

AG
T

*Algebraic & Geometric
Topology*

Volume 26 (2026)

Extensions of finitely generated Veech groups

ELIOT BONGIOVANNI



Extensions of finitely generated Veech groups

ELIOT BONGIOVANNI

Given a closed surface S with finitely generated Veech group G and its $\pi_1(S)$ -extension Γ , there exists a hyperbolic space \widehat{E} on which Γ acts isometrically and cocompactly. The space \widehat{E} is obtained by collapsing some regions of the surface bundle over the convex hull of the limit set of G . Using the nice action of Γ on the hyperbolic space \widehat{E} , it is shown that Γ is hierarchically hyperbolic. These are generalizations of Dowdall–Durham–Leininger–Sisto, who assume in addition that G is a lattice. Because finitely generated Veech groups are among the most basic examples of subgroups of mapping class groups which are expected to qualify as geometrically finite, this result is evidence for the development of a broader theory of geometric finiteness.

1. Introduction	989
2. Background concepts	991
3. Generalized Veech dichotomy	998
4. Construction	1000
5. Guessing geodesics	1005
6. Slim triangles of collapsed preferred paths	1010
7. Hierarchical hyperbolicity	1031
List of symbols	1032
Acknowledgements	1034
References	1034

1 Introduction

In the context of Kleinian groups, there is a well-defined notion of “convex cocompactness” and a generalization known as “geometric finiteness”. For subgroups of mapping class groups, there is an analogous notion of convex cocompactness, but it is unclear what might be meant by geometric finiteness in this context. Finitely generated Veech groups are of interest here because they are subgroups of the mapping class group (those which stabilize Teichmüller disks) and Kleinian groups (since each of those Teichmüller disks is isometric to the hyperbolic plane). In general, a finitely generated Veech group is not convex cocompact—neither as a subgroup of a mapping class group nor as a Kleinian group—but it is geometrically finite as a Kleinian group. It is widely agreed upon that finitely generated Veech groups should qualify as geometrically finite as subgroups of mapping class groups; see Mosher’s discussion in [23, Section 6]. For example, Tang [28, Theorem 1.4] shows that finitely generated Veech groups are parabolically geometrically finite in the sense of Dowdall–Durham–Leininger–Sisto [9], so these serve as a fundamental example for developing a theory of geometric finiteness in the context of mapping

MSC2020: primary 20F65, 20F67, 57M60; secondary 30F40, 30F60.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

class groups. It is also known from Farb–Mosher and Hamenstädt that a subgroup of the mapping class group is convex cocompact if and only if its extension group is hyperbolic [11; 15]. Literature including that of Mosher [23, Problem 6.2], Dowdall–Durham–Leininger–Sisto [9, Section 1.4], and Russell [26, Section 1.1] suggests that some notion of geometric finiteness of subgroups of mapping class groups corresponds to hierarchical hyperbolicity of the extension group. This paper provides further evidence toward that conclusion.

The main result is a generalization of Dowdall–Durham–Leininger–Sisto [8, Theorem 1.1]. The following statement is identical, except that the term “lattice” has been generalized to “finitely generated”. The $\pi_1(S)$ extension of G is precisely the group Γ that fits into the short exact sequence

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

and the vertex subgroups of Γ are those that stabilize the vertices of particular Bass–Serre trees upon which Γ acts isometrically.

Theorem 5.8 *Suppose $G < \text{MCG}(S)$ is a finitely generated, nonelementary 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 Gromov hyperbolic space \widehat{E} , quasi-isometric to the Cayley graph of Γ coned off along the cosets of $\Upsilon_1, \dots, \Upsilon_k$.*

Although this paper follows the basic construction and argument outline from [8], the details differ substantially. First, in both cases, the space \widehat{E} is constructed from the hyperbolic plane bundle over the convex hull of the limit set of the associated Teichmüller disk. However, when G was also assumed to be a lattice, the convex hull of the limit set coincided with the Teichmüller disk. When G is only finitely generated, the convex hull of the limit set is generally some (strict) subset of the Teichmüller disk. As a consequence, the proof of the “fan lemma” (Lemma 6.22, which generalizes [8, Lemma 4.12]) is reformulated entirely. The fan lemma is the cornerstone for constructing sets that form slim triangles, which are used to prove hyperbolicity of \widehat{E} via the guessing geodesics criterion. (The guessing geodesics criterion essentially says that the existence of *paths* that form slim triangles is sufficient evidence for the existence of *geodesics* that form slim triangles; see Proposition 5.1. The version used here is due to Masur–Schleimer [21] and Bowditch [5].) As in [8], this paper constructs the “guessed geodesics” by concatenating hyperbolic geodesic segments orthogonal to the fibers with saddle connections within fibers. When G was also assumed to be a lattice, the Veech dichotomy ensured that every saddle connection arose as a boundary component of a cylinder decomposition. It is therefore necessary to generalize the Veech dichotomy to the case of finitely generated Veech groups. Although this generalization is known to experts and follows quickly from other long-established results, it appears to be missing from the literature.

Theorem 3.3 (generalized Veech dichotomy) *Let (S, X, q) be a flat surface with finitely generated Veech group G . Every direction $\alpha \in \Lambda(G) \subset \partial D$ is either minimal and uniquely ergodic or completely periodic and invariant by a parabolic element of the maximal Veech group.*

Section 3, therefore, can be read independently of the rest of the paper. Its relevance here is to arrive at Corollary 3.9, which states that saddle connections either arise as boundary components of cylinder decompositions (as in the classical Veech dichotomy) or are associated to directions not in the limit set of G . Therefore, when G is only finitely generated, it is possible for the “guessed geodesics” to include saddle connections which are arbitrarily long — an issue not encountered when G was also assumed to be a lattice — and this complicates the process of proving that these paths form slim triangles.

It is important to note that some arguments in this paper require that the finitely generated Veech group G has at least one parabolic element. If G has no parabolic elements (equivalently, if G is convex cocompact) then it is already known that the extension group is hyperbolic by the aforementioned results of [11; 15]. However, the arguments in this paper provide a new proof in this special case with little extra work.

Theorem 6.29 (special case of [11; 15]) *Let $G < \text{MCG}(S)$ be a finitely generated Veech group with extension group Γ . If G has no parabolic elements, then Γ is Gromov hyperbolic.*

After establishing the main result of Theorem 5.8, hierarchical hyperbolicity of the extension group is a natural next step.

Theorem 7.1 (hierarchical hyperbolicity) *Let $G < \text{MCG}(S)$ be a finitely generated Veech group with extension group Γ . Then Γ is a hierarchically hyperbolic group.*

The proof of this theorem proceeds almost exactly as in [9], after patching one subcase.

One might also ask whether an extension of a finitely generated Veech group is quasi-isometrically rigid. Indeed, this is a result of Margolis [19]. Though [9] gives a stronger result for extensions of lattice Veech groups, their version does not generalize to extensions of finitely generated Veech groups.

Outline

Section 2 surveys the background information needed to approach this problem and refers to resources that explore these concepts in more depth. Section 3 is dedicated to a generalization of the classical Veech dichotomy, which lays the foundation for the rest of the paper. Section 4 constructs the space \widehat{E} featured in Theorem 5.8. Section 5 proves that \widehat{E} is hyperbolic (Theorem 5.7), and shows how the rest of the statement of Theorem 5.8 follows. A key result for the proof of hyperbolicity of \widehat{E} , Theorem 6.28, is deferred to and comprises the entirety of Section 6. Finally, Section 7 proves that an extension of a finitely generated Veech group is hierarchically hyperbolic. For the reader’s convenience, a list of symbols and references to the pages on which they first appear begins on page 1032.

2 Background concepts

This section surveys concepts used throughout the paper. References for deeper reading are included at the beginning of each subsection.

2.1 Fundamental geometry

2.1.1 Paths in metric spaces See [6, Chapter I.1].

Let (X, d) be a metric space. A *path* from x to y in X is a continuous map c from an interval $[0, \ell] \subset \mathbb{R}$ to X with $c(0) = x$ and $c(\ell) = y$. The *length* of the path c is

$$\sup_P \sum_{i=1}^{nP} d(c(t_{i-1}), c(t_i)),$$

where the supremum is taken over all partitions P of $[0, \ell]$ with $t_0 = 0$ and $t_{nP} = \ell$. If the length of c is finite, then c is called *rectifiable*. The space X is called *rectifiably path connected* if any two points are connected by a rectifiable path. If for any $x, y \in X$ the distance $d(x, y)$ is precisely the infimum of the lengths of all paths from x to y , then (X, d) is called a *length space*.

A path c is called a *geodesic* if for all $t, t' \in [0, \ell]$,

$$d(c(t), c(t')) = |t - t'|.$$

The image of the map c is also sometimes referred to as a geodesic. The space X is called a *geodesic space* if any two points are connected by a geodesic. A geodesic between x and y in a geodesic space is typically denoted $[x, y]$, which is the notation used throughout this section.

2.1.2 Quotient spaces See [6, pp. 2, 64–70].

Let (X, d) be a metric space, and let \sim be an equivalence relation on elements of X . A *chain* from x to y in X is a sequence $\{x_1, y_1, x_2, y_2, \dots, x_n, y_n\}$ with $y_i \sim x_{i+1}$. For $x', y' \in X' = X / \sim$ representing $x, y \in X$, respectively, the *quotient pseudometric* is defined by

$$d'(x', y') := \inf_C \sum_{i=1}^n d(x_i, y_i),$$

where the infimum is taken over all chains joining x to y . A *pseudometric* meets all of the requirements to be a metric except that it is perhaps not *positive-definite*, meaning that $d'(x', y') = 0$ need not imply $x' = y'$. The next lemma is used several times in this paper to verify that a quotient space is in fact a length space.

Lemma 2.1 [6, Lemma I.5.20] *Let (X, d) be a length space, let \sim be an equivalence relation on X and let d' be the quotient pseudometric on $X' = X / \sim$. If d' is a metric then (X', d') is a length space.*

2.1.3 Gromov hyperbolicity See [6, Chapter III.H.1; 31].

Let (X, d) be a metric space with $x \in X$. The *Gromov product* of $y, z \in X$ with respect to x is

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

For $\delta \geq 0$, X is called δ -*hyperbolic* (or *Gromov hyperbolic with hyperbolicity constant δ*) if

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

for all $w, x, y, z \in X$. When X is a geodesic space, δ -hyperbolicity is equivalent to every geodesic triangle in X being δ' -slim for some $\delta' \geq 0$ depending on δ : A *geodesic triangle* consists of three points

$x, y, z \in X$ and choices of geodesics $[x, y]$, $[y, z]$, and $[x, z]$ and is δ -*slim* if each side is contained in the δ -neighborhood of the other two, i.e.,

$$\begin{aligned} [x, z] &\subset N_\delta([x, y] \cup [y, z]), \\ [x, y] &\subset N_\delta([y, z] \cup [x, z]), \\ [y, z] &\subset N_\delta([x, z] \cup [x, y]), \end{aligned}$$

where N_δ denotes the δ -neighborhood.

In this paper, any mention of hyperbolicity is Gromov hyperbolicity (with some hyperbolicity constant).

2.1.4 Coarse geometry See [6, pp. 138–144; 31].

Let (X, d_X) and (Y, d_Y) be metric spaces. Given $\lambda \geq 1$ and $\varepsilon \geq 0$, a map $f : X \rightarrow Y$ is called a (λ, ε) -*quasi-isometric embedding* if for all $x, y \in X$,

$$\frac{1}{\lambda}d_X(x, y) - \varepsilon \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y) + \varepsilon.$$

When the domain X is an interval in \mathbb{R} or \mathbb{Z} , the map f is called a *quasi-geodesic*; the image of the map is also sometimes referred to as a quasi-geodesic. If in addition there is a constant $K \geq 0$ so that $N_K(f(X)) = Y$ — that is, if f is *coarsely surjective* — then f is called a *quasi-isometry*, and the spaces X and Y are said to be quasi-isometric. If f is a quasi-isometry, then there exists a *quasi-inverse* of f , which is a quasi-isometry $g : Y \rightarrow X$ so that $d_X(x, g(f(x)))$ and $d_Y(y, f(g(y)))$ are uniformly bounded for all $x \in X$ and $y \in Y$.

Gromov hyperbolicity of length spaces is an invariant of quasi-isometry: If (X, d_X) and (Y, d_Y) are length spaces, X is hyperbolic, and $f : X \rightarrow Y$ is a quasi-isometry, then Y is hyperbolic (perhaps with a different hyperbolicity constant from that of X) [31, p. 16].

2.1.5 Cayley graphs Given a finitely generated group G with generating set \mathcal{A} , the *Cayley graph of G* , denoted $\text{Cay}(G)$, is the graph whose vertices are elements of G and whose edges connect g to ga for any $g \in G, a \in \mathcal{A}$. Defining all edges to be unit length makes the Cayley graph a metric space, where the induced metric is precisely the word metric with respect to the set \mathcal{A} . The group G equipped with the word metric is quasi-isometric to its Cayley graph, and any two Cayley graphs for G (obtained from two different generating sets) are quasi-isometric (see [6, pp. 139–141]).

Given a finite family of subgroups $\mathcal{H} = \{H_1, H_2, \dots, H_n\}$ of G , the *Cayley graph of G with respect to \mathcal{H}* or *coned-off Cayley graph*, denoted $\text{Cay}(G, \mathcal{H})$ is the graph obtained from $\text{Cay}(G)$ by adding a vertex V_{gH} for each left coset gH (with $g \in G$ and $H \in \mathcal{H}$) and attaching V_{gH} by an edge of length $\frac{1}{2}$ to each vertex of $\text{Cay}(G)$ which is an element of the coset gH . The isometric left action of G on $\text{Cay}(G)$ extends to an isometric action of G on $\text{Cay}(G, \mathcal{H})$: For all $g, g' \in G$ and $H \in \mathcal{H}$, set $gV_{g'H} = V_{gg'H}$. The nontrivial vertex stabilizers of the action are conjugate to subgroups of \mathcal{H} . See [7], including the following variant of the Schwarz–Milnor lemma.

Theorem 2.2 [7, Theorem 5.1] *Let G be a finitely generated group and suppose that G admits a discontinuous (that is, with discrete orbits), cocompact, isometric action on a length space X . Let \mathcal{H} denote a*

collection of subgroups of G consisting of exactly one representative of each conjugacy class of maximal isotropy subgroups for the action of G on X . Then \mathcal{H} is finite and, for any finite generating set \mathcal{A} of G , the coned-off Cayley graph $\text{Cay}(G, \mathcal{H})$ is quasi-isometric to X . In particular, if X is a hyperbolic space then the coned-off Cayley graph is hyperbolic.

2.1.6 Mapping class groups See [10].

Let S be a closed, orientable surface of genus at least 2. The *mapping class group* of S is

$$\text{MCG}(S) := \pi_0(\text{Homeo}^+(S)).$$

That is, the mapping class group is the group of isotopy classes of orientation-preserving homeomorphisms on S .

Mark a point $*$ on S and denote the marked surface by \dot{S} . The mapping class group of the marked surface is then

$$\text{MCG}(\dot{S}) := \pi_0(\text{Homeo}^+(S, *)),$$

i.e., the isotopy classes of elements of $\text{Homeo}^+(S)$ which fix the marked point $*$. These mapping class groups fit into the *Birman exact sequence*,

$$(1) \quad 1 \rightarrow \pi_1(S) \rightarrow \text{MCG}(\dot{S}) \rightarrow \text{MCG}(S) \rightarrow 1,$$

where $\text{MCG}(\dot{S}) \rightarrow \text{MCG}(S)$ is the map which forgets the marked point on S [10; 4]. For a subgroup $G < \text{MCG}(S)$, denote its preimage by $\Gamma_G < \text{MCG}(\dot{S})$ and observe that it fits into a short exact sequence

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

which includes into the Birman exact sequence above. The group Γ_G is called the $\pi_1(S)$ -*extension* of G (or simply the extension of G), and it is the fundamental group of an S -bundle with monodromy an isomorphism onto G .

2.2 Flat surfaces

See [12; 27].

Let S be a closed, connected, and oriented surface of genus at least 2. Equip S with a *complex structure* X , which is an atlas of charts $\{z_\alpha : U_\alpha \rightarrow \mathbb{C}\}$ whose transition functions $z_\beta^{-1} \circ z_\alpha$ are biholomorphic wherever the composition is defined. A *quadratic differential* for X , denoted q_X or more simply q , is a nonzero holomorphic section of the square of the canonical line bundle over (S, X) . Note in particular that any nonzero quadratic differential has finitely many zeroes. The pair (X, q) are a *flat structure* on S , described as follows.

In a small disk neighborhood of a nonzero point p , choose a coordinate chart z so that p corresponds to $z(p) = 0$ and pick a branch of $q^{1/2}(z)$. The *natural coordinate* or *preferred coordinate* in a neighborhood of p is given by

$$\zeta(z) = \int_0^z q^{1/2}(u) du.$$

In this coordinate, q is given by $q(z) dz^2 = d\xi^2$. In the neighborhood of a zero of order $k \geq 1$ there are natural coordinates such that $q(z) dz^2 = \zeta^k d\xi^2$.

Away from the zeroes of q , the transition functions for overlapping preferred coordinates of q are locally given by $z \mapsto \pm z + c$ for some $c \in \mathbb{C}$. Because the Euclidean metric is invariant under these transition functions, the Euclidean metric pulls back to a metric on S minus the zeroes of q . To complete the pullback metric, each zero of order $k - 2$ is filled back in so that a neighborhood of the point is isometric to the image of k Euclidean half planes glued together in a cyclic pattern by identifying the positive real axis of one half plane with the negative real axis of another (and only one other) half plane. A point of S corresponding to a zero of q — that is, a singular point — of order $k - 2$ is called a *cone point* with *cone angle* $k\pi$. The resulting metric is called a *flat metric*, which is also denoted by q . (Using the same notation for both the quadratic differential and the flat metric is a bit imprecise: The flat metric determines the quadratic differential up to multiplication by a nonzero complex number.)

Given a complex structure X on S and an associated flat metric q , the triple (S, X, q) is a *flat surface*. Where the structure and metric are implied, a flat surface is often denoted simply as S .

2.2.1 Directions Because the transition functions for overlapping preferred coordinates for q are locally given by $z \mapsto \pm z + c$, a line in the tangent space at any nonsingular point can be parallel translated (everywhere except the cone points) to produce a smooth line field, which corresponds to a line in the projective tangent space at any nonsingular point. In other words, any tangent line to the surface has a distinguishable direction, up to a rotation by π , which is consistent everywhere on the surface away from the cone points. This is known as the *space of directions*. It is sometimes denoted by $\mathfrak{P}^1(q)$, but in this paper the notation ∂D (introduced in Section 2.3) is used instead.

2.2.2 Geodesics Because the flat metric q is Euclidean away from the cone points, (local) geodesics on S minus the cone points are straight lines. A geodesic containing a cone point locally consists of two straight line segments meeting at the cone point and forming angles of at least π on both sides. A geodesic between two cone points and with no cone points on its interior is called a *saddle connection*. In particular, a saddle connection σ determines a line in the tangent space at any of its interior points, and therefore determines a point in the space of directions which is denoted by $[\sigma]$. The length of a saddle connection σ is denoted by $\ell(\sigma)$.

2.2.3 Foliations The line field obtained by parallel translating a tangent line around S minus the cone points also corresponds to a foliation of S minus the cone points by geodesics, which extends to a singular foliation over all of S when the cone points are included back in. So for any direction α there is a corresponding (singular) foliation $\mathcal{F}(\alpha)$ in direction α .

It is possible that for particular α the foliation $\mathcal{F}(\alpha)$ defines a *cylinder decomposition* in which every nonsingular leaf is a closed geodesic, and the singular leaves are concatenations of saddle connections which separate S into a union of Euclidean cylinders. A description of some of the directions α for which this occurs is given by Theorem 3.3.

2.2.4 The universal cover Let \tilde{S} denote the universal cover of S . The complex structure X and quadratic differential q on S can be pulled back to \tilde{S} , and the covering also gives a canonical identification of the directions on S with those on \tilde{S} . The covering map from \tilde{S} to S sends cone points to cone points (and saddle connections to saddle connections) and is a local isometry. For simplicity, the same notation is used for the structure, quadratic differential (and associated metric), and space of directions for \tilde{S} .

