-1$, we have

$$\gamma_i^+, \gamma_{i+1}^- \in D_{z_i}, \quad x, \gamma_0^- \in D_{z_0}, \quad \text{and} \quad y, \gamma_k^+ \in D_{z_k}.$$

Then for $i = 1, \dots, k-1$, h_i is the geodesic from γ_i^+ to γ_{i+1}^- in D_{z_i} , and h_0 and h_k are the geodesic segments from x to γ_0^- and γ_k^+ to y in D_{z_0} and D_{z_k} , respectively. We call the segments γ_i the *saddle pieces* and the h_i the *horizontal pieces*. We note that when γ_i and γ_{i+1} are in the same direction, then h_i is degenerate. Also observe that the construction is symmetric: $\zeta(x, y)$ and $\zeta(y, x)$ are the same paths with opposite orientations.

We call $\zeta(x, y)$ the *preferred path* from x to y . We can push preferred paths forward via the 1-Lipschitz map $P: E \rightarrow \bar{E}$ (Lemma 3.3) and thus consider the image $\hat{\zeta}(x, y) = P(\zeta(x, y))$. We will call $\hat{\zeta}(x, y)$ a *collapsed preferred path*. Note that collapsed preferred paths are combinatorial paths. The key fact about these paths needed to prove that \hat{E} is hyperbolic is the following.

Theorem 4.2. *There exists $\delta > 0$ so that collapsed preferred paths form δ -slim triangles. That is, for any $x, y, z \in \Sigma$, we have*

$$\hat{\zeta}(x, y) \subset N_\delta(\hat{\zeta}(x, z) \cup \hat{\zeta}(y, z)).$$

We divide the proof of this theorem into a sequence of lemmas, which requires some further setup and notation, and occupies the bulk of this section. Before we do that, however, we assume Theorem 4.2 and use it to prove Theorem 4.1.

Although the preferred paths depend on the choice of basepoints X_α for each $\alpha \in \mathcal{P}$, all choices produce collapsed preferred paths that have uniformly bounded Hausdorff distance to geodesics (see Proposition 2.2), and hence each other. The fact that that all choices are uniformly bounded Hausdorff distance can also be shown more directly via the arguments that follow. In any case, we do not need this fact to apply Proposition 2.2.

4.2. Hyperbolicity. Given $u, v \in \mathcal{V}$, define

$$L(u, v) = \bigcup \hat{\zeta}(x, y),$$

where the union is taken over all $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$ (see §3.7 for notation). Since $\zeta(x, y)$ is a finite length path connecting x and y and P is Lipschitz, it follows that $L(u, v)$ is a rectifiably path connected set containing both u and v .

Lemma 4.3. *Suppose $u, v \in \mathcal{V}$. For any $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$ we have*

$$(5) \quad L(u, v) \subset N_{2\delta}(\hat{\zeta}(x, y))$$

where $\delta > 0$ is the constant from Theorem 4.2.

Proof. Let $x' \in \theta^u \cap \Sigma$ and $y' \in \theta^v \cap \Sigma$ be any points. Since $[f(x), f(x')]$ is a concatenation of saddle connections all in direction $\alpha(u)$ and since $x, x' \in \mathcal{B}_{\alpha(u)}$, we have $\hat{\zeta}(x, x') = \{u\}$. Similarly, $\hat{\zeta}(y, y') = \{v\}$. Now by Theorem 4.2, we have

$$\hat{\zeta}(x', y) \subset N_{\delta}(\hat{\zeta}(x, y) \cup \hat{\zeta}(x, x')) = N_{\delta}(\hat{\zeta}(x, y)),$$

and

$$\hat{\zeta}(x', y') \subset N_{\delta}(\hat{\zeta}(x', y) \cup \hat{\zeta}(y, y')) = N_{\delta}(\hat{\zeta}(x', y)) \subset N_{2\delta}(\hat{\zeta}(x, y)).$$

Since x', y' were arbitrary, the result follows. \square

Recall from Lemma 3.6 that \mathcal{V} is R_0 dense in \hat{E} .

Lemma 4.4. *There exists a constant $C > 0$ so that if $u, v \in \mathcal{V}$ with $\hat{d}(u, v) \leq 3R_0$, then $\text{diam}(L(u, v)) \leq C$.*

Proof. Let $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$. By Lemma 3.8, there is a bounded length combinatorial path connecting $P(x) = u$ to $P(y) = v$ which is a concatenation of n horizontal jumps, where the bound on the length and the number n depends only on R_0 . Each such horizontal jump is a collapsed preferred path.

By iterated application of Theorem 4.2 one can show that $\hat{\zeta}(x, y)$ is contained in the $n\delta$ -neighborhood of this combinatorial path, and hence has uniformly bounded diameter. \square

The proof of Theorem 4.1 is now an easy consequence:

Proof of Theorem 4.1. Since $\mathcal{V} \subset \hat{E}$ is R_0 dense by Lemma 3.6, we need only verify the sets $L(u, v)$ satisfy the two conditions of Proposition 2.2. Theorem 4.2 and Lemma 4.3 immediately imply the sets $L(u, v)$ form 3δ -slim triangles, verifying condition (1), and Lemma 4.4 precisely gives condition (2). \square

We now record a corollary of Theorem 4.1:

Corollary 4.5. *Let $\Upsilon_1, \dots, \Upsilon_k < \Gamma$ be representatives of the conjugacy classes of vertex subgroups, and let \mathcal{S} be any finite generating set for Γ . Then the Cayley graph $\text{Cay}(\Gamma, \mathcal{S} \cup \bigcup \Upsilon_i)$ is Γ -equivariantly quasi-isometric to \hat{E} , and in particular it is hyperbolic.*

Proof. This follows from Theorem 4.1 and the fact that the Cayley graph described in the statement is quasi-isometric to \hat{E} by the version of the Schwartz-Milnor Lemma given by [CC07, Theorem 5.1].

In fact, the action of Γ on \hat{E} is cocompact since the action of Γ on \bar{E} is, and the natural map from \bar{E} to \hat{E} is continuous. Point-stabilizers for the action on \hat{E} are either trivial or conjugate into one of the subgroups described in the statement. Finally, discreteness of orbits follows considering separately points outside of the T_{α} (which have neighborhoods isometric to open sets of \bar{E}), vertices of the T_{α} , and points along edges of the T_{α} , in the latter two cases appealing to Lemma 3.2. \square

We now move on to the analysis of triangles of collapsed preferred paths necessary to prove Theorem 4.2.

4.3. (Non)degenerate triangles. The proof of Theorem 4.1 relies on an analysis of geodesic triangles in the base fiber E_0 with cone point vertices. Given any $x, y, z \in \Sigma$, write

$$\Delta(x, y, z) = \Delta(f(x), f(y), f(z)) = [f(x), f(y)] \cup [f(y), f(z)] \cup [f(z), f(x)]$$

for the associated geodesic *reference triangle* in E_0 . Such a triangle may have some *degenerate corners* where the initial geodesic segments emanating from a vertex share a number of saddle connections. See Figure 2. As illustrated in the figure, there is a subtriangle $\Delta' \subset \Delta$, any two distinct edges of which intersect in exactly one point (an endpoint). The fact that geodesic triangles in E_0 actually fit this description is easily deduced from the fact that the flat metric on E_0 is uniquely geodesic, as it is CAT(0). If $\Delta = \Delta'$, then we say that Δ is *nondegenerate*, and otherwise it is *degenerate*.

Let $\Delta^\varsigma(x, y, z) \subset E$ be the triangle of preferred paths and $\Delta^\hat{\varsigma}(x, y, z) \subset \hat{E}$ its image in the collapsed space. We say that these triangles are degenerate or nondegenerate according to whether $\Delta(x, y, z)$ is.




FIGURE 2. A general triangle Δ in E_0 . The nondegenerate subtriangle $\Delta' \subset \Delta$ is shown in bold.

Lemma 4.6. *If there exists $\delta > 0$ so that every nondegenerate triangle $\Delta^\hat{\varsigma}(x, y, z)$ is δ -slim, then the same is true for degenerate triangles of collapsed preferred paths.*

In particular, to prove Theorem 4.2, it suffices to prove that nondegenerate triangles of collapsed preferred paths are δ -slim.

Proof. Assume all nondegenerate triangles of collapsed preferred paths are δ -slim for some δ , and let $x, y, z \in \Sigma$ be any points. If $[f(x), f(y)] \cap [f(x), f(z)] = \sigma_1 \cdots \sigma_r$, where σ_i is a saddle connection, then the preferred paths from x to y and from x to z agree on the first r horizontal pieces and r saddle pieces. Likewise for the segments starting at y and z . Points along the common paths are within distance 0 of another side. After excising these common paths, we are left with a nondegenerate triangle of preferred paths, which is δ -slim by assumption. In particular, the distance from any point on one side of the original triangle either has distance 0 or distance at

most δ from some point on one of the other sides, thus the original (arbitrary) triangle is δ -slim. \square

4.4. Decomposing triangles. To prove Theorem 4.2, we can restrict our attention to points x, y, z for which the reference triangle $\Delta(x, y, z)$ in the base fiber E_0 is nondegenerate, by Lemma 4.6. Our analysis of the collapsed preferred paths of nondegenerate triangles is then carried out in steps by analyzing increasingly complicated reference triangles. We now prove two lemmas which will be useful for decomposing more complicated triangles in the base fiber into simpler pieces.

Lemma 4.7. *Given cone points $x, y, z \in \Sigma_0$, if $\Delta(x, y, z)$ is nondegenerate, then it bounds a topological 2-simplex (also denoted $\Delta(x, y, z)$) which is convex and has no cone points in its interior.*

Proof. To see that $\Delta(x, y, z)$ bounds a convex topological 2-simplex, we notice that we can obtain the simplex by prolonging the sides to geodesic lines, and intersecting half-spaces bounded by these lines.

Let $N \geq 0$ denote the number of cone points in the interior of $\Delta = \Delta(x, y, z)$, each of which has cone angle at least 3π . The double Δ' of Δ along its boundary is a topological sphere. If $a, b, c > 0$ are the interior angles of the vertices of $\Delta(x, y, z)$, then Δ' has cone angles $2a, 2b, 2c$ at these points. Every other point of Δ' has cone angle at least 2π ; this is because all interior points of $\Delta(x, y, z)$ have angle at least 2π and all non-vertex points on the boundary of Δ have interior angle at least π since the sides are geodesics. The curvature at a point is 2π minus the cone angle at that point, and each of the $2N$ cone points in Δ' coming from the N cone points in the interior of Δ have curvature $\leq -\pi$. The Gauss-Bonnet theorem implies the total curvature of Δ' (the sum of the curvatures over all the cone points) is $2\pi\chi(\Delta') = 4\pi$, and therefore

$$4\pi \leq (2\pi - 2a) + (2\pi - 2b) + (2\pi - 2c) - 2N\pi < 6\pi - 2N\pi.$$

Thus $2N\pi < 2\pi$, which is only possible if $N = 0$. \square

The next lemma will be our main tool for decomposing triangles in the base fiber.

Lemma 4.8. *Let $x, y, z \in \Sigma_0 \subset E_0$ be cone points defining a nondegenerate triangle $\Delta(x, y, z)$. Let γ be the flat geodesic from x to a cone point w in the interior of the opposite side $[y, z]$ in $\Delta(x, y, z)$. Then γ is the union of a proper subsegment ν of either $[x, y]$ or $[x, z]$ and a single saddle μ in the interior of $\Delta(x, y, z)$ ending at w .*

Proof. The geodesic γ must lie in $\Delta(x, y, z)$ by convexity (Lemma 4.7). Starting from x , if γ does not immediately enter the interior of $\Delta(x, y, z)$, then it must run along one of the edges for some time, then leave that edge. Note that once it leaves the edge it cannot return to it as this would produce two geodesic segments between a pair of points, contradicting the CAT(0) condition. The segment cannot follow that edge all the way to one of the other vertices, for then it can be continued beyond w to the third vertex, giving two distinct geodesics between x and the third vertex, a contradiction again (because $\Delta(x, y, z)$ is nondegenerate). It must therefore leave the edge it is following before reaching the vertex. By a similar reasoning, the segment cannot run along the side opposite x for any of its length, for we could then again continue to another vertex of the triangle and produce two




FIGURE 3. Possible configurations of geodesic segments from x to a cone point w in the interior of $[y, z]$ (shown as thickened lines). The portion in the interior of $\Delta(x, y, z)$ is a single saddle connection. In the left-most picture ν is just the point x .

distinct geodesics between a pair of points, a contradiction once again. Finally, the fact that the part in the interior of $\Delta(x, y, z)$ is a single saddle connection comes from the fact that there are no cone points in the interior by Lemma 4.7. \square

4.5. Euclidean triangles. We now begin by analyzing the simplest type of non-degenerate triangles. Given $x, y, z \in \Sigma_0$, the triangle $\Delta(x, y, z)$ is called *Euclidean*, if it is nondegenerate, and each side consists of a single saddle connection.

