University of Bristol - Explore Bristol Research
General rights
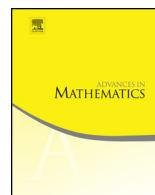
Hyperbolic cone metrics and billiards

Viveka Erlandsson [a,b,*], Christopher J. Leininger [c],
Chandrika Sadanand [d]

[a] *School of Mathematics, University of Bristol, Bristol, UK*
[b] *Department of Mathematics and Statistics, UiT The Arctic University of Norway, Norway*
[c] *Department of Mathematics, Rice University, Houston, TX, USA*
[d] *Department of Mathematics, Bowdoin College, Brunswick, ME, USA*

A B S T R A C T

A negatively curved hyperbolic cone metric is called *rigid* if it is determined (up to isotopy) by the support of its Liouville current, and *flexible* otherwise. We provide a complete characterization of rigidity and flexibility, prove that rigidity is a generic property, and parameterize the associated deformation space for any flexible metric. As an application, we parameterize the space of hyperbolic polygons with the same symbolic coding for their billiard dynamics, and prove that generically this parameter space is a point.

© 2022 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Let S be a closed, oriented surface of genus at least two and let $\varphi \in \mathrm{Hyp}_c(S)$ be the isotopy class of a negatively curved hyperbolic cone metric on S; see §2.1. Hersonsky and Paulin [16] proved that $\varphi \in \mathrm{Hyp}_c(S)$ is determined by its marked length spec-

trum. The proof follows Otal's approach [18], associating to φ a *Liouville* (or *Möbius*) *geodesic current*, L_φ which determines, and is determined by, the marked length spectrum, then proving that L_φ determines φ; see also [7,5,12]. The analogous result for unit area nonpositively curved Euclidean cone metrics—also called, *flat metrics*—was proved by the second author with Bankovic in [3] extending a special case with Duchin and Rafi [10], and subsequently extended to all nonpositively curved Riemannian cone metrics by Constantine [6].

In our earlier paper with Duchin [9] we proved a stronger theorem for flat metrics: they are typically determined by the *support* of their Liouville current; see §2.3. This result starkly contrasts the case of nonpositively curved *Riemannian* metrics for which the Liouville current always has full support, and thus no two such metrics can be distinguished by supports. In general, we proved that two flat metrics have the same support for their Liouville currents if and only if the metrics differ by an affine deformation (up to isotopy), and this is possible precisely when the holonomy of the metrics have order at most 2; see [9] for more details.

For $\varphi \in \mathrm{Hyp}_c(S)$, there is no notion of an "affine deformation" which preserves $\mathcal{G}_{\tilde\varphi}$, the support of the Liouville current L_φ. On the other hand, in Theorem 3.1 below, we show that if we are in the (highly non-generic) situation that (S, φ) admits a locally isometric branched covering of a hyperbolic orbifold $(S, \varphi) \to \mathcal{O}$ *and* every cone point maps to an orbifold point of even order, then any deformation of $\mathcal{O}$ to another hyperbolic orbifold $\mathcal{O}'$ lifts to a deformation φ' of φ without disturbing the support of the Liouville current; i.e. so that $\mathcal{G}_{\tilde\varphi} = \mathcal{G}_{\tilde\varphi'}$. Our main theorem states that this is the only way two metrics can have the same support for their currents.

Current Support Theorem. *Suppose $\varphi_1, \varphi_2 \in \mathrm{Hyp}_c(S)$ and $\mathcal{G}_{\tilde\varphi_1} = \mathcal{G}_{\tilde\varphi_2}$. Then $\varphi_1 = \varphi_2$ or else φ_1, φ_2 arise precisely as described above (that is, as in Theorem 3.1).*

We reiterate that for either a flat or hyperbolic cone metric, φ, the generic situation is that $\mathcal{G}_{\tilde\varphi}$ determines φ. In the non-generic case, the dimension of the deformation space for flat cone metrics preserving the support is precisely 2, while for hyperbolic cone metrics, the dimension is still finite but can be arbitrarily large: indeed, it is parameterized by the Teichmüller space of the quotient orbifold produced in the proof of the Current Support Theorem. See the remark at the end of §6.2.

An interesting special case is described by the following, which states that when a cone metric has "too many cone points", then it is necessarily rigid.

Corollary 1.1. *If $\varphi \in \mathrm{Hyp}_c(S)$ has at least $32(g-1)$ cone points (where g is the genus of S) then φ is rigid.*

See §6.3 for the proof. We do not claim that the bound $32(g-1)$ is sharp, but one cannot do better than a bound which is linear in g. Indeed, by taking unbranched covers of a genus two surface by a surface of genus $g \geq 3$ with the pull-back metric from the one

constructed in Example 3.2, we find flexible metrics on the surface of genus g with $g-1$ cone points.

1.1. Billiard rigidity

Suppose $P \subset \mathbb{H}$ is a compact n–gon in the hyperbolic plane (assumed simply connected, but not necessarily convex). We assume that P comes equipped with a labeling of the side by elements of $\mathcal{A} = \{1, 2, \ldots, n\}$, starting with 1 on some edge and proceeding in counterclockwise order. A biinfinite billiard trajectory determines a *bounce sequence* in $\mathcal{A}^{\mathbb{Z}}$ by recording the labels of the sides encountered in order, and the *bounce spectrum of* P is the set $\mathcal{B}(P) \subset \mathcal{A}^{\mathbb{Z}}$ consisting of all bounce sequences. We say that P is *billiard rigid* if $\mathcal{B}(P)$ determines P up to label-preserving isometry, and *billiard flexible* otherwise.

As an application of the Current Support Theorem, we prove a sharp symbolic-dynamical rigidity theorem characterizing rigidity and flexibility for hyperbolic polygons. Our theorem is analogous to, but with quite a different conclusion from, the "Bounce Theorem" of our paper with Duchin, [9] where we proved that a Euclidean polygon is billiard flexible if and only if all its interior angles are either $\frac{\pi}{2}$ or $\frac{3\pi}{2}$. For hyperbolic polygons, the statement is slightly more technical, and requires an additional definition.

A polygon P is *reflectively tiled* if it can be tiled by isometric copies of a single polygon, called the *tiles*, so that any two adjacent tiles differ by reflection in their shared edge. We further require that each interior angle of a tile is of the form $\frac{\pi}{k}$, for some $k \in \mathbb{Z}$, where k is even if the vertex of the tile is also a vertex of P. See §7 for a detailed discussion.

Billiard Rigidity Theorem. *Given hyperbolic polygons P_1, P_2, we have $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ if and only if*

1. P_1 *is isometric to* P_2 *by a label preserving isometry, or*
2. P_1, P_2 *are reflectively tiled and there exists a label-preserving homeomorphism* $H \colon P_1 \to P_2$ *that maps tiles to tiles, preserving their interior angles.*

The characterization of billiard flexibility is easily derived from the theorem; see §7.

Corollary 1.2. *A hyperbolic polygon P is billiard flexible if and only if it is reflectively tiled with a non-triangular tile.*

Because the statement of Billiard Rigidity Theorem is slightly technical, we note a couple of special cases that help to illustrate the conclusion. First, rigidity is generic; for example, Corollary 7.7 says that any polygon with at least one interior angle that is not in $\pi\mathbb{Q}$ is rigid. Moreover, Proposition 7.8 shows that an n–gon for which no interior angle is an even submultiple of π (i.e. of the form $\frac{\pi}{2k}$ for some $k \in \mathbb{Z}$), is rigid, while Corollary 7.9 shows that if all angles are even submultiples of π, then there is an $(n-3)$–dimensional space of n–gons with the same bounce spectrum; namely those having the same corresponding interior angles. See §7.4

For any polygon P, the bounce spectrum $\mathcal{B}(P)$ consists of uncountably many biinfinite sequences and it is natural to wonder if one can draw the same conclusions with less information. The answer is yes (as in the Euclidean case) and one needs only consider the bounce sequences of *generalized diagonals*, which are *compact* billiard trajectories in P that start and end at a vertex of P. The bounce sequence of a generalized diagonal is a finite sequence and the set of all such, $\mathcal{B}_\Delta(P)$, is countable. This is enough information to determine $\mathcal{B}(P)$ and thus we have the following.