Because every cone point on S has cone angle greater than 2π , both S and \tilde{S} are nonpositively curved via Gromov’s link condition (see, for instance, [6, Chapter II.5]). Therefore the pulled back metric on \tilde{S} is CAT(0), and so \tilde{S} is uniquely geodesic. The geodesics are analogous to those in S , consisting of concatenations of saddle connections (with perhaps a straight line segment at the beginning or end of the path). Given a Euclidean cylinder in S , the preimage is a union of strips, and the covering map restricts to a universal covering of the cylinder on each strip.

2.3 The Teichmüller space

See [13, Section 8.2; 20].

Two complex structures X and Y on S are called *equivalent* if there is a map $f : (S, X) \rightarrow (S, Y)$, biholomorphic in the coordinate charts, which is 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 . The notation X is used both for a particular complex structure and its isotopy class $X \in \mathcal{T}(S)$. The Teichmüller space of S comes equipped with a metric known as the Teichmüller metric, which will not appear explicitly in this paper, but whose relevant features are described below.

Denoting by \dot{S} the surface S with a marked point, the space $\mathcal{T}(\dot{S})$ is the space of isotopy classes of complex structures in which isotopies are also required to fix the marked point. A fibration of Teichmüller spaces called the *Bers fibration* is given by

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

obtained by forgetting the marked point, where the fiber over a point $X \in \mathcal{T}(S)$ is canonically identified with \tilde{S} [2; 18].

2.3.1 The Teichmüller disk Given a flat surface (S, X, q) , where q has preferred coordinates ζ_i , a new complex structure can be obtained from any $A \in \text{SL}_2(\mathbb{R})$ by applying A to the given atlas — that is, a new atlas $\{A \circ \zeta_i\}$, where A acts as a linear transformation of $\mathbb{R}^2 \cong \mathbb{C}$. The new complex structure is denoted $A \cdot (X, q) = (A \cdot X, A \cdot q)$. Note that this deformation preserves the zeroes (including their orders) of the original structure (X, q) . The map $\text{SO}(2)A \mapsto A \cdot X$ gives a homeomorphism from $\text{SO}(2)\backslash\text{SL}_2(\mathbb{R})$ to the image of the orbit $D \subset \mathcal{T}(S)$ of (X, q) , since $\text{SO}(2)$ preserves the underlying complex structure. The disk D is called the *Teichmüller disk* of q , on which the Teichmüller metric is the push-forward of the Poincaré metric (by Teichmüller’s theorem). As a consequence,

$$\text{Isom}^+(D) \cong \text{PSL}_2(\mathbb{R}),$$

where the latter is the group of orientation-preserving isometries of the hyperbolic plane. It is often convenient to think of D as the Poincaré disk model of the hyperbolic plane, where each point represents a complex structure on S and traversing a geodesic between two points in D corresponds to varying the underlying flat structure by affine deformations.

More specifically, a geodesic through $X \in D$ is the map $t \mapsto A_t \cdot (X, q)$, where $\{A_t\}_{t \in \mathbb{R}}$ is a symmetric, 1-parameter hyperbolic subgroup of $SL_2(\mathbb{R})$ — that is, given by a matrix conjugate to

$$\pm \begin{bmatrix} e^t & 0 \\ 0 & e^{-t} \end{bmatrix}$$

by an element of $SO(2)$. All geodesics in D can be obtained by $SL_2(\mathbb{R})$ conjugates of the same 1-parameter family. The unit eigenvectors of such a transformation are orthogonal, leading to an identification of the boundary circle ∂D with the space of directions $\mathbb{P}^1(q)$ by associating the endpoint of the positive ray with the direction of the contracting eigenvector. As the ray approaches the boundary point associated to direction α , the length of any saddle connection in direction α (if one exists) shrinks exponentially. Since this occurs for a ray based at any point and ending at the boundary point associated to direction α , any horocycle based at the boundary point is a level set for the length of a saddle connection in direction α .

2.4 Veech groups

See [13, Chapters 7–8; 16; 30].

For a flat surface (S, X, q) and its associated Teichmüller disk $D \subset \mathcal{T}(S)$, a *Veech group* of q is a subgroup of the stabilizer of D . (Note also that the stabilizer of D is a subgroup of the mapping class group of S .) In this paper, a choice of Veech group of q is denoted by G . Equivalently, a Veech group is a subgroup of the affine homeomorphisms of (S, X, q) in preferred coordinates for q which fix the cone points, projected into the mapping class group.

Any affine homeomorphism in the Veech group has a derivative in preferred coordinates which is well defined up to sign and determines an element of $PSL_2(\mathbb{R})$, the group of orientation-preserving isometries of D — or, equivalently, of the hyperbolic plane. The action of the Veech group G on D is conjugate to an action on \mathbb{H} via the derivative (an element of $PSL_2(\mathbb{R})$). Since these two actions are essentially the same, both perspectives are employed throughout this paper. Elements of G are referred to as *parabolic*, *hyperbolic*, or *elliptic* according to whether their images under the derivative homomorphism are parabolic, hyperbolic, or elliptic isometries of \mathbb{H} , respectively. Finally, because $MCG(S)$ acts properly discontinuously on $\mathcal{T}(S)$, a Veech group G acts properly discontinuously on \mathbb{H} . Therefore the image of G under the derivative homomorphism is a *Fuchsian group*, that is, a discrete subgroup of $PSL_2(\mathbb{R})$.

In this paper, all Veech groups are assumed to be finitely generated.

2.4.1 Limit points The *limit set* of G , denoted $\Lambda(G)$, is the set of all possible limit points of some G -orbit, Gz , for $z \in D$. A Fuchsian group G is called *nonelementary* if $|\Lambda(G)| > 2$, which implies that G is not virtually cyclic. The limit set $\Lambda(G)$ is contained in the boundary circle ∂D [17, Corollary 2.2.7].

By identifying the boundary circle ∂D with $\mathbb{P}^1(q)$, each $\alpha \in \Lambda(G)$ corresponds to a direction $\alpha \in \mathbb{P}^1(q)$. Throughout this paper, limit points and their associated directions are referred to interchangeably.

A limit point is called a *parabolic fixed point* if it is the fixed point of a parabolic element of G . A limit point $\alpha \in \partial D$ is a *conical limit point* if for any ray R ending at α there is a point $Y \in D$, a sequence $\{g_i\}_{i=1}^\infty$ of elements of G , and an $\epsilon > 0$ such that $\{g_i Y\}_{i=1}^\infty$ converges to α within the ϵ -neighborhood of R in D [25, p. 617]. Since a finitely generated Veech group is Fuchsian, the following characterization of limit points applies.

Theorem 2.3 [1, Theorem 10.2.5] *A Fuchsian group is finitely generated if and only if each limit point is either a parabolic fixed point or a conical limit point.*

The precise definition of a conical limit point is only necessary for a proof of the generalized Veech dichotomy in Section 3.

2.4.2 Convex hull and convex core The *convex hull* of G , denoted $\text{hull}(G)$, is the intersection of all hyperbolic half planes whose closures in $D \cup \partial D$ contain $\Lambda(G)$. The convex hull is the minimal closed, convex, G -invariant subset of D [25, p. 637]. Consequently, the *convex core* of D/G , denoted $\text{core}(D/G) := \text{hull}(G)/G$, is the smallest closed convex subset of D/G for which the inclusion is a homotopy equivalence. When G is finitely generated, $|\Lambda(G)| \leq 1$ implies that $\text{core}(D/G) = \emptyset$, otherwise $\text{core}(D/G)$ is nonempty [25, p. 637]. Moreover, when the group G is finitely generated, the convex core has finite area, and truncating the convex core along its cusps results in a compact space [1, Theorem 10.1.2; 29, Proposition 8.4.3].

2.4.3 Lattices If the quotient space of D under the action of G has finite area, then the Veech group G is called a *lattice*, and the associated flat surface (S, X, q) is called a *lattice surface*. In particular, if G is a lattice then $\text{hull}(G) = D$.

3 Generalized Veech dichotomy

The Veech dichotomy characterizes foliations on a lattice surface. A generalization to foliations in certain directions on any flat surface follows quickly from several well-known results, but appears to be missing from the literature. The result most relevant to this paper is Corollary 3.9 at the end of this section. Otherwise, this section is self-contained.

Definition 3.1 Let \mathcal{F}_α denote a foliation of (S, X, q) in the direction α . If each of the leaves of \mathcal{F}_α is dense in (S, X, q) , then \mathcal{F}_α is called *minimal*; if in addition the transverse measure is unique up to scalar multiplication, then it is called *uniquely ergodic*. If, on the other hand, each of the leaves of \mathcal{F}_α is either closed or a saddle connection, then \mathcal{F}_α is called *completely periodic*.

A saddle connection is considered a leaf, so when \mathcal{F}_α is minimal, there are no saddle connections on S in the direction α . When \mathcal{F}_α is completely periodic, this means that there is a cylinder decomposition of S in the direction α .

Theorem 3.2 (classical Veech dichotomy¹) *Let (S, X, q) be a lattice surface. Every direction $\alpha \in \partial D$ is either minimal and uniquely ergodic or completely periodic and invariant by a parabolic element of the maximal Veech group.*

When the maximal Veech group G is not a lattice, there may be directions on S which do not correspond to points in the limit set of G . The general statement considers only directions corresponding to limit points for a finitely generated Veech group.

Theorem 3.3 (generalized Veech dichotomy) *Let (S, X, q) be a flat surface with finitely generated Veech group G . Every direction $\alpha \in \Lambda(G) \subset \partial D$ is either minimal and uniquely ergodic or completely periodic and invariant by a parabolic element of the maximal Veech group.*

Remark 3.4 When the Veech group is a lattice, the limit set is all of ∂D . Therefore the classical Veech dichotomy appears as a special case of this theorem.

Definition 3.5 [22, p. 1033] The *Teichmüller geodesic flow* is the one parameter subgroup of $\mathrm{SL}_2(\mathbb{R})$ given by

$$g_t = \begin{bmatrix} e^t & 0 \\ 0 & e^{-t} \end{bmatrix}$$

acting on the space of all quadratic differentials. A flat structure (X, q) is called *divergent* if $g_t \cdot (X, q)$ eventually exits every compact set in the moduli space $\mathcal{T}(S)/\mathrm{MCG}(S)$ as $t \rightarrow \infty$.

Similar to the proof of the classical Veech dichotomy, the proof of the generalized Veech dichotomy will require Masur's criterion.

Theorem 3.6 (Masur's criterion; [22, Theorem 3.8]) *Let α be the vertical direction. If the foliation \mathcal{F}_α of (S, X, q) is minimal but not uniquely ergodic, then (X, q) is divergent.*

Lemma 3.7 (from proof of [22, Theorem 1.8]) *If there is no saddle connection in direction α , then the foliation \mathcal{F}_α is minimal.*

Lemma 3.8 *Let $\alpha \in \Lambda(G) \subset \partial D$ and let R be a ray in D ending at α . If the image of R under $D \rightarrow D/G$ returns to a compact set infinitely often, there is no saddle connection on (S, X, q) in direction α .*

Proof Note that for all $g \in G$ and $(S, X, q) \in D$, $g \cdot (S, X, q)$ is isometric to (S, X, q) —so points of D/G are flat surfaces up to isometry. The length of the shortest saddle connection in direction α is a continuous function on D/G , so over any compact subset of D/G there is a lower bound on the length of the shortest saddle connection in direction α .

Suppose that there is a saddle connection on (S, X, q) in direction α . Without loss of generality assume that α is the vertical direction, so that R is the image of the Teichmüller geodesic flow g_t . Then the length of the shortest saddle connection in the vertical direction tends to zero. Therefore the image of R in D/G cannot return to a compact set of D/G infinitely often. \square

¹This statement is primarily based on [22, Theorem 5.10]. An alternate statement and proof (in the language of flows rather than foliations) is [16, Theorem 1]. Both papers note the proof of [33, Theorem 3.4] as another resource. The original statement is due to [32].

Proof of generalized Veech dichotomy Because G is finitely generated, α is either a conical limit point or parabolic fixed point (Theorem 2.3). When α is a parabolic direction the foliation in direction α is completely periodic; see, for example, [30, Section 6].

Suppose that α is conical, and let R be a ray based at a point $X \in \text{hull}(G)$ and ending at α . Then there exist a sequence $\{g_i\}_{i=1}^{\infty}$ of elements of G and an $\epsilon > 0$ large enough that $\{g_i X\}_{i=1}^{\infty}$ converges to α within the ϵ -neighborhood of R in D . Then R intersects the closed ϵ -neighborhood of X and its images under each of the isometries g_i (i.e., the translates of the neighborhood around X), so R must return to a compact set in D/G infinitely often. By Lemma 3.8 there is no saddle connection in direction α , so the foliation \mathcal{F}_α is minimal (Lemma 3.7). Finally, because \mathcal{F}_α is minimal and X is not divergent, the contrapositive to Masur's criterion (Theorem 3.6) implies that \mathcal{F}_α must be uniquely ergodic. \square

Corollary 3.9 *For any saddle connection σ , the associated direction $[\sigma] \in \partial D$ is either parabolic or lies outside of $\Lambda(G)$.*

4 Construction

Let S be a closed, connected, oriented surface of genus at least 2. Fix a complex structure X_0 on S and a flat metric q_0 for X_0 so that (S, X_0, q_0) is a flat surface. Let D be the Teichmüller disk of q_0 , let ρ be the Poincaré metric on D , and let G be a finitely generated Veech group of q_0 .

See [8] for the motivating construction in the case where G is not only finitely generated but also a lattice. Their construction of analogous spaces E and \widehat{E} differs from the construction given here. However, several of their arguments continue to apply in this more general setting essentially verbatim, and their work is cited wherever this is the case.

The construction is illustrated in Figure 2 on page 1005.

4.1 Base spaces

4.1.1 Horoballs and horopoints Denote by \mathcal{D} the set of all directions of all saddle connections on (S, X_0, q_0) . Consider \mathcal{D} to be a subset of ∂D as described in Section 2.2.1. For a saddle connection σ , denote its direction by $[\sigma] \in \mathcal{D}$. The notation $\alpha \in \mathcal{D}$ is sometimes used to refer to a direction without a specific choice of saddle connection.

Because G is assumed to be finitely generated and not necessarily a lattice, the direction of any saddle connection is associated to either a parabolic limit point or to a point outside the limit set of G (Corollary 3.9), which inspires the definitions below.

Definition 4.1 Let σ be a saddle connection with direction $[\sigma] \in \mathcal{D}$.

(i) If $[\sigma]$ is a parabolic limit point, then it is called a *parabolic direction* and σ is called a *parabolic saddle connection*. For each such $[\sigma]$, make a choice of closed horoball which is invariant by the maximal parabolic subgroup of G corresponding to $[\sigma]$. Choose the horoball to be small enough so that in any fiber over a point in its interior, the length of a saddle connection in direction $[\sigma]$ is no more than one third the

length of a saddle connection in any other direction. (This ensures that the horoballs are 1-separated.) Also choose the horoball to be small enough that its distance to $\partial \text{hull}(G)$ is at least one. Finally, choose the horoballs so that the set of all horoballs (for all parabolic directions) is G -invariant. For each parabolic $[\sigma]$, fix this choice of horoball and call it the *horoball* for $[\sigma]$, denoted $B_{[\sigma]}$.

(ii) If $[\sigma]$ is not a parabolic limit point, then it is called a *nonparabolic direction* and σ is called a *nonparabolic saddle connection*. For each such $[\sigma]$, define the *horopoint* for $[\sigma]$, denoted $B_{[\sigma]}$, to be the ρ -closest orthogonal projection of $[\sigma]$ to $\partial \text{hull}(G)$ — that is, the unique point $x \in \partial \text{hull}(G)$ such that the (unique, hyperbolic) geodesic through x and $[\sigma]$ meets $\partial \text{hull}(G)$ orthogonally. Note that the set of all horopoints (for all nonparabolic directions) is necessarily G -invariant.

See Figure 1 (blue) for examples. The notation B_* is used in both definitions because these objects function analogously later, when it is sometimes helpful to refer to them interchangeably.

The set of horopoints is not necessarily pairwise separated by a fixed constant, although by construction the horopoints are separated by a fixed constant from the set of horoballs associated to parabolic saddle connections. For each $\alpha \in \mathcal{D}$ fix a point $X_\alpha \in \partial B_\alpha$; for nonparabolic directions the only possible choice is $X_\alpha = B_\alpha$, the horopoint for α .

4.1.2 Convex hull and truncated convex hull $(\bar{D}, \bar{\rho})$ Let $\text{hull}(G)$ denote the convex hull of the limit set of G as defined in Section 2.4.2. Because $\text{hull}(G)$ is a subset of D and (D, ρ) is a length space, the space $(\text{hull}(G), \rho|_{\text{hull}(G)})$ is also a length space.

The *truncated convex hull* of G , denoted \bar{D} , is obtained from $\text{hull}(G)$ by deleting the interiors of the horoballs B_α . That is,

$$\bar{D} := \text{hull}(G) \setminus \bigcup_{\alpha \in \mathcal{D}} B_\alpha^\circ,$$

where B_α° denotes the interior of B_α . (Note that for nonparabolic $\alpha \in \mathcal{D}$, B_α is a point and $B_\alpha^\circ = \emptyset$.) By [6, Theorem II.11.27], the truncated disk $D \setminus \bigcup_{\alpha \in \mathcal{D}} B_\alpha^\circ$ with the induced path metric is CAT(0). Therefore the restriction to \bar{D} with the associated restricted metric $\bar{\rho}$ is also a CAT(0) space. Because G is finitely generated, it acts cocompactly on \bar{D} (refer to Section 2.4.2).

Denote by $\text{proj}_{\bar{D}} : D \rightarrow \bar{D}$ the ρ -closest-point projection of D onto \bar{D} .

4.1.3 Collapsed convex hull $(\hat{D}, \hat{\rho})$ Let \hat{D} be the *collapsed convex hull*, the quotient space obtained from \bar{D} by collapsing the boundary of each horoball B_α to a point. The quotient pseudometric $\hat{\rho}$ is positive-definite, therefore a metric, and so the space $(\hat{D}, \hat{\rho})$ is a length space (Lemma 2.1).

Analogous to the map $\bar{p} : \bar{D} \rightarrow \hat{D}$ which collapses each horoball to a point, there is also a map $p : \text{hull}(G) \rightarrow \hat{D}$ which collapses the *interior* of each horoball to a point.

4.2 Bundles

4.2.1 Total space E Let $\pi : E \rightarrow \text{hull}(G)$ be the pullback bundle of the Bers fibration (2) via the inclusion $\text{hull}(G) \subset D \subset \mathcal{T}(S)$, identifying $E \subset \mathcal{T}(\dot{S})$. Let $\Gamma < \text{MCG}(\dot{S})$ be the $\pi_1 S$ -extension of G (that is, the group fitting into the Birman exact sequence (1)), which acts on E .

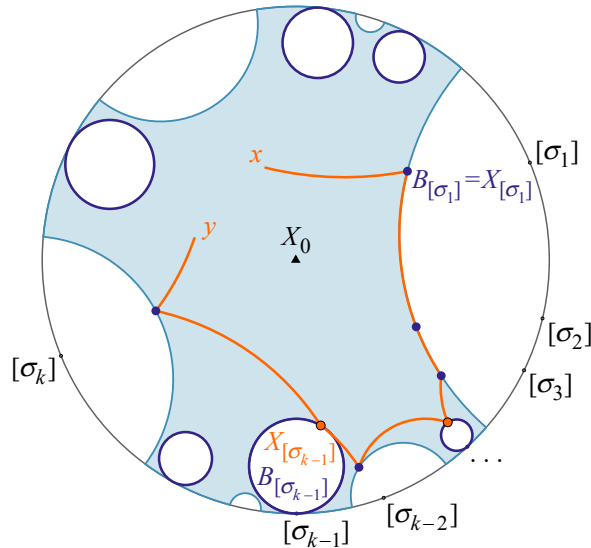


Figure 1: Examples of the constructions from Definitions 4.1 and 5.2, projected to D . The region shaded in blue, including its boundary in D , is the truncated convex hull \bar{D} . Some saddle connection directions are labeled along ∂D . For nonparabolic directions $[\sigma_*]$ —that is, those outside hull (G) —the associated “horopoint” $B_{[\sigma_*]}$ (blue dot) is the ρ -closest projection of $[\sigma_*]$ onto $\partial \text{hull}(G)$. For the remaining parabolic directions $[\sigma_*]$, the associated “horoball” $B_{[\sigma_*]}$ (blue circle and its interior) is a closed horoball in hull (G) based at $[\sigma_*]$, with the choice of horoball made according to the technical requirements of Definition 4.1. Finally, for each σ_* there is a choice of point $X_{[\sigma_*]}$: When σ_* is nonparabolic $X_{[\sigma_*]} := B_{[\sigma_*]}$ (blue dot), and when σ_* is parabolic $X_{[\sigma_*]} \in \partial B_{[\sigma_*]}$ (orange dot).