Lemma 4.9. *For any Euclidean triangle $\Delta(x, y, z)$, there exists a unique point $X \in D$ such that $f_X(\Delta(x, y, z))$ is an equilateral triangle.*

We call the point X in this lemma the *balance point* of $\Delta(x, y, z)$.

Proof. Recall that the points of D parameterize all affine deformations of the flat metric q . Since $\Delta(x, y, z)$ bounds a 2-simplex with no cone points in its interior (Lemma 4.7), we may develop the entire 2-simplex into \mathbb{C} via preferred coordinates for q . The image is a triangle in \mathbb{C} , and we choose any $A \in \mathrm{SL}_2(\mathbb{R})$ so that the associated real linear transformation sends the triangle to an equilateral triangle. Setting $(X, q_X) = A \cdot (X_0, q)$, we have $f_X(\Delta(x, y, z))$ is equilateral since the preferred coordinates for $A \cdot q$ differ from those for q by composing with A ; see §2.8. If $Y \in D$ and $f_Y(\Delta(x, y, z))$ is equilateral, then the affine map $f_{Y, X}$ maps the equilateral triangle $f_X(\Delta(x, y, z))$ to the equilateral triangle $f_Y(\Delta(x, y, z))$. Therefore $f_{X, Y}$ is conformal, and hence $X = Y$. \square

The following lemma is an immediate consequence of [SW10, Theorem 1.2] which states that there are only finitely many triangles with area above any given number $\epsilon > 0$. We will use this lemma repeatedly to bound the distance between certain pairs of horoballs in the proof of Lemma 4.12 below.

Lemma 4.10. *There exists $A > 0$ so that the area of any Euclidean triangle $\Delta(x, y, z)$ is at most A . Consequently, if X is the balance point of $\Delta(x, y, z)$, then each side of $f_X(\Delta(x, y, z))$ has length at most $2\sqrt{A}$. In particular, the balance point lies uniformly close to all 3 horoballs corresponding to the directions of the sides of $\Delta(x, y, z)$. \square*

4.6. Fans. The next simplest type of nondegenerate triangles are those built from Euclidean triangles in the pattern of a “fan”. The analysis of the associated collapsed preferred paths of triangles in this case is the key technical result behind the proof of Theorem 4.2. We now proceed to the precise description of these triangles and their analysis.

Given $x, y, z \in \Sigma$, the triangle $\Delta(x, y, z)$ in E_0 is called a *fan* if it is nondegenerate and at least two sides consist of a single saddle connection. By Lemma 4.8, each geodesic between a cone point in the interior of the third side to the vertex opposite that side is a single saddle connection. It follows that fans decompose into unions of Euclidean triangles, all of which have a common vertex. See Figure 4.

We say that two saddle connections in a fiber E_X *span a triangle* if they share an endpoint and the flat geodesic joining their other endpoints is a single (possibly degenerate) saddle connection. Given a saddle connection σ in some fiber E_X , let us write $\mathcal{P}(\sigma) \subset \mathcal{P}$ for the set of directions of all saddle connections that span a triangle with σ . Notice that the direction of σ is contained in $\mathcal{P}(\sigma)$ since σ spans a (degenerate) triangle with itself. Denote

$$B(\sigma) = \bigcup_{\alpha \in \mathcal{P}(\sigma)} B_\alpha$$

for the union of horoballs at these directions. Similarly, taking horoball preimages, we denote $\mathcal{B}(\sigma) = \bigcup_{\alpha \in \mathcal{P}(\sigma)} \mathcal{B}_\alpha$.

The utility of this notion comes from the following corollary of Lemma 4.10.

Corollary 4.11. *There exists $A' > 0$ so that for each saddle connection σ in the direction α the following holds. For any $\beta \in \mathcal{P}(\sigma)$ we have $d(B_\alpha, B_\beta) < A'$ in D .*

We now show that fans give rise to slim triangles of collapsed preferred paths, and we moreover give a criterion, in terms of the notions that we just defined, for the triangle to be coarsely degenerate.

Lemma 4.12 (Fan lemma). *There exists a constant $\delta' > 0$ such that if the geodesic triangle $\Delta(x, y, z)$ in E_0 is a fan, where $x, y, z \in \Sigma$, then the triangle $\Delta^{\hat{\zeta}}(x, y, z)$ of collapsed preferred paths is δ' -slim. If, furthermore, $[f(x), f(y)] = \sigma_z$ and $[f(y), f(z)] = \sigma_x$ are each a single saddle connection and y lies on a geodesic in D_y with respective endpoints in $\mathcal{B}(\sigma_z)$ and $\mathcal{B}(\sigma_x)$, then the collapsed preferred paths satisfy*

$$\hat{\zeta}(x, y), \hat{\zeta}(y, z) \subset N_{\delta'}(\hat{\zeta}(x, z)) \quad \text{and} \quad \hat{\zeta}(x, z) \subset N_{\delta'}(\hat{\zeta}(x, y) \cup \hat{\zeta}(y, z)).$$

We will make frequent (sometimes implicit) use of the following fact about the geometry of the hyperbolic plane; see e.g. [GM08, Lemma 4.5].

Lemma 4.13. *If points $v, w \in D$ respectively lie within bounded distance of horoballs B_α, B_β , for $\alpha, \beta \in \mathcal{P}$, then the geodesic joining v to w lies within a uniform neighborhood of $B_\alpha \cup B_\beta$ and the shortest geodesic joining these horoballs. \square*

Hence, if h_1, h_2 are horizontal geodesics in some Teichmüller disc D_x whose endpoints lie within bounded distance of the same horoball preimages \mathcal{B}_α and \mathcal{B}_β , then $P(h_1)$ and $P(h_2)$ have bounded Hausdorff distance in \hat{E} .

Proof of Lemma 4.12. In what follows, we will often lift objects (points, paths, etc) from D to a horizontal disk D_w , through a point $w \in E$. When we can do so without confusion, we will use the same name for the object in D and in D_w , to avoid introducing even more notation than is already necessary.

As depicted in Figure 4, we can express the three sides of $\Delta(x, y, z) \subset E_0$ as concatenations of saddle connections:

$$[f(x), f(y)] = \sigma_z, \quad [f(y), f(z)] = \sigma_x, \quad \text{and} \quad [f(x), f(z)] = \sigma_1 \sigma_2 \cdots \sigma_k.$$




FIGURE 4. A fan $\Delta(x, y, z) \subset E_0$, in which $[f(x), f(y)]$ and $[f(y), f(z)]$ are both a single saddle connection.

For each index $*$ $\in \{z, x, 1, \dots, k\}$, write $\alpha_* \in \mathcal{P}$ for the direction of the saddle connection σ_* and set $\gamma_* = f_{X_{\alpha_*}}(\sigma_*)$. We may then write

$$\varsigma(x, y) = h_z \gamma_z h'_z \quad \text{and} \quad \varsigma(y, z) = h'_x \gamma_x h_x,$$

where h_z is the horizontal geodesic from x to X_{α_z} in D_x , h'_z the geodesic from X_{α_z} to y in D_y , h'_x the geodesic from y to X_{α_x} in D_y , and h_x the geodesic from X_{α_x} to z in D_z . We also have

$$\varsigma(x, z) = h_0 \gamma_1 h_1 \gamma_2 h_2 \cdots \gamma_k h_k$$

where h_0 is the geodesic from x to X_{α_1} in D_x , h_k the geodesic from X_{α_k} to z in D_z , and h_i the geodesic from X_{α_i} to $X_{\alpha_{i+1}}$ in $D_{\gamma_i^+} = D_{\gamma_{i+1}^-}$ for each $1 \leq i < k$.

Let h'_y be the horizontal geodesic from X_{α_z} to X_{α_x} in D_y . The three paths h'_z , h'_y , h'_x form a geodesic triangle in the 2-hyperbolic space D_y ; thus we have

$$h'_z \subset N_2(h'_x \cup h'_y), \quad h'_x \subset N_2(h'_z \cup h'_y), \quad \text{and} \quad h'_y \subset N_2(h'_z \cup h'_x).$$

Since $P: E \rightarrow \hat{E}$ is 1-Lipschitz; the collapsed horizontal segments $P(h'_z)$, $P(h'_x)$, and $P(h'_y)$ also form a 2-slim triangle. Therefore, the first conclusion of the lemma follows from the following claim, which we prove below.

Claim 4.14. *The collapsed paths $P(h_z \gamma_z h'_y \gamma_x h_x)$ and $P(h_0 \gamma_1 h_1 \gamma_2 h_2 \cdots \gamma_k h_k) = P(\varsigma(x, z))$ have uniformly bounded Hausdorff distance.*

Now we argue that the claim also implies the ‘furthermore’ conclusion of the lemma. Assume that y lies on a geodesic g in D_y with endpoints in $\mathcal{B}(\sigma_z)$ and $\mathcal{B}(\sigma_x)$. This means that there are saddle connections σ'_z and σ'_x spanning triangles T_z and T_x with σ_z and σ_x , respectively, such that the endpoints of g lie over the horoballs B'_z and B'_x for the directions of σ'_z and σ'_x (meaning that π of the endpoints lie on the prescribed horoballs). By Corollary 4.11, the endpoints of h'_y lie over horoballs close to B'_z and B'_x , respectively. In particular, $P(g)$ and $P(h'_y)$ lie within uniformly bounded Hausdorff distance (Lemma 4.13).

We now see that the sets

$$P(h'_z) \cup P(h'_x) \quad \text{and} \quad P(h'_y)$$

have uniformly bounded Hausdorff distance, since h'_z is a geodesic sharing an endpoint with h'_y , while the P -image of the other endpoint y is close to $P(h'_y)$, and similarly for h'_x . Therefore, the ‘furthermore’ conclusion of the lemma follows from Claim 4.14 as well. This concludes the proof of Lemma 4.12, assuming Claim 4.14. \square

Before we prove Claim 4.14, we need a little more notation and another claim. We let $T_i \subset E_0$ be the Euclidean triangle determined by $f(y)$ and the saddle connection σ_i , that is, with vertices $f(y)$, σ_i^- , and σ_i^+ . We call σ_i the base of T_i and $\tau_i^\pm = [\sigma_i^\pm, f(y)]$ the sides of T_i . Let λ_i and ρ_i denote the angles of T_i between σ_i and the sides τ_i^- and τ_i^+ , respectively; see Figure 4. For any $X \in D$, we write $T_i(X)$, $\sigma_i(X)$, $\lambda_i(X)$, etc., for the corresponding pieces of the image fan $\Delta_X = f_X(\Delta(x, y, z))$ in the fiber E_X . Let $b_i \in D$ denote the balance point for the Euclidean triangle T_i , and h the geodesic in D from b_1 to b_k .

The next claim is the key to proving Claim 4.14 (and is perhaps the biggest miracle in the proof of hyperbolicity of \hat{E}). Roughly, it says that as we traverse h from b_1 to b_k in D , each of the saddle connections $\sigma_1, \dots, \sigma_k$ become bounded in length when viewed in the fibers over h , and that the times when they are bounded occur in order.

Claim 4.15. *There are points t_1, \dots, t_k appearing in order along the geodesic h that respectively lie within uniformly bounded distance of the horoballs $B_{\alpha_1}, \dots, B_{\alpha_k}$.*

Proof of Claim 4.15. Denote lengths of saddle connections by $\text{len}(\cdot)$. We find points t_i as in the statement with $\text{len}(\sigma_i(t_i))$ uniformly bounded for each $i = 1, \dots, k$. By Lemma 4.10, we may take $t_1 = b_1$ and $t_k = b_k$.

First observe that for any $X \in D$ we have

$$\lambda_1(X) < \lambda_2(X) < \dots < \lambda_k(X) \quad \text{and} \quad \rho_1(X) > \rho_2(X) > \dots > \rho_k(X).$$

Indeed, $\lambda_i(X) + \rho_i(X) < \pi$ by the fact that these are two angles of the Euclidean triangle $T_i(X)$, whereas $\rho_i(X) + \lambda_{i+1}(X) \geq \pi$ by the fact that $\sigma_1(X) \cdots \sigma_k(X)$ is a geodesic in E_X . Thus $\lambda_i(X) < \lambda_{i+1}(X)$, proving the claimed inequalities for the λ_i ; the ones for the ρ_i can be proven similarly.