Theorem 1.3. *Given polygons $P_1, P_2 \subset \mathbb{H}$, we have $\mathcal{B}_\Delta(P_1) = \mathcal{B}_\Delta(P_2)$ if and only if one of the two conclusions of the Billiard Rigidity Theorem holds.*

The bounce spectrum of a hyperbolic polygon was previously studied by Ullmo and Giannoni [23] in the special case that all interior angles are acute (see also [13] and Nagar–Singh [17]). The primary objective in [23] was to find a set of "grammar rules" that completely describe $\mathcal{B}(P)$ for a given polygon. There are some interesting connections with the Billiard Rigidity Theorem that we explain in §7.6.

1.2. Outline

Here we explain the key ideas in the proofs of the two theorems above, focusing primarily on the Current Support Theorem. First, given a metric $\varphi \in \mathrm{Hyp}_c(S)$ consider the pull back metric $\tilde{\varphi}$ on the universal cover $\tilde{S}$ of S. The set $\mathcal{G}_{\tilde{\varphi}}$ can be alternatively described as the set of pairs of endpoints at infinity of the *basic $\tilde{\varphi}$–geodesics*, that is those which are limits of *nonsingular* (disjoint from cone points) $\tilde{\varphi}$–geodesics; see §2.3.

Now, fix $\varphi_1, \varphi_2 \in \mathrm{Hyp}_c(S)$ and suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. Having the same set of endpoints translates into a correspondence between basic $\tilde{\varphi}_1$–geodesics and basic $\tilde{\varphi}_2$-geodesics since they must be bounded Hausdorff distance apart. The proof of the Current Support Theorem begins by successively adjusting (a representative of) φ_2 by an isotopy until it and φ_1 share some useful common features. This starts by first showing that the cone points are naturally "aligned" by the condition $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ using *chains* at infinity: after adjusting φ_2 by an isotopy, the two metrics have the same set of cone points; see §4.1. Analyzing the way the basic geodesics partition the cone points in the universal covering, we explain how to adjust φ_2 and construct a triangulation with vertices at the cone points so that in *both metrics* the triangles are isometric to triangles in $\mathbb{H}$; see §4.3.

Up to this point the outline of the proof is similar to that of the analogous theorem for Euclidean cone metrics in [9], though there are various technical differences we must account for throughout. In contrast, the remainder of the proof substantially diverges from the Euclidean case. To prove the theorem we must **either** show that the interior angles of the triangles of the triangulation are the same in the two metrics, ensuring the triangles are isometric and thus $\varphi_1 = \varphi_2$, **or** produce a locally isometric branched cover of an orbifold $(S, \varphi_1) \to \mathcal{O}_1$ (with cone points mapping to even order orbifold points) and prove that φ_2 is obtained by lifting a deformation of $\mathcal{O}_1$.

This analysis starts by removing the common set of cone points to produce a punctured surface $\dot{S}$, and taking (the metric completion of) the universal cover $\hat{S}$ of $\dot{S}$. Using the triangulation, we show that the basic geodesics of the pulled back metrics $\hat{\varphi}_1$ and $\hat{\varphi}_2$ also have the same endpoints at infinity. Associated to $\hat{\varphi}_i$, for each $i = 1, 2$, there is a locally isometric developing map $D_i \colon (\hat{S}, \hat{\varphi}_i) \to \mathbb{H}$ which is equivariant with respect to its associated holonomy homomorphism $\rho_i \colon \pi_1(\dot{S}) \to \mathrm{PSL}_2(\mathbb{R})$. A delicate argument analyzing intersection patterns of basic geodesics is used to construct an orientation preserving homomorphism $h \colon \partial\mathbb{H} \to \partial\mathbb{H}$ that conjugates ρ_1 to ρ_2 within the homeomorphism group of the circle $\partial\mathbb{H}$. Furthermore, h sends the endpoints of any D_1–image of a basic $\hat{\varphi}_1$–geodesic to the endpoints of the D_2–image of the associated basic $\hat{\varphi}_2$–geodesic; see §5.2. It is by investigating the holonomy homomorphism image $\rho_1(\pi_1(\dot{S}))$ (and a slightly enlarged group) and using the existence of the conjugating homeomorphism that we can determine whether the two metrics are the same or differ by an orbifold deformation.

Indeed, if $\rho_1(\pi_1(\dot{S}))$ is indiscrete we show that h must in fact lie in $\mathrm{PSL}_2(\mathbb{R})$ by an analysis of the closed, cocompact subgroups of $\mathrm{PSL}_2(\mathbb{R})$; see §2.5. This implies that the angle between a pair of D_1–developed basic $\hat{\varphi}_1$–geodesics and the corresponding D_2–developed basic $\hat{\varphi}_2$–geodesics must agree. Applying this to extensions of the sides of the triangles implies that the interior angles for each triangle in the triangulation are the same in the two metrics, and hence we have that $\varphi_1 = \varphi_2$.

One can push this argument further using an observation from earlier in the proof: during the first adjustment stage, it is shown that all basic $\tilde{\varphi}_1$–geodesics through a cone point in $(\widetilde{S}, \tilde{\varphi}_1)$ correspond to $\tilde{\varphi}_2$–geodesics that pass through a cone point in $(\widetilde{S}, \tilde{\varphi}_2)$ (after an adjustment, these become the same point). Passing to $\hat{S}$, we deduce that for every cone point ζ of $\hat{S}$, h "witnesses" the concurrence of the D_1– and D_2–images of the respective basic geodesics through ζ. This manifests itself in the fact that h conjugates the order two elliptic isometry fixing $D_1(\zeta)$ to the order two elliptic isometry fixing $D_2(\zeta)$. Taking the groups $\Gamma_i^0 < \mathrm{PSL}_2(\mathbb{R})$ generated by $\rho_i(\pi_1(\dot{S}))$ together with these order two elliptic elements about $D_i(\zeta)$, over all cone points ζ of $\hat{S}$, we see that h actually conjugates Γ_1^0 to Γ_2^0; see §6.1. Consequently, if Γ_1^0 is indiscrete we again deduce, by the same argument as above, that h must lie in $\mathrm{PSL}_2(\mathbb{R})$ and so $\varphi_1 = \varphi_2$ as before.

The only option left is that Γ_1^0 and Γ_2^0 are *discrete*. Then the quotient $\mathcal{O}_i^0 = \mathbb{H}/\Gamma_i^0$ is hyperbolic orbifolds, for $i = 1, 2$. By construction each D_i is equivariant and hence descend to branched covers from $(S, \varphi_i) \to \mathcal{O}_i^0$ under which the cone points map to even order orbifold points. The homeomorphism h conjugates Γ_1^0 to Γ_2^0 and hence determines homeomorphisms $\mathcal{O}_1^0 \to \mathcal{O}_2^0$. Chasing through the diagram shows that (after another adjustment) φ_2 differs from φ_1 by lifting this homeomorphism, completing the proof; see §6.2.

The Billiard Rigidity Theorem follows from the Current Support Theorem by an unfolding procedure. Roughly speaking, given a hyperbolic n–gon P, we consider a negatively curved hyperbolic cone surfaces tiled by copies of P. Then "folding the tiles up" defines a map from the surface to P that sends each tile isometrically to P; consequently, we call such a surface and metric an unfolding. A key property of an unfolding is that

nonsingular geodesics in the surface project to billiard trajectories, and vice versa billiard trajectories lift to nonsingular geodesics.