To construct the preferred path from x to y , $\zeta(x, y)$, as in Definition 5.2, first observe that the geodesic between $f(x)$ and $f(y)$ in the fiber over X_0 (black triangle) consists of saddle connections σ_1 through σ_k , in order. The fact that their associated directions appear in a nice order along ∂D is not a coincidence; see Lemma 6.19. The preferred path is constructed from concatenations of hyperbolic geodesics through the horizontal fibers of \bar{E} (orange curves) with saddle connections traversed in the fibers over the chosen points X_* (blue dots or orange dots). Because this diagram shows only the projection of $\zeta(x, y)$ to D , it is important to keep in mind that the horizontal pieces (orange curves) all belong to different horizontal fibers in \bar{E} and that the saddle connections (blue dots or orange dots) in fact have positive length. The saddle connections associated to parabolic directions (orange dots only) will have zero length after collapsing the bundle as in Section 4.3.2.

The space E is a surface bundle over hull (G) : By construction the fiber over $X \in \text{hull}(G)$, denoted $E_X := \pi^{-1}(X)$, is canonically identified with \tilde{S} equipped with the pulled back complex structure X and flat metric q_X . Recall that there was a fixed choice of a complex structure X_0 (which is necessarily in hull (G)), and denote the fiber over X_0 by E_0 .

It will be necessary to have maps between fibers. For $X, Y \in \text{hull}(G)$, let $f_{X,Y} : E_Y \rightarrow E_X$ be the lift of the Teichmüller map, i.e., the map that sends $y \in E_Y$ to the unique point $f(y) \in E_X$ along the lift of the geodesic in D connecting $\pi(x)$ and $\pi(y)$. This map is affine with respect to the flat metrics q_X and q_Y

and is $e^{\rho(X,Y)}$ -bilipschitz by construction. For any $X, Y, Z \in \text{hull}(G)$, the composition $f_{X,Y} \circ f_{Y,Z}$ agrees with $f_{X,Z}$. For any $X \in \text{hull}(G)$, define $f_X : E \rightarrow E_X$ by $f_X|_{E_Y} = f_{X,Y}$ for $Y \in \text{hull}(G)$. When $X = X_0$, the map is denoted $f = f_{X_0} : E \rightarrow E_0$.

In fact, E is a product: It admits a product structure $E \cong \text{hull}(G) \times \tilde{S}$. The map π is the projection onto the first factor $\text{hull}(G)$. The projection onto \tilde{S} , the universal cover of S , is any map f_X for $X \in \text{hull}(G)$. Define a metric d on E as the orthogonal direct sum of the Poincaré metric ρ in each horizontal fiber D_x and the pulled back flat metric q_X in each fiber $E_X \cong \tilde{S}$.

For any $X \in \text{hull}(G)$, denote by $\Sigma_X \subset E_X$ the set of cone points of the flat structure q_X on E_X , and define

$$\Sigma = \bigcup_{X \in \text{hull}(G)} \Sigma_X.$$

The quotient E/Γ is generally a noncompact S -bundle over the convex core $\text{hull}(G)/G$. If the quotient $\text{hull}(G)/G$ is compact and hence E/Γ is compact, then this represents a special case in which there are no parabolic directions for G —that is, the Veech group G is convex cocompact and therefore the extension group Γ is hyperbolic (see [11; 15] or Theorem 6.29). In the general case, a Γ -equivariant quotient of E is constructed from the G -equivariant quotient $\text{hull}(G) \rightarrow \hat{D}$ as described in the following several subsections.

4.2.2 Horizontal fibers D_x For any $X \in \text{hull}(G)$ and $x \in E_X$, the horizontal fiber $D_x = f_X^{-1}(x)$ is the unique lift of $\text{hull}(G)$ to the Teichmüller disk in $\mathcal{T}(\tilde{S})$ through x that covers $\text{hull}(G)$ via the projection π .

Remark 4.2 A more appropriate choice of notation might be $\text{hull}(G)_x$, to emphasize that this is a lift of only the convex hull rather than the entire Teichmüller disk D . Because there will not be any objects lifted to the complement of the hull, the notation D_x is used for simplicity.

4.2.3 Pullback bundle over truncated convex hull \bar{E} Define the pullback bundle over the truncated convex hull by

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

where $\mathcal{B}_\alpha^\circ = \pi^{-1}(B_\alpha^\circ)$ is the interior of the horoball preimage $B_\alpha = \pi^{-1}(B_\alpha)$. (As before, if α is nonparabolic then B_α is a point and $B_\alpha^\circ = \emptyset$, so $\mathcal{B}_\alpha^\circ = \emptyset$ also.)

The truncated bundle \bar{E} is equipped with the length metric \bar{d} induced from (E, d) . Because G acts isometrically and cocompactly on $(\bar{D}, \bar{\rho})$ and $\pi_1(S)$ acts isometrically and cocompactly on (\tilde{S}, \tilde{q}_0) , Γ acts isometrically and cocompactly on (\bar{E}, \bar{d}) . Therefore the space (\bar{E}, \bar{d}) is quasi-isometric to the group Γ equipped with the word metric. Then the quotient \bar{E}/Γ is a compact S -bundle over \bar{D}/G ,

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

As in [8], compactness of \bar{E}/Γ gives that Σ is r -dense in \bar{E} for some $r > 0$. By an application of the Arzelà–Ascoli theorem, (\bar{E}, \bar{d}) is a geodesic space.

4.3 Electrified space \widehat{E}

4.3.1 Bass–Serre trees Let α be a parabolic direction. The \mathbb{R} -tree dual to the foliation of E_{X_α} in the direction α , denoted T_α , is a weighted Bass–Serre tree when equipped with the metric defined by the transverse measure on the foliation of E_{X_α} in direction α . (See, for instance, the construction in [3, Subection 4.4].) Define a map

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

to be the composition of $f_{X_\alpha} : E \rightarrow E_{X_\alpha}$ followed by the projection $E_{X_\alpha} \rightarrow T_\alpha$. Recall from Section 4.1.1 that there was a choice of fixed point $X_\alpha \in \partial B_\alpha$, and note that the tree T_α is independent of the choice of $X_\alpha \in \partial B_\alpha$. Therefore the map t_α is also independent of this choice. Because there are only finitely many Γ -orbits of edges of these trees, every edge of every tree has length uniformly bounded above and below.

4.3.2 Collapsed “bundle” \widehat{E} Define \widehat{E} to be the quotient of E obtained by collapsing each B_α onto T_α via the restriction $t_\alpha|_{B_\alpha}$, and denote this map by $P : E \rightarrow \widehat{E}$. Note that the quotient is Γ -equivariant because it is constructed from the G -equivariant quotient hull $(G) \rightarrow \widehat{D}$. Finally, \widehat{E} is equipped with the quotient pseudometric \widehat{d} obtained from \bar{d} by the restriction $\bar{P} := P|_{\bar{E}}$. As in [8, Lemma 3.2], \widehat{d} is positive-definite, therefore a metric. Then by Lemma 2.1, the space $(\widehat{E}, \widehat{d})$ is a length space. The following fact is used throughout the paper.

Lemma 4.3 [8, Lemma 3.3] *The map P is 1-Lipschitz.*

4.4 Illustration of the construction

Figure 2 is a modified copy of [8, Figure 1], which illustrates the construction of \widehat{E} when the Veech group G is not a lattice. The base is a copy of the Teichmüller disk for the flat structure (S, X_0) , which is isometric to the hyperbolic plane. Over each point in the base disk is a copy of the universal cover of S with the lift of the associated flat metric. The truncated convex hull \bar{D} is shaded. The structures illustrated are:

- (a) Some points of ∂D associated to nonparabolic directions. For each nonparabolic direction η , the horopoint B_η is the nearest-point projection of η to hull (G) . The point X_η is chosen to be identical to B_η . The fiber over X_η is quasi-isometric to E_0 , the fiber over X_0 .
- (b) Some horoballs based at points of ∂D associated to parabolic directions. For each parabolic direction α , the associated horoball B_α is a sublevel set for the length of a saddle connection in direction α in a fiber over a point of B_α . The point X_α is a fixed choice in ∂B_α . By construction, the horoballs are pairwise 1-separated and also 1-separated from the set of horopoints associated to nonparabolic directions.
- (c) For any parabolic direction β , there is a multitwist $\tau_\beta \in G$ which takes X_β to other points in ∂B_β . As τ_β moves X_β along ∂B_β in D , it induces shearing in the fibers, which can be seen in E_{X_β} and $E_{\tau_\beta(X_\beta)}$.
- (d) The fiber E_0 over the basepoint X_0 is the universal cover of (S, X_0) . The highlighted path γ is a flat geodesic between cone points, which consists of concatenated saddle connections.

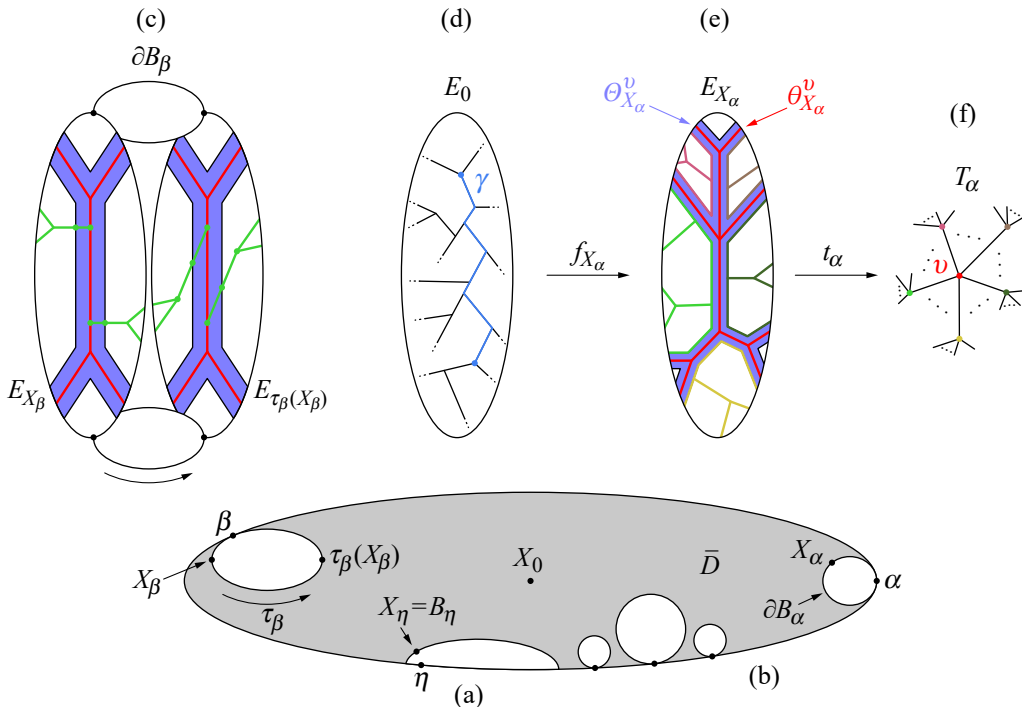


Figure 2: A copy of [8, Figure 1] which has been modified to reflect the case in which the Veech group G is not a lattice.

(e) The fiber E_{X_α} over X_α is the universal cover of (S, X_α) . It can also be obtained from E_0 by mapping via f_{X_α} , the lift of the Teichmüller map between X_0 and X_α in D . Each collection of colored segments, for example $\theta_{X_\alpha}^v$ in red, represents a connected set of saddle connections in direction α . These are called spines; see Section 5.2. (The lavender region $\Theta_{X_\alpha}^v$ is the “thickened spine neighborhood”, which only appears in this paper once when referring to results from [9].) The complement of the saddle connections is a cylinder decomposition of E_{X_α} in the direction α .

(f) The Bass–Serre tree T_α associated to the direction α . Each vertex is associated to a spine in E_{X_α} . If two spines in E_{X_α} bound a cylinder, their associated vertices in T_α are connected by an edge with length corresponding to the transverse measure of the cylinder.

5 Guessing geodesics

Hyperbolicity of \widehat{E} will be proven using the “guessing geodesics” criterion of Bowditch [5, Proposition 3.1] and Masur–Schleimer [21, Theorem 3.15], as formulated by [8].

Proposition 5.1 (guessing geodesics; [8, Proposition 2.2]) *Suppose Ω is a length 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 these properties:*

- (1) The $L(x, y)$ form δ -slim triangles.
- (2) If $x, y \in \Upsilon$ have $\hat{d}(x, y) \leq 3R$, then the diameter of $L(x, y)$ is at most δ .

Then Ω is hyperbolic.

Saying the $L(x, y)$ form δ -slim triangles means that for all $x, y, z \in \Upsilon$,

$$L(x, y) \subset N_\delta (L(x, z) \cup L(z, y)),$$

where N_δ denotes the δ -neighborhood. This is similar to the usual slim triangles condition for hyperbolicity from Section 2.1.3, except that the $L(x, y)$ need not be geodesics.

This section will verify most of the hypotheses required to apply the guessing geodesics criterion. Many of the arguments made by [8] continue to hold in this case. However, the proof that the sets $L(x, y)$ (constructed in Section 5.1) form slim triangles differs considerably and composes the entirety of Section 6.

5.1 Preferred paths and collapsed preferred paths

The sets $L(u, v)$ in the guessing geodesics criterion will be constructed from the following preferred paths in E which connect cone points by concatenating geodesics in the horizontal and vertical fibers.

Definition 5.2 For two cone points $x, y \in \Sigma$ the *preferred path* from x to y , denoted $\zeta(x, y)$, is constructed as follows.

- (1) In E_0 , $f(x)$ and $f(y)$ are cone points with respect to the flat metric q_0 , and they are connected by a geodesic segment which is a concatenation of saddle connections in E_0 . That is,

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

Each saddle connection σ_i is considered to be oriented with initial point $\sigma_i^- := \sigma_i \cap \sigma_{i-1}$ for $i = 2, \dots, k$ (and $\sigma_1^- = f(x)$) and terminal point $\sigma_i^+ := \sigma_i \cap \sigma_{i+1}$ for $i = 1, \dots, k-1$ (and $\sigma_k^+ = f(y)$). In the preferred path, each saddle connection σ_i will be traversed in the fiber over the associated point $X_{[\sigma_i]} \in B_{[\sigma_i]}$ as fixed in Section 4.1.1. More precisely, define the i -th *saddle piece* to be

$$\gamma_i := f_{X_{[\sigma_i]}}(\sigma_i),$$

which is a segment of the preferred path for all $1 \leq i \leq k$. Each γ_i is considered to be oriented with initial point

$$\gamma_i^- := f_{X_{[\sigma_i]}}(\sigma_i^-)$$

and terminal point

$$\gamma_i^+ := f_{X_{[\sigma_i]}}(\sigma_i^+).$$

- (2) The *horizontal pieces* h_i are chosen to make the preferred path continuous. Specifically, for $1 \leq i \leq k-1$, h_i is the geodesic in $D_{\gamma_i^+} = D_{\gamma_{i+1}^-}$ which connects γ_i^+ to γ_{i+1}^- . The first horizontal piece, denoted h_0 , is the geodesic in $D_x = D_{\gamma_1^-}$ connecting x to γ_1^- . The last horizontal piece, denoted h_k , is the geodesic in $D_{\gamma_k^+} = D_y$ connecting γ_k^+ to y .

(3) Define

$$\zeta(x, y) := h_0 \gamma_1 h_1 \gamma_2 h_2 \cdots \gamma_k h_k.$$

See Figure 1 for an example. The *collapsed preferred paths*, denoted $\hat{\zeta}(x, y)$, are the images of the preferred paths $\zeta(x, y)$ under the map $P : E \rightarrow \hat{E}$; see Section 4.3.2.

5.2 The sets $L(u, v)$

Let \mathcal{V} be the set of all vertices of all Bass–Serre trees in \hat{E} . Given $X \in \text{hull}(G)$ and a saddle connection σ , the union of the saddle connections in E_X with direction $[\sigma]$ is precisely the preimage of the vertices of the Bass–Serre tree $T_{[\sigma]}$ under the map

$$(3) \quad E_X \xrightarrow{f_{X_{[\sigma]}, X}} E_{X_{[\sigma]}} \xrightarrow{P} T_{[\sigma]},$$

which is simply the intersection of E_X with the preimage of $T_{[\sigma]}$ under $t_{[\sigma]}$ (see Section 4.3.1). For any vertex $v \in \mathcal{V}$ and $X \in \text{hull}(G)$, the preimage of v under this composition (3) is called the *v -spine in E_X* and is denoted θ_X^v . Note that a choice of $v \in \mathcal{V}$ determines the Bass–Serre tree $T_{[\sigma]}$ to which v belongs, and denote the union of the v -spines over all $X \in \partial B_{[\sigma]}$ by θ^v .

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 $y \in \theta^v \cap \Sigma$. Since $\zeta(x, y)$ is a finite length path (by construction, each saddle piece and horizontal piece has finite length), P is 1-Lipschitz (Lemma 4.3), and all $\hat{\zeta}(x, y)$ in $L(u, v)$ connect u to v , it follows that the set $L(u, v)$ is a path connected, rectifiable set containing both u and v .

5.3 Hyperbolicity

The statement of the following theorem is identical to [8, Theorem 4.2], but the proof depends on results which require considerable reformulation in the case that G is only finitely generated.

Theorem 6.28 (collapsed preferred paths form slim triangles) *There exists $\delta > 0$ so that collapsed preferred paths form δ -slim triangles. That is, for any $x, y, z \in \Sigma$,*

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

As the numbering suggests, the entirety of Section 6 is dedicated to the proof of this theorem. For now, Theorem 6.28 is assumed and used to prove hyperbolicity of \hat{E} in Theorem 5.7. Besides this, the last result needed is Corollary 5.6, which is stated after a few more constructions below.

Definition 5.3, Lemma 5.4, and Claim 5.5 do not reappear in this paper until Section 7; since the proofs are purely technical rather than geometrically intuitive, the reader may want to skip these on a first pass.

Definition 5.3 A *horizontal jump* in \hat{E} is the image under \bar{P} of a geodesic in \bar{D}_z , for some $z \in \Sigma$, that either is a subsegment of a component of $\partial \bar{D}_z$ or connects two components of $\partial \bar{D}_z$ with interior

disjoint from $\partial\bar{D}_z$. A *combinatorial path* in \widehat{E} is a concatenation of horizontal jumps and (perhaps zero) nonparabolic saddle connections; in particular, it is a concatenation of the collapsed preferred paths from Definition 5.2.

When the Veech group is a lattice, horizontal jumps and combinatorial paths coincide with the objects of the same names in [8, Definition 3.7]. Both definitions have been changed in this paper to fit the next statement, which is adapted from [8, Lemma 3.8]; their argument is sketched here and supplemented as necessary.

Lemma 5.4 *There is a 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 $C\hat{d}(x, y)$. In particular, there is a constant $C' > 0$ such that this combinatorial path consists of at most $C'\hat{d}(x, y)$ horizontal jumps and saddle connections.*

Proof sketch Define $\Sigma_\alpha := \Sigma \cap \partial\mathcal{B}_\alpha$. There exists some $K \geq 30$ so that for each $\alpha \in \mathcal{D}$ and $z \in \partial\mathcal{B}_\alpha$ there is $w \in \Sigma_\alpha$ with $\bar{d}(z, w) \leq K/30$ (in particular, K is chosen to be $30M$, where M is from [8, Lemma 3.4(3)]). The next step in the argument is to describe how to traverse between points of $\bar{P}(\Sigma_\alpha)$. Recall from Section 4.3.1 that T_α is the \mathbb{R} -tree dual to the foliation of E_{X_α} , and denote the length metric on T_α by ℓ_α . The statement of the following claim is identical to [8, Claim 3.9]; this is the only part of the argument for which the proof does not extend verbatim to the nonlattice case.