Since $T_1(b_1)$ and $T_k(b_k)$ are equilateral, for each $1 < i < k$ we thus have

$$\rho_i(b_1) < \rho_1(b_1) = \frac{\pi}{3} = \lambda_k(b_k) > \lambda_i(b_k).$$

On the other hand, since $\sigma_1(X) \cdots \sigma_k(X)$ forms a geodesic in any fiber E_X , we also have that $\rho_{k-1}(b_k) + \lambda_k(b_k) \geq \pi$ and $\rho_1(b_1) + \lambda_2(b_1) \geq \pi$. Hence, again for each $1 < i < k$, we have

$$\rho_i(b_k) \geq \rho_{k-1}(b_k) \geq \frac{2\pi}{3} \leq \lambda_2(b_1) \leq \lambda_i(b_1).$$

Using the intermediate value theorem, for each $1 < i < k$ we define $l_i, r_i \in D$ to be the first points along the geodesic h such that $\lambda_i(l_i) = \pi/2$ and $\rho_i(r_i) = \pi/2$.

Let us use \leq to denote the natural order along h (oriented from b_1 to b_k). Notice that $\lambda_i(r_i) < \pi/2$ (since $\rho_i(r_i) + \lambda_i(r_i) < \pi$). As we also have $\lambda_i(b_1) > \pi/2$, the definition of l_i implies that $l_i \leq r_i$. The fact $\rho_i(l_{i+1}) + \lambda_{i+1}(l_{i+1}) \geq \pi$ further

implies $\rho_i(l_{i+1}) \geq \pi/2$. Since $\rho_i(b_1) < \pi/2$, the definition of r_i gives $r_i \leq l_{i+1}$. Thus our points along the geodesic $h = [b_1, b_k]$ are ordered as

$$b_1 = t_1 \leq l_2 \leq r_2 \leq \dots \leq l_{k-1} \leq r_{k-1} \leq t_k = b_k.$$

To complete the proof, it now suffices to show, for each $1 < i < k$, that there is a point $t_i \in [l_i, r_i]$ with $\sigma_i(t_i)$ of uniformly bounded length. Then by Lemma 4.10 we have

$$\begin{aligned} \text{len}(\tau_i^-(l_i)) \cdot \text{len}(\sigma_i(l_i)) &= 2\text{area}(T_i(l_i)) \leq 2A, \quad \text{and} \\ \text{len}(\tau_i^+(r_i)) \cdot \text{len}(\sigma_i(r_i)) &= 2\text{area}(T_i(r_i)) \leq 2A. \end{aligned}$$

If $\text{len}(\sigma_i(l_i)) \leq 2\sqrt{A}$ or $\text{len}(\sigma_i(r_i)) \leq 2\sqrt{A}$ we are done. Otherwise $\text{len}(\tau_i^-(l_i)) \leq \sqrt{A}$ and $\text{len}(\tau_i^+(r_i)) \leq \sqrt{A}$. Thus l_i lies within bounded distance of the horoball for the direction τ_i^- , and r_i lies within bounded distance of the horoball for the direction τ_i^+ . Since the balance point b_i for T_i lies close to both of these horoballs by Lemma 4.10, the balance point lies near the shortest geodesic joining the horoballs. It follows from Lemma 4.13 that the geodesic $[l_i, r_i]$ passes near b_i , at which point σ_i has bounded length, again by Lemma 4.10. \square

We now commence with the proof of Claim 4.14.

Proof of Claim 4.14. Let h_y denote the horizontal geodesic in D_y from b_1 to b_k , that is, h_y is the horizontal lift of h to D_y . By Lemma 4.10, b_1 and b_k respectively lie within bounded distance of the horoballs B_z and B_x . Since h'_y runs between these horoballs as well, we see that $P(h'_y)$ and $P(h_y)$ have bounded Hausdorff distance. The horizontal paths h_z and h_0 in D_x go from x to the respective horoballs B_{α_z} and B_{α_1} which lie bounded distance from each other (they are near the balance point b_1 for T_1). Thus we similarly conclude that $P(h_z)$ and $P(h_0)$ have bounded Hausdorff distance. Symmetrically, the same holds for $P(h_x)$ and $P(h_k)$.

By the above remarks, and using the facts that P is Lipschitz and the vertical paths $\gamma_z, \gamma_x, \gamma_1, \dots, \gamma_k$ each have uniformly bounded length once projected to \hat{E} by P , it suffices to show that the sets

$$P(h_y) \quad \text{and} \quad P(h_1 \cup \dots \cup h_{k-1})$$

have bounded Hausdorff distance.

For $1 \leq i < k$, let g_i denote the horizontal geodesic in D_y between the balance points b_i and b_{i+1} of the triangles T_i and T_{i+1} . Similarly, let h'_i be the horizontal geodesic between these points b_i and b_{i+1} in the disc $D_{\gamma_i^+} = D_{\gamma_{i+1}^-}$. Since the saddle connection $\tau_i^+ = \tau_{i+1}^-$ from $\gamma_i^+ = \gamma_{i+1}^-$ to $f(y)$ is common to both triangles T_i and T_{i+1} , Lemma 4.10 implies that b_i and b_{i+1} both lie within bounded distance of the horoball for the direction of τ_i^+ . Thus the saddle connection τ_i^+ has bounded length over the whole geodesic $[b_i, b_{i+1}]$, and we conclude that the paths g_i and h'_i have bounded Hausdorff distance in E .

Using Claim 4.15, let s_i denote the lift of $t_i \in h$ to h_y . This gives a decomposition $h_y = h_y^1 \cdots h_y^{k-1}$ of h_y into segments

$$h_y^i = [s_i, s_{i+1}], \quad \text{for } i = 1, \dots, k-1.$$

By construction of t_i , we see that h_y^i begins and ends near the horoballs B_{α_i} and $B_{\alpha_{i+1}}$ in $D_{\gamma_i^+}$. Since the balance points b_i and b_{i+1} lie near these horoballs as well, it follows that $P(h_y^i)$ and $P(g_i)$ are at bounded Hausdorff distance. By the previous

paragraph $P(g_i)$ and $P(h'_i)$ are at bounded Hausdorff distance, and hence so are $P(h'_y)$ and $P(h'_i)$. Therefore, $P(h_y)$ and $P(h'_1 \cup \dots \cup h'_{k-1})$ have bounded Hausdorff distance.

To complete the proof of Claim 4.14, it now suffices to show that

$$P(h'_i) \quad \text{and} \quad P(h_i)$$

have uniformly bounded Hausdorff distance for each $1 \leq i < k$. But this is clear: The endpoints b_i and b_{i+1} of h'_i lie uniformly close to the horoballs B_{α_i} and $B_{\alpha_{i+1}}$ containing the endpoints X_{α_i} and $X_{\alpha_{i+1}}$ of h_i . Since both paths h'_i and h_i lie in the same disc $D_{\gamma_i^+} = D_{\gamma_{i+1}^-}$, it follows that $P(h'_i)$ and $P(h_i)$ have uniformly bounded Hausdorff distance. This establishes Claim 4.14 and concludes the proof of Lemma 4.12. \square

4.7. Decomposing triangles into fans. We next consider the case that $\Delta(x, y, z)$ is a triangle where one side is a single saddle connection. The main idea is that such triangles decompose into an alternating union of fans at certain ‘‘pivot’’ cone points along the sides with more than one saddle connection, as in Figure 5.

Lemma 4.16. *There exists $\delta'' > 0$ so that if $x, y, z \in \Sigma$ and $\Delta(x, y, z)$ is a nondegenerate triangle so that one side is a saddle connection, then $\Delta^\xi(x, y, z)$ is δ'' -slim.*

Proof. Let $\delta' > 0$ be the constant provided by Lemma 4.12. We may assume that the side $[f(x), f(y)]$ consists of a single saddle, which we denote η . To economize notation, set $x' = y$. Write $[f(x), f(z)] = \sigma_1 \cdots \sigma_m$ as a concatenation of saddle connections, and let $f(x) = a_0, a_1, \dots, a_m = f(z)$ be the sequence of endpoints of these saddle connections. Similarly write $[f(x'), f(z)] = \sigma'_1 \cdots \sigma'_n$ with $f(x') = a'_0, a'_1, \dots, a'_n = f(z)$ the corresponding sequence of cone points.

Letting $\alpha, \alpha_i, \alpha'_j \in \mathcal{P}$ be the directions of η, σ_i and σ'_j , respectively, we then have

$$(6) \quad \varsigma(x, z) = h_0 \gamma_1 h_1 \gamma_2 \cdots \gamma_m h_m \quad \text{and} \quad \varsigma(x', z) = h'_0 \gamma'_1 h'_1 \gamma'_2 \cdots \gamma'_n h'_n,$$

where $\gamma_i = f_{X_{\alpha_i}}(\sigma_i)$ and each $h_i \subset D_{\alpha_i}$ is a horizontal geodesic making the concatenation into a path, and similarly for $\gamma'_j = f_{X_{\alpha'_j}}(\sigma'_j)$ and $h'_j \subset D_{\alpha'_j}$.

We now explain how to decompose $\Delta(x, y, z)$ into a union of fans as in Figure 5, where the base points (called pivots) of the fans alternate sides of the triangle.

First, we claim that there is a single saddle connection joining $f(x) = a_0$ to a'_1 or else there is a single saddle connection joining $f(x') = a'_0$ to a_1 (or both). If not, then the geodesics $[a_0, a'_1]$ and $[a'_0, a_1]$ are both nontrivial concatenations of saddle connections joining cone points in $\Delta(x, x', z) \subset E_0$.

Since $[a_0, a'_0]$ is a single saddle connection, Lemma 4.8 implies that $\sigma_1 = [a_0, a_1]$ is the first saddle connection in the geodesic $[a_0, a'_1]$. In particular, $[a_1, a'_1]$ is a subpath of $[a_0, a'_1]$. Similarly, $\sigma'_1 = [a'_0, a'_1]$ is the first saddle connection of the geodesic $[a'_0, a_1]$, which contains $[a'_1, a_1]$ as a subpath. These observations imply that $\sigma_1[a_1, a'_1]\sigma'_1$ is a concatenation of saddle connections making angle at least π on both sides of each cone point encountered. Therefore it is the geodesic in E_0 connecting $f(x)$ to $f(x')$. But we are assuming that this geodesic consists of a single saddle connection, a contradiction.

By the previous paragraph, we may without loss of generality assume that $[a_0, a'_1]$ is a single saddle connection. Let $j_1 \in \{1, \dots, n\}$ be the largest index such that $[a_0, a'_{j_1}]$ is a single saddle connection.

If $j_1 = n$, then $\Delta(x, x', z)$ is evidently a fan and the conclusion follows from Lemma 4.12. We may therefore assume $j_1 < n$, in which case the cone points $a_0, a'_{j_1}, f(z)$ span a nondegenerate geodesic triangle with the side $[a_0, a'_{j_1}]$ consisting of a single saddle connection. The above observations now imply that $[a'_{j_1}, a_1]$ is a single saddle connection, and we are justified in letting $i_1 \in \{1, \dots, m\}$ be the largest index such that $[a'_{j_1}, a_{i_1}]$ is a single saddle connection.

Continuing in this manner, we recursively choose indices $j_1, i_1, j_2, i_2, \dots$, terminating when some $j_k = n$ or some $i_k = m$, with the defining property that j_{k+1} is the largest index in $\{j_k + 1, \dots, n\}$ such that $[a_{i_k}, a'_{j_{k+1}}]$ is a single saddle connection and i_{k+1} is the largest index in $\{i_k + 1, \dots, m\}$ such that $[a'_{j_{k+1}}, a_{i_{k+1}}]$ is a single saddle connection. In this way, we decompose the triangle $\Delta(x, x', z)$ into fans based at the *pivot vertices* $a_0, a'_{j_1}, a_{i_1}, \dots$. For the sake of argument, we may assume some $i_k = n$ so that the situation is as depicted in Figure 5.




FIGURE 5. Decomposing a triangle into fans based at pivot vertices.