Given two hyperbolic n–gons P_1 and P_2, we find a *common unfolding*, which is a surface with two different hyperbolic cone metrics φ_1 and φ_2 so that the tilings define the same cell structure (the 2-cells are the tiles) and so the labelings from the polygons agree. When $\mathcal{B}(P_1) = \mathcal{B}(P_2)$, we lift a common unfolding to the universal cover of the surface, and observe that the bounce sequence of a billiard trajectory also specifies (from a given starting tile) a symbolic coding of a nonsingular geodesic. In this way, $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ naturally implies $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ since a $\tilde{\varphi}_1$–geodesic and $\tilde{\varphi}_2$–geodesic passing through the same set of tiles forces them to remain a bounded distance apart; see §7.2.

This first part of the proof is similar to the Euclidean billiard case, but again, at this point the proofs diverge. Since $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, we find a homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ and show that it conjugates the group generated by reflections in the faces of P_1 to the group generated by reflections in the faces of P_2. If $h \in \mathrm{PSL}_2(\mathbb{R})$, by looking at the extensions of the geodesic sides to biinfinite geodesics and the effect of h on the endpoints, we see that the extension of h to an isometry of $\mathbb{H}$ maps P_1 to P_2.

Now, to determine whether h belongs to $\mathrm{PSL}_2(\mathbb{R})$ we proceed similarly to the argument above. If the group $R_{P_1}^0$ generated by reflections in the sides of P_1 together with the order two elliptic elements fixing the vertices is indiscrete, then h is necessarily in $\mathrm{PSL}_2(\mathbb{R})$, and we are in case (1) of the theorem; see §7.3. If on the other hand $R_{P_1}^0$ is discrete, one can check that this group is also generated by reflections in a (typically) smaller polygon inside of P_1. The union of the lines of reflection for this group defines a tiling of $\mathbb{H}$ for which P_1 is a union of tiles. The map h sends the endpoints of all these lines of reflection to the endpoints of the lines of reflection for the associated group determined by P_2. This is enough to construct a homeomorphism $\mathbb{H} \to \mathbb{H}$ sending tiles of one tiling to tiles of the other and sending P_1 to P_2. The interior angles of the tiles are preserved by this homeomorphism since h conjugates the dihedral vertex stabilizer in one group to an isomorphic group in the other. This essentially puts us in case (2) of the theorem, and completes the proof; see §7.4.

The paper is organized as follows. In §2 we introduce the notation and terminology used throughout the paper, describe the geodesics in $\widetilde{S}$ and $\hat{S}$, and discuss topological conjugates of subgroups of $\mathrm{PSL}_2(\mathbb{R})$. §3 is devoted to proving Theorem 3.1 mentioned above, showing that two hyperbolic cone metrics φ_1, φ_2 which branch cover hyperbolic orbifolds (with cone points projecting to even order orbifold points) and differ by a lift of an orbifold deformation satisfy $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. The normalization procedures of the metric φ_2 alluded to in the outline above are explained in §4. In §5 we construct the homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ which conjugates the holonomies and is a key tool in the proof of the Current Support Theorem, which we prove in §6. Finally, in §7 we discuss hyperbolic billiards and prove the Billiard Rigidity Theorem.

In particular, the argument involving triangle groups in Lemma 7.6 is due to Alan Reid. The authors would like to thank the referee for helpful comments on the initial version of the paper. Leininger was partially supported by NSF grant DMS-1811518 and DMS-2106419. Erlandsson was partially supported by EPSRC grant EP/T015926/1.

2. Preliminaries

2.1. Hyperbolic cone metrics

We write $\widetilde{\mathrm{Hyp}}_c(S)$ for the space of negatively curved hyperbolic cone metrics on S. These are metrics, locally isometric to $\mathbb{H}$ away from a finite, positive number of cone singularities with cone angles greater than 2π. We refer to the non-cone points as *regular points*. The group $\mathrm{Homeo}_0(S)$ of homeomorphisms isotopic to the identity acts (on the right) on $\widetilde{\mathrm{Hyp}}_c(S)$ by pullback, and we let $\mathrm{Hyp}_c(S)$ be the quotient by this action. That is, given $\varphi, \varphi' \in \widetilde{\mathrm{Hyp}}_c(S)$, declaring $\varphi \sim \varphi'$ if there exists a homeomorphism $f \colon S \to S$ isotopic to the identity, such that $f^*(\varphi') = \varphi$, then $\mathrm{Hyp}_c(S) = \widetilde{\mathrm{Hyp}}_c(S)/\sim$ is the quotient of $\widetilde{\mathrm{Hyp}}_c(S)$ by the equivalence relation $\sim$.

Remark. While we have assumed that the cone point set is nonempty, we note that the Current Support Theorem is trivially valid without this assumption. Indeed, for any nonsingular hyperbolic metric, its Liouville current has full support. Thus, any two nonsingular hyperbolic metrics have the same support for their Liouville currents, and do in fact differ by the lift of a deformation from the trivial orbifold covering of itself.

We fix a universal covering $p \colon \widetilde{S} \to S$. For any reference negatively curved metric, we let S^1_∞ be the boundary at infinity of $\widetilde{S}$ with respect to the pull back metric. This defines a compactification of $\widetilde{S}$ which is independent of the metric, up to homeomorphism that is the identity on $\widetilde{S}$. In particular, S^1_∞ is independent of the choice of negatively curved reference metric (indeed, S^1_∞ is the Gromov boundary, which can be defined by any geodesic metric on S pulled back to $\widetilde{S}$). For any $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$, the metric $\tilde{\varphi} = p^*(\varphi)$ is CAT(-1), by Gromov's link condition and the Cartan-Hadamard Theorem; see [2].

2.2. Puncturing surfaces

Suppose $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$, let $C_0 = \mathrm{cone}(\varphi)$ be the set of cone points for φ, and let $\dot{S} = S \setminus C_0$ denote the punctured surface obtained by deleting the cone points. Let $\hat{p} \colon \widetilde{\dot{S}} \to \dot{S}$ be the universal cover and $\hat{\varphi} = \hat{p}^*(\varphi)$ the pull back metric. Since $\hat{p}$ is a local isometry, it extends over the metric completions, and we denote the map by the same name:

$$\hat{p} \colon \hat{S} \to S$$

where $\hat{S}$ denotes the completion of $\widetilde{\dot{S}}$. To describe the local behavior of this map and the space $\hat{S}$ near the completion points, suppose that $\epsilon > 0$ is sufficiently small so that the

ϵ–neighborhood of C_0 is a disjoint union of topological disks, $\Delta_1 \cup \ldots \cup \Delta_k$, with each Δ_j containing exactly one point $\zeta_j \in C_0$. For any component $U \subset \hat{p}^{-1}(\Delta_j - \{\zeta_j\})$,

$$\hat{p}|_U : U \to \Delta_j - \{\zeta_j\}$$

is a universal cover of $\Delta_j - \{\zeta_j\}$. For such U, there is a single completion point $\hat{\zeta}_U$ that projects to ζ_j.

In fact, any non-convergent Cauchy sequence in $\widetilde{S}$ projects to such a sequence in $\dot{S}$, which is necessarily eventually contained in some Δ_j. By taking $\epsilon > 0$ sufficiently small, we can assume that the sequence is contained in a single component U, and hence converges to the point $\hat{\zeta}_U$. That is, the points $\hat{\zeta}_U$ account for all completion points. We write $\hat{C}_0 = \hat{p}^{-1}(C_0) \subset \hat{S}$ for the completion points, which (by an abuse of terminology) we also refer to as cone points.