Claim 5.5 *There exists $R' > 0$ such that any pair of points $v_1, v_2 \in \bar{P}(\Sigma_\alpha) \subset \mathcal{V}$ may be connected by a combinatorial path of length at most $R'\ell_\alpha(v_1, v_2)$.*

Proof It suffices to assume that v_1 and v_2 are adjacent vertices of T_α . Choosing any $X \in \partial\mathcal{B}_\alpha$, the preimages of v_1 and v_2 under the restriction $\bar{P}|_{E_X} : E_X \rightarrow T_\alpha$ are adjacent spines $\theta_X^{v_1}$ and $\theta_X^{v_2}$ which are separated by a strip of uniformly bounded width [8, Lemma 3.4]. Then there exists a saddle connection $\sigma \subset E_X$ of bounded length which joins cone points $y_1 \in \theta_X^{v_1}$ and $y_2 \in \theta_X^{v_2}$. Denote the direction of σ by $[\sigma] \in \mathcal{D}$; unlike the case of [8], this direction may be nonparabolic.

Because σ has bounded length, $X \in \partial\mathcal{B}_\alpha$ lies within a bounded distance of $Y \in \partial\mathcal{B}_{[\sigma]}$. Let $z_i := f_{Y,X}(y_i)$. Let h_i be the horizontal geodesic in D_{y_i} from y_i to z_i , and note that the \bar{P} -image of each component of $h_i \cap \bar{D}_{y_i}$ is a horizontal jump in \widehat{E} . In particular, each $h_i \cap \bar{D}_{y_i}$ must have total length at most $\rho(X, Y)$, which is uniformly bounded. The saddle connection $f_{Y,X}(\sigma) \subset E_Y$ has length bounded by $e^{\rho(X,Y)}$ times the length of σ , which are both bounded. Therefore following the horizontal jumps along h_1 , traversing $f_{Y,X}(\sigma)$, and then following the horizontal jumps along h_2 gives a bounded length combinatorial path in \widehat{E} from v_1 to v_2 . The claim follows from the fact that an edge in T_α must be at least some minimal length, so $\ell_\alpha(v_1, v_2)$ is uniformly bounded below. □

Note that there exists a constant R_0 so that \mathcal{V} is R_0 -dense in \widehat{E} because \mathcal{V} is Γ -invariant and \widehat{E}/Γ is compact (since Γ acts cocompactly on \bar{E} and \widehat{E}/Γ is the continuous image of \bar{E}/Γ under the descent of $P : \bar{E} \rightarrow \widehat{E}$; [8, Lemma 3.6]).

Now it can be shown that 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 $\leq 4R'e^{Kr}r$ by [8, Claim 3.10]. Set $C := 36R'e^{3KR_0}$. Then, if $r := \hat{d}(x, y) \leq 3R_0$,

there is a combinatorial path joining x and y of length at most

$$4R'e^{3KR_0}\hat{d}(x, y) = \frac{1}{9}C\hat{d}(x, y),$$

satisfying the lemma. Otherwise if $r > 3R_0$, x and y can be joined by a path γ of length at most $2r$, which can be subdivided into $n = \lceil \text{length}(\gamma)/R_0 \rceil$ equal-length subsegments of length at most R_0 . Since \mathcal{V} is R_0 -dense, there is a sequence $\{x_i\} \subset \mathcal{V}$ with $x_0 = x$, $x_n = y$, and $\hat{d}(x_i, x_{i+1})$. Each pair x_i, x_{i+1} can then be connected by a combinatorial path of length at most $CR_0/3$, and so there is a combinatorial path from x to y of length at most $CR_0n/3$. Finally, because $R_0n < \text{length}(\gamma) + R_0 \leq 3\hat{d}(x, y)$, the first statement of the lemma holds.

The number of horizontal jumps and saddle connections in the combinatorial path joining x and y can be bounded as follows. Each horizontal jump between different components of the hull boundary (see Definition 5.3) has length at least 1 because the horoballs associated to parabolic directions were constructed to be 1-separated from each other and from $\partial \text{hull}(G)$ (see Definition 4.1). Therefore the combinatorial path may contain at most $C\hat{d}(x, y)$ horizontal jumps between different components of the hull boundary. By Proposition 6.3, there are no short saddle connections in the fibers over \bar{D} ; supposing that $A' < 1$ is a uniform positive lower bound on the lengths of the saddle connections, the combinatorial path may contain at most $2C\hat{d}(x, y)/A'$ nonparabolic saddle connections and horizontal jumps which are subsegments of components of the hull boundary. Then the combinatorial path joining x and y contains at most $3C\hat{d}(x, y)/A'$ horizontal jumps and saddle connections. Write $C' := 3C/A'$ to simplify the notation. □

Combined with Theorem 6.28, this bound on $\hat{d}(x, y)$ allows for a bound on the diameter of the sets $L(u, v)$. The following statement is identical to [8, Lemma 4.4]; the proof is nearly identical except for some changes in the language regarding combinatorial paths, which were given a different definition in this paper (see Definition 5.3 above).

Corollary 5.6 *There exists a constant $C > 0$ so that if $u, v \in \mathcal{V}$ with $\hat{d}(u, v) \leq 3R_0$ (where R_0 is a constant so that \mathcal{V} is R_0 -dense in \hat{E} , as in the proof of Claim 5.5), then $\text{diam}(L(u, v)) \leq C$.*

Proof Pick $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$. By Lemma 5.4 there is a combinatorial path from $P(x) = u$ to $P(y) = v$ of bounded length which is a concatenation of n horizontal jumps and saddle connections, where the bound on the length and the number n depend only on R_0 . By repeatedly applying Theorem 6.28, $\hat{\zeta}(x, y)$ is in the $n\delta$ -neighborhood of the combinatorial path joining $P(x) = u$ to $P(y) = v$. Therefore $\hat{\zeta}(x, y)$ has uniformly bounded diameter. As in [8, Lemma 4.3], $L(u, v)$ is contained in the 2δ -neighborhood of $\hat{\zeta}(x, y)$, where $\delta > 0$ is the constant from Theorem 6.28. Therefore $L(u, v)$ has uniformly bounded diameter also. □

Finally, it is possible to establish hyperbolicity of \hat{E} .

Theorem 5.7 *The space (\hat{E}, \hat{d}) is hyperbolic when the limit set $\Lambda(G)$ contains at least one parabolic point.*

Proof Hyperbolicity of \widehat{E} will be proven by the guessing geodesics criterion (Proposition 5.1). It was established during construction that (\widehat{E}, \hat{d}) is a length space (see Section 4.3.2 and [8, Lemma 3.2]). Let $\mathcal{V} \subset \widehat{E}$ be the collection of all vertices of Bass–Serre trees in \widehat{E} . Let $R_0 > 0$ be as in the proof of Claim 5.5, so that \mathcal{V} is R_0 -dense in \widehat{E} . For any $u, v \in \mathcal{V}$, $L(u, v)$ is a rectifiably path-connected set containing u and v (see Section 5.2). Let $\delta > 0$ be the constant from Theorem 6.28. As in [8, Lemma 4.3], each $L(u, v)$ is contained in the 2δ -neighborhood of $\hat{\zeta}(x, y)$ for any $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$; because triangles of collapsed preferred paths in \widehat{E} are δ -slim (Theorem 6.28), the $L(u, v)$ also form 3δ -slim triangles, satisfying condition (1) of the guessing geodesics criterion. Finally, by Corollary 5.6 the sets $L(u, v)$ have diameter bounded by a constant $C > 0$ whenever $\hat{d}(u, v) \leq 3R_0$, satisfying condition (2) of the guessing geodesics criterion. Therefore (\widehat{E}, \hat{d}) is hyperbolic. \square

The action of Γ on the total space E is isometric by construction but noncocompact in general. The truncated convex hull \bar{E} was constructed as a subspace on which Γ acts cocompactly in addition to isometrically. By a direct application of the Schwarz–Milnor lemma, \bar{E} and Γ are quasi-isometric. However, neither is hyperbolic unless there are no parabolic directions on S . To construct a hyperbolic space from E , for parabolic saddle connections σ the sets $\mathcal{B}_{[\sigma]}$ are collapsed to Bass–Serre trees. In Γ , this corresponds to coning off by the *vertex subgroups*, the stabilizers of the vertices of the Bass–Serre trees under the isometric action of Γ . Then when the limit set $\Lambda(G)$ contains at least one parabolic point, the main theorem follows. The proof is essentially identical to that of [8, Corollary 4.5], but it is included here for the sake of completeness.

Theorem 5.8 *Suppose $G < \text{MCG}(S)$ is a finitely generated, nonelementary 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 Gromov hyperbolic space \widehat{E} , quasi-isometric to the Cayley graph of Γ coned off along the cosets of $\Upsilon_1, \dots, \Upsilon_k$.*

Proof Recall from Section 4.3.2 that \widehat{E} is a length space. By construction, the action of Γ on \bar{E} is isometric and cocompact, and therefore so is the action of Γ on \widehat{E} . Let $\Upsilon_1, \dots, \Upsilon_k < \Gamma$ be representatives of the conjugacy classes of vertex subgroups stabilizing the vertices of the Bass–Serre trees. Any point-stabilizers for the action of Γ on \widehat{E} are trivial or conjugate into one of the Υ_* , and therefore any point in \widehat{E} has a discrete orbit under Γ . By Theorem 2.2 it follows that, for any finite generating set \mathcal{S} of Γ , the coned-off Cayley graph $\text{Cay}(\Gamma, \mathcal{S} \cup \bigcup \Upsilon_i)$ is quasi-isometric to \widehat{E} . Finally, because \widehat{E} is hyperbolic by Theorem 5.7, so too is the coned-off Cayley graph of Γ . \square

6 Slim triangles of collapsed preferred paths

Section 5 verified all but one condition for the guessing geodesics criterion (Proposition 5.1). This section will verify the remaining condition, which is that all triangles of collapsed preferred paths in \widehat{E} are slim (Theorem 6.28).

The following definitions are comparable to those in [8, Section 4.3].

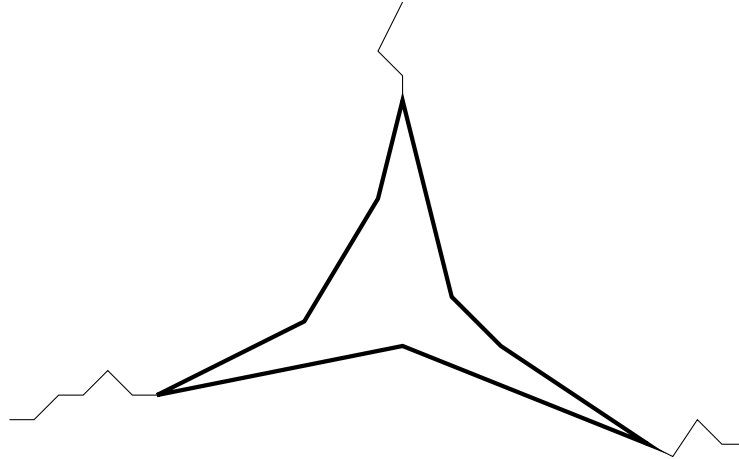


Figure 3: Illustration from [8, Figure 2] of a degenerate triangle in E_0 with the nondegenerate subtriangle in bold.

Definition 6.1 Given $x, y, z \in \Sigma$, define the associated geodesic triangle in E_0 by

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

the triangle of preferred paths in E by

$$\Delta^\zeta(x, y, z) := [\zeta(x, y)] \cup [\zeta(y, z)] \cup [\zeta(z, x)],$$

and the triangle of collapsed preferred paths in \widehat{E} by

$$\Delta^{\hat{\zeta}}(x, y, z) := [\hat{\zeta}(x, y)] \cup [\hat{\zeta}(y, z)] \cup [\hat{\zeta}(z, x)].$$

Equivalently, $\Delta^{\hat{\zeta}}(x, y, z)$ is the image of $\Delta^\zeta(x, y, z)$ under the map P . If any pair of the sides of $\Delta(x, y, z)$ intersect at one or more (nontrivial) saddle connections, then $\Delta(x, y, z)$ is called *degenerate*. Otherwise, if each pair of sides of $\Delta(x, y, z)$ intersect only at the vertices, then $\Delta(x, y, z)$ is called *nondegenerate*. The triangles $\Delta^\zeta(x, y, z)$ and $\Delta^{\hat{\zeta}}(x, y, z)$ are described as degenerate (or nondegenerate) if the triangle $\Delta(x, y, z)$ is degenerate (or nondegenerate).

Recall from the construction that $E_0 \cong \widetilde{S}$ (equipped with the lift of the flat metric q_0). Since E_0 is complete, simply connected, and nonpositively curved, it is CAT(0) (as a result of the Cartan–Hadamard theorem) and therefore it is uniquely geodesic. Then the nontrivial intersection of two sides of a degenerate triangle must be a concatenation of saddle connections based at the shared vertex. As a consequence, every degenerate triangle — whether in E_0 , E , or \widehat{E} — contains a subtriangle which is nondegenerate. See Figure 3.

To prove that all triangles are slim, it suffices to prove that all nondegenerate triangles are slim (see [8, Lemma 4.6] for a detailed proof). Therefore the rest of this section assumes that all triangles are nondegenerate.

6.1 Euclidean triangles

Definition 6.2 Given $x, y, z \in \Sigma \cap E_0$, the triangle $\Delta(x, y, z)$ is called *Euclidean* if it is nondegenerate and each side consists of a single saddle connection. When the information about x, y, z is clear from context, a Euclidean triangle is sometimes denoted simply by T .

Proposition 6.3 (no short saddle connections condition) *For any $X \in \bar{D}$, there is a positive lower bound on the lengths of all saddle connections in (S, X) .*

Proof Recall from Definition 4.1 that for parabolic saddle connections σ the horoballs $B_{[\sigma]}$ were chosen so that over $\partial B_{[\sigma]}$ the saddle connection σ would have (positive) length less than any saddle connection in a different direction. (The lengths of the nonparabolic saddle connections are necessarily positive over the convex hull.) Because the length of the shortest saddle connection on (S, X) is a continuous function on D/G and \bar{D}/G is compact, there must be a positive lower bound on the lengths of all saddle connections in the fibers over \bar{D} . \square

It is also true that Euclidean triangles cannot be too large. This is known for lattice surfaces (see, for instance, [33]), but requires a new proof for general flat surfaces.

Proposition 6.4 (no large triangles condition) *There exists $A > 0$ so that for any $X \in D$ the area of any Euclidean triangle in (S, X, q) is at most A .*

Proof Any Euclidean triangle minus its vertices is embedded in the flat surface (S, X, q) [24, Lemma 2.1]. Choose A to be the area of (S, X, q) . Then any Euclidean triangle in (S, X, q) has area at most A . Recalling a fact from Section 2.3, the choice of A is uniform for all $X \in D$ because the area of (S, X, q) is preserved by the action of $\mathrm{SL}_2(\mathbb{R})$. \square

Corollary 6.5 *Every Euclidean triangle is uniformly slim.*

Proof The inradius of any Euclidean triangle is bounded above by the inradius of an equilateral triangle of area A (where A is the constant from Proposition 6.4), so every Euclidean triangle is $2\sqrt{A}/3^{3/4}$ -slim. \square

Throughout this section, it will be necessary to relate Euclidean triangles in E_0 to particular ideal triangles in D , described below.

Definition 6.6 Let $x, y, z \in \Sigma$ so that $T = \Delta(x, y, z)$ is a Euclidean triangle consisting of saddle connections σ_x, σ_y , and σ_z . The *ideal triangle associated to T* , denoted \tilde{T} , is the ideal triangle in D with vertices $[\sigma_x]$, $[\sigma_y]$, and $[\sigma_z]$.

These are studied in depth by [24].

Definition 6.7 Let $T = \Delta(x, y, z)$ be a Euclidean triangle. The *balance point of T* (or of its associated ideal triangle \tilde{T}) is the unique point $b \in D$ such that $f_b(T)$ is an equilateral triangle. The *adjusted balance point of T* (or of its associated ideal triangle \tilde{T}), denoted \bar{b} , is the projection of b onto $\mathrm{hull}(G)$. If $b \in \mathrm{hull}(G)$, then b and \bar{b} coincide.

Remark 6.8 Existence and uniqueness of the balance point is given by [8, Lemma 4.9]. Uniqueness of the adjusted balance point follows from uniqueness of the closest-point projection.

Proposition 6.9 Let $T = \Delta(x, y, z)$ be a Euclidean triangle with sides consisting of saddle connections σ_x , σ_y , and σ_z . Then the balance point b is the incenter of the ideal triangle \tilde{T} (that is, the center of the largest circle inscribed in \tilde{T}).

Proof Without loss of generality assume that $[\sigma_x]$ is the vertical direction (perhaps after acting on D with an element of $SO(2)$; see Section 2.3). Changing the coordinate chart by a rotation of $\pm\pi/3$ makes $[\sigma_y]$ or $[\sigma_z]$ vertical. Therefore the three geodesic rays in D starting at b and ending at $[\sigma_x]$, $[\sigma_y]$, or $[\sigma_z]$ diverge from b at equal angles. In particular, changing the coordinate chart by a rotation of $\pm\pi/3$ preserves $\tilde{T} \subset D$. It also preserves the three geodesic segments from b to the nearest point on each side of \tilde{T} , so b must be the incenter of \tilde{T} . \square

Corollary 6.10 Let T be a Euclidean triangle, and let b be the balance point of T . Because the area of $f_b(T)$ is bounded above by A (Proposition 6.4), each side of $f_b(T)$ has length at most $2\sqrt{A}$.

For many of the following proofs it is important to recall from Section 2 that given a saddle connection σ with direction $[\sigma] \in \partial D$, the horoballs based at $[\sigma] \in D$ are sublevel sets for the length of σ , where the length of σ shrinks exponentially along any geodesic ending at $[\sigma]$. The following technical lemma states this more precisely.

Throughout Lemmas 6.11 and 6.12, a geodesic between $[\sigma_x]$ and $[\sigma_y]$ is denoted $\overleftrightarrow{[\sigma_x], [\sigma_y]}$.

Lemma 6.11 Let $T = \Delta(x, y, z)$ be a Euclidean triangle with sides consisting of saddle connections σ_x , σ_y , and σ_z . Then for any segment $c \subset \overleftrightarrow{[\sigma_x], [\sigma_y]}$ of length $L > 0$,

$$\min_{t \in c} (\min \{ \ell(f_t(\sigma_x)), \ell(f_t(\sigma_y)) \}) \leq 2\sqrt{3}Ae^{-L/2},$$

where A is the constant from Proposition 6.4.

Proof At the balance point b for T , the lengths of $f_b(\sigma_x)$ and $f_b(\sigma_y)$ are at most $2\sqrt{A}$ (Corollary 6.10). Denote by t' the point of intersection of $\overleftrightarrow{[\sigma_x], [\sigma_y]}$ with the inscribed circle of \tilde{T} . The distance between b and t' is $\ln \sqrt{3}$ (as this is the inradius of an ideal triangle — see Proposition 6.9), so the lengths of $f_{t'}(\sigma_x)$ and $f_{t'}(\sigma_y)$ are at most $2\sqrt{A}e^{\ln \sqrt{3}} = 2\sqrt{3A}$. Traversing $\overleftrightarrow{[\sigma_x], [\sigma_y]}$ towards $[\sigma_x]$ (or towards $[\sigma_y]$) shrinks $\ell(\sigma_x)$ (or $\ell(\sigma_y)$) exponentially as a function of the distance traversed. Since one of the endpoints has distance at least $L/2$ from t' , it follows that $\min(\ell(\sigma_x), \ell(\sigma_y)) \leq 2\sqrt{3A}e^{-L/2}$ on c . \square

Lemma 6.12 Let $T = \Delta(x, y, z)$ be a Euclidean triangle consisting of saddle connections σ_x , σ_y , and σ_z ; let $\tilde{T} \subset D$ be the ideal triangle associated to directions $[\sigma_x]$, $[\sigma_y]$, and $[\sigma_z]$; and let $\text{proj}_{\bar{D}}(\tilde{T})$ denote the closest-point projection of \tilde{T} onto \bar{D} (where this map is the same as in Section 4). Then the perimeter of $\text{proj}_{\bar{D}}(\tilde{T})$ is uniformly bounded. As a consequence, the diameter of $\text{proj}_{\bar{D}}(\tilde{T})$ is uniformly bounded.