For each of the pivot vertices, $a_0, a'_{j_1}, a_{i_1}, \dots, a_{i_{k-1}}, a'_{j_k}$, we choose a point in the corresponding Teichmüller disk as follows: Recall our notation (6) for the preferred paths $\varsigma(x, z)$ and $\varsigma(x', z)$. Let $x_0 \in D_{a_0}$ be the terminal endpoint of the initial horizontal segment h_0 of $\varsigma(x, z)$. That is, x_0 is the intersection of h_0 and γ_1 , and thus lies in the designated fiber $E_{X_{\alpha_1}}$ over the the boundary of the horoball ∂B_{α_1} for the direction α_1 of the saddle connection σ_1 . For each $s \in \{1, \dots, k\}$, let $x'_s \in D_{a'_{j_s}}$ be any point on the horizontal segment h'_{j_s} of $\varsigma(x', z)$ corresponding to the vertex a'_{j_s} . Similarly, for any $r \in \{1, \dots, k-1\}$, let $x_r \in D_{a_{i_r}}$ be any point on the horizontal segment h_{i_r} of $\varsigma(x, z)$ corresponding to the vertex a_{i_r} . To round out the notation, we also set $x_k = z$. Notice that these choices decompose the preferred paths $\varsigma(x, z)$ and $\varsigma(x', z)$ into concatenations of preferred paths:

$$\begin{aligned} \varsigma(x, z) &= \varsigma(x, x_0)\varsigma(x_0, x_1)\varsigma(x_1, x_2)\cdots\varsigma(x_{k-1}, z), \text{ and} \\ \varsigma(x', z) &= \varsigma(x', x'_1)\varsigma(x'_1, x'_2)\cdots\varsigma(x'_k, z). \end{aligned}$$

Now, for each $r \in \{1, \dots, k-1\}$, the three points $x'_r, x_r, x'_{r+1} \in \Sigma$ satisfy all of the hypotheses of the Fan Lemma 4.12, including the furthermore hypothesis on the location of the pivot vertex with respect to the horoballs for the adjacent saddle

connections. Therefore, Lemma 4.12 implies the sets

$$\hat{\zeta}(x'_r, x_r) \cup \hat{\zeta}(x_r, x'_{r+1}) \quad \text{and} \quad \hat{\zeta}(x'_r, x'_{r+1})$$

have Hausdorff distance at most δ' . Similarly, for each $s \in \{1, \dots, k\}$, we may apply Lemma 4.12 to the points $x_{s-1}, x'_s, x_s \in \Sigma$ to bound the Hausdorff distance between

$$\hat{\zeta}(x_{s-1}, x'_s) \cup \hat{\zeta}(x'_s, x_s) \quad \text{and} \quad \hat{\zeta}(x_{s-1}, x_s)$$

by δ' . Finally, observe that Lemma 4.12 implies $\Delta^\zeta(x, x', x_0)$ is δ' -slim.

We now show that the triangle $\Delta^\zeta(x, x', z)$ is $2\delta'$ -slim. The above shows that the projected path $\hat{\zeta}(x', z) = \hat{\zeta}(x', x'_1) \cup \dots \cup \hat{\zeta}(x'_k, z)$ is contained in the δ' -neighborhood of the set

$$\left(\hat{\zeta}(x', x_0) \right) \cup \left(\hat{\zeta}(x_0, x'_1) \cup \hat{\zeta}(x'_1, x_1) \cup \dots \cup \hat{\zeta}(x_{k-1}, x'_k) \cup \hat{\zeta}(x'_k, z) \right),$$

which is in turn contained in the δ' neighborhood of

$$\left(\hat{\zeta}(x', x) \cup \hat{\zeta}(x, x_0) \right) \cup \left(\hat{\zeta}(x_0, x_1) \cup \dots \cup \hat{\zeta}(x_{k-1}, z) \right) = \hat{\zeta}(x', x) \cup \hat{\zeta}(x, z).$$

Thus $\hat{\zeta}(x', z)$ is contained in the $2\delta'$ -neighborhood of $\hat{\zeta}(x', x) \cup \hat{\zeta}(x, z)$.

A similar argument shows that the projected preferred path

$$\hat{\zeta}(x, z) = \left(\hat{\zeta}(x, x_0) \right) \cup \left(\hat{\zeta}(x_0, x_1) \cup \dots \cup \hat{\zeta}(x_{k-1}, z) \right)$$

is contained in the 2δ -neighborhood of $\hat{\zeta}(x, x') \cup \hat{\zeta}(x', z)$

Finally, we see that $\hat{\zeta}(x, x')$ is contained in the δ' -neighborhood of

$$\hat{\zeta}(x, x_0) \cup \hat{\zeta}(x_0, x')$$

which itself is contained in the δ' -neighborhood of

$$\hat{\zeta}(x, x_0) \cup \hat{\zeta}(x', x'_1) \subset \hat{\zeta}(x, z) \cup \hat{\zeta}(x', z).$$

Thus each side of $\Delta^\zeta(x, x', z) = \Delta^\zeta(x, y, z)$ is contained in the $(2\delta' + 2)$ -neighborhood of the union of the other two, which proves the Lemma. \square

4.8. Proving that general triangles are thin. We are now ready to prove Theorem 4.2, which again we do by decomposing a general triangle into simpler ones, this time of the previous type where one side of the reference triangle in the base fiber is a single saddle connection.

Proof of Theorem 4.2. Let $\delta'' > 0$ be as provided by Lemma 4.16 and set $\delta = 3\delta''$. Let $x, y, z \in \Sigma$ be any three points, and we must prove that $\Delta^\zeta(x, y, z)$ is δ -slim. By Lemma 4.6, we may assume that $\Delta(x, y, z)$ is nondegenerate. We show that $\hat{\zeta}(x, y)$ is in $N_\delta(\hat{\zeta}(x, z) \cup \hat{\zeta}(y, z))$. If any side of $\Delta(x, y, z)$ is a saddle connection, this follows from Lemma 4.16, so we suppose that all sides have a least two saddle connections.

We appeal to Lemma 4.8 to decompose $\Delta^\zeta(x, y, z)$ into two or three triangles, depending on the configuration of geodesics from $f(x)$ to points in $[f(y), f(z)]$. To describe the decomposition, recall that for each cone point w on the side $[f(y), f(z)]$ of $\Delta(x, y, z) = \Delta(f(x), f(y), f(z))$, other than $f(y)$ and $f(z)$, the geodesic segment from $f(x)$ to w is the concatenation of two geodesic segments $\nu\mu$, where:

- (1) ν is a geodesic subsegment of either $[f(x), f(y)]$ or $[f(x), f(z)]$, possibly a single point, and

- (2) μ is a single saddle connection with interior contained in the interior of $\Delta(x, y, z)$.

See Figure 3 above.

If for some cone point w in the interior of $[f(y), f(z)]$, ν degenerates to the point $f(x)$ (as in the left-most picture in Figure 3), we subdivide $\Delta(x, y, z)$ into two triangles along the segment $[f(x), w]$. We use this to subdivide $\Delta^\zeta(x, y, z)$ into two triangles by choosing some point \tilde{w} along the horizontal piece of the preferred path $\zeta(y, z)$ with $f(\tilde{w}) = w$, and subdividing into $\Delta^\zeta(x, y, \tilde{w})$ and $\Delta^\zeta(x, z, \tilde{w})$. Since \tilde{w} lies on the preferred path $\zeta(y, z)$, we note that $\zeta(y, z) = \zeta(y, \tilde{w}) \cup \zeta(\tilde{w}, z)$. Noting also that the two triangles $\Delta(x, y, \tilde{w}) = \Delta(f(x), f(y), w)$ and $\Delta(f(x), f(z), w)$ in the subdivision of $\Delta(x, y, z)$ are both nondegenerate with one side consisting of a single saddle connection, Lemma 4.16 implies that $\Delta^\zeta(x, y, \tilde{w})$ and $\Delta^\zeta(x, z, \tilde{w})$ are δ'' -slim. Thus two applications of Lemma 4.16 give the desired containment:

$$\begin{aligned} \hat{\zeta}(x, y) &\subset N_{\delta''}(\hat{\zeta}(y, \tilde{w}) \cup \hat{\zeta}(\tilde{w}, x)) \\ &\subset N_{\delta''}(\hat{\zeta}(y, \tilde{w})) \cup N_{2\delta''}(\hat{\zeta}(\tilde{w}, z) \cup \hat{\zeta}(z, x)) \\ &\subset N_{2\delta''}(\hat{\zeta}(y, z) \cup \hat{\zeta}(z, x)). \end{aligned}$$

It remains to consider the case that for every cone point w in the interior of $[f(y), f(z)]$, the geodesic does *not* consist of a single saddle connection as in the left-most picture of Figure 3. Let $f(y) = w_0, w_1, \dots, w_k = f(z)$ be the sequence of cone points along the $[f(y), f(z)]$. By hypothesis, for each $i \in \{0, \dots, k\}$, the geodesic $[f(x), w_i]$ has nondegenerate intersection ν with either $[f(x), f(y)]$ or $[f(x), f(z)]$. We may thus partition the index set

$$\{0, \dots, k\} = Y \sqcup Z$$

according to whether this initial segment ν of $[f(x), w_i]$ lies in $[f(x), f(y)]$ (for $i \in Y$) or in $[f(x), f(z)]$ (for $i \in Z$). Notice $0 \in Y$ and $k \in Z$ by construction. Furthermore, if $i \in Y$ then $j \in Y$ for all $j \leq i$; indeed, this follows from the fact that the geodesic $[f(x), w_j]$ must lie in the convex subset $\Delta(f(x), f(y), w_i)$ of $\Delta(f(x), f(y), f(z))$. Therefore, there exists an index $0 \leq i < k$ such that $Y = \{0, \dots, i\}$ and $Z = \{i+1, \dots, k\}$.

There are three cases to consider: Firstly, if $0 < i < k-1$, we choose points $\tilde{w}_i, \tilde{w}_{i+1} \in \zeta(y, z)$ along the corresponding horizontal pieces of the preferred path so that $f(\tilde{w}_i) = w_i$ and $f(\tilde{w}_{i+1}) = w_{i+1}$. The preferred paths $\zeta(x, y)$ and $\zeta(x, \tilde{w}_i)$ then share a degenerate initial segment, and we let \tilde{v}_i be the endpoint of this common initial segment; that is, $\zeta(x, y) \cap \zeta(x, \tilde{w}_i) = \zeta(x, \tilde{v}_i)$. Similarly choose \tilde{v}_{i+1} so that $\zeta(x, z) \cap \zeta(x, \tilde{w}_{i+1}) = \zeta(x, \tilde{v}_{i+1})$. This decomposes $\Delta(x, y, z)$ into three nondegenerate triangles $\Delta(\tilde{v}_i, y, \tilde{w}_i)$, $\Delta(x, \tilde{w}_i, \tilde{w}_{i+1})$, and $\Delta(\tilde{v}_{i+1}, \tilde{w}_{i+1}, z)$, each of which having one side a single saddle connection. Therefore we may apply Lemma 4.16

three times to conclude the containment

$$\begin{aligned}
\hat{\zeta}(y, x) &= \hat{\zeta}(y, \tilde{v}_i) \cup \hat{\zeta}(\tilde{v}_i, x) \\
&\subset N_{\delta''}(\hat{\zeta}(y, \tilde{w}_i) \cup \hat{\zeta}(\tilde{w}_i, \tilde{v}_i)) \cup \hat{\zeta}(\tilde{v}_i, x) \\
&\subset N_{\delta''}(\hat{\zeta}(y, \tilde{w}_i) \cup \hat{\zeta}(\tilde{w}_i, x)) \\
&\subset N_{2\delta''}(\hat{\zeta}(y, \tilde{w}_i) \cup \hat{\zeta}(\tilde{w}_i, \tilde{w}_{i+1}) \cup \hat{\zeta}(\tilde{w}_{i+1}, x)) \\
&= N_{2\delta''}(\hat{\zeta}(y, \tilde{w}_{i+1}) \cup \hat{\zeta}(\tilde{w}_{i+1}, \tilde{v}_{i+1}) \cup \hat{\zeta}(\tilde{v}_{i+1}, x)) \\
&\subset N_{3\delta''}(\hat{\zeta}(y, \tilde{w}_{i+1}) \cup \hat{\zeta}(\tilde{w}_{i+1}, z) \cup \hat{\zeta}(z, \tilde{v}_{i+1}) \cup \hat{\zeta}(\tilde{v}_{i+1}, x)) \\
&= N_{3\delta''}(\hat{\zeta}(y, z) \cup \hat{\zeta}(z, x)).
\end{aligned}$$

Secondly, if $i = 0$, we similarly choose a point $\tilde{w}_1 \in \zeta(y, z)$ along a horizontal piece such that $f(\tilde{w}_1) = w_1$. Since $1 \in Z$ by hypothesis, the paths $\zeta(x, z)$ and $\zeta(x, \tilde{w}_1)$ share an initial segment and we again choose $\tilde{v}_1 \in \zeta(x, z)$ so that $\zeta(x, z) \cap \zeta(x, \tilde{w}_1) = \zeta(x, \tilde{v}_1)$. This decomposes $\Delta(x, y, z)$ into two nondegenerate triangles $\Delta(x, y, \tilde{w}_1)$ and $\Delta(\tilde{v}_1, \tilde{w}_1, z)$ which each have a side consisting of a single saddle connection. Thus we may apply Lemma 4.16 twice, as above, to obtain

$$\begin{aligned}
\hat{\zeta}(y, z) &\subset N_{\delta''}(\hat{\zeta}(y, \tilde{w}_1) \cup \hat{\zeta}(\tilde{w}_1, x)) \\
&= N_{\delta''}(\hat{\zeta}(y, \tilde{w}_1) \cup \hat{\zeta}(\tilde{w}_1, \tilde{v}_1) \cup \hat{\zeta}(\tilde{v}_1, x)) \\
&\subset N_{2\delta''}(\hat{\zeta}(y, \tilde{w}_1) \cup \hat{\zeta}(\tilde{w}_1, z) \cup \hat{\zeta}(z, \tilde{v}_1) \cup \hat{\zeta}(\tilde{v}_1, x)) \\
&= N_{2\delta''}(\hat{\zeta}(y, z) \cup \hat{\zeta}(z, x)),
\end{aligned}$$

showing that $\Delta^{\hat{\zeta}}(x, y, z)$ is $2\delta''$ -slim. The remaining case $i = k - 1$ is handled symmetrically by choosing $\tilde{w}_{k-1} \in \zeta(y, z)$ with $f(\tilde{w}_{k-1}) = w_{k-1}$, and choosing \tilde{v}_{k-1} so that $\zeta(x, y) \cap \zeta(x, \tilde{w}_{k-1}) = \zeta(x, \tilde{v}_{k-1})$. This decomposes $\Delta(x, y, z)$ into two triangles $\Delta(y, \tilde{v}_{k-1}, \tilde{w}_{k-1})$ and $\Delta(x, \tilde{w}_{k-1}, z)$ to which we may apply Lemma 4.16 and conclude $\hat{\zeta}(x, y) \subset N_{2\delta''}(\hat{\zeta}(y, z) \cup \hat{\zeta}(z, x))$ as above. This covers all the cases and completes the proof of Theorem 4.2. \square