Note that $\hat{p}$ factors through a map $\tilde{p} \colon \hat{S} \to \widetilde{S}$. Furthermore, there are covering actions of $\pi_1 \dot{S}$ and $\pi_1 S$ on $\hat{S}$ and $\widetilde{S}$ respectively, and a homomorphism $\pi_1 \dot{S} \to \pi_1 S$ (induced by inclusion $\dot{S} \to S$, up to conjugation) so that $\tilde{p}$ is equivariant:

$$
\begin{array}{ccc}
\pi_1 \dot{S} & \longrightarrow & \pi_1 S \\
\circlearrowleft & & \circlearrowleft \\
\hat{S} & \xrightarrow{\ \tilde{p}\ } & \widetilde{S} \\
& \searrow{\scriptstyle \hat{p}} & \downarrow{\scriptstyle p} \\
& & S.
\end{array}
$$

In what follows, when considering two metrics $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$, we will only consider the space $\hat{S}$ when φ_1 and φ_2 have common cone point set $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2)$. Under this assumption, we note that the discussion above implies that the metric completion of $\widetilde{S}$ is the same for either metric $\hat{\varphi}_1$ or $\hat{\varphi}_2$. Also, with respect to either metric $\hat{\varphi}_1$ or $\hat{\varphi}_2$, $\hat{S}$ is CAT(-1), and in particular is Gromov hyperbolic. The identity $(\hat{S}, \hat{\varphi}_1) \to (\hat{S}, \hat{\varphi}_2)$ is a quasi-isometry (in fact, after adjusting φ_2 by an arbitrarily small homeomorphism (fixing C_0), we may assume φ_1 and φ_2 are biLipschitz equivalent, and hence the same is true of $\hat{\varphi}_1$ and $\hat{\varphi}_2$). The boundary at infinity of $\hat{S}$, denoted $\hat{S}^1_\infty$, is thus well-defined and independent of the metric (in the same sense as for $\widetilde{S}$). Moreover, any two points of the boundary are connected by a biinfinite geodesic. Although $\hat{S}^1_\infty$ is not a circle (it is not even compact), it is homeomorphic to a subset of the circle in a $\pi_1 \dot{S}$–invariant way, since $\widetilde{S}$ is a surface.

2.3. Geodesics in $\widetilde{S}$ and $\hat{S}$

We suppose $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$ throughout and let $C_0 = \mathrm{cone}(\varphi)$. We refer to the sets $\widetilde{C}_0 = p^{-1}(C_0)$ and $\hat{C}_0 = \hat{p}^{-1}(C_0)$ as the cone point sets of $(\widetilde{S}, \tilde{\varphi})$ and $(\hat{S}, \hat{\varphi})$. By a $\tilde{\varphi}$–geodesic or $\hat{\varphi}$–geodesic we will always mean a biinfinite geodesic in $(\widetilde{S}, \tilde{\varphi})$ or $(\hat{S}, \hat{\varphi})$,

respectively. Moreover, we consider the geodesics as unparameterized, and we equip the space of geodesics with the quotient topology coming from the compact-open topology on the set of parametrized geodesics, where we forget the parametrization.

A $\tilde{\varphi}$–geodesic segment is a subsegment of a $\tilde{\varphi}$–geodesic and can be finite, half-infinite (i.e. a ray), or bi-infinite, and we make the analogous convention for $\hat{\varphi}$–geodesic segments. A $\tilde{\varphi}$–geodesic (resp. $\hat{\varphi}$–geodesic) segment is *singular* if its interior intersects $\widetilde{C}_0$ (resp. $\hat{C}_0$) and *nonsingular* otherwise. Note that nonsingular geodesic segments are hyperbolic geodesic segments. If a $\tilde{\varphi}$–geodesic segment η meets a cone point $\zeta \in \widetilde{C}_0$ in the interior of η, then any sufficiently small disk neighborhood of ζ is divided into two half-disks by η which we call the *sides of η near ζ* (nested disks determine nested sides), and we note that being geodesic, the angle on each side is at least π. If a cone point $\hat{\zeta} \in \hat{C}_0$ is encountered in the interior of a singular $\hat{\varphi}$–geodesic segment, then any sufficiently small "disk" neighborhood (i.e. an open set U as above) is divided into three sides; on only one side does it make sense to measure the angle, but again that angle must be at least π. Finally, a nonsingular $\tilde{\varphi}$- or $\hat{\varphi}$–geodesic segment between two cone points is called a *saddle connection*.

Observe that a $\tilde{\varphi}$–geodesic η has two *globally* defined sides since it divides $\widetilde{S}$ into two components. Orienting η, and appealing to the orientation on $\widetilde{S}$, we can refer to the two sides as the *positive* (left-hand) and *negative* (right-hand) sides of η. A $\hat{\varphi}$–geodesic can have infinitely many globally defined sides, but choosing an orientation they can still be prescribed a sign, positive or negative. Of course, orienting a geodesic segment also gives rise to a well-defined sign for the sides at cone points encountered (locally), as well.

The set of distinct pairs of points in S^1_∞ is denoted by

$$\mathcal{G}(\widetilde{S}) = S^1_\infty \times S^1_\infty \setminus \Delta / (x,y) \sim (y,x)$$

where Δ is the diagonal in $S^1_\infty \times S^1_\infty$. Given a $\tilde{\varphi}$–geodesic η, we let $\partial_{\tilde{\varphi}}(\eta) \in \mathcal{G}(\widetilde{S})$ be the set of endpoints on S^1_∞ of γ. Since $\tilde{\varphi}$ is CAT(-1), this determines a homeomorphism

$$\partial_{\tilde{\varphi}} : \{\tilde{\varphi}\text{–geodesics in } \widetilde{S}\} \to \mathcal{G}(\widetilde{S}).$$

Note that two $\tilde{\varphi}$–geodesics are either disjoint, meet at exactly one point where they transversely intersect, or coincide either in a single cone point or in a (unique maximal) geodesic segment. If they coincide in a segment, this is either a concatenation of saddle connections or an infinite ray emanating from a cone point (and the two geodesics share an endpoint at infinity). In the latter case we say the two geodesics are *cone point asymptotic*. We say that $\{x_1, y_1\}, \{x_2, y_2\} \in \mathcal{G}(\widetilde{S})$ *link* if $x_2, y_2 \notin \{x_1, y_1\}$ and x_2 and y_2 lie in different components of $S^1_\infty \setminus \{x_1, y_1\}$. We say that two $\tilde{\varphi}$–geodesics η_1 and η_2 *cross* if their endpoints $\partial_{\tilde{\varphi}}(\eta_1)$ and $\partial_{\tilde{\varphi}}(\eta_2)$ link (meaning, $S^1_\infty \setminus \partial_{\tilde{\varphi}}(\eta_1)$ contains one point of $\partial_{\tilde{\varphi}}(\eta_2)$ in each component). When η_1 and η_2 cross, they either intersect transversely at a point or they coincide either in a single cone point or in a concatenation of saddle connections, see Fig. 1. Moreover, if a pair of geodesics intersect transversely, then they
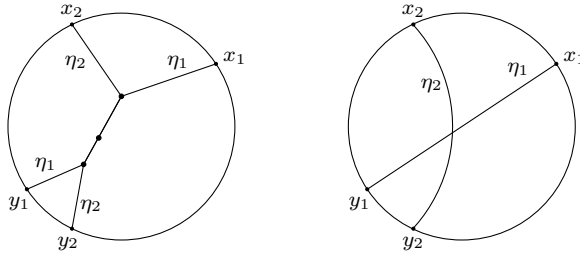
Fig. 1. Two examples of crossing geodesics η_1 and η_2 with their linking endpoints $\partial_{\tilde{\varphi}}(\eta_1) = \{x_1, y_1\}$ and $\partial_{\tilde{\varphi}}(\eta_2) = \{x_2, y_2\}$.

must cross. On the other hand, a pair of geodesics may coincide in a nontrivial segment or cone point without crossing.