Proof Denote the balance point for \tilde{T} by b and recall that the lengths of σ_x , σ_y , and σ_z are all uniformly bounded at b . They are also uniformly bounded at the point closest to b on the side of \tilde{T} opposite (without

loss of generality) $[\sigma_x]$ (as in the proof of Lemma 6.11). Then, traversing the geodesic from this point toward (without loss of generality) $[\sigma_y]$, $\ell(\sigma_y) \rightarrow 0$. Because there are no short saddle connections in \bar{D} (Proposition 6.3), there is a bound on the length of the intersection of each side of \tilde{T} with \bar{D} .

If $[\sigma]$ is a parabolic vertex of \tilde{T} , then $\text{proj}_{\bar{D}}(\tilde{T} \cap B_{[\sigma]}) = \tilde{T} \cap \partial B_{[\sigma]}$, which has uniformly bounded length due to the choice of horoballs (see Definition 4.1).

If the distance between \tilde{T} and $\text{hull}(G)$ is at least 1, then the projection of \tilde{T} onto any component of $\partial \text{hull}(G)$ is bounded. More generally, any component of \tilde{T} which lies outside of the 1-neighborhood of $\text{hull}(G)$ in D projects to a bounded subsegment of $\partial \text{hull}(G)$.

It only remains to consider the portions of \tilde{T} contained in the 1-neighborhood of $\text{hull}(G)$. By convexity of $\text{hull}(G)$, the intersection of a side of \tilde{T} with the 1-neighborhood of $\text{hull}(G)$ consists of at most one component. Fix some $L > 0$ and assume (without loss of generality) that there is such a component $c \subset [\sigma_x], [\sigma_y]$ of length $L > 0$ contained in the 1-neighborhood of a component of $\partial \text{hull}(G)$. Then the lengths of σ_x and σ_y in the fibers over points of $\text{proj}_{\bar{D}}(c)$ are bounded by $2\sqrt{3}Ae^{-L/2+1}$ (Lemma 6.11). Again, because there are no short saddle connections in \bar{D} (Proposition 6.3), this implies that the projection of c to $\partial \text{hull}(G)$ has bounded length.

Because each of these cases can occur at most three times for any \tilde{T} , there is a bound on the perimeter of $\text{proj}_{\bar{D}}(\tilde{T}) \cap \bar{D}$. Therefore the diameter of $\text{proj}_{\bar{D}}(\tilde{T}) \cap \bar{D}$ is also bounded. □

The next lemma shows that for any ideal triangle that intersects $\text{hull}(G)$, the balance point and adjusted balance point are close to each other, to the horopoints associated to any nonparabolic vertices, and to the horoballs associated to any parabolic vertices. This statement functions similarly to [8, Lemma 4.10], which is a consequence of [33]; both assume in addition that G is a lattice, which is not necessary for the result below.

Lemma 6.13 *Let $T = \Delta(x, y, z)$ be a Euclidean triangle consisting of saddle connections σ_x, σ_y , and σ_z so that $\tilde{T} \cap \text{hull}(G) \neq \emptyset$. Then the balance point b of \tilde{T} , the adjusted balance point \bar{b} of \tilde{T} , the horoballs associated to any parabolic vertices of \tilde{T} , and the horopoints associated to any nonparabolic vertices of \tilde{T} are all uniformly close to one another. See Figure 4.*

Proof Because the projection of \tilde{T} to $\text{hull}(G)$ has bounded diameter (Lemma 6.12), the adjusted balance point \bar{b} and the horopoints associated to any nonparabolic vertices of \tilde{T} must be uniformly close. Also, the adjusted balance point \bar{b} must be uniformly close to the horoballs associated to any parabolic vertices of \tilde{T} because G acts cocompactly on the truncated hull \bar{D} , meaning that \bar{b} is in some uniformly thick part of $\text{hull}(G)$. Lastly, if $b \in \text{hull}(G)$, then $b \in \bar{D}$ (by the choice of horoballs made during construction; see Definition 4.1) and therefore $b = \bar{b}$.

It only remains to show that b and \bar{b} are uniformly close even when b lies outside of $\text{hull}(G)$. Choose a saddle connection σ in T for which $[\sigma] \in \partial D$ and b are in different components of the closure of $D - \text{hull}(G)$. Denote by c the geodesic from b to $[\sigma]$. The length of σ decreases along c as it approaches $[\sigma]$. The length of σ is uniformly bounded above at b (Corollary 6.10) and uniformly bounded below at a point $t \in c \cap \partial \text{hull}(G)$ (Proposition 6.3). Therefore the subsegment of c between b and t has

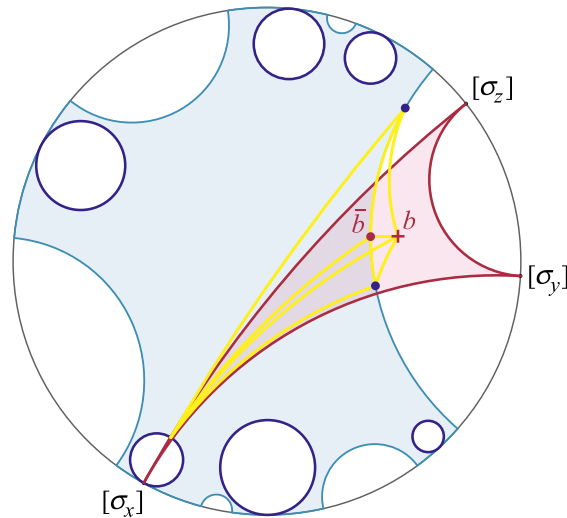


Figure 4: An example of Lemma 6.13: When \tilde{T} (red) intersects hull (G) (blue), there is a uniform bound on the distances (yellow curves) between the balance point b of \tilde{T} (red +), the adjusted balance point \bar{b} of \tilde{T} (red dot), the horoball B_x associated to the parabolic direction $[\sigma_x]$ (blue circle), and the horopoints B_y and B_z associated to the nonparabolic directions $[\sigma_y]$ and $[\sigma_z]$ (blue dots).

uniformly bounded length. Since \bar{b} is the closest point projection of b to $\partial \text{hull}(G)$, the distance between b and \bar{b} is no greater than the distance between b and t , so b and \bar{b} are uniformly close. \square

Corollary 6.14 Let $T = \Delta(x, y, z)$ be a Euclidean triangle consisting of saddle connections σ_x, σ_y , and σ_z so that $\tilde{T} \cap \text{hull}(G) \neq \emptyset$. If σ_x is nonparabolic, then it has uniformly bounded length at its associated horopoint. That is, the length of

$$f_{X_{[\sigma_x]}}(\sigma_x)$$

is uniformly bounded.

The next definition is from [8, Section 4.6]. The following corollary is comparable to [8, Corollary 4.11], which is given a new proof.

Definition 6.15 Two saddle connections in a common fiber span a triangle if they share an endpoint and the geodesic joining their other endpoints is a single (possibly degenerate) saddle connection. Write $\mathcal{P}(\sigma)$ for the set of saddle connections that span a triangle with σ . Denote

$$B(\sigma) = \bigcup_{\sigma' \in \mathcal{P}(\sigma)} B_{[\sigma']}$$

for the union of horoballs or horopoints associated to the saddle connections that span a triangle with σ . Taking preimages of the horoballs or horopoints, denote

$$\mathcal{B}(\sigma) = \bigcup_{\sigma' \in \mathcal{P}(\sigma)} \mathcal{B}_{[\sigma']}$$

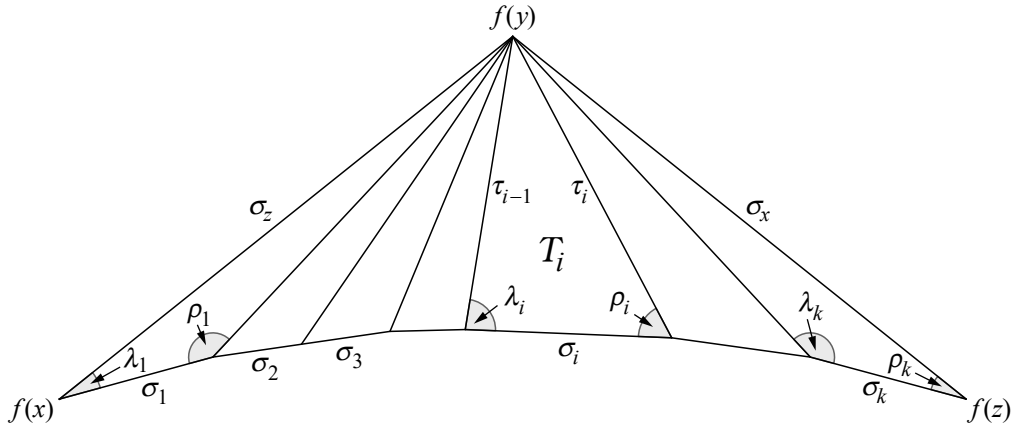


Figure 5: An example of a fan $\Delta(x, y, z)$ in E_0 in which $[f(x), f(y)]$ and $[f(y), f(z)]$ are each a single saddle connection, labeled with the notation used in the proof of Lemma 6.19 and throughout the rest of the paper. In this example, the saddle connections σ_2 and σ_3 are parallel, i.e., $[\sigma_2] = [\sigma_3]$.

Corollary 6.16 For any $\sigma' \in \mathcal{P}(\sigma)$, the sets $B_{[\sigma]}$ and $B_{[\sigma']}$ are uniformly close in D .

Proof Suppose saddle connections σ and σ' span a triangle T , and let \tilde{T} be the associated ideal triangle. Then $B_{[\sigma]}$ and $B_{[\sigma']}$ are close by Lemma 6.12. □

6.2 Fans

Definition 6.17 Let $x, y, z \in \Sigma$. The triangle $\Delta(x, y, z) \subset E_0$ is called a *fan* if it is nondegenerate and at least two sides consist of a single saddle connection. When the information about x, y, z is clear from context, a fan is sometimes denoted simply by F .

A fan canonically decomposes into a finite union of Euclidean triangles sharing a common vertex. See Figure 5 for an illustration and notation used throughout this section.

Definition 6.18 Let $F = \Delta(x, y, z) \subset E_0$ be a fan as in Figure 5. The *ideal fan associated to $\Delta(x, y, z)$* , denoted \tilde{F} , is the union of the ideal triangles $\tilde{T}_i \subset D$ associated to the Euclidean triangles $T_i \subset \Delta(x, y, z)$.

Figure 6 shows the ideal fan associated to the fan from Figure 5. For a fan F consisting of k Euclidean triangles, the associated ideal fan \tilde{F} consists of k ideal triangles with disjoint interiors formed from at most $2k + 1$ vertices; of these vertices, $k - 1$ vertices each belong to exactly two ideal triangles. The following lemma describes that these $k - 1$ vertices are contained in a connected component of the boundary at infinity.

Lemma 6.19 (structure lemma) Let $F = \Delta(x, y, z)$ be a fan consisting of k Euclidean triangles $\{T_i\}_{1 \leq i \leq k}$ as in Figure 5. Then the vertices of the associated ideal fan \tilde{F} appear in cyclic order (counterclockwise) in ∂D :

$$[\sigma_z] < [\tau_1] < \cdots < [\tau_{k-1}] < [\sigma_x] < [\sigma_k] \leq \cdots \leq [\sigma_1] < [\sigma_z].$$

Moreover, the ideal triangles $\{\tilde{T}_i\}_{1 \leq i \leq k}$ have disjoint interiors. See Figure 6.

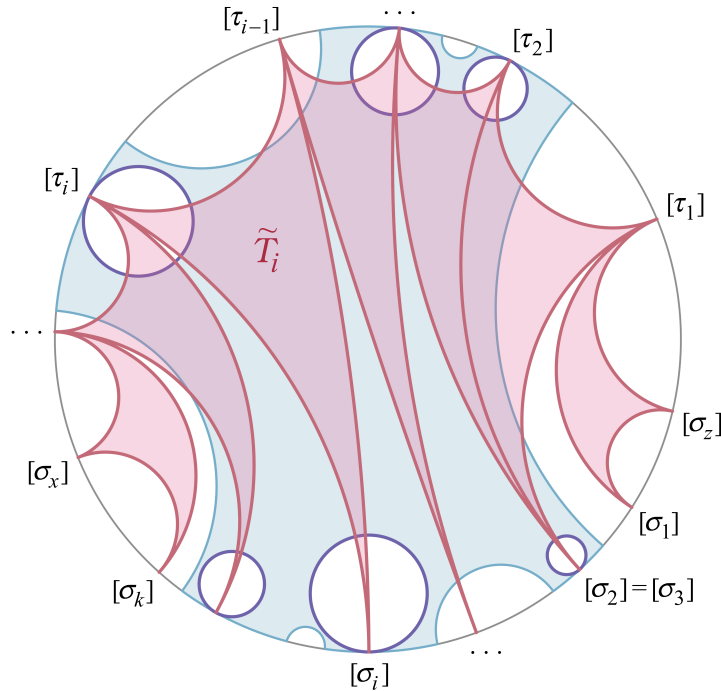


Figure 6: An example of an ideal fan (red) in D (with ∂D represented by the gray circle) corresponding to the Euclidean fan from Figure 5. An example of hull (G) is in blue. As given by the structure lemma (Lemma 6.19) each vertex $[\tau_*$] belongs to exactly two ideal triangles, and $[\sigma_x]$ and $[\sigma_z]$ each belong to exactly one ideal triangle. The remaining vertices may belong to one or more triangles: In this example, the vertex $[\sigma_2] = [\sigma_3]$ belongs to two triangles because σ_2 and σ_3 were parallel in the Euclidean fan. Most importantly, the interiors of the ideal triangles are pairwise disjoint. Notice that some triangles in the ideal fan do not intersect hull (G) , which is important for Corollary 6.20.

Proof Refer to Figures 5 and 6 for a Euclidean fan, an associated ideal fan, and related notation. Develop F into the Euclidean plane; it can be assumed that, after rotating, σ_1 is in the horizontal direction. Because $\sigma_1\sigma_2\cdots\sigma_k$ is a geodesic in E_0 , $\rho_i + \lambda_{i+1} \geq \pi$, and therefore

$$[\sigma_k] \leq [\sigma_{k-1}] \leq \cdots \leq [\sigma_1].$$

Note that for each T_i

$$[\tau_{i-1}] < [\tau_i] < [\sigma_i],$$

and therefore

$$[\sigma_z] = [\tau_0] < [\tau_1] < [\tau_2] < \cdots < [\tau_{k-2}] < [\tau_{k-1}] < [\tau_k] = [\sigma_x] < [\sigma_k] \leq [\sigma_{k-1}] \leq \cdots \leq [\sigma_1] < [\sigma_z].$$

Notice that the inequalities concerning the vertices of \tilde{T}_1 are “nested” inside those concerning \tilde{T}_2 — that is, the inequalities $[\sigma_1] < [\sigma_z] < [\tau_1]$ relating the vertices of \tilde{T}_1 are inserted (as $*$) into the inequalities $[\sigma_2] \leq * < [\tau_1] < [\tau_2]$ relating the vertices of \tilde{T}_2 — the inequalities concerning the vertices of \tilde{T}_2 are

nested inside those concerning \tilde{T}_3 , and so on. This demonstrates that the ideal triangles \tilde{T}_i , $1 \leq i \leq k$, have disjoint interiors and proves the lemma. \square

The following corollaries are used in the proofs of the fan lemma (Lemmas 6.24 and 6.22).

Corollary 6.20 *Any ideal fan in D can be split into at most three subfans: $\{\tilde{T}_1, \dots, \tilde{T}_n\}$, which is disjoint from $\text{hull}(G)$; $\{\tilde{T}_{n+1}, \dots, \tilde{T}_{m-1}\}$, for which all triangles intersect $\text{hull}(G)$; and $\{\tilde{T}_m, \dots, \tilde{T}_k\}$, which is also disjoint from $\text{hull}(G)$. (See, for example, Figure 6.)*

Recall that $\text{proj}_{\bar{D}} : D \rightarrow \bar{D}$ is the ρ -closest point projection.

Corollary 6.21 *Let $F = \Delta(x, y, z)$ be a fan, and let $\tilde{F} \subset D$ be the ideal fan associated to F . If $\tilde{F} \cap \text{hull}(G) = \emptyset$, then there is an upper bound on the diameter of $\text{proj}_{\bar{D}}(\tilde{F}) \subset \partial \text{hull}(G)$.*

Proof By Lemma 6.19, $\text{proj}_{\bar{D}}(\tilde{F})$ is contained in the image of the projection of one or two ideal triangles in \tilde{F} . In the former case, $\text{proj}_{\bar{D}}(\tilde{F})$ lies in the image of the projection of a single ideal triangle, which is uniformly bounded (Lemma 6.12). In the latter case these two ideal triangles share a common vertex, so an upper bound on the lengths of the projections of both triangles implies an upper bound on the diameter of $\text{proj}_{\bar{D}}(\tilde{F})$. \square

6.3 The fan lemma

This section presents a statement and proof of the fan lemma (Lemma 6.22), which says that the triangle of collapsed preferred paths associated to any Euclidean fan is uniformly slim. It is used later to demonstrate uniform slimness of general triangles of collapsed preferred paths (Theorem 6.28). Slimness of triangles of collapsed preferred paths is the only remaining condition needed to satisfy all hypotheses of the guessing geodesics criterion (Proposition 5.1), which gives hyperbolicity of \hat{E} (Theorem 5.7).

The statement of the fan lemma is identical to that of [8, Lemma 4.12], but the proof requires significant reformulation in order to accommodate the circumstance in which the Veech group G is not a lattice.

Lemma 6.22 (fan lemma) *There is $\delta' > 0$ so that if $x, y, z \in \Sigma$ and the geodesic triangle $\Delta(x, y, z)$ in E_0 is a fan with $[f(x), f(y)] = \sigma_z$ and $[f(y), f(z)] = \sigma_x$, then the triangle $\Delta^{\hat{\zeta}}(x, y, z)$ of collapsed preferred paths is δ' -slim. Furthermore, if y lies on a geodesic in D_y with endpoints in $\mathcal{B}(\sigma_z)$ and $\mathcal{B}(\sigma_x)$, then the collapsed preferred paths satisfy*

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

The “furthermore” statement is necessary for getting a uniform slimness constant for general triangles in later arguments. A Euclidean fan fitting the description in the fan lemma is shown in Figure 5. The (uncollapsed) preferred paths joining the vertices $x, y, z \in \Sigma$ are described in Tables 1 and 2.

The fan lemma is proved by [8] under the additional hypothesis that G is a lattice. In their case, all of the saddle connections are parabolic and therefore have length zero after being projected to \hat{E} by the map P . This places most of the burden of the proof on showing that the horizontal pieces of the paths along the top and bottom of the fan are sufficiently close. Here G is only required to be finitely

h'_z	horizontal geodesic in D_x from x to $X_{[\sigma_z]}$
γ'_z	saddle connection σ_z traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_z]}}(\sigma_z)$
\bar{h}'_z	horizontal geodesic in D_y from $X_{[\sigma_z]}$ to y
\bar{h}'_x	horizontal geodesic in D_y from y to $X_{[\sigma_x]}$
γ'_x	saddle connection σ_x traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_x]}}(\sigma_x)$
h'_x	horizontal geodesic in D_z from $X_{[\sigma_x]}$ to z

Table 1: In order, the pieces forming the concatenation of the (uncollapsed) preferred paths forming the “top” of the fan referred to in Lemma 6.22. The primed notation on the paths in this table and in Table 2 serves a practical purpose: The unprimed notation is reserved for paths which play a much more active role in the proof of the fan lemma.

generated, so preferred paths may contain nonparabolic saddle connections whose lengths are arbitrarily long even after collapsing. In the proof of the fan lemma that follows, the preferred path along the bottom of the fan is replaced with a substitute path to which it is uniformly close. The substitute path is chosen specifically to allow nonparabolic saddle connections to be traversed in a select set of fibers where

h'_0	horizontal geodesic in D_x from x to $X_{[\sigma_1]}$
γ'_1	saddle connection σ_1 traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_1]}}(\sigma_1)$
h'_1	horizontal geodesic in $D_{\gamma'^+_1} = D_{\gamma'^-_2}$ from $X_{[\sigma_1]}$ to $X_{[\sigma_2]}$
...	...
γ'_i	saddle connection σ_i traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_i]}}(\sigma_i)$
h'_i	horizontal geodesic in $D_{\gamma'^+_i} = D_{\gamma'^-_{i+1}}$ from $X_{[\sigma_i]}$ to $X_{[\sigma_{i+1}]}$
...	...
γ'_k	saddle connection σ_k traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_k]}}(\sigma_k)$
h'_k	horizontal geodesic in D_z from $X_{[\sigma_k]}$ to z

Table 2: In order, the pieces forming the (uncollapsed) preferred path forming the “bottom” of the fan referred to in Lemma 6.22. The primed notation on the paths in this table and in Table 1 serves a practical purpose: The unprimed notation is reserved for paths which play a much more active role in the proof of the fan lemma. The superscripts + and – denote endpoints of geodesics, as described in Section 5.1.

slimness of Euclidean fans can be employed. The horizontal pieces of the substitute path are handled mostly in the same manner as in the lattice case presented by [8], except for a few horizontal pieces possibly requiring special treatment.