4.9. Loxodromic elements. Recall that an isometry ϕ of a hyperbolic space Y is *loxodromic* if its *translation length*

$$t(\phi) = \lim_{n \rightarrow \infty} \frac{d(y, \phi^n(y))}{n} \quad \text{for any } y \in Y$$

is positive. Now that we know \hat{E} is hyperbolic, we can characterize which elements of Γ act loxodromically on \hat{E} . The characterization is quite simple to state:

Proposition 4.17. *An element $\gamma \in \Gamma$ acts loxodromically on \hat{E} if and only if γ does not stabilize any vertex $v \in \mathcal{V}$.*

Proof. Obviously γ cannot be loxodromic if it fixes a point. Conversely, suppose γ does not fix any point of \mathcal{V} and let $g \in G$ be the image of γ in the quotient. If g is a pseudo-Anosov element of $G \leq \text{Mod}(S)$, then g acts with positive translation length on \hat{D} and therefore γ acts with positive translation length on \hat{E} (since $\hat{\pi}: \hat{E} \rightarrow \hat{D}$ is Lipschitz). Hence, after passing to a power if necessary, we may assume that g is parabolic or trivial in G . Either way, we may choose some $\alpha \in \mathcal{P}$ that is fixed by g .

It follows that γ preserves the subsets $\partial\mathcal{B}_\alpha \subset \bar{E}$ and $T_\alpha \subset \hat{E}$. Since γ does not stabilize any vertex of T_α , it restricts to a loxodromic isometry of the tree T_α . Let

$$\omega = \cdots \omega_{-2} \omega_{-1} \omega_0 \omega_1 \omega_2 \cdots$$

be the bi-infinite axis of γ in T_α , viewed as an edge path in which each ω_i denotes an edge of T_α between, say, vertices $v_{i-1}, v_i \in T_\alpha^{(0)} \subset \mathcal{V}$.

For any $X \in \partial B_\alpha$ and $i \in \mathbb{Z}$, let $\Omega_X^i \subset E_X$ be the closure of the preimage of the interior of the edge ω_i under $P|_{E_X}$. Thus Ω_X^i is a closed strip in direction α that is foliated by lines in direction α and separates E_X into two pieces. Note that the strip Ω_X^i intersects each of the trees $\theta_X^{v_{i-1}}$ and $\theta_X^{v_i}$ in a line, namely, one of the boundary components of the strip. Hence, two consecutive Ω_X^i and Ω_X^{i+1} have intersection which is either empty, a saddle connection of $\theta_X^{v_i}$, or a single cone-point of $\theta_X^{v_i}$. Let

$$\Omega_X = \bigcup_i \Omega_X^i$$

be the union of these strips. Each connected component Ω' of Ω_X is a union of strips $\Omega_X^j \cup \Omega_X^{j+1} \cup \dots \cup \Omega_X^k$ where consecutive ones $\Omega_X^i, \Omega_X^{i+1}$ have non-empty intersection. Note that each such Ω' is convex, and that the connected components of its boundary are each contained in a spine $\theta_X^{v_i}$ for some $i \in \mathbb{Z}$.

Let $r: \partial B_\alpha \rightarrow \omega$ be the composition of $P|_{\partial B_\alpha}: \partial B_\alpha \rightarrow T_\alpha$ with the closest-point-projection $T_\alpha \rightarrow \omega$. Since P is Γ -equivariant and γ is an ω -preserving isometry of T_α , this map r is equivariant with respect to the action of the cyclic group $\langle \gamma \rangle$; that is, $r(\gamma x) = \gamma r(x)$. We note that $r|_{E_X} = r \circ f_{X,Y}$ for any $X, Y \in \partial B_\alpha$, which follows from the fact that in this case $P|_{E_Y} = P \circ f_{X,Y}$. Next define a map

$$r_\omega: \mathcal{V} \rightarrow 2^\omega \quad \text{given by} \quad r_\omega(v) = r(\theta_X^v) \subset \omega \quad \text{for any} \quad X \in \partial B_\alpha,$$

which is well-defined since $\theta_X^v = f_{X,Y}(\theta_Y^v)$. Furthermore, since $\gamma\theta_X^v = \theta_{gX}^{\gamma v}$, we have

$$\gamma r_\omega(v) = \gamma r(\theta_X^v) = r(\gamma\theta_X^v) = r(\theta_{gX}^{\gamma v}) = r_\omega(\gamma v).$$

That is, r_ω is equivariant with respect to the actions of γ on \mathcal{V} and ω .

The heart of the proposition is captured in the following claim:

Claim 4.18. *There exists $M > 0$ such that $\text{diam}_{T_\alpha}(r_\omega(v)) \leq M$ for all $v \in \mathcal{V}$.*

Assuming the claim, let us deduce the proposition. For $u, v \in \mathcal{V}$, define

$$d_\omega(u, v) = \text{diam}_{T_\alpha}(r_\omega(u) \cup r_\omega(v)).$$

If u, v are connected by a horizontal jump in \hat{E} , then the spines θ_X^u and θ_X^v intersect for any $X \in D$. Thus by the claim $d_\omega(u, v) \leq 2M$. By Lemma 3.8, there is some C such that any two vertices $u, v \in \mathcal{V}$ may be connected by a combinatorial path consisting of at most $C\hat{d}(u, v)$ jumps. Therefore

$$d_\omega(u, v) \leq 2MC\hat{d}(u, v) \quad \text{for all} \quad u, v \in \mathcal{V}.$$

Finally, if $t(\gamma) > 0$ denotes the translation length of γ acting on T_α , then by definition for all $v \in \mathcal{V}$ and $n \geq 0$ the equivariance $r_\omega(\gamma^n v) = \gamma^n r_\omega(v)$ implies that

$$2MC\hat{d}(v, \gamma^n v) \geq d_\omega(v, \gamma^n v) = \text{diam}_{T_\alpha}(r_\omega(v) \cup \gamma^n r_\omega(v)) \geq nt(\gamma).$$

Thus γ indeed acts loxodromically on \hat{E} as claimed.

It remains to prove Claim 4.18. Let us first describe the restriction of r to E_X for $X \in \partial B_\alpha$. Since P maps Ω_X to ω , we see that r agrees with P on Ω_X . For each component U of $E_X \setminus \Omega_X$, there exists some $i \in \mathbb{Z}$ such that $r(\bar{U}) = \{v_i\}$ with $\partial\bar{U} \subset \theta_X^{v_i}$. From this it follows that if $r_\omega(v)$ is not a single point, then $\theta_X^v \cap \Omega_X \neq \emptyset$ and moreover that $r(\theta_X^v \cap \bar{U}) = r(\theta_X^v \cap \partial\bar{U})$ for every such component U . Since

$\partial\bar{U} \subset \Omega_X$, we conclude $r_\omega(v) = r(\theta_X^v \cap \Omega_X)$ in this case. Hence, in Claim 4.18 it suffices to bound $\text{diam}(r(\theta_X^v \cap \Omega_X))$.

We also note that when $r_\omega(v)$ is not a singleton, $\theta_X^v \cap \Omega_X$ is contained in a convex component of Ω_X . Indeed, the convex set $(P|_{E_X})^{-1}(\omega)$ is the union of Ω_X with $\cup_{i \in \mathbb{Z}} \theta_X^{v_i}$. Thus if θ_X^v were to intersect distinct components of Ω_X , then θ_X^v must intersect some $\theta_X^{v_i}$ in a saddle connection. But that implies $\alpha(v) = \alpha(v_i) = \alpha$ and therefore that $P(\theta_X^v)$ is a single point, contrary to our assumption on v . \square

We now complete the proposition by proving Claim 4.18 in two cases:

Proof of Claim 4.18 when $g = 1$. Fix any $X \in \partial B_\alpha$. In this case $\gamma\Omega_X = \Omega_{gX} = \Omega_X$, so that γ gives an isometry of E_X that preserves the convex set Ω_X .

The strips Ω_X^i limit, as $i \rightarrow \pm\infty$, on two points in the CAT(0) boundary of E_X . Let ℓ denote a geodesic joining these boundary points, which is a γ -invariant, bi-infinite geodesic that crosses each strip Ω_X^i . Note that ℓ cannot lie in any thickened spine Θ_X^v , since then γ would fix $v \in \mathcal{V}$. Therefore the saddle connections comprising ℓ do not all have the same direction. Write

$$\ell = \dots \sigma_{-2}\sigma_{-1}\sigma_0\sigma_1\sigma_2 \dots$$

where each σ_i is a maximal concatenation of saddle connections in a single direction. There exists $\mu > \pi$ so that σ_i and σ_{i+1} form an angle of least μ on both sides.

Consider any vertex $v \in \mathcal{V}$. By the above discussion, it suffices to take any points $y, z \in \theta_X^v \cap \Omega_X$ with $r(y) \neq r(z)$. Let U be the convex component of Ω_X containing $\theta_X^v \cap \Omega_X$. Choose $i \in \mathbb{Z}$ so that $y \in \Omega_X^i$, and let y' be the unique point of $\ell \cap \Omega_X^i$ with $r(y') = y$. That is, y' is the intersection of ℓ with the leaf \mathcal{F}_y through y of the foliation of Ω_X^i in direction α (note that $\ell \cap \Omega_X^i$ is a segment in a direction distinct from α). Choose $z' \in \ell$ and \mathcal{F}_z similarly. Finally, let $\eta \subset \theta_X^v$ be the geodesic from y to z and $\ell_0 \subset \ell$ the geodesic from y' to z' .

If ℓ_0 and η are disjoint, then the geodesics $\mathcal{F}_y, \eta, \mathcal{F}_z, \ell_0$ form a geodesic quadrilateral in U with no cone points in its interior (since there are no cone points in the interior of Ω_X). Otherwise ℓ_0 and η intersect and we get two geodesic triangles connected by the (possibly degenerate) segment $\ell_0 \cap \eta$. By Gauss–Bonnet, doubling this picture produces an object W that is either a sphere with total curvature 4π (in the quadrilateral case) or two spheres (when $\ell_0 \cap \eta \neq \emptyset$) with total curvature 8π . The only points of positive curvature come from the four corners y, y', z, z' and the two endpoints of $\ell_0 \cap \eta$ (when it exists). Further, each of these points contributes positive curvature at most 2π . On the other hand, each cone point juncture $\sigma_j \sigma_{j+1}$ in the interior of $\ell_0 \setminus \eta$ contributes negative curvature at most $2(\pi - \mu) < 0$ to W . Note also that $\ell_0 \cap \eta$ is necessarily contained in a single segment σ_j by construction, since all of η lies in the single direction $\alpha(v)$. Therefore, if m denotes the number of segments σ_j that have nondegenerate intersection with ℓ_0 we conclude that

$$4\pi \leq 2\pi\chi(W) \leq 6(2\pi) + 2(\pi - \mu)(m - 3) \implies m \leq \frac{4\pi}{\mu - \pi} + 3.$$

Since ℓ is γ -invariant, there is some maximum number N of strips crossed by any segment σ_i . Thus ℓ_0 crosses at most $3N + 4N\pi/(\mu - \pi)$ strips, and we are done. \square

Proof of Claim 4.18 when $g \neq 1$. In this case g is a parabolic element of the Veech group $G = G_q$ fixing the direction $\alpha \in \mathbb{P}^1(q)$. Moreover, the iterates g^m converge uniformly on compact subsets of $\mathbb{P}^1(q) \setminus \{\alpha\}$ to the constant map $\mathbb{P}^1(q) \setminus \{\alpha\} \mapsto \{\alpha\}$.