We let $\mathcal{G}(\tilde{\varphi})$ denote the closure of the set of nonsingular $\tilde{\varphi}$–geodesics and call the geodesics in $\mathcal{G}(\tilde{\varphi})$ the *basic $\tilde{\varphi}$–geodesics*. We define

$$\mathcal{G}_{\tilde{\varphi}} = \partial_{\tilde{\varphi}}(\mathcal{G}(\tilde{\varphi})) \subset \mathcal{G}(\widetilde{S})$$

to be the set of all endpoints of basic geodesics. As mentioned in the introduction, the set $\mathcal{G}_{\tilde{\varphi}}$ is exactly the support of the Liouville current L_φ (see [3,6]). As a notational reminder, when some metric is involved, we will typically include the metric as a subscript to denote endpoints of geodesics while it will be included in parentheses to denote the actual geodesics.

Similarly, we let $\mathcal{G}(\hat{\varphi})$ be the closure of the set of nonsingular $\hat{\varphi}$–geodesics which we call *basic $\hat{\varphi}$–geodesics*, denoting the set of their endpoints by $\mathcal{G}_{\hat{\varphi}} \subset \mathcal{G}(\hat{S}) = \hat{S}^1_\infty \times \hat{S}^1_\infty \setminus \Delta / (x, y) \sim (y, x)$.

The following lemma describes a few elementary properties of basic geodesics in $\widetilde{S}$ and $\hat{S}$. For $\widetilde{S}$, the analogous statements in the case of Euclidean cone metrics were proved in [3, Section 3] and for non-positively curved Riemannian cone metrics in [6, Section 5].

Lemma 2.1 *(Basic geodesics). Suppose $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$. For any basic geodesic in $(\widetilde{S}, \tilde{\varphi})$ that meets a cone point, it does so making angle exactly π on one side at that point. If a basic $\tilde{\varphi}$–geodesic encounters more than one cone point, then after choosing an orientation and ordering the cone points it meets accordingly, the sides on which the angle is π can change sign at most once. Consequently, the set $\mathcal{G}^2_{\tilde{\varphi}}$ of basic geodesics meeting at least 2 cone points is countable. The same statements are true for basic geodesics in $(\hat{S}, \hat{\varphi})$: at each point $\zeta \in \hat{C}_0$ a basic $\hat{\varphi}$–geodesic encountering it makes angle exactly π, and if a basic $\hat{\varphi}$–geodesic encounters more than one point of $\hat{C}_0$, the signs of the sides on which it makes angle π switch at most once.*

Proof. We prove the statements for basic $\tilde{\varphi}$–geodesics, then describe any modifications necessary for the corresponding statement for basic $\hat{\varphi}$–geodesics.

Consider any $\eta \in \mathcal{G}(\tilde{\varphi})$ and suppose that ζ is a cone point on η. We first prove that the angle is π on one side of η at ζ. Let $\{\eta_i\}$ be a sequence of nonsingular geodesics converging to η, and observe that for each sufficiently large i, there is a small disk neighborhood of ζ which η_i intersects, and lies on one side or the other of ζ (since η_i does not contain ζ). Passing to a subsequence if necessary, we can assume all $\{\eta_i\}$ are on the same side of η in each of these neighborhoods. Since each η_i make angle π on each side of every point it passes through, it follows that η must make angle π at ζ on the side containing all η_i. For $\eta \in \mathcal{G}(\hat{\varphi})$, the argument is similar, noting that in this case (in sufficiently small neighborhoods) the approximating geodesics η_i must lie on the side of η making finite angle, which is then necessarily equal to π.

Next we assume $\eta \in \mathcal{G}(\tilde{\varphi})$, and show that the side one which it makes angle π can change at most once. To this end, assume η passes through at least three points of $\widetilde{C}_0$, say $\zeta_1, \zeta_2, \zeta_3$, and that these points appear in order along η, with respect to some orientation on η. We suppose η makes angle π on the positive side at ζ_1 and ζ_3 but at the negative side at ζ_2, and will arrive at a contradiction. As above, let $\{\eta_i\}$ be an approximating sequence of nonsingular geodesics. Note that for large i we must have that η_i is on the positive side of η in neighborhoods of ζ_1 and ζ_3 while on the negative side of η in a neighborhood of ζ_2. It follows that for large i the geodesics η_i and η contain segments that bound a bigon, which is impossible since the metric is CAT(-1). Since the metric $\hat{\varphi}$ on $\hat{S}$ is also CAT(-1), the exact same argument works in $\hat{S}$ as well.

Finally, we explain why these facts imply that the set of basic geodesics meeting at least 2 cone points is countable. Every $\tilde{\varphi}$–geodesic meeting at least two cone points contains a saddle connection. Since a saddle connection is determined by its endpoints in $\widetilde{C}_0$, and since $\widetilde{C}_0$ is countable, there are only countably many saddle connections. For each saddle connection, there are only countably many basic $\tilde{\varphi}$–geodesics containing it: one that always makes angle π on the positive side, one that always makes angle π on the negative side, and a countable number that switch. To see that this last set is countable, divide into two cases, depending on whether it switches from π on the positive side to π on the negative side, or vice-versa, then index the cone points encountered by any such geodesic with an interval in $\mathbb{Z}$ (with one of the endpoints of the saddle connections being 0) and record the index of the last cone point before the switch of sides occurs. Therefore, the set of basic $\tilde{\varphi}$–geodesics containing at least two cone points is a countable union of countable sets, hence countable. □

We say that a geodesic segment (in either $\widetilde{S}$ or $\hat{S}$) is a *basic segment* if it is either nonsingular or it meets each cone point in its interior at an angle exactly π on one side and this side switches at most once. This name is justified because we will show below that a segment is basic if and only if it can be extended to a basic geodesic. In fact, we show that the description of basic geodesics in Lemma 2.1 actually characterizes basic geodesics:

Lemma 2.2. *If η is a $\tilde{\varphi}$-geodesic (or $\hat{\varphi}$–geodesic) which makes angle π on one side at every cone point it meets, and this side switches at most once, then η belongs to $\mathcal{G}(\tilde{\varphi})$*
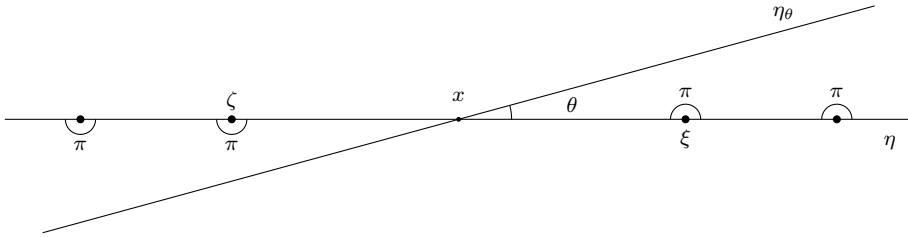
Fig. 2. Constructing an approximation of η by η_θ, as $\theta \to 0$.

(respectively, $\mathcal{G}(\hat{\varphi})$). Moreover, a geodesic segment is a basic segment if and only if it can be extended to a basic geodesic.

Proof. Let η be a $\tilde{\varphi}$–geodesic as in the statement (the case of $\hat{\varphi}$–geodesics follow by almost identical arguments). We will construct a sequence of nonsingular geodesics $\{\eta_i\}$ converging to η, showing that η is a basic geodesic.