The following lemma will be necessary on a few occasions.

Lemma 6.23 (tube lemma; adapted from [14, Lemma 4.5]) *Let $[\sigma], [\sigma'] \in \mathcal{D}$, and let v be a geodesic in D with endpoints in $\partial B_{[\sigma]}$ and $\partial B_{[\sigma']}$. If w is a geodesic in D with endpoints each within a bounded distance of $B_{[\sigma]}$ and $B_{[\sigma']}$, then w lies within a uniform neighborhood of $v \cup B_{[\sigma]} \cup B_{[\sigma']}$.*

To simplify the proof of the fan lemma and to demonstrate the utility of substituting preferred paths, the case in which the ideal fan is disjoint from hull (G) is presented separately.

Lemma 6.24 (disjoint fan lemma) *The fan lemma (Lemma 6.22) holds under the additional assumption that the ideal fan in D associated to $\Delta(x, y, z)$ does not intersect hull (G).*

Proof Note that objects are sometimes lifted from D to various horizontal disks. For the sake of simplicity, the lifts are referred to by the same notation wherever it does not cause confusion.

Because the ideal fan never intersects the hull, all of the saddle connections in the preferred paths are traversed in the fibers over their associated horopoints, which lie in a uniformly bounded subsegment of $\partial \text{hull}(G)$ (Corollary 6.21). The general trajectory of the proof is to replace the preferred paths with uniformly close substitute paths in which all of the saddle connections are traversed in the fiber over $X_{[\sigma_z]}$, then prove that the triangle of substitute paths is slim.

The substitute path along the “top” of the fan is the concatenation of the following paths, in order:

h_z	horizontal geodesic in D_x from x to $X_{[\sigma_z]}$
γ_z	saddle connection σ_z traversed in the fiber over $X_{[\sigma_z]}$, i.e., $f_{X_{[\sigma_z]}}(\sigma_z)$
\bar{h}_z	horizontal geodesic in D_y from $X_{[\sigma_z]}$ to y
\bar{h}_x	horizontal geodesic in D_y from y to $X_{[\sigma_x]}$
γ_x	saddle connection σ_x traversed in the fiber over $X_{[\sigma_x]}$, i.e., $f_{X_{[\sigma_x]}}(\sigma_x)$
h_x	horizontal geodesic in D_z from $X_{[\sigma_x]}$ to z

This substitute path is uniformly close to the path in Table 1, which can be seen by comparing each pair of analogous pieces as follows.

- Since $h_z = h'_z$, these are uniformly close.
- Since $\gamma_z = \gamma'_z$, these are uniformly close.
- Since $\bar{h}_z = \bar{h}'_z$, these are uniformly close.
- Since $\bar{h}_x, \bar{h}'_x \subset D_y$ share an initial point and have terminal points ($X_{[\sigma_z]}$ and $X_{[\sigma_x]}$, respectively) which are a uniformly bounded distance apart by Corollary 6.21, \bar{h}_x and \bar{h}'_x are uniformly close.

- The saddle connections γ_x and γ'_x are saddle connections in the fibers over $X_{[\sigma_z]}$ and $X_{[\sigma_x]}$, respectively. The map $f_{X_{[\sigma_z]}, X_{[\sigma_x]}}$ maps γ'_x to γ_x , and so the distance between them is exactly the distance between $X_{[\sigma_x]}$ and $X_{[\sigma_z]}$, which is uniformly bounded.
- Since $h_x, h'_x \subset D_z$ share a terminal point and have initial points ($X_{[\sigma_z]}$ and $X_{[\sigma_x]}$, respectively) a uniformly bounded distance apart, h_x and h'_x are uniformly close.

The substitute path along the “bottom” of the fan is the concatenation of the following paths, in order:

h_0	horizontal geodesic in D_x from x to $X_{[\sigma_z]}$
γ_1	saddle connection σ_1 traversed in the fiber over $X_{[\sigma_z]}$, i.e., $f_{X_{[\sigma_z]}}(\sigma_1)$
h_1	point $\gamma_1^+ = \gamma_2^-$
...	...
γ_i	saddle connection σ_i traversed in the fiber over $X_{[\sigma_z]}$, i.e., $f_{X_{[\sigma_z]}}(\sigma_i)$
h_i	point $\gamma_i^+ = \gamma_{i+1}^-$
...	...
γ_k	saddle connection σ_k traversed in the fiber over $X_{[\sigma_z]}$, i.e., $f_{X_{[\sigma_z]}}(\sigma_k)$
h_k	horizontal geodesic in D_z from $X_{[\sigma_z]}$ to z

All but the first and last horizontal paths are points because the concatenation of saddle connections γ_i in the fiber over $X_{[\sigma_z]}$ is already continuous. Again, this substitute path is uniformly close to the path in Table 2, which can be seen by comparing each pair of analogous pieces as follows.

- Since $h_0, h'_0 \subset D_x$ share an initial point and have terminal points ($X_{[\sigma_z]}$ and $X_{[\sigma_1]}$, respectively) a uniformly bounded distance apart, h_0 and h'_0 are uniformly close.
- For all i , the saddle connections γ_i and γ'_i are saddle connections in the fibers over $X_{[\sigma_z]}$ and $X_{[\sigma_i]}$, respectively. The map $f_{X_{[\sigma_z]}, X_{[\sigma_i]}}$ maps γ'_i to γ_i , and so the distance between them is exactly the distance between $X_{[\sigma_i]}$ and $X_{[\sigma_z]}$, which is uniformly bounded.
- For each $1 \leq i \leq k - 1$, h_i and h'_i have initial points which are the terminal points of γ_i and γ'_i , respectively, and terminal points which are the initial points of γ_{i+1} and γ'_{i+1} , respectively. Since for all i the pair of γ_i, γ'_i are uniformly close and h_i is contained in a hyperbolic space (with hyperbolicity constant uniform for all i), the pair h_i, h'_i must also be uniformly close.
- Since $h_k, h'_k \subset D_z$ share a terminal point and have initial points ($X_{[\sigma_z]}$ and $X_{[\sigma_k]}$, respectively) a uniformly bounded distance apart, h_k and h'_k are uniformly close.

Since h_i is degenerate for each $1 \leq i \leq k - 1$, the substitute path along the “bottom” of the fan is simply $h_0\gamma_1\gamma_2 \cdots \gamma_k h_k$. Then because $h_z = h_0, \bar{h}_z = \bar{h}_x$ (ignoring orientation), and $h_x = h_k$, to show that the triangle of substitute paths is slim it only remains to show that the saddle connections γ_* form a slim triangle.

Claim 6.25 *The Euclidean fan formed by the saddle connections $\gamma_x, \gamma_z, \gamma_1, \dots, \gamma_k$ in the fiber over $X_{[\sigma_z]}$ is uniformly slim.*

Proof For any $X \in D$, the fiber E_X is quasi-isometric to (\tilde{S}, \tilde{X}_0) , the universal cover of the closed surface (S, X_0) , which is hyperbolic. Given any X_1 and X_2 in an ϵ -neighborhood of X_0 , the map $f_{X_2, X_1} : E_{X_1} \rightarrow E_{X_2}$ is e^ϵ -bilipschitz (Section 4.2.1), so the fibers E_{X_1} and E_{X_2} are uniformly hyperbolic. Finally, because G acts cocompactly on the truncated convex hull \bar{D} , it holds that all fibers E_X with $X \in \bar{D}$ are uniformly hyperbolic. Therefore the Euclidean fan formed by the saddle connections $\gamma_x, \gamma_z, \gamma_1, \dots, \gamma_k$ in the fiber over $X_{[\sigma_z]}$ is uniformly slim. \square

Therefore, the original triangle of preferred paths is slim.

To prove the “furthermore” statement, note that since \bar{h}'_z and \bar{h}'_x are both in D_y , they form a slim triangle with the geodesic joining their endpoints, denoted h'_y . When y is assumed to lie on a geodesic with endpoints in $\mathcal{B}(\sigma_z)$ and $\mathcal{B}(\sigma_x)$, the concatenation of \bar{h}'_z and \bar{h}'_x lies uniformly close to h'_y (Lemma 6.23). By construction, h'_y is a uniformly bounded subsegment of $\partial \text{hull}(G)$. Therefore h'_y — and the concatenation of \bar{h}'_z and \bar{h}'_x — is uniformly close to the shared endpoint of γ_z and γ_x , which results in the containments from the lemma. \square

When the ideal fan is not disjoint from $\text{hull}(G)$, the proof is more complex. First the fan is decomposed into subfans as in Corollary 6.20: One subfan consists of all of the ideal triangles that intersect $\text{hull}(G)$, and the other subfans (up to two, if any exist) are disjoint from $\text{hull}(G)$. Thanks to the tools in Lemma 6.13, the paths associated to the ideal subfan which intersects $\text{hull}(G)$ can be handled in a manner similar to that of [8], with a few technicalities to consider for the horizontal paths at the beginning and end. A subfan which is disjoint from $\text{hull}(G)$ is handled in a similar manner to the previous proof; however, some additional work is needed since slimness of the associated Euclidean subfan in any particular fiber over $\partial \text{hull}(G)$ is no longer sufficient to prove slimness of the whole fan.

Proof of fan lemma Like the disjoint case, note that objects are sometimes lifted from D to other horizontal disks, and the lifts are referred to by the same notation wherever it does not cause confusion. Finally, note that some arguments are made by showing that the uncollapsed pieces (i.e., before applying P) are close, which is sufficient because P is 1-Lipschitz (Lemma 4.3).

Substituting the concatenated preferred paths along the top of the fan. Compare to the concatenated preferred paths described in Table 1. The substitute path along the top of the fan is the concatenation of the following paths, in order:

$h_z := h'_z$	horizontal geodesic in D_x from x to $X_{[\sigma_z]}$
$\gamma_z := \gamma'_z$	saddle connection σ_z traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_z]}}(\sigma_z)$
h'_y	horizontal geodesic in D_y joining $\bar{h}'_z{}^- \in X_{[\sigma_z]}$ to $\bar{h}'_x{}^+ \in X_{[\sigma_x]}$
$\gamma_x := \gamma'_x$	saddle connection σ_x traversed in the fiber over the associated horoball or horopoint, i.e., $f_{X_{[\sigma_x]}}(\sigma_x)$
$h_x := h'_x$	horizontal geodesic in D_z from $X_{[\sigma_x]}$ to z

Since \bar{h}'_z and \bar{h}'_x are both in D_y , they form a slim triangle with the geodesic joining their endpoints, denoted h'_y . In particular, h'_y lies in a uniform neighborhood of \bar{h}'_z and \bar{h}'_x . Because the rest of the pieces of the substitute path along the top of the fan are identical to those appearing in the (uncollapsed) preferred paths, the substitute path along the top of the fan lies in a uniform neighborhood of $\zeta(x, y) \cup \zeta(y, z)$ (while the reverse is not necessarily true). If in addition y lies on a geodesic with endpoints in $B(\sigma_z)$ and $B(\sigma_x)$, then the concatenation of \bar{h}'_z with \bar{h}'_x is in fact *uniformly* close to $h'_y \cup B(\sigma_z) \cup B(\sigma_x)$ (Lemma 6.23); therefore the image under P of the substitute path along the top of the fan is uniformly close to $\hat{\zeta}(x, y) \cup \hat{\zeta}(y, z)$.

Substituting the preferred path along the bottom of the fan. Compare to the preferred path described in Table 2. There are at most two maximal subfans of the ideal fan which do not intersect hull (G), and any such subfan must be at the beginning or end of the fan (Corollary 6.20). Let n be the largest index for which $\tilde{T}_1, \dots, \tilde{T}_n$ do not intersect hull (G). If no such n exists, set $n = 0$. Assume that $n < k$; the case when $n = k$ (that is, when no ideal triangle in the fan intersects hull (G)) is proven in Lemma 6.24. Similarly, let m be the smallest index for which $\tilde{T}_m, \dots, \tilde{T}_k$ do not intersect hull (G). If no such m exists, set $m = k + 1$. Note that as a result of the structure lemma, $m - n \geq 2$ (Corollary 6.20). Define

$$\gamma_i := \begin{cases} f_{X_{[\sigma_z]}}(\sigma_i), & 1 \leq i \leq n, \\ \gamma'_i, & n < i < m, \\ f_{X_{[\sigma_x]}}(\sigma_i), & m \leq i \leq k. \end{cases}$$

The horizontal paths are chosen to make the substitute path along the bottom of the fan continuous. These are explicitly described as follows. Because the saddle connections associated to the triangles T_{n+1}, \dots, T_{m-1} remain unchanged, the horizontal paths connecting them are also unchanged. That is, define

$$h_i := h'_i \quad \text{for } n + 1 \leq i \leq m - 2.$$

The paths $\gamma_1 \gamma_2 \cdots \gamma_n$ and $\gamma_m \gamma_{m+1} \cdots \gamma_k$ are the concatenation of the saddle connections $\sigma_1, \dots, \sigma_n$ in the fiber over $X_{[\sigma_z]}$ and the concatenation of the saddle connections $\sigma_m, \dots, \sigma_k$ in the fiber over $X_{[\sigma_x]}$, respectively. Note in particular that these paths are continuous without the presence of horizontal paths, so the horizontal “paths” are chosen to be the endpoints of saddle connections,

$$h_i := \gamma_i^+ = \gamma_{i+1}^- \quad \text{for } 1 \leq i \leq n - 1 \text{ or } m \leq i \leq k - 1.$$

When $n \geq 1$ and $m \leq k$, define the remaining horizontal paths as follows:

- $h_0 := h_z$ horizontal geodesic in D_x from x to $X_{[\sigma_z]}$
- h_n horizontal geodesic in $D_{\gamma_n^+}$ from γ_n^+ to γ_{n+1}^-
- h_{m-1} horizontal geodesic in $D_{\gamma_{m-1}^+}$ from γ_{m-1}^+ to γ_m^-
- $h_k := h_x$ horizontal geodesic in D_z from $X_{[\sigma_x]}$ to z

If instead $n = 0$, define $h_0 = h_n := h'_0$, and if instead $m = k + 1$, define $h_k = h_{m-1} := h'_k$. In other words, if there is no ideal subfan $\widetilde{T}_1, \dots, \widetilde{T}_n$ (or $\widetilde{T}_m, \dots, \widetilde{T}_k$) disjoint from $\text{hull}(G)$, then the beginning (or end) of the substitute path along the bottom of the fan is chosen to be identical to the original preferred path.

The reader will gain the most insight from this proof by assuming in addition that $n \geq 1$ and $m \leq k$, since this represents the most novel circumstance in which the ideal fan can be split into three subfans depending on their intersection with $\text{hull}(G)$ (see Corollary 6.20). The following claim, for instance, becomes tautological otherwise.

Claim 6.26 *The substitute path along the bottom of the fan, $h_0\gamma_1h_1 \cdots \gamma_ih_i \cdots \gamma_kh_k$, is uniformly close to the preferred path along the bottom of the fan,*

$$\zeta(x, z) = h'_0\gamma'_1h'_1 \cdots \gamma'_ih'_i \cdots \gamma'_kh'_k.$$

Proof of claim If $n = 0$ and $m = k + 1$ — that is, if every triangle in the ideal fan intersects $\text{hull}(G)$ — then the preferred path and the substitute path are identical. Otherwise if there are triangles in the ideal fan which do not intersect $\text{hull}(G)$, then only the initial and/or terminal portions of the preferred and substitute paths are distinct. This proof will demonstrate the case when $n \geq 1$; the proof for the case when $m \leq k$ is analogous, and then the case when both $n \geq 1$ and $m \leq k$ follows immediately.

For any $i > n$, $h_i = h'_i$ and $\gamma_i = \gamma'_i$. So it remains to demonstrate that

$$h'_0\gamma'_1h'_1 \cdots \gamma'_nh'_n \quad \text{and} \quad h_0\gamma_1h_1 \cdots \gamma_nh_n$$

are uniformly close. This comes as a consequence of Corollary 6.21, which states that the projection of the ideal subfan $\widetilde{T}_1, \dots, \widetilde{T}_n$ to $\text{hull}(G)$ is a uniformly bounded subsegment of $\partial \text{hull}(G)$. Therefore the initial segments h'_0 and h_0 share an endpoint at x , and their other endpoints are uniformly close, so h'_0 and h_0 are uniformly close since they lie in a horizontal disk which is uniformly hyperbolic. Similarly, the final segments h'_n and h_n share an endpoint at γ_{n+1}^- , and their other endpoints are uniformly close, so h'_n and h_n are uniformly close. The point $\gamma_i^- = h_{i-1}^+$ is uniformly close to $\gamma_i'^- = h_{i-1}^+$, and $\gamma_i^+ = h_i^-$ is uniformly close to $\gamma_i'^+ = h_i^-$, so γ_i' is uniformly close to γ_i (as in the proof of Lemma 6.24). Finally, the endpoints of each h'_i , $1 \leq i \leq n - 1$, must be within a uniformly bounded distance of each h_i . \square

To prove the fan lemma, it now suffices to show that the substitute paths along the top and bottom of the fan are uniformly close to each other.

Horizontal pieces at x and z are close. When $n \geq 1$, $h_z = h_0$ by construction. Otherwise for $n = 0$, h_z and h_0 share an endpoint at x , and their other endpoints are uniformly close because the ideal triangle \widetilde{T}_1 intersects $\text{hull}(G)$ (Lemma 6.13). So, in either case, it follows that h_z and h_0 are uniformly close.

Analogous arguments show that h_x and h_k are uniformly close.

Saddle connections γ_i for $i \leq n$ or $i \geq m$. Recall that $\gamma_1h_1\gamma_2h_2 \cdots h_{n-1}\gamma_n = \gamma_1\gamma_2 \cdots \gamma_n$ is a continuous path in the fiber over $X_{[\sigma_z]}$. It forms one side of a Euclidean fan whose other sides are γ_z and τ_n — that is, the fan formed by triangles T_1, \dots, T_n , which is slim (Claim 6.25). By construction, the saddle connection τ_n is nonparabolic. Because the ideal triangle \widetilde{T}_{n+1} intersects $\text{hull}(G)$, the saddle connection τ_n

has uniformly bounded length in the fiber over $X_{[\tau_n]}$ (Corollary 6.14). Since $X_{[\tau_n]}$ is within a uniformly bounded distance of $X_{[\sigma_z]}$ (Corollary 6.21), the saddle connection τ_n has uniformly bounded length in the fiber over $X_{[\sigma_z]}$. Therefore $\gamma_1 \cdots \gamma_n$ is uniformly close to γ_z because they are in $E_{X_{[\sigma_z]}}$ which is uniformly hyperbolic.

Analogous arguments show that for $i \geq m$ the saddle connection γ_x in the top of the fan is uniformly close to the paths $\gamma_m \gamma_{m+1} \cdots \gamma_k$ in the substitute path along the bottom of the fan.