Note that any geodesic connecting the two boundary components of a strip Ω_X^i is a straight segment in a single direction with no cone points in its interior. Thus for $i \in \mathbb{Z}$ we are justified in defining $A_X^i \subset \mathbb{P}^1(q)$ to be the set of directions of all segments of the form $\ell \cap \Omega_X^i$, where ℓ is any geodesic segment from Ω_X^{i-1} to Ω_X^{i+1} . Note that such a segment $\ell \cap \Omega_X^i$ has its endpoints in the closest-point-projections of $\Omega_X^{i\pm 1}$ to Ω_X^i . Since such a closest-point-projection is equal to a saddle connection or cone point in $\partial\Omega_X^i$, we see that A_X^i is a compact subset of $\mathbb{P}^1(q) \setminus \{\alpha\}$.

Let $k \in \mathbb{N}$ be such that $\gamma\omega_i = \omega_{i+k}$ for all $i \in \mathbb{Z}$. Thus for any $X \in \partial B_\alpha$ we have $\gamma\Omega_X^i = \Omega_{gX}^{i+k}$. By definition, it follows that $gA_X^i = A_{gX}^{i+k}$. Since the affine maps $f_{Y,Z}: E_Z \rightarrow E_Y$ fix the directions of lines (viewed in $\mathbb{P}^1(q)$), we conclude that

$$A_X^{i+k} = f_{X,gX}(A_{gX}^{i+k}) = gA_X^i.$$

By this and the uniform convergence of the maps $\{g^m\}$ on $\mathbb{P}^1(q) \setminus \{\alpha\}$, we may choose $m \geq 1$ such that $A_X^{i+km} = g^m A_X^i$ and A_X^i are disjoint for all $i \in \mathbb{Z}$.

To prove the claim, consider any $v \in \mathcal{V}$. Note that all saddle connections of θ_X^v lie in one fixed direction $\alpha(v)$. Suppose θ_X^v intersects strips Ω_X^{i-1} and Ω_X^{i+km+1} . It follows that θ_X^v intersects every strip Ω_X^j with $i \leq j \leq i+km$, since these each separate Ω_X^{i-1} from Ω_X^{i+km+1} . Therefore the direction $\alpha(v)$ lies in both A_X^i and A_X^{i+km} , contradicting our choice of m . This proves that θ_X^v can intersect at most $km + 2$ strips and bounds $\text{diam}_{T_\alpha}(r_\omega(v))$. \square

5. THE LOCALLY HOMOGENEOUS GEOMETRY OF E

In this section, we explain in more detail the homogeneous geometry on which E (and \bar{E}) is modeled, away from the singular locus $\Sigma \subset E$. This is one of the 4-dimensional geometries (in the sense of Thurston [Thu97, §3.8]) studied in the thesis of Filipkiewicz [Fil54, Chapter 5] (see also [Hil02, Section 7.1] and [Wal85] where the geometry is called \mathbb{F}^4). In the sequel [DDLS21] we will use the fact that isometries of E are necessarily fiber preserving. The proof of this fact requires a description of local isometries of the homogeneous geometry that we could not find in the literature, and so we carry out the required calculations here. Along the way, we compute the geometry's full isometry group (see Proposition 5.1). This seems to be known (see [Wal85]), though we could only find a computation of the component containing the identity (see [Fil54]).

To describe the geometry, consider the area-preserving affine group of \mathbb{R}^2

$$\mathfrak{G} = \text{SAff}^\pm(\mathbb{R}^2).$$

There is a homomorphism $\mathfrak{G} \rightarrow \text{SL}_2^\pm(\mathbb{R})$ sending an affine map to its derivative, expressed as a matrix of determinant ± 1 , with respect to the standard basis. The kernel is the group of translations isomorphic to \mathbb{R}^2 , and the linear action of $\text{SL}_2^\pm(\mathbb{R})$ on \mathbb{R}^2 defines a splitting as a semi-direct product,

$$\mathfrak{G} = \mathbb{R}^2 \rtimes \text{SL}_2^\pm(\mathbb{R}).$$

We compose the derivative homomorphism above with the surjective homomorphism $\text{SL}_2^\pm(\mathbb{R}) \rightarrow \text{PGL}_2(\mathbb{R})$, given by taking the quotient by the center. This in turn defines an action of \mathfrak{G} on the hyperbolic plane \mathbb{H}^2 by isometries (acting as Möbius transformations). Combining these actions, we obtain a transitive action of \mathfrak{G} on the space $\mathfrak{X} = \mathbb{H}^2 \times \mathbb{R}^2$, which is by hyperbolic isometries in the first factor and affine transformation in the second factor.

In the discussion above, it is most natural to consider the upper-half plane model of the hyperbolic plane, $\mathbb{H}^2 = \{(x, y) \in \mathbb{R}^2 \mid y > 0\}$. For any $\kappa > 0$, \mathbb{H}^2 can be equipped with the $\mathrm{PGL}_2(\mathbb{R})$ -invariant metric $ds^2 = \frac{1}{\kappa y^2}(dx^2 + dy^2)$ of constant curvature $-\kappa < 0$ ($\kappa = 4$ is the Poincaré metric and $\kappa = 1$ the hyperbolic metric). These coordinates (x, y) on \mathbb{H}^2 , together with standard coordinates (s, t) on \mathbb{R}^2 define global coordinates (x, y, s, t) on \mathfrak{X} . The stabilizer of $(0, 1, 0, 0)$ is $\mathfrak{K} = \mathrm{O}(2) < \mathrm{SL}_2^\pm(\mathbb{R})$, exhibiting $\mathfrak{X} = \mathfrak{G}/\mathfrak{K}$ as a homogeneous space for \mathfrak{G} . Since \mathfrak{K} is a compact subgroup, we can construct an invariant Riemannian metric. In the next section, we explicitly describe such a metric.

5.1. Computations. Here we define the metric g on \mathfrak{X} , and compute its expression in coordinates. Since the metric of constant curvature $-\kappa$ on each slice $\mathbb{H}^2 \times \{*\}$ is already \mathfrak{G} -invariant, we will choose our metric to agree with this metric on these slices. Projecting onto the first factor, \mathfrak{X} fibers over \mathbb{H}^2 in a \mathfrak{G} -equivariant way. We declare the \mathbb{R}^2 -fibers to be orthogonal to the slices $\mathbb{H}^2 \times \{*\}$ and define the metric on the fiber over $(0, 1)$ to be the standard Euclidean metric. Since the stabilizer in \mathfrak{G} of this fiber is $\mathbb{R}^2 \rtimes \mathrm{O}(2)$, and since this acts transitively by isometries on this metric on the fiber, we can compute this metric at $(0, 1, 0, 0)$ and push it around by a subgroup of \mathfrak{G} acting transitively on \mathfrak{X} . Finally, we note that we may *globally* scale this metric by any positive constant without affecting any of the above properties of this metric (i.e. the fact that the fibers are Euclidean, orthogonality to the \mathbb{H}^2 -slices, etc), *except* that the curvature scales. In particular, without loss of generality we may assume $\kappa = 1$ for the remainder of this section, and we do so. See the last paragraph of §5.2 for more on the choices of \mathfrak{G} -invariant metrics.

Now, at the point $(0, 1, 0, 0)$ the metric is then given by

$$g(0, 1, 0, 0) = \begin{pmatrix} \frac{1}{y^2} & 0 & 0 & 0 \\ 0 & \frac{1}{y^2} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

In terms of coordinates (x, y, s, t) on \mathfrak{X} , we want to write the metric $g(x, y, s, t)$. We use the subgroup

$$\mathfrak{H} = \left\{ \begin{pmatrix} a & b \\ 0 & \frac{1}{a} \end{pmatrix} \mid a, b \in \mathbb{R}, a > 0 \right\},$$

to translate $(0, 1, 0, 0)$ to any point $(x, y, 0, 0)$. We note that since $\mathbb{R}^2 < \mathfrak{G}$ acts transitively on each \mathbb{R}^2 -fiber with trivial derivative in these coordinates, the metric is independent of (s, t) and hence $g(x, y, s, t) = g(x, y, 0, 0)$, so pushing the metric around by the subgroup \mathfrak{H} suffices (in particular, the metric is Euclidean on each \mathbb{R}^2 -fiber).

Now the subgroup \mathfrak{H} acts as follows:

$$\begin{pmatrix} a & b \\ 0 & \frac{1}{a} \end{pmatrix} \cdot (x, y, s, t) = (a^2x + ab, a^2y, as + bt, \frac{t}{a}).$$

Given $a, b \in \mathbb{R}$, $a > 0$, set $F_{a,b}(x, y, s, t) = (a^2x + ab, a^2y, as + bt, \frac{t}{a})$, and observe that the derivative of $F_{a,b}$ at any point (x, y, s, t) is given by

$$(dF_{a,b})_{(x,y,s,t)} = \begin{pmatrix} a^2 & 0 & 0 & 0 \\ 0 & a^2 & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & 0 & \frac{1}{a} \end{pmatrix}.$$

We drop the subscript (x, y, s, t) since the derivative is constant.

Next, note that for $a = \sqrt{y}$ and $b = \frac{x}{\sqrt{y}}$ we have

$$F_{\sqrt{y}, \frac{x}{\sqrt{y}}}(0, 1, 0, 0) = (x, y, 0, 0).$$

This is a bit cumbersome, so we write $\phi_{x,y} = F_{\sqrt{y}, \frac{x}{\sqrt{y}}}$. Thus, to compute the metric at $(x, y, 0, 0)$, we can push forward the metric $g(0, 1, 0, 0)$ by the derivative of $\phi_{x,y}$. Applying $d\phi_{x,y}$ to the basis vectors, we get

$$(7) \quad \begin{aligned} d\phi_{x,y}(e_1) &= ye_1, & d\phi_{x,y}(e_2) &= ye_2, \\ d\phi_{x,y}(e_3) &= \sqrt{y}e_3, & d\phi_{x,y}(e_4) &= \frac{x}{\sqrt{y}}e_3 + \frac{1}{\sqrt{y}}e_4. \end{aligned}$$

We want to compute $g_{ij}(x, y, 0, 0) = g(x, y, 0, 0)(e_i, e_j)$, and by definition, we have

$$g(x, y, 0, 0)(d\phi_{x,y}(e_i), d\phi_{x,y}(e_j)) = g(0, 1, 0, 0)(e_i, e_j).$$

From this, the bilinearity of the metric, and (7), we obtain a system of equations implicitly defining $g(x, y, 0, 0)$. Solving this system we find

$$(8) \quad g(x, y, s, t) = g(x, y, 0, 0) = \begin{pmatrix} \frac{1}{y^2} & 0 & 0 & 0 \\ 0 & \frac{1}{y^2} & 0 & 0 \\ 0 & 0 & \frac{1}{y} & -\frac{x}{y} \\ 0 & 0 & -\frac{x}{y} & y + \frac{x^2}{y} \end{pmatrix}.$$

5.2. Isometries. We are now ready to describe the isometry group.

Proposition 5.1. *For the metric g defined by (8) we have $\text{Isom}(\mathfrak{X}, g) = \mathfrak{G}$.*

The main ingredient—and really the point of interest for us—is the following. To state it, let $\mathcal{F}_{\mathbb{H}}$ and $\mathcal{F}_{\mathbb{R}}$ denote the foliations of \mathfrak{X} whose leaves are the slices $\mathbb{H}^2 \times \{*\}$ and $\{*\} \times \mathbb{R}^2$, respectively. We also use the same name for the foliations restricted to any open subset.

Proposition 5.2. *Any local isometry between open subsets of (\mathfrak{X}, g) with g as in (8) preserves $\mathcal{F}_{\mathbb{H}}$ and $\mathcal{F}_{\mathbb{R}}$.*

Proof. This will boil down to a computation in Riemannian geometry. Specifically, we claim that the 2–planes tangent to the leaves $\{*\} \times \mathbb{R}^2$ of $\mathcal{F}_{\mathbb{R}}$ have the property that the sectional curvatures in these 2–planes are uniquely maximal among all 2–planes. More precisely, the sectional curvature in any one of these 2–planes is $\frac{1}{2}$, and in any other 2–plane it is strictly smaller.

To carry out such a computation, we can appeal to the *Mathematica* code detailed in [FOH20]. In particular, plugging the metric above in one can compute the Riemannian curvatures $R_{ijkl}(x, y, s, t)$ for $1 \leq i, j, k, \ell \leq 4$, and verify that the curvature of the plane P_0 spanned by $\frac{\partial}{\partial s}$ and $\frac{\partial}{\partial t}$ at $(0, 1, 0, 0)$ is indeed $\frac{1}{2}$. To prove

that these planes uniquely maximize the sectional curvature, we note that for any other 2-plane $P \subset T_{(0,1,0,0)}\mathfrak{X}$, the orthogonal projection to the plane P_1 spanned of $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ is nontrivial. We consider the two cases that the projection is an isomorphism, and when it maps onto a line separately.