First suppose that η makes angle π on the same side (say the positive side) at every cone point it encounters. Pick any regular point x on η and let δ be a short geodesic segment emanating from x on the positive side of η and meeting η orthogonally. By taking δ short enough we can assume it contains no cone points. For every point y on δ let η_y be a basic geodesic through y and orthogonal to δ. Note that $\{\eta_y\}_{y\in\delta}$ is a set of pairwise disjoint geodesics, since if any two intersected this would result in a geodesic triangle with angle sum at least π, which is impossible in a CAT(-1) metric. Hence for each $y \in \delta - \{x\}$, the geodesic η_y lies on the positive side of $\eta = \eta_x$. Moreover, only countably many y gives rise to a singular η_y (since there are only countably many points in $\widetilde{C}_0$) and hence picking any sequence of points $\{y_i\}$ on δ that avoids this countable set and converging to x we get a sequence of nonsingular geodesics $\eta_i = \eta_{y_i}$ limiting to η.

Now suppose η passes through at least two cone points and that the side at which it makes angle π at the cone points switches once. Assume η makes angle π on, say, the negative side at cone point ζ and on the positive side at cone point ξ and that hat η encounters ζ immediately before ξ. Let $[\zeta, \xi]$ denote the segment of η between the two points and let x be any regular point on $[\xi, \zeta]$. For each $\theta \in (0, \pi)$ let η_θ be a geodesic that transversely intersects η at x making an angle as shown in Fig. 2. Then η_θ is nonsingular for all but countably many directions θ, and hence picking any sequence $\{\theta_i\}$, avoiding this countable set, such that $\theta_i \to 0$ and setting $\eta_i = \eta_{\theta_i}$ results in a sequence of nonsingular geodesics $\{\eta_i\}$ converging to η.

For the second claim, let σ be a basic geodesic segment. First suppose its interior makes angle π on the positive side at every cone point it meets. Extend σ to a geodesic by making angle exactly π on the positive side at every cone point it meets. By the above, the resulting geodesic is a basic geodesic. Now, if σ switches once from making angle π on the positive side to making angle π on the negative side, extend σ forward by a geodesic ray making angle π on the negative side, and backward by a geodesic ray

making angle π on the positive side. Again by the above, the result is a basic geodesic. The converse follows from Lemma 2.1 and the definition of a basic geodesic segment. □

We note that, since in a neighborhood of any regular point the map $\tilde{p} : \hat{S} \to \widetilde{S}$ is an isometry, and since cone points are sent to cone points, any nonsingular geodesic in $\hat{S}$ projects to a nonsingular geodesic in $\widetilde{S}$. For the same reason, every nonsingular geodesic in $\widetilde{S}$ can be lifted to a nonsingular geodesic in $\hat{S}$. In fact, the same is true for basic geodesics.

Lemma 2.3. *The map $\tilde{p} : \hat{S} \to \widetilde{S}$ sends basic geodesics to basic geodesics, and every basic geodesic in $\widetilde{S}$ is the $\tilde{p}$–image of a basic geodesic in $\hat{S}$.*

Proof. Given a basic geodesic η in $\hat{S}$, there is a sequence of nonsingular geodesics $\{\eta_i\}$ that converges to η. Then the sequence $\{\tilde{p}(\eta_i)\}$ of nonsingular geodesics must converge to $\tilde{p}(\eta)$, and so $\tilde{p}(\eta)$ must be a basic geodesic in $\widetilde{S}$.

Now let λ be a basic geodesic in $\widetilde{S}$ that is parametrized at unit speed, and $\{\lambda_i\}$ a sequence of nonsingular unit speed geodesics converging to λ. We can assume $\{\lambda_i\}$ is one of the sequences constructed in the proof of Lemma 2.2 above. In particular, it is a sequence of nonsingular geodesics that are either pairwise disjoint or all coincide in a unique (regular) point. In the latter case, let x be this point (which by construction is a point on λ) and in the former case, there is a common perpendicular geodesic segment to $\{\lambda_i\}$ and λ and x is its intersection with λ.

Let U be an evenly covered open disk containing x and no cone points. Consider the tail of the sequence $\{\lambda_i\}$ that intersects U. Dropping everything but the tail and re-indexing, every geodesic in the sequence $\{\lambda_i\}$ intersects U. Choosing base points in U, lift the geodesics $\{\lambda_i\}$ to nonsingular geodesics $\{\eta_i\}$ in $\hat{S}$. We claim that $\{\eta_i\}$ converges to a geodesic η in $\hat{S}$. If this were true, η would be a basic geodesic since it is the limit of a sequence of nonsingular geodesics, and $\tilde{p}(\eta) = \lambda$. To show that $\{\eta_i\}$ converges, we first recall that convergence in the compact open topology is equivalent to uniform convergence on compact sets.

Since $\{\lambda_i\}$ converges in the compact open topology, $\{\lambda_i\}$ must converge uniformly on compact sets. Given a compact set $K \subseteq \mathbb{R}$, there must be a ball B_K containing U that contains $\cup_i \lambda_i(K)$. If B_K contains cone points, there are geodesic rays emanating from the cone points to ∂B_K that do not intersect $\cup_i \lambda_i$ and do not intersect each other.

In the case where λ_i intersect in a point x, these rays can be chosen as follows. Connect x to each cone point by a geodesic segment, and then take the rays to be a geodesic continuation of these segments emanating from the cone points. No two can intersect because if they did, a geodesic bigon would be created with vertices x and the intersection point.

In the case where λ_i are disjoint, let γ be the common perpendicular geodesic segment to $\{\lambda_i\}$ and recall $x = \lambda \cap \gamma$. The rays can be found as follows. Take the perpendicular geodesic segments from γ to each cone point, and then take the rays to be a geodesic

continuation of these segments emanating from the cone points. No two can intersect because if they did, a geodesic triangle would be formed by the intersecting geodesics and γ with angle sum larger than π.

Let the union of the rays be called R. Then let $B := B_K \setminus R$. We have chosen R so that B contains no cone points, and so that B is simply connected. Therefore, B can be lifted homeomorphically and intrinsically isometrically to a bounded compact set in $\hat{S}$ containing $\cup_i \eta_i(K)$. Therefore $\{\eta_i\}$ converges uniformly on K as desired. $\quad\square$

We end this subsection with an observation about the sets $\mathcal{G}_{\tilde{\varphi}}$ and $\mathcal{G}_{\hat{\varphi}}$ of endpoints of basic geodesics in $\widetilde{S}$ and $\hat{S}$. To prove the Current Support Theorem we will be interested in metrics φ_1, φ_2 for which $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, and we will see that this is equivalent to the assumption that $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ (up to replacing φ_2 with an equivalent metric). We prove one implication here; see Lemma 5.1 for the converse.