Horizontal pieces h_i for $n + 1 \leq i \leq m - 2$. The ideal triangles \tilde{T}_{n+1} and \tilde{T}_{m-1} are the first and last ideal triangles intersecting the hull, respectively. Define h to be the geodesic in D between the adjusted balance points \bar{b}_{n+1} and \bar{b}_{m-1} , and let h_y denote its lift to D_y .² Because the ideal triangle \tilde{T}_{n+1} intersects hull (G) , there is a uniformly bounded distance from the endpoint $h_y^- = \bar{b}_{n+1}$ to $B_{[\tau_n]}$ (Lemma 6.13), which lies a uniformly bounded distance from $B_{[\sigma_z]}$ (indeed if $n = 0$, $B_{[\tau_n]} = B_{[\sigma_z]}$; otherwise this follows from Corollary 6.21). Similarly, there is a uniformly bounded distance from the endpoint $h_y^+ = \bar{b}_{m-1}$ to $B_{[\sigma_x]}$. Because $h_y^- \in \partial B_{[\sigma_z]}$ and $h_y^+ \in \partial B_{[\sigma_x]}$, the tube lemma (Lemma 6.23) ensures that h_y lies uniformly close to $h'_y \cup B_{[\sigma_z]} \cup B_{[\sigma_x]}$. Therefore $P(h_y)$ and $P(h'_y)$ must have uniformly bounded Hausdorff distance.

The aim is to show that $P(h_y)$ and $P(h_{n+1} \cup \cdots \cup h_{m-1})$ have uniformly bounded Hausdorff distance. This is accomplished by breaking h_y into segments, each of which has uniformly bounded Hausdorff distance to its image under P .

Recall that h_i is in the disk $D_{\mathcal{V}_i^+}$ connecting the horoballs (or horopoints) $B_{[\sigma_i]}$ and $B_{[\sigma_{i+1}]}$. Define h''_i to be the geodesic in $D_{\mathcal{V}_i^+}$ between (the lifts of) \bar{b}_i and \bar{b}_{i+1} , the adjusted balance points of ideal triangles \tilde{T}_i and \tilde{T}_{i+1} . Because $\tilde{T}_i \cap \text{hull}(G) \neq \emptyset$ and $\tilde{T}_{i+1} \cap \text{hull}(G) \neq \emptyset$, there is a uniform bound on the distance from \bar{b}_i to $B_{[\sigma_i]}$ and from \bar{b}_{i+1} to $B_{[\sigma_{i+1}]}$ (Lemma 6.13). (Note that if there are multiple saddle connections in the same direction, then h_i might be a point which is uniformly close to both endpoints of h''_i .) Then by the tube lemma (Lemma 6.23), h''_i lies uniformly close (in $D_{\mathcal{V}_i^+}$) to $h_i \cup B_{[\sigma_i]} \cup B_{[\sigma_{i+1}]}$. Therefore $P(h''_i)$ and $P(h_i)$ have uniformly bounded Hausdorff distance.

Define g_i to be the geodesic in D_y joining \bar{b}_i and \bar{b}_{i+1} . Since \bar{b}_i and \bar{b}_{i+1} are both within a uniformly bounded distance of $B_{[\tau_i]}$ (Lemma 6.13), the saddle connection τ_i has bounded length over the geodesic $[\bar{b}_i, \bar{b}_{i+1}]$ and therefore h''_i and g_i have bounded Hausdorff distance in E (and in \bar{E}).

The following claim will allow h_y to be broken into segments whose images under P will be shown to be uniformly close to each of the $P(g_i)$. Recall that $p : \text{hull}(G) \rightarrow \widehat{D}$ is the map that collapses horoballs in \bar{D} to points.

Claim 6.27 *There are points t_{n+1}, \dots, t_{m-1} appearing in order along the geodesic h whose images under $\bar{p} : D \rightarrow \widehat{D}$ respectively lie within uniformly bounded distance of the collapsed horoballs (or horopoints) $\bar{p}(B_{[\sigma_{n+1}]}) , \dots , \bar{p}(B_{[\sigma_{m-1}]})$.*

²Recall from the beginning of the proof that for the sake of simplicity, objects in D and their lifts to various horizontal disks are generally referred to by the same notation. This is the first (and only) instance when an object in D , h , is *notationally* distinguished from one of its lifts, $h_y \subset D_y$.

This statement is functionally analogous to [8, Claim 4.15]. Their proof relies on defining h between two balance points, where the associated Euclidean triangles are equilateral, and using the intermediate value theorem to identify points on h which must be close to the balance points of each of the intermediate triangles in the fan. The proof below relies instead on the structure lemma (Lemma 6.19) to force the same conclusion. Importantly, the approach presented here works in the more general case where h is defined between *adjusted* balance points, where the associated Euclidean triangles might not be equilateral. A key advantage to this construction is that all of the t_i are forced to be in $\text{hull}(G)$, which avoids the issue of later needing to project paths back onto the bundle over the hull.

Proof of Claim 6.27 Recall that h is the geodesic joining adjusted balance points \bar{b}_{n+1} and \bar{b}_{m-1} . Recall also that for each $i \in \{n + 1, \dots, m - 1\}$ the ideal triangle \tilde{T}_i intersects $\text{hull}(G)$ by construction. Therefore each adjusted balance point \bar{b}_i is within a uniformly bounded distance of all three horoballs or horopoints associated to \tilde{T}_i (Lemma 6.13). Choose $t_{n+1} := \bar{b}_{n+1}$ and $t_{m-1} := \bar{b}_{m-1}$, which are close to $B_{[\sigma_{n+1}]}$ and $B_{[\sigma_{m-1}]}$, respectively.

By the structure lemma (Lemma 6.19), h must pass through all of the ideal triangles \tilde{T}_i for $n + 1 \leq i \leq m - 1$; and by convexity of $\text{hull}(G)$, h is entirely contained in $\text{hull}(G)$. An example is illustrated in Figure 7. Therefore if h intersects any of the geodesics joining \bar{b}_i to the horoballs or horopoints associated to \tilde{T}_i , then setting t_i to be the intersection point satisfies the claim. (If h intersects more than one of these geodesics, then any choice of intersection point will suffice.) If $h \cap \tilde{T}_i \subset B_{[\sigma_i]}$, then any choice of $t_i \in h \cap \tilde{T}_i$ will satisfy the claim.

Finally, consider the case in which h avoids intersecting any of the three bounded length geodesics joining \bar{b}_i to the horoballs or horopoints associated to \tilde{T}_i by passing through both of the horoballs $B_{[\tau_{i-1}]}$ and $B_{[\tau_i]}$. An example is illustrated in Figure 8. By Lemma 6.13 and the tube lemma (Lemma 6.23), the point \bar{b}_i lies within a uniform neighborhood of $B_{[\tau_{i-1}]} \cup B_{[\tau_i]}$ and the subsegment of h between these two horoballs,

$$h_i^\# := h \cap \tilde{T}_i \cap \left(B_{[\sigma_i]}^\circ \cup B_{[\tau_{i-1}]}^\circ \cup B_{[\tau_i]}^\circ \right)^C.$$

Choose t_i to be the closest point on $h_i^\#$ to \bar{b}_i , which may be in $\partial B_{[\tau_{i-1}]}$ or $\partial B_{[\tau_i]}$. Then $\bar{p}(t_i)$ is uniformly close to \bar{b}_i , which is uniformly close to $\bar{p}(B_{[\sigma_i]})$. □

Let s_i be the lift of t_i to D_y , and break h_y into segments

$$h_y^i = [s_i, s_{i+1}], \quad i = n + 1, \dots, m - 2.$$

By construction, in D_y each $P(s_i)$ is within uniformly bounded distance of $P(B_{[\sigma_i]})$, which is within uniformly bounded distance of \bar{b}_i (note that $P(\bar{b}_i) = \bar{b}_i$). Recalling that g_i is the geodesic in D_y joining \bar{b}_i and \bar{b}_{i+1} , this implies that $P(h_y^i)$ is uniformly close to $P(g_i)$. Also recalling that $P(g_i)$ and $P(h_i)$ have uniformly bounded Hausdorff distance, each $P(h_i)$ is uniformly close to each $P(h_y^i)$, so $P(h_{n+1} \cup \dots \cup h_{m-2})$ is uniformly close to $P(h_y^{n+1} \cup \dots \cup h_y^{m-2}) = P(h_y)$. Finally, recalling that $P(h_y)$ and $P(h'_y)$ have uniformly bounded Hausdorff distance, $P(h_{n+1} \cup \dots \cup h_{m-1})$ and $P(h'_y)$ have

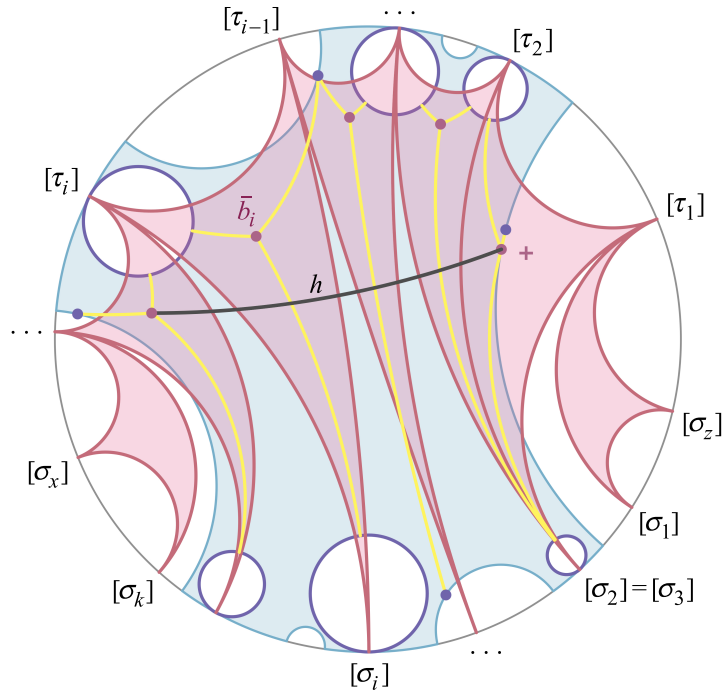


Figure 7: An example of a construction of the $\{t_i\}$ in the proof of Claim 6.27. The ideal fan from Figure 6 (red) has been superimposed on an example of an approximation to a truncated convex hull \bar{D} (blue). Note that in this drawing, $n = 1$ and $m = 1$ since there is one triangle disjoint from the hull at the beginning and end of the fan, so the initial and final triangles are not considered by the claim. For the vertices of the fan at parabolic points, the dark blue circles are the boundaries of the horoballs B_* . For the vertices of the fan at nonparabolic points, the dark blue dots are the horopoints B_* . (Refer to Definition 4.1.) Then for each ideal triangle the adjusted balance point \bar{b}_* is represented by a red dot; in this example, the only triangle with $\bar{b}_* \neq b_*$ is on the right, where the balance point is represented by a red +. Each of the adjusted balance points is a uniformly bounded distance from the horoballs and/or horopoints associated to the same triangle (Lemma 6.13), and the shortest distance to each of these is represented by a yellow curve. The geodesic h is drawn in black between adjusted balance points \bar{b}_{n+1} and \bar{b}_{m-1} . In this example, all of the t_* are chosen to be the points of intersection of h with the yellow curves.

uniformly bounded Hausdorff distance. Since h'_y was a piece of the substitute path along the top of the fan and the h_i , $n + 1 \leq i \leq m - 2$ were pieces of the substitute path along the bottom of the fan, this piece of the argument is complete.

Note that in the case where $m - n = 2$ — that is, where only one ideal triangle intersects hull (G) — this argument continues to work despite that the geodesic h is in fact not a path but rather the point $\bar{b}_{n+1} = \bar{b}_{m-1}$. In fact, Claim 6.27 follows immediately from Lemma 6.13.

Horizontal pieces h_n ($n \geq 1$) and h_{m-1} ($m \leq k$). When $n \geq 1$, the horizontal path h_n still requires consideration. Conveniently, $P(h_n)$ must have uniformly bounded length. The endpoint γ_n^+ of h_n is joined to γ_z^+ by the saddle connection τ_n in the fiber over $X_{[\sigma_z]}$, where it has uniformly bounded length because

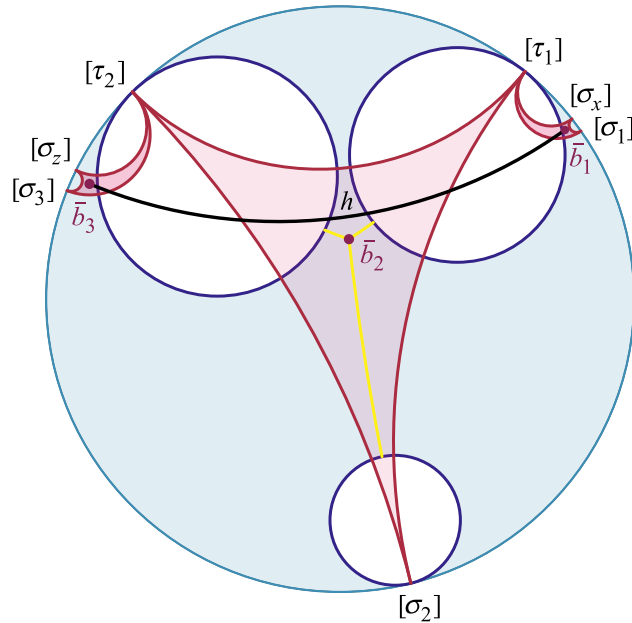


Figure 8: An example of a case in the proof of Claim 6.27. The colors and labels are analogous to those in Figure 7; for simplicity, in this example hull (G) (blue) is assumed to be all of D . Because h (black geodesic) does not intersect any of the yellow curves, t_2 will be chosen to be the point on h and in \tilde{T}_2 (red) which is closest to b_2 ; although t_2 and b_2 may not be uniformly close, their images under \tilde{p} must be uniformly close.

the ideal triangle \tilde{T}_{n+1} intersects hull (G) and because $X_{[\tau_n]}$ is uniformly close to $X_{[\sigma_z]}$ (Corollaries 6.14 and 6.21). Since $\gamma_z^+ = h_y'^-$, the arguments in the previous part of this proof give that $P(\gamma_z^+) = P(h_y'^-)$ is within a uniformly bounded distance of $P(h_y^-) = P(\bar{b}_{n+1})$ in $P(D_y)$, which is within a uniformly bounded distance of the P -image of $B_{[\sigma_{n+1}]} \ni \gamma_{n+1}^+ = h_{n+1}^-$ (Claim 6.27). Lastly, γ_{n+1} has uniformly bounded length in the fiber over $X_{[\sigma_{n+1}]}$ —either it is a nonparabolic saddle connection associated to an ideal triangle which intersects hull (G) (Corollary 6.14), or it is a parabolic saddle connection with bounded length by construction—and $\gamma_{n+1}^- = h_n^+$, so $P(h_n^-)$ and $P(h_n^+)$ must be uniformly close.

When $m \leq k$, a similar argument shows that $P(h_{m-1})$ has uniformly bounded length.

Remaining saddle connections. The remaining saddle connections must have uniformly bounded length because their endpoints—which are endpoints of the horizontal paths—have been shown to be uniformly close. However, it is insightful to see these remaining saddle connections handled explicitly.

Any remaining nonparabolic saddle connections must be associated to ideal triangles which intersect the hull, so each of these saddle connections has uniformly bounded length at its associated horopoint (Corollary 6.14). Therefore any isolated saddle connection—that is, any σ_i for which $[\sigma_{i-1}] \neq [\sigma_i] \neq [\sigma_{i+1}]$ —will not prevent slimness of the collapsed preferred paths. (Note that this also resolves any obstruction to slimness posed by the saddle connections γ_z and γ_x in the cases where $n = 0$ and $m = k + 1$, respectively.)

Suppose instead that there are consecutive indices

$$I = \{i_1, \dots, i_j\} \subset \{n + 1, n + 2, \dots, m - 1\}$$

such that $[\sigma_{i_1}] = \dots = [\sigma_{i_j}]$. Denote this common direction by $[\sigma_I]$. Since the ideal triangles $\widetilde{T}_{i \in I}$ intersect hull (G) , all the saddle connections in $T_{i \in I}$ have uniformly bounded length at their associated (adjusted) balance points \bar{b}_i (Corollary 6.10), which are all uniformly close to $X_{[\sigma_I]}$ (Lemma 6.13). Therefore all of the saddle connections associated to $T_{i \in I}$ have uniformly bounded length in the fiber over $X_{[\sigma_I]}$. In particular, because all of the $\gamma_{i \in I}$ form one side of the geodesic triangle $f_{X_{[\sigma_I]}}(T_{i_1} \cup \dots \cup T_{i_j})$, the triangle inequality gives that

$$\sum_{i \in I} \ell(\gamma_i) \leq \ell\left(f_{X_{[\sigma_I]}}(\tau_{i_1-1})\right) + \ell\left(f_{X_{[\sigma_I]}}(\tau_{i_j})\right),$$

which is uniformly bounded above. Therefore any set of consecutive nonparabolic saddle connections does not prevent slimness of the collapsed preferred paths. □

6.4 General triangles

It remains to show that any triangle of collapsed preferred paths is uniformly slim. Equipped with this new version of the fan lemma for finitely generated Veech groups, the subsequent arguments of [8] continue to hold for proving that general triangles of preferred paths are slim. As in Section 4, their results are outlined here and cited wherever their arguments apply essentially verbatim.

Theorem 6.28 (collapsed preferred paths form slim triangles) *There exists $\delta > 0$ so that collapsed preferred paths form δ -slim triangles. That is, for any $x, y, z \in \Sigma$,*

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

Proof of Theorem 6.28 Let δ' be the constant from the fan lemma (Lemma 6.22). Any triangle in $\Delta(x, y, z)$ which has at least one side consisting of exactly one saddle connection can be decomposed into a union of fans; see Figure 9.

Slimness of each fan and the “furthermore” statement of the fan lemma (Lemma 6.22) ensure that the associated triangle of collapsed preferred paths $\Delta^{\hat{\zeta}}(x, y, z)$ is $\delta'' = (2\delta' + 2)$ -slim [8, Lemma 4.16]. General triangles $\Delta(x, y, z)$ may be similarly decomposed into fans, and after a careful treatment of all cases [8, Theorem 4.2] concludes that general triangles of collapsed preferred paths $\Delta^{\hat{\zeta}}(x, y, z)$ must be $\delta = 3\delta'' = 3(2\delta' + 2)$ -slim. □

Coincidentally, this gives an alternate proof for special case of a theorem by [11] (generalized by [15]) which states that a virtually free subgroup of the mapping class group is convex cocompact if and only if its extension group is hyperbolic.

Theorem 6.29 (special case of [11; 15]) *Let $G < \text{MCG}(S)$ be a finitely generated Veech group with extension group Γ . If G has no parabolic elements, then Γ is Gromov hyperbolic.*

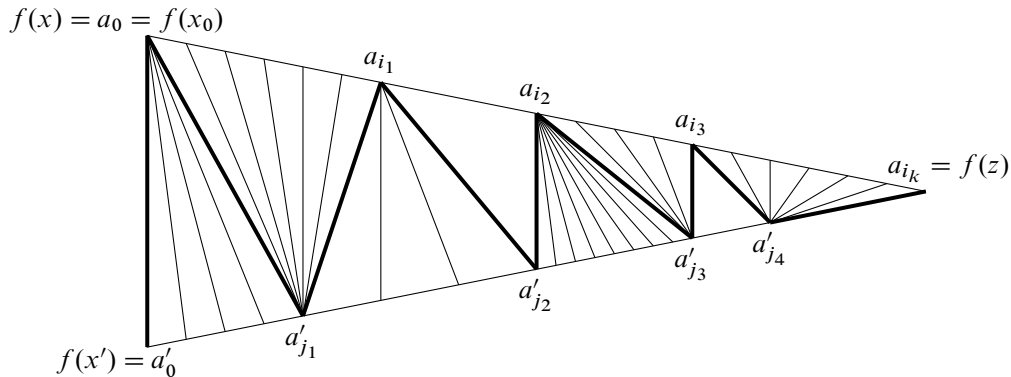


Figure 9: Illustration from [8, Figure 6] of a triangle with exactly one side being a single saddle connection, showing how it can be decomposed into fans. The labeled points are the vertices of the fans.

In the sense of [11], G is convex cocompact as a discrete subgroup of $\text{Isom}(D)$ if it acts cocompactly on the convex hull of its limit set—that is, if the quotient of $\text{hull}(G)$ by the action of G is compact. Therefore if G has no parabolic elements, it is convex cocompact as a Fuchsian group. A consequence is that the G -orbit of $X \in \text{hull}(G)$ is quasiconvex.

Proof Let \bar{E} be the space constructed in Section 4. Since there are no horoball preimages to collapse, $\bar{E} = \widehat{E}$ and Γ acts isometrically and cocompactly on \bar{E} by construction. It only remains to show that \bar{E} (equipped with the metric \bar{d}) is hyperbolic, which is accomplished by applying the guessing geodesics criterion (Proposition 5.1). Here, the sets $L(x, y)$ in the statement of the guessing geodesics criterion are precisely the preferred paths $\zeta(x, y) = \hat{\zeta}(x, y)$.