When P projects isomorphically to P_1 , the plane is given by

$$P = \left\langle \frac{\partial}{\partial x} + \eta \frac{\partial}{\partial t} + \rho \frac{\partial}{\partial s}, \frac{\partial}{\partial y} + \lambda \frac{\partial}{\partial s} + \mu \frac{\partial}{\partial t} \right\rangle,$$

for some $\lambda, \mu, \eta, \rho \in \mathbb{R}$. From *Mathematica* computations, one obtains that the sectional curvature in the plane P is given by

$$K(\eta, \rho, \lambda, \mu) = \frac{-4 - \eta^2 - \rho^2 - \mu^2 - \lambda^2 + 6(\eta\mu - \rho\lambda) + 2(\eta\mu - \rho\lambda)^2}{4(1 + \eta^2 + \rho^2 + \lambda^2 + \mu^2 + (\eta\mu - \rho\lambda)^2)}.$$

Set $\eta = r \cos(\theta)$, $\rho = r \sin(\theta)$, $\lambda = u \cos(\phi)$, and $\mu = u \sin(\phi)$. Then observe that

$$\eta\mu - \rho\lambda = ru(\sin(\phi) \cos(\theta) - \cos(\phi) \sin(\theta)) = ru \sin(\phi - \theta),$$

so setting $\alpha = \phi - \theta$ we get

$$\begin{aligned} K(\eta, \rho, \lambda, \mu) &= \frac{-4 - r^2 - u^2 + 6ru \sin(\alpha) + 2r^2u^2 \sin^2(\alpha)}{4(1 + r^2 + u^2 + r^2u^2 \sin^2(\alpha))} \\ &= \frac{2 + 2r^2 + 2u^2 + 2r^2u^2 \sin^2(\alpha) - 6 - 3r^2 - 3u^2 + 6ru \sin(\alpha)}{4(1 + r^2 + u^2 + r^2u^2 \sin^2(\alpha))} \\ &= \frac{1}{2} - \frac{3(2 + r^2 + u^2 - 2ru \sin(\alpha))}{4(1 + r^2 + u^2 + r^2u^2 \sin^2(\alpha))} \end{aligned}$$

We want to see that the curvature above is less than $\frac{1}{2}$. It suffices to show that the numerator in the second term of the last line is positive; that is, we must show that $2 + r^2 + u^2 - 2ru \sin(\alpha) > 0$. For this, observe that $2 + r^2 + u^2 - 2ru \sin(\alpha)$ is minimized when $ru \sin(\alpha)$ is maximized, hence equal to ru (which occurs when $\alpha = \frac{\pi}{2}$). In this situation, observe that we have $2 + r^2 + u^2 - 2ru \sin(\alpha) = 2 + (r-u)^2$, and this is clearly positive. Therefore $K(\eta, \rho, \lambda, \mu) < \frac{1}{2}$.

When P projects to a line, after applying an isometry fixing $(0, 1)$ if necessary, we may assume that this line is spanned by $\frac{\partial}{\partial x}$. Then P is given by

$$P = \left\langle \frac{\partial}{\partial x} + \eta \frac{\partial}{\partial s} + \rho \frac{\partial}{\partial t}, \frac{\partial}{\partial s} + \lambda \frac{\partial}{\partial t} \right\rangle,$$

for some $\eta, \rho, \lambda \in \mathbb{R}$ or else

$$P = \left\langle \frac{\partial}{\partial x} + \eta \frac{\partial}{\partial s}, \frac{\partial}{\partial t} \right\rangle,$$

for some $\eta \in \mathbb{R}$. The sectional curvatures are then respectively given by

$$K(\eta, \rho, \lambda) = \frac{-1 - \lambda^2 + 2(\eta\lambda - \rho)^2}{4(1 + \lambda^2 + (\eta\lambda - \rho)^2)} = \frac{1}{2} - \frac{3 + 3\lambda^2}{4(1 + \lambda^2 + (\eta\lambda - \rho)^2)},$$

and

$$K(\eta) = \frac{2\eta^2 - 1}{4(1 + \eta^2)} = \frac{1}{2} - \frac{3}{4(1 + \eta^2)}.$$

These curvatures are all clearly less than $\frac{1}{2}$.

Since the sectional curvature of the tangent planes to the leaves of $\mathcal{F}_{\mathbb{R}}$ uniquely maximize curvature, any isometry must preserve this 2-plane field, and hence the foliation $\mathcal{F}_{\mathbb{R}}$. Since the tangent planes to the leaves of $\mathcal{F}_{\mathbb{H}}$ are pointwise orthogonal to those of $\mathcal{F}_{\mathbb{R}}$, this plane field must also be preserved, and hence so is $\mathcal{F}_{\mathbb{H}}$. \square

Proof of Proposition 5.1. The stabilizer of $(0, 1, 0, 0)$ in \mathfrak{G} is $O(2)$, and so it suffices to show that this is the stabilizer in $\text{Isom}(\mathfrak{X})$. Fix any isometry $T \in \text{Isom}(\mathfrak{X})$ in the stabilizer of $(0, 1, 0, 0)$. According to Proposition 5.2, T preserves $\mathcal{F}_{\mathbb{H}}$ and $\mathcal{F}_{\mathbb{R}}$, it follows that $T((x, y), (s, t)) = (T_1(x, y), T_2(s, t))$, where $T_1 \in \text{Isom}(\mathbb{H}^2)$ preserves $(0, 1)$ and $T_2 \in \text{Isom}(\mathbb{R}^2)$ preserves $(0, 0)$. Therefore, after composing with an element of $O(2) < \mathfrak{G}$, we can assume that T_1 is the identity on $\mathbb{H}^2 \times \{(0, 0)\}$, and thus $T((x, y), (s, t)) = ((x, y), T_2(s, t))$. It suffices to show that $T_2(s, t) = \pm(s, t)$. For this, note that for every $(x, y) \in \mathbb{H}^2$, the map $(s, t) \mapsto T_2(s, t)$ has to be an isometry from $\mathbb{R}^2 \cong \{(x, y)\} \times \mathbb{R}^2$ to itself, with the induced metric from \mathfrak{X} . Varying over all $(x, y) \in \mathbb{H}^2$, the metrics on \mathbb{R}^2 vary over all unit area affine deformations of the standard metric, and thus T_2 has to be an isometry with respect to all such metrics. This readily implies $T_2(s, t) = \pm(s, t)$, as required, completing the proof. \square

We note in passing that all \mathfrak{G} -invariant Riemannian metrics on \mathfrak{X} are obtained by applying independent, global scaling factors to the leaves of $\mathcal{F}_{\mathbb{R}}$ and the leaves of $\mathcal{F}_{\mathbb{H}}$ (thus producing a 2-parameter family of invariant metrics). To see this, we note that \mathfrak{G} acts transitively on the projective tangent bundles to each of the foliations $\mathcal{F}_{\mathbb{R}}$ and $\mathcal{F}_{\mathbb{H}}$, and so a \mathfrak{G} -invariant metric on the leaves of each of these two foliations are determined by the norm of a single non-zero tangent vector to a leaf at any point. Furthermore, the stabilizer of a point contains an involution that fixes the leaf of $\mathcal{F}_{\mathbb{H}}$ through the point, and acts as $-\text{Id}$ on the leaf of $\mathcal{F}_{\mathbb{R}}$ through that point, and thus the foliations must be orthogonal.

5.3. Application to E . The geometry $(\mathfrak{G}, \mathfrak{X})$ locally models the geometry of E as described in §3.6, away from the singular locus (including any orbifold locus). The next proposition immediately implies Theorem 1.3.

Proposition 5.3. *Every nonsingular point of E has a neighborhood which is isometric to an open set in \mathfrak{X} (with $\kappa = 4$). Moreover, such an isometry sends intersections with horizontal disks of E into leaves of $\mathcal{F}_{\mathbb{H}}$ and intersections with fibers $E_X \subset E$ into leaves of $\mathcal{F}_{\mathbb{R}}$.*

Proof. Write $E = D \times E_0$. Consider any nonsingular point $x \in E_0$ and consider a coordinate chart $\zeta: U \rightarrow \mathbb{R}^2$ about x for q . Define $\phi: D \times U \rightarrow \mathbb{H}^2 \times \mathbb{R}^2 = \mathfrak{X}$ by

$$(A \cdot (X, q), y) \mapsto (A^{-1}(0, 1), \zeta(y)).$$

We claim that ϕ is a local isometry. To see this, note that the restriction to each $D \times \{y\}$, for $y \in U$ is an isometry to the leaf $\mathbb{H}^2 \times \{\zeta(y)\}$ of $\mathcal{F}_{\mathbb{H}}$. We therefore need to show that for all $A \in \text{SL}_2(\mathbb{R})$, we have $(A \cdot (X, q), y) \mapsto \zeta(y)$ maps isometrically into the leaf $\{A^{-1}(0, 1)\} \times \mathbb{R}^2$ of $\mathcal{F}_{\mathbb{R}}$. For this, first note that the metric on $\{A \cdot (X, q)\} \times U$ is such that $A \circ \zeta: U \rightarrow \mathbb{R}^2$ is an isometry onto its image (with respect to the standard metric on \mathbb{R}^2). On the other hand, the metric on $\{A^{-1}(0, 1)\} \times \mathbb{R}^2 \subset \mathfrak{X}$ is such that $(A^{-1}(0, 1), s, t) \mapsto A(s, t)$ is an isometry to \mathbb{R}^2 . Precomposing this isometry with ϕ we have

$$(A \cdot (X, q), y) \xrightarrow{\phi} (A^{-1}(0, 1), \zeta(y)) \mapsto A(\zeta(y)) = A \circ \zeta(y),$$

which as noted above, is an isometry. Therefore ϕ is an isometry. \square

In the sequel [DDLS21], we study the isometry groups of E and \bar{E} , which we denote $\text{Isom}(E)$ and $\text{Isom}(\bar{E})$, respectively. Since these spaces fiber over D and

\bar{D} , the isometry groups contain subgroups consisting of the fiber-preserving isometries which we denote $\text{Isom}_{\text{fib}}(E)$ and $\text{Isom}_{\text{fib}}(\bar{E})$. In particular, we will need the following, which is an immediate corollary of Proposition 5.2 and Proposition 5.3.

Corollary 5.4. *We have $\text{Isom}(E) = \text{Isom}_{\text{fib}}(E)$ and $\text{Isom}(\bar{E}) = \text{Isom}_{\text{fib}}(\bar{E})$.*

We will also use the following fact.

Proposition 5.5. *The isometry group $\text{Isom}(E)$ acts properly discontinuously on E , and likewise for $\text{Isom}(\bar{E})$ on \bar{E} .*

Proof. Let $H = \text{Isom}(E) = \text{Isom}_{\text{fib}}(E)$ or $\text{Isom}(\bar{E}) = \text{Isom}_{\text{fib}}(\bar{E})$, and $Z = E$ or \bar{E} , respectively. Any isometry of Z preserves the singular locus, and we can thus view H as the isometry group of the Riemannian manifold Z_0 , obtained by removing this set from Z . It follows that H , with the topology of pointwise convergence is in fact a Lie group; in fact, for any five points of Z_0 sufficiently close together (and not contained in a lower dimensional geodesic submanifold), the orbit map from H to Z_0^5 is an embedding onto a closed submanifold; see [MS39]. From this it is straightforward to see that H acts properly discontinuously on Z if and only if it is a discrete group; that is, if and only if the identity is isolated.

Now suppose $\{T_n\} \subset H$ is any sequence converging to the identity and we must show that T_n is eventually the identity. Since the T_n are all fiber preserving, they descend to isometries of D (or \bar{D} , if $Z = \bar{E}$) and restrict to isometries on each fiber. The set of lengths of saddle connections on a given fiber is a closed discrete subset of \mathbb{R} which locally determines the fiber. Thus, since T_n is converging to the identity, the descent to D (or \bar{D}) must eventually be equal to the identity. It follows that for n sufficiently large, T_n restricts to an isometry of E_0 to itself. Since isometries of E_0 must send the discrete set of cone points to itself, it must eventually be the identity on three cone points that span a Euclidean triangle, and hence the identity on all of E_0 . Finally, once T_n is the identity on D (or \bar{D}) and on the fiber E_0 , then it must be the identity on a neighborhood of E_0 , and by the discussion of the topology above, it is the identity on all of Z . \square

A consequence of the proof above is the following, which we also record for use in [DDLS21].

Corollary 5.6. *Any isometry which is the identity on some fiber E_X and some disk D_x is the identity.*

5.4. Relationship with Sol. The geometry $(\mathfrak{X}, \mathfrak{G})$ is closely related to Thurston's *solvgeometry*, **Sol**. More precisely, \mathfrak{X} fibers over \mathbb{H}^2 , and the bundle over any geodesic line in \mathbb{H}^2 is a model for **Sol**; see [Thu97, §3.8]. To be a geometry (in the sense of Thurston), one needs a finite volume quotient, which is provided by the lattice $\mathbb{Z}^2 \rtimes \text{SL}_2(\mathbb{Z}) < \mathfrak{G}$. The quotient is the universal (orbifold) torus bundle over the moduli space $\mathcal{M}_1 = \mathbb{H}^2/\text{PSL}_2(\mathbb{Z})$ of genus 1 Riemann surfaces, and the compact, orientable **Sol** torus bundles all arise as bundles over closed geodesics into \mathcal{M}_1 ; see [Thu97, Example 3.8.9]. The space E/Γ is a singular $(\mathfrak{X}, \mathfrak{G})$ -space that fibers over D/G , and one similarly obtains a singular **Sol** structure on the bundles over closed geodesics in D/G .