Lemma 2.4. *Let $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ with $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2)$. If $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ then $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$.*

Proof. Let $\eta_1 \in \mathcal{G}(\tilde{\varphi}_1)$ be a basic $\tilde{\varphi}_1$–geodesic. Then there exists a basic $\hat{\varphi}_1$–geodesic $\hat{\eta}_1$ such that $\tilde{p}(\hat{\eta}_1) = \eta_1$. Since $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ there exists a basic $\hat{\varphi}_2$–geodesic $\hat{\eta}_1$ such that the Hausdorff distance (with respect to, say, $\hat{\varphi}_1$) between $\hat{\eta}_1$ and $\hat{\eta}_2$ is bounded. Since $\tilde{p}$ maps basic geodesics in $\hat{S}$ to basic geodesics in $\widetilde{S}$ we have that $\eta_2 = \tilde{p}(\hat{\eta}_2)$ is a basic $\tilde{\varphi}_2$–geodesic. Moreover, since $\tilde{p}$ is 1-Lipschitz, we must have that $\eta_1 = \tilde{p}(\hat{\eta}_1)$ and $\eta_2 = \tilde{p}(\hat{\eta}_2)$ stay bounded Hausdorff distance apart, and hence must have the same endpoints. We have shown that $\mathcal{G}_{\tilde{\varphi}_1} \subset \mathcal{G}_{\tilde{\varphi}_2}$. The opposite inclusion follows by a symmetric argument. $\quad\square$

2.4. Developing and holonomy

As all along, suppose $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$ and let $C_0 = \mathrm{cone}(\varphi) \subset S$ be the set of cone points, $\hat{\varphi}$ the pull back metric $\hat{p}^*(\varphi)$ and $\hat{C}_0 = \hat{p}^{-1}(C_0) \subset \hat{S}$. Since $\widetilde{\hat{S}}$ is simply connected and $\hat{\varphi}$ is locally isometric to the hyperbolic plane, there is a *developing map* $\widetilde{\hat{S}} \to \mathbb{H}$: an orientation preserving local isometry, unique up to post-composing by an element of $\mathrm{PSL}_2(\mathbb{R})$. This map also extends to the metric completion, which we continue to call a developing map and denote it by

$$D \colon \hat{S} \to \mathbb{H}.$$

Observe that for every basic geodesic $\eta \in \mathcal{G}(\hat{\varphi})$, $D(\eta)$ is a bi-infinite geodesic in $\mathbb{H}$.

Lemma 2.5. *Fix any $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$. Then for any point $\zeta \in \hat{S}$, any geodesic in $\mathbb{H}$ through $D(\zeta)$ is the D–image of a basic $\hat{\varphi}$–geodesic in $\hat{S}$ through ζ. In particular, D is surjective and every geodesic in $\mathbb{H}$ is the D–image of a basic $\hat{\varphi}$–geodesic.*

Proof. First suppose ζ is any point in $\hat{S} \smallsetminus \hat{C}_0$. For every unit tangent vector in $\mathbf{v} \in T_\zeta^1 \hat{S}$ we can find a basic $\hat{\varphi}$–geodesic $\eta_{\mathbf{v}}$ through ζ tangent to $\mathbf{v}$ (see Lemma 2.2). Observe that the derivative dD_ζ is an isometry, and that $D(\eta_{\mathbf{v}})$ is a geodesic through $D(\zeta)$ tangent to $dD_\zeta(\mathbf{v})$. Since dD_ζ is an isometry, every geodesic in $\mathbb{H}$ through $D(\zeta)$ is the image of some basic $\hat{\varphi}$–geodesic in $\hat{S}$. Essentially the same argument works if $\zeta \in \hat{C}_0$: in this case, however, the derivative dD_ζ restricted to the unit vectors $T_\zeta^1 \hat{S}$ in the tangent cone $T_\zeta \hat{S}$ at ζ maps onto $T_{D(\zeta)}^1(\mathbb{H})$ (in fact, the restriction of the derivative $dD_\zeta \colon T_\zeta^1 \hat{S} \to T_{D(\zeta)}^1 \mathbb{H}$ is the universal cover). This proves the first claim of the lemma. The second claim is immediate from the first. $\square$

Let $\mathcal{K} \subseteq \hat{S}_\infty^1$ be the subset of the Gromov boundary of $\hat{S}$ consisting of end points of a basic $\hat{\varphi}$–geodesics. We extend the map D to the union

$$\bar{D} \colon \hat{S} \cup \mathcal{K} \to \overline{\mathbb{H}},$$

so that $\bar{D}(\mathcal{K}) \subset \partial \mathbb{H}$. To see how this is done, let $k \in \mathcal{K}$, and define $\bar{D}(k)$ to be the forward end point of $D(\eta)$ where η is any oriented basic $\hat{\varphi}$–geodesic whose forward end point is k. This is well-defined because any two forward asymptotic basic $\hat{\varphi}$–geodesics will have forward asymptotic D–images.

Corollary 2.6. *The restriction $\bar{D}|_\mathcal{K} \colon \mathcal{K} \to \partial \mathbb{H}$ is surjective.*

Proof. For every $y \in \partial \mathbb{H}$ one can choose an oriented geodesic η' that has y as its forward end point. By Lemma 2.5, there is a basic $\hat{\varphi}$–geodesic η with $D(\eta) = \eta'$. Then $\bar{D}(k) = y$ where k is the forward end point of η. $\square$

Precomposing D with any covering transformation $\gamma \in \pi_1 \dot{S}$ acting on $\hat{S}$, gives another developing map to $\mathbb{H}$ that must therefore differ from D by post-composing with an element in $\mathrm{PSL}_2(\mathbb{R})$ denoted $\rho(\gamma)$. The assignment $\gamma \mapsto \rho(\gamma)$ thus defines the *holonomy homomorphism* $\rho \colon \pi_1 \dot{S} \to \mathrm{PSL}_2(\mathbb{R})$. That is, ρ is defined by

$$D(\gamma \cdot x) = \rho(\gamma) \cdot D(x),$$

for all $x \in \hat{S}$ and $\gamma \in \pi_1 \dot{S}$. We let $\Gamma = \rho(\pi_1 \dot{S})$, and note that since D is ρ–equivariant, it descends to a continuous map $q \colon S \to \mathcal{O} = \mathbb{H}/\Gamma$ making the following diagram commute:

$$
\begin{array}{ccc}
\hat{S} & \xrightarrow{\;D\;} & \mathbb{H} \\
{\scriptstyle \hat{p}} \downarrow & & \downarrow \\
S & \xrightarrow{\;q\;} & \mathcal{O}.
\end{array}
$$

Here $\mathcal{O} = \mathbb{H}/\Gamma$ is just the quotient topological space, *unless* Γ is discrete, in which case it admits the structure of a hyperbolic 2–orbifold. In any case, we note that since S is compact and q is continuous, $\mathcal{O}$ is compact. That is, Γ is a cocompact subgroup.

2.5. Topological conjugacy

Recall that the action of $\text{Isom}(\mathbb{H})$ on $\mathbb{H}$ extends uniquely to an action on $\overline{\overline{\mathbb{H}}} = \mathbb{H} \cup \partial\mathbb{H}$. We will implicitly use this extension without further comment, and likewise for the orientation preserving index two subgroup $\text{PSL}_2(\mathbb{R}) < \text{Isom}(\mathbb{H})$.

Given two subgroups $G_1, G_2 < \text{Isom}(\mathbb{H})$, we say that the homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ *topologically conjugates G_1 to G_2* if for all $\gamma_1 \in G_1$ there exists $\gamma_2 \in G_2$ such that

$$h\gamma_1 h^{-1}(x) = \gamma_2(x)$$

for all $x \in \partial\mathbb{H}$. Given two hyperbolic cone metrics φ_1 and φ_2 with $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and the corresponding holonomy homomorphisms ρ_1 and ρ_2, we will see in the course of the proof that $\Gamma_1 = \rho_1(\pi_1\dot{S})$ and $\Gamma_2 = \rho_2(\pi_1\dot{S})$ are topologically conjugate. In fact, we will prove that the homomorphisms are topologically conjugate (see Proposition 5.4), meaning that there is a homeomorphism $h: \partial\mathbb{H} \to \partial\mathbb{H}$ so that for all $x \in \partial\mathbb{H}$ and $\gamma \in \pi_1\dot{S}$,

$$h(\rho_1(\gamma) \cdot x) = \rho_2(\gamma) \cdot h(x).$$

In this case we will say that *h topologically conjugates ρ_1 to ρ_2*.