It was shown in the construction that (\bar{E}, \bar{d}) is a length space. Because the collection Σ of all cone points is Γ -invariant and \bar{E}/Γ is compact, there exists some constant $R > 0$ so that Σ is R -dense in \bar{E} . For $x, y \in \Sigma$ there is a preferred path $\zeta(x, y) \subset \bar{E}$, and the preferred paths form slim triangles (Theorem 6.28), satisfying condition (1) of the guessing geodesics criterion.

To verify condition (2) of the guessing geodesics criterion, suppose that $\bar{d}(x, y) \leq 3R$ for some $x \in E_X$ and $y \in E_Y$. Then $\rho(X, Y) \leq 3R$ in D . Because $f_{X,Y} = f_X|_{E_Y}$ is $e^{\rho(X,Y)}$ -bilipschitz (Section 4.2.1), the length of $f_X(\zeta(x, y))$ is bounded by $3Re^{3R}$. The geodesic in E_X joining x to $f_X(y)$ also has bounded length and in particular is a concatenation of saddle connections each with length bounded above (by $3Re^{3R}$) and below (uniformly, due to Proposition 6.3). So there are at most n saddle connections in $\zeta(x, y)$, where n depends only on R . Each of these saddle connections σ has length bounded above in the fiber over X and bounded below in the fiber at its associated horopoint $X_{[\sigma]}$, so X and $X_{[\sigma]}$ are uniformly close for all saddle connections σ appearing in $\zeta(x, y)$. Finally, since $\zeta(x, y)$ consists of at most $2n + 1$ pieces (hyperbolic geodesics in the horizontal fibers and saddle connections in the vertical fibers), where n depends only on R , and each of those pieces has uniformly bounded length, it follows that $L(x, y) = \zeta(x, y)$ has bounded length—and therefore bounded diameter, as required. \square

7 Hierarchical hyperbolicity

After establishing a nice action of Γ on a hyperbolic space \widehat{E} , a natural next step is to prove that Γ is hierarchically hyperbolic.

Theorem 7.1 (hierarchical hyperbolicity) *Let $G < \text{MCG}(S)$ be a finitely generated Veech group with extension group Γ . Then Γ is a hierarchically hyperbolic group.*

Nearly all of [9] continues to hold essentially verbatim when G is finitely generated but not necessarily a lattice: Constructing a hierarchically hyperbolic space structure once \widehat{E} is shown to be hyperbolic is almost solely concerned with what happens with the Bass–Serre trees T_α for parabolic directions α for G , so the existence of nonparabolic directions does not interfere with the construction of the HHS. For this same reason, Γ is already a hierarchically hyperbolic group whenever G has no parabolic limit points; indeed, Γ is hyperbolic by [11; 15] (or by Theorem 6.29). This section is deliberately kept brief; the reader is referred to [9] for more comprehensive coverage of the HHS structure and background information.

Let \mathcal{X} be the flag simplicial complex with 1-skeleton $\mathcal{X}^{(1)}$ as given in [9, Section 4.2]. Let \mathcal{W} be the \mathcal{X} -graph whose vertices are the maximal simplices of \mathcal{X} and whose edges are as given in [9, Section 4.2]. Let $\mathcal{X}^{+\mathcal{W}}$ be the \mathcal{W} -augmented dual graph as given in [9, Definition 4.2]. The following statement is comparable to [9, Lemma 4.17]. While this proof is nearly identical to the original, it includes consideration for when nontrivial saddle connections might appear in the case that G is not a lattice.

Lemma 7.2 (empty simplex) *The graph $\mathcal{X}^{+\mathcal{W}}$ is quasi-isometric to \widehat{E} .*

Proof Starting with the map $Z : \mathcal{X}^{(1)} \rightarrow \sqcup T_\alpha$ defined in [9, Section 4.2], extend this to a map $Z' : \mathcal{X}^{+\mathcal{W}} \rightarrow \widehat{E}$ as described in the proof of [9, Lemma 4.17]. As in that proof, Z is a one-sided inverse of Z' . It remains to show that Z is coarsely Lipschitz by demonstrating that for any edge $e = [x, y]$ of $\mathcal{X}^{+\mathcal{W}}$ with $v = Z(x)$ and $w = Z(y)$ in \mathcal{V} , there exists a path of uniformly bounded length in $\mathcal{X}^{+\mathcal{W}}$ joining v and w .

Now by Lemma 5.4, any $v \in T_\alpha$, $w \in T_\beta$ are joined by a combinatorial path of length proportional to $\hat{d}(v, w)$ which, as the proof shows, is in fact a concatenation of combinatorial subpaths of uniformly bounded length connecting points of \mathcal{V} . Moreover, by the lemma, the number of horizontal jumps and saddle connections is bounded by $C'\hat{d}(v, w)$. Therefore it suffices to prove the lemma in the cases that v and w are joined by a bounded length combinatorial path which is either

- (a) a single horizontal jump between two horoballs, or
- (b) two horizontal jumps from horoballs to a common component of the hull boundary concatenated with nonparabolic saddle connections (which are joined by horizontal jumps in lifts of the hull boundary).

Regardless of the presence of nonparabolic saddle connections, the proof may proceed as in [9] because this combinatorial path is the image of a path whose interior is in \overline{E} and connects the “thickened spine neighborhoods” (see Figure 2(e)) associated to v and w . \square

Proof of Theorem 7.1 The proof of [9, Theorem 4.16] holds verbatim after substituting Lemma 7.2 for [9, Lemma 4.17]. Therefore the pair $(\mathcal{X}, \mathcal{W})$ is a combinatorial HHS by the definition found in [9, Definition 4.8]. The rest follows from [9, Theorem 4.11]. \square

List of symbols

(S, X, q)	flat surface determined by a closed surface S of genus at least 2, a choice of complex structure X , and its associated quadratic differential q	page 995
$[\sigma]$	direction associated to the saddle connection σ	page 1000
$\overleftrightarrow{[\sigma_x], [\sigma_y]}$	hyperbolic geodesic segment between $[\sigma_x]$ and $[\sigma_y]$	page 1013
$[x, y]$	geodesic segment between x and y	page 992
α	an element of \mathcal{D} (without specific reference to an associated saddle connection)	page 1000
B_α°	interior of the horoball preimage B_α	page 1003
B_α	preimage $\pi^{-1}(B_\alpha)$	page 1003
B_α°	interior of B_α	page 1003
B_α	if α is parabolic, closed horoball in D that is invariant by the maximal parabolic subgroup of G corresponding to α ; if α is nonparabolic, orthogonal projection of the point $\alpha \in \partial D$ onto hull (G)	page 1001
\bar{b}	adjusted balance point of a Euclidean triangle T (or of its associated ideal triangle \tilde{T})—that is, the projection of the balance point b to hull (G)	page 1012
b	balance point of a Euclidean triangle T (or of its associated ideal triangle \tilde{T})	page 1012
\bar{D}	truncated convex hull, obtained from hull (G) by deleting the interiors of the horoballs B_α	page 1001
\bar{d}	metric on \bar{E}	page 1003
\hat{D}	collapsed convex hull, quotient space obtained from \bar{D} by collapsing each horoball B_α to a point	page 1001
\hat{d}	metric on \hat{E}	page 1004
\mathcal{D}	set of all directions of all saddle connections on (S, X_0, q_0) , considered as a subset of ∂D	page 1000
d	metric on E	page 1003
D	Teichmüller disk (determined by a flat structure (X, q) , although this is often dropped from the notation)	page 996
$\Delta^\mathcal{S}(x, y, z)$	triangle of preferred paths in E with vertices $x, y, z \in \Sigma$	page 1011
$\hat{\Delta}^\mathcal{S}(x, y, z)$	triangle of collapsed preferred paths in \hat{E} with vertices $x, y, z \in \Sigma$	page 1011

$\Delta(x, y, z)$	triangle of geodesics in E_0 with vertices $f(x), f(y), f(z)$ for $x, y, z \in \Sigma$	page 1011
D_x	preimage $f_X^{-1}(x)$, i.e., the unique lift of $\text{hull}(G)$ to the Teichmüller disk in $\mathcal{T}(S)$ through x that covers $\text{hull}(G)$ via the projection π	page 1003
\widehat{E}	space obtained from \overline{E} by collapsing horoballs to Bass–Serre trees via the map P	page 1004
E	surface bundle over $\text{hull}(G)$ given by π ; also admits a product structure $E \cong \text{hull}(G) \times \widetilde{S}$	page 1001
E_0	fiber of E over X_0	page 1002
E_X	fiber of E over X , i.e., $\pi^{-1}(X)$	page 1002
$f : E \rightarrow E_0$	alternative notation for f_{X_0}	page 1003
$f_X : E \rightarrow E_X$	map determined by varying $f_{X,Y}$ over all $Y \in D$	page 1003
$f_{X,Y} : E_Y \rightarrow E_X$	lift of the Teichmüller map to the universal cover	page 1002
G	finitely generated Veech group of q_0	page 1000
γ_i	i -th saddle piece of $\zeta(x, y)$	page 1006
Γ	extension group of G	page 1001
h_i	i -th horizontal piece of $\zeta(x, y)$	page 1006
$\text{hull}(G)$	convex hull of G	page 998
$\text{hull}(G)$	convex hull of the limit set of G	page 1001
ℓ_α	length metric on T_α	page 1008
$\ell()$	length of a saddle connection σ	page 995
$L(u, v)$	union of $\hat{\zeta}(x, y)$ over all $x \in \theta^u \cap \Sigma$ and $y \in \theta^v \cap \Sigma$	page 1007
$N_\delta()$	δ -neighborhood	page 1006
\overline{P}	restriction of $P : E \rightarrow \widehat{E}$ to \overline{E}	page 1004
\bar{p}	map $\overline{D} \rightarrow \widehat{D}$ which collapses each B_α to a point	page 1001
p	map $\text{hull}(G) \rightarrow \widehat{D}$ which collapses the interior of each B_α to a point	page 1001
$P : E \rightarrow \widehat{E}$	quotient of E obtained by collapsing each B_α onto T_α via the restriction $t_\alpha _{B_\alpha}$	page 1004
$\mathcal{P}(\sigma)$	set of saddle connections that span a triangle with σ	page 1015
$\pi : E \rightarrow \text{hull}(G)$	pullback bundle of the Bers fibration via the inclusion $\text{hull}(G) \subset \mathcal{T}(S)$	page 1001
$\text{proj}_{\overline{D}}$	ρ -closest-point projection of D onto \overline{D}	page 1001
q	quadratic differential or flat metric for a complex structure in the Teichmüller space of a closed surface of genus at least 2	page 994
q_0	flat metric for the complex structure X_0	page 1000
q_X	flat metric associated to X , i.e., terminal flat metric from the Teichmüller mapping $(S, X_0, q_0) \rightarrow (S, X, q_X)$	page 1002
$\bar{\rho}$	induced path metric on the truncated convex hull \overline{D}	page 1001
$\hat{\rho}$	quotient pseudometric on \widehat{D}	page 1001

ρ	Poincaré metric on the Teichmüller disk D	page 1000
\tilde{S}	universal cover of S	page 1003
S	closed, connected, oriented surface of genus at least 2	page 1000
Σ	the set of all cone points in E , i.e., $\bigcup_{X \in \text{hull}(G)} \Sigma_X$	page 1003
$\hat{\zeta}(x, y)$	collapsed preferred path between $x, y \in \Sigma$	page 1007
$\zeta(x, y)$	preferred path between $x, y \in \Sigma$	page 1006
Σ_X	for any $X \in \text{hull}(G)$, the set of cone points of the flat structure q_X on E_X	page 1003
T_α	weighted Bass–Serre tree for a parabolic direction α , i.e., the \mathbb{R} -tree dual to the foliation of E_{X_α} in the direction α	page 1004
$t_\alpha : E \rightarrow T_\alpha$	composition of f_{X_α} followed by projection to T_α	page 1004
θ_X^v	v -spine in E_X	page 1007
θ^v	union of v -spines over all fibers in the horoball associated to v	page 1007
\tilde{T}_i	i -th triangle in an ideal fan, often the ideal triangle associated to the Euclidean triangle T_i	page 1016
T_i	i -th triangle in a Euclidean fan	page 1016
\mathcal{V}	set of all vertices of all Bass–Serre trees T_α	page 1007
X_α	choice of fixed point in ∂B_α ; if α is nonparabolic, identical to B_α	page 1001
X_0	fixed choice of complex structure on S	page 1000

Acknowledgements

I would like to thank my PhD advisor, Chris Leininger, for being extraordinarily generous with his time and attention as I prepared this paper. Chris provided valuable insights into existing research and ensured that the final manuscript is as sensible as possible, though neither of us could come up with a less silly term for the “horopoints” described in Definition 4.1 and appearing throughout the paper.

I would like to thank the reviewer for thoughtful comments on the preprint.

Finally, I would like to thank my friends and colleagues for their support (in every sense) throughout the final months of preparing this paper, during which I was without a home after my apartment caught fire. This paper could not have been completed in a timely manner — and, more generally, the last several months of my life would have been immeasurably more difficult — if not for their generosity. I hope to be able to pay it forward someday.

I am financially supported in part by NSF-1842494 and NSF DMS-1745670.

References

- [1] **A F Beardon**, *The geometry of discrete groups*, Graduate Texts in Mathematics 91, Springer (1983) MR
- [2] **L Bers**, *Fiber spaces over Teichmüller spaces*, Acta Math. 130 (1973) 89–126 MR
- [3] **M Bestvina**, *Real trees in topology, geometry, and group theory*, preprint (1997) arXiv math/9712210
- [4] **J S Birman**, *Braids, links, and mapping class groups*, Annals of Mathematics Studies 82, Princeton Univ. Press (1974) MR

- [5] **B H Bowditch**, *Uniform hyperbolicity of the curve graphs*, Pacific J. Math. 269:2 (2014) 269–280 MR
- [6] **M R Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle Math. Wissen. 319, Springer (1999) MR
- [7] **R Charney, J Crisp**, *Relative hyperbolicity and Artin groups*, Geom. Dedicata 129 (2007) 1–13 MR
- [8] **S Dowdall, M G Durham, C J Leininger, A Sisto**, *Extensions of Veech groups, I: A hyperbolic action*, J. Topol. 16:2 (2023) 757–805 MR
- [9] **S Dowdall, M G Durham, C J Leininger, A Sisto**, *Extensions of Veech groups, II: Hierarchical hyperbolicity and quasi-isometric rigidity*, Comment. Math. Helv. 99:1 (2024) 149–228 MR
- [10] **B Farb, D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton Univ. Press (2012) MR
- [11] **B Farb, L Mosher**, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. 6 (2002) 91–152 MR
- [12] **F P Gardiner**, *Teichmüller theory and quadratic differentials*, Wiley, New York (1987) MR
- [13] **F P Gardiner, N Lakic**, *Quasiconformal Teichmüller theory*, Mathematical Surveys and Monographs 76, Amer. Math. Soc., Providence, RI (2000)
- [14] **D Groves, J F Manning**, *Dehn filling in relatively hyperbolic groups*, Israel J. Math. 168 (2008) 317–429 MR
- [15] **U Hamenstaedt**, *Word hyperbolic extensions of surface groups*, preprint (2005) arXiv math/0505244
- [16] **P Hubert, T A Schmidt**, *An introduction to Veech surfaces*, from “Handbook of Dynamical Systems” (B Hasselblatt, A Katok, editors), volume 1, Elsevier (2006) 501–526
- [17] **S Katok**, *Fuchsian groups*, University of Chicago Press (1992) MR
- [18] **C J Leininger, S Schleimer**, *Hyperbolic spaces in Teichmüller spaces*, J. Eur. Math. Soc. 16:12 (2014) 2669–2692 MR
- [19] **A Margolis**, *The geometry of groups containing almost normal subgroups*, Geom. Topol. 25:5 (2021) 2405–2468 MR
- [20] **H Masur**, *Geometry of Teichmüller space with the Teichmüller metric*, from “Geometry of Riemann surfaces and their moduli spaces” (L Ji, S A Wolpert, S-T Yau, editors), Surv. Differ. Geom. 14, International, Somerville, MA (2009) 295–313 MR
- [21] **H Masur, S Schleimer**, *The geometry of the disk complex*, J. Amer. Math. Soc. 26:1 (2013) 1–62 MR
- [22] **H Masur, S Tabachnikov**, *Rational billiards and flat structures*, from “Handbook of Dynamical Systems” (B Hasselblatt, A Katok, editors), volume 1, Elsevier (2002) 1015–1089
- [23] **L Mosher**, *Problems in the geometry of surface group extensions*, from “Problems on mapping class groups and related topics” (B Farb, editor), Proc. Sympos. Pure Math. 74, Amer. Math. Soc., Providence, RI (2006) 245–256 MR
- [24] **D-M Nguyen**, *Topological Veech dichotomy and tessellations of the hyperbolic plane*, Israel J. Math. 249:2 (2022) 577–616 MR
- [25] **J G Ratcliffe**, *Foundations of hyperbolic manifolds*, 2nd edition, Graduate Texts in Mathematics 149, Springer (2006) MR
- [26] **J Russell**, *Extensions of multicurve stabilizers are hierarchically hyperbolic*, Geom. Topol. 29:6 (2025) 3187–3240
- [27] **K Strebel**, *Quadratic differentials*, Ergebnisse der Math. 5, Springer (1984) MR
- [28] **R Tang**, *Affine diffeomorphism groups are undistorted*, J. Lond. Math. Soc. (2) 104:2 (2021) 747–769 MR
- [29] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes (1980) Available at <http://library.msri.org/books/gt3m/>
- [30] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19:2 (1988) 417–431 MR
- [31] **J Väisälä**, *Gromov hyperbolic spaces*, Expo. Math. 23:3 (2005) 187–231 MR
- [32] **W A Veech**, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Invent. Math. 97:3 (1989) 553–583 MR
- [33] **Y B Vorobets**, *Plane structures and billiards in rational polygons: the Veech alternative*, Uspekhi Mat. Nauk 51:5(311) (1996) 3–42 MR In Russian; translated in Russian Math. Surveys 51 (1996), 779–817

ELIOT BONGIOVANNI bongio@umich.edu

Department of Mathematics, University of Michigan, Ann Arbor, MI, United States

Received: July 9, 2024 Revised: January 15, 2025

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Vesna Stojanoska
vesna@illinois.edu
University of Illinois at Urbana-Champaign

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Daniel Isaksen	Wayne State University isaksen@math.wayne.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Markus Land	LMU München markus.land@math.lmu.de
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Octav Cornea	Université de Montreal cornea@dms.umontreal.ca	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futер	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Kathryn Hess	École Polytechnique Féd. de Lausanne kathryn.hess@epfl.ch		


See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2026 is US \$795/year for the electronic version, and \$1170/year (+\$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<https://msp.org/>

© 2026 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 26 Issue 3 (pages 825–1227) 2026

Standard position for surfaces in link complements in arbitrary 3-manifolds	825
JESSICA S. PURCELL and ANASTASIIA TSVIETKOVA	
Geometric rigidity of quasi-isometries in horospherical products	863
TOM FERRAGUT	
Large volume fibered knots in 3-manifolds	955
J ROBERT OAKLEY	
The primitive curve complex for a handlebody	973
SANGBUM CHO and JUNG HOON LEE	
Extensions of finitely generated Veech groups	989
ELIOT BONGIOVANNI	
Primitive Feynman diagrams and the rational Goussarov–Habiro Lie algebra of string links	1037
BRUNO DULAR	
Cusp-transitive 4-manifolds with every cusp section	1077
JACOPO GUOYI CHEN and EDOARDO RIZZI	
Finiteness conjecture for 3-manifolds obtained from handlebodies by attaching 2-handles	1095
HIROAKI KARUO and ZHIHAO WANG	
L -spaces, taut foliations and fibred hyperbolic two-bridge links	1115
DIEGO SANTORO	
New results on tilings via cup products and Chern characters on tiling spaces	1155
JIANLONG LIU, JONATHAN ROSENBERG and RODRIGO TREVIÑO	
Relative bounded cohomology on groups with contracting elements	1195
ZHENGUO HUANGFU and RENXING WAN	