APPENDIX A. NOTATION INDEX

- X_0 , the fixed complex structures on S , and base point in the Teichmüller disk D , §3, page 11.

- q , fixed flat metric/quadratic differential, §3, page 11.
- $D = D_q$, the Teichmüller disk defined by q , with Poincaré metric ρ , §3, page 11.
- X, q_X , generic point in D and associated flat metric, §3, page 11.
- $G = G_q < \text{Mod}(S)$, the lattice Veech group of q , §3, page 11.
- $\mathcal{P} = \mathcal{P}(q) \subset \mathbb{P}^1(q)$, the set of parabolic directions on q , §3, page 11.
- $\{B_\alpha\}_{\alpha \in \mathcal{P}}$ the 1-separated horoballs, each invariant by the associated parabolic subgroup, §3.1, page 11.
- $c_\alpha: D \rightarrow B_\alpha$, the ρ -closest-point projection map, §3.1, page 11.
- $p: D \rightarrow \hat{D}$, ρ , the quotient by collapsing horoballs and its quotient metric, §3.1, page 11.
- $\bar{D} \subset D$, $\bar{\rho}$, the truncated Teichmüller disk obtained by removing all B_α and its induced path metric §3.1, 11. See also diagram 3.
- $\pi: E \rightarrow D$, the \tilde{S} -bundle over D , §3.2, page 11. See also diagram 3, page 14.
- Γ , the extension group of G , §3.2, page 11. See also diagram 3, page 14.
- $\mathcal{B}_\alpha = \pi^{-1}(B_\alpha)$, for $\alpha \in \mathcal{P}$, §3.2, page 11.
- $\bar{E} = \pi^{-1}(\bar{D})$, the truncated bundle, §3.2, page 11.
- $E_X \subset E$, the fiber over $X \in D$, §3.3, page 12.
- $E_0 = E_{X_0}$, the “base fiber” in E , §3.3, page 12.
- $f_{X,Y}: E_Y \rightarrow E_X$, the map between fibers over Y and X , §3.3, page 12.
- $f_X: E \rightarrow E_X$, $f_0: E \rightarrow E_0$, assembled maps to the fiber, §3.3, page 12.
- $D_x = f_X^{-1}(x)$, $\bar{D}_x = D_x \cap \bar{E}$, the horizontal disk, and truncated disk, through $x \in E_X$, §3.3, page 12.
- $f_\alpha: \bar{E} \rightarrow \partial B_\alpha$, horizontal closest point projection, §3.3, page 12.
- T_α , the tree dual to foliation of E_0 in direction $\alpha \in \mathcal{P}$, §3.4, page 12.
- $t_\alpha: E \rightarrow T_\alpha$, the projection to the tree, §3.4, page 12.
- $\Sigma_X \subset E_X$, $\Sigma \subset E$, $\bar{\Sigma} = \Sigma \cap \bar{E}$, cone points in a fiber, and their union in E and \bar{E} , §3.5, page 12.
- $P: E \rightarrow \hat{E}$, $\bar{P}: \bar{E} \rightarrow \hat{E}$, quotient maps collapsing \mathcal{B}_α to T_α , §3.6, page 13, see also diagram 3, page 14.
- $d_\alpha, d, \bar{d}, \hat{d}$, metrics on $T_\alpha, E, \bar{E}, \hat{E}$, respectively, §3.6, page 13.
- $\mathcal{V} \subset \hat{E}$, the union of all vertices of all trees, §3.6, page 13.
- $B_v = B_{\alpha(v)}$, $\mathcal{B}_v = \mathcal{B}_{\alpha(v)}$, for $v \in \mathcal{V}$ in the tree $T_{\alpha(v)}$, §3.6, page 13.
- $\bar{\pi}: \bar{E} \rightarrow \bar{D}$, $\hat{\pi}: \hat{E} \rightarrow \hat{D}$, associated projection maps, §3.6, page 13, see also diagram 3, page 14.
- $\theta_X^v \subset \Theta_X^v \subset E_X$ the v -spine and thickened v -spine, for $v \in \mathcal{V}$, §3.7, page 16.
- $\theta^v, \Theta^v \subset \partial \mathcal{B}_v$, the v -spine bundle and thickened v -spine bundle over ∂B_v , for $v \in \mathcal{V}$, §3.7, page 16.
- $\zeta(x, y)$, $\hat{\zeta}(x, y) = P(\zeta(x, y))$, preferred path and collapsed preferred path between $x, y \in \Sigma \subset E$, §4.1, Equation 4, page 21.
- $\Delta(x, y, z) = \Delta(f(x), f(y), f(z)) = [f(x), f(y)] \cup [f(y), f(z)] \cup [f(z), f(x)] \subset E_0$, reference triangle for $x, y, z \in \Sigma$, §4.3, page 23.
- $\Delta^\zeta(x, y, z)$, $\Delta^{\hat{\zeta}}(x, y, z)$, the preferred path and collapsed preferred path triangles for $x, y, z \in \Sigma$, §4.3, page 23.
- $\mathcal{P}(\sigma) \subset \mathcal{P}$ are directions of all saddle connections that span a non-degenerate triangle with the saddle connection σ , §4.6, page 26.
- $B(\sigma), \mathcal{B}(\sigma)$, the horoballs and bundles over them, associated to the directions in $\mathcal{P}(\sigma)$, §4.6, page 26.

- $\mathfrak{G} = \text{SAff}^{\pm}(\mathbb{R}^2) = \mathbb{R}^2 \rtimes \text{SL}_2^{\pm}(\mathbb{R})$ and $\mathfrak{X} = \mathfrak{G}/\mathfrak{K}$, area preserving affine group of \mathbb{R}^2 and associated homogeneous space, §5, page 37.

REFERENCES

- [Ber60] Lipman Bers. Quasiconformal mappings and Teichmüller’s theorem. In *Analytic functions*, pages 89–119. Princeton Univ. Press, Princeton, N.J., 1960.
- [Ber73] Lipman Bers. Fiber spaces over Teichmüller spaces. *Acta Math.*, 130:89–126, 1973.
- [BH99] M.R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [Bir69] Joan S. Birman. Mapping class groups and their relationship to braid groups. *Comm. Pure Appl. Math.*, 22:213–238, 1969.
- [Bir74] Joan S. Birman. *Braids, links, and mapping class groups*. Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1974. *Annals of Mathematics Studies*, No. 82.
- [Bow13] B. H. Bowditch. Stacks of hyperbolic spaces and ends of 3-manifolds. In *Geometry and topology down under*, volume 597 of *Contemp. Math.*, pages 65–138. Amer. Math. Soc., Providence, RI, 2013.
- [Bow14] Brian H. Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math.*, 269(2):269–280, 2014.
- [CC07] Ruth Charney and John Crisp. Relative hyperbolicity and Artin groups. *Geom. Dedicata*, 129:1–13, 2007.
- [CT07] James W. Cannon and William P. Thurston. Group invariant Peano curves. *Geom. Topol.*, 11:1315–1355, 2007.
- [DDL21] Spencer Dowdall, Matthew G. Durham, Christopher J. Leininger, and Allesandro Sisto. Extensions of Veech groups II: Hierarchical hyperbolicity and quasi-isometric rigidity. *preprint, to appear in Comment. Math. Helv.*, 2021.
- [Fil83] R. P. Filipkiewicz. *Four dimensional geometries*. PhD thesis, University of Warwick, December 1983, <http://go.warwick.ac.uk/wrap/954>.
- [FM02] Benson Farb and Lee Mosher. Convex cocompact subgroups of mapping class groups. *Geom. Topol.*, 6:91–152, 2002.
- [FM10] B. Farb and D. Margalit. *A primer on mapping class groups*. Princeton Univ. Press, Princeton, N.J., 2010.
- [FOH20] E. Fairchild, F. Owen, and B. B. Healy. Sectional curvature in Riemannian manifolds. *The Mathematica Journal*, 22, 2020.
- [Gar87] Frederick P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). John Wiley & Sons, Inc., New York, 1987. A Wiley-Interscience Publication.
- [GL00] F. P. Gardiner and N. Lakic. *Quasiconformal Teichmüller theory*, volume 76 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2000.
- [GM08] Daniel Groves and Jason Fox Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [Ham] Ursula Hamenstädt. Word hyperbolic extensions of surface groups. Preprint, [arXiv:math.GT/0505244](https://arxiv.org/abs/math/0505244).
- [Hil02] J. A. Hillman. *Four-manifolds, geometries and knots*, volume 5 of *Geometry & Topology Monographs*. Geometry & Topology Publications, Coventry, 2002.
- [IT92] Y. Iwayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. Translated and revised from the Japanese by the authors.
- [Koz16] Kenji Kozai. Hyperbolic structures from Sol on pseudo-Anosov mapping tori. *Geom. Topol.*, 20(1):437–468, 2016.
- [Min94] Yair N. Minsky. On rigidity, limit sets, and end invariants of hyperbolic 3-manifolds. *J. Amer. Math. Soc.*, 7(3):539–588, 1994.
- [Mor01] Shigeyuki Morita. *Geometry of characteristic classes*, volume 199 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 2001. Translated from the 1999 Japanese original, Iwanami Series in Modern Mathematics.

- [Mos03] Lee Mosher. Stable Teichmüller quasigeodesics and ending laminations. *Geom. Topol.*, 7:33–90, 2003.
- [Mos06] Lee Mosher. Problems in the geometry of surface group extensions. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 245–256. Amer. Math. Soc., Providence, RI, 2006.
- [MS39] S. B. Myers and N. E. Steenrod. The group of isometries of a Riemannian manifold. *Ann. of Math. (2)*, 40(2):400–416, 1939.
- [MS91] Howard Masur and John Smillie. Hausdorff dimension of sets of nonergodic measured foliations. *Ann. of Math. (2)*, 134(3):455–543, 1991.
- [MS12] Mahan Mj and Pranab Sardar. A combination theorem for metric bundles. *Geom. Funct. Anal.*, 22(6):1636–1707, 2012.
- [MS13] Howard Masur and Saul Schleimer. The geometry of the disk complex. *J. Amer. Math. Soc.*, 26(1):1–62, 2013.
- [MT02] Howard Masur and Serge Tabachnikov. Rational billiards and flat structures. In *Handbook of dynamical systems, Vol. 1A*, pages 1015–1089. North-Holland, Amsterdam, 2002.
- [Ota98] Jean-Pierre Otal. Thurston’s hyperbolization of Haken manifolds. In *Surveys in differential geometry, Vol. III (Cambridge, MA, 1996)*, pages 77–194. Int. Press, Boston, MA, 1998.
- [ST20] Nick Salter and Bena Tshishiku. Surface bundles in topology, algebraic geometry, and group theory. *Notices Amer. Math. Soc.*, 67(2):146–154, 2020.
- [Str80] K. Strebel. *Quadratic differentials*, volume 5 of *A series of Modern Surveys in Mathematics*. Springer-Verlag, Berlin, 1980.
- [SW10] John Smillie and Barak Weiss. Characterizations of lattice surfaces. *Invent. Math.*, 180(3):535–557, 2010.
- [Thu88] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 19(2):417–431, 1988.
- [Thu97] William P. Thurston. *Three-dimensional geometry and topology. Vol. 1*, volume 35 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1997. Edited by Silvio Levy.
- [Väi05] Jussi Väisälä. Gromov hyperbolic spaces. *Expo. Math.*, 23(3):187–231, 2005.
- [Wal85] C. T. C. Wall. Geometries and geometric structures in real dimension 4 and complex dimension 2. In *Geometry and topology (College Park, Md., 1983/84)*, volume 1167 of *Lecture Notes in Math.*, pages 268–292. Springer, Berlin, 1985.
- [Zor02] Anton Zorich. Square tiled surfaces and Teichmüller volumes of the moduli spaces of abelian differentials. In *Rigidity in dynamics and geometry (Cambridge, 2000)*, pages 459–471. Springer, Berlin, 2002.

DEPARTMENT OF MATHEMATICS, VANDERBILT UNIVERSITY, NASHVILLE, TN
E-mail address: `spencer.dowdall@vanderbilt.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, RIVERSIDE, CA
E-mail address: `mdurham@ucr.edu`

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY, HOUSTON, TX
E-mail address: `cjl12@rice.edu`

MAXWELL INSTITUTE AND DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK
E-mail address: `a.sisto@hw.ac.uk`