Part of the proof of the Current Support Theorem involves proving that quite often the homeomorphism h which topologically conjugates ρ_1 to ρ_2 is in fact in $\text{PSL}_2(\mathbb{R})$. The following will be an important ingredient.

Proposition 2.7. *Suppose G_1, G_2 are two cocompact subgroups of $\text{Isom}(\mathbb{H})$. If there exists an orientation preserving homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ topologically conjugating G_1 to G_2, then either the groups are discrete or $h \in \text{PSL}_2(\mathbb{R})$.*

Proof. First note that h must topologically conjugate the orientation preserving subgroup of G_1 to the orientation preserving subgroup of G_2. Since discreteness of a group is equivalent to discreteness of its orientation preserving subgroup, we may assume that G_1 and G_2 are in $\text{PSL}_2(\mathbb{R})$.

For $i = 1, 2$ let $\bar{G}_i$ denote the closure of G_i and let $\bar{G}_i^{id}$ be the connected component of $\bar{G}_i$ containing the identity. Since h conjugates G_1 to G_2 it also topologically conjugates $\bar{G}_1$ to $\bar{G}_2$ and thus $\bar{G}_1^{id}$ to $\bar{G}_2^{id}$. Also, since $\bar{G}_i^{id}$ is a closed subgroup of $\text{PSL}_2(\mathbb{R})$ it is a Lie subgroup by the closed subgroup theorem. It follows from general Lie theory (by the bijection between closed connected subgroups of $\text{PSL}_2(\mathbb{R})$ and subalgebras of $\mathfrak{sl}_2(\mathbb{R})$) that $\bar{G}_1^{id}$, as well as $\bar{G}_2^{id}$, must be one of the following groups up to conjugation in $\text{PSL}_2(\mathbb{R})$:

- $T = \left\{ \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \middle| t \in \mathbb{R} \right\}$ $\bullet$ PSO(2)

- $U = \left\{ \begin{pmatrix} s & t \\ 0 & \frac{1}{s} \end{pmatrix} \middle| s,t \in \mathbb{R}, s > 0 \right\} = \mathrm{Stab}\{\infty\}$ $\bullet$ $\mathrm{PSL}_2(\mathbb{R})$

- $D = \left\{ \begin{pmatrix} s & 0 \\ 0 & \frac{1}{s} \end{pmatrix} \middle| s > 0 \right\} = \mathrm{Stab}\{0\} \cap \mathrm{Stab}\{\infty\}$ $\bullet$ $\{id\}$

Note that if $\bar{G}_1^{id} = \{id\}$ then G_1, and therefore also G_2, is discrete. So suppose $\bar{G}_1^{id}$, $\bar{G}_2^{id}$ are non-trivial. Using the facts that $\gamma \in \bar{G}_1^{id}$ fixes $p \in \partial\mathbb{H}$ if and only if $h\gamma h^{-1} \in \bar{G}_2^{id}$ fixes $h(p)$ and that the five non-trivial groups in the above list can be distinguished by considering fixed points sets and point stabilizers, we see that $\bar{G}_1^{id}$ and $\bar{G}_2^{id}$ are $\mathrm{PSL}_2(\mathbb{R})$–conjugate to the same group. Hence, by composing h with an element in $\mathrm{PSL}_2(\mathbb{R})$ we can assume that $\bar{G}_1^{id} = \bar{G}_2^{id}$ and is one of the above non-trivial groups. Now, conjugation by h defines an orientation preserving continuous automorphism of $\bar{G}_1^{id}$:

$$\rho \colon \bar{G}_1^{id} \to \bar{G}_1^{id}, \quad \gamma \mapsto h\gamma h^{-1}.$$

Below we will see that in all cases except $\bar{G}_1^{id} = D$ this forces h to be an element of $\mathrm{PSL}_2(\mathbb{R})$.

Case 1: Suppose that $\bar{G}_1^{id} \in \{T, U, \mathrm{PSL}_2(\mathbb{R})\}$ and identify $\partial\mathbb{H}$ with $\mathbb{R} \cup \{\infty\}$. We can assume h fixes infinity: this is necessarily the case if $\bar{G}_1^{id} \in \{T, U\}$ and if $\bar{G}_1^{id} = \mathrm{PSL}_2(\mathbb{R})$ we can further compose h with an element of $\mathrm{PSL}_2(\mathbb{R})$ so that it is true. Note that in all cases $T < \bar{G}_1^{id}$ and h must conjugate T to itself (again by considering the fixed point sets). That is, $\rho|_T$ defines an orientation preserving continuous automorphism of T. Each element in T corresponds to a translation $\gamma_t \colon x \mapsto x + t$ for some $t \in \mathbb{R}$ and, by identifying T with $\mathbb{R}$ through the identification $\gamma_t \mapsto t$, ρ gives rise to an orientation preserving continuous automorphism of $\mathbb{R}$. Any such automorphism is of the form $t \mapsto \lambda t$ for some fixed $\lambda > 0$. Hence there exists $\lambda \neq 0$ such that $\rho(\gamma_t) = \gamma_{\lambda t}$ for all $\gamma_t \in T$ and letting $p = h(0) \neq \infty$ we have for each $t \in \mathbb{R}$,

$$h(t) = h(\gamma_t(0)) = h\gamma_t h^{-1}(p) = \rho(\gamma_t)(p) = \gamma_{\lambda t}(p) = \lambda t + p.$$

Hence h is in $\mathrm{PSL}_2(\mathbb{R})$.

Case 2: Suppose $\bar{G}_1^{id} = \mathrm{PSO}(2)$ and identify $\partial\mathbb{H}$ with the circle S^1. By also identifying $\mathrm{PSO}(2)$ with S^1 (by sending each element to its angle of rotation) we have that $\rho \colon \mathrm{PSO}(2) \to \mathrm{PSO}(2)$ induces an orientation preserving continuous automorphism of S^1. The only such automorphism is the identity. It follows that $h\gamma h^{-1} = \gamma$ for all $\gamma \in \mathrm{PSO}(2)$, i.e. h commutes with every rotation. Therefore h is a rotation itself, and in particular, h is in $\mathrm{PSL}_2(\mathbb{R})$.

Case 3: Suppose $\bar{G}_1^{id} = D$. In fact, we will show that this cannot happen due to cocompactness of G_1. Suppose $\bar{G}_1 \setminus D \neq \emptyset$ and let $\gamma \in \bar{G}_1 \setminus D$. Then $\gamma \notin \mathrm{Stab}\{0\} \cap \mathrm{Stab}\{\infty\}$, but $\gamma \bar{G}_1^{id} \gamma^{-1} = \bar{G}_1^{id} = D = \mathrm{Stab}\{0\} \cap \mathrm{Stab}\{\infty\}$. Hence we must have that $\gamma(0) = \infty$ and $\gamma(\infty) = 0$, i.e. $\gamma \in \mathrm{Stab}\{0, \infty\}$. It follows that $\bar{G}_1 = \mathrm{Stab}\{0, \infty\} = \left\langle D, \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \right\rangle = D'$. Hence either $\bar{G}_1 = D$ or $\bar{G}_1 = D'$. However, neither D nor D' is cocompact. To see this, set

$$W_n = \left\{ (x, y) \mid y > \frac{|x|}{n} \right\}$$

and note that $\{W_n \mid n \in \mathbb{N}\}$ is an open cover of $\mathbb{H}$ (in the upper half space model) which descends to an open cover of both $\mathbb{H}/D$ and $\mathbb{H}/D'$ with no finite subcover. However G_1, and hence also $\bar{G}_1$, is cocompact giving a contradiction. $\quad \square$

We note that Proposition 2.7 above gives no information about the conjugating homeomorphism h when the groups are discrete. However, in a very special case of discrete groups we can also conclude that h is in $\mathrm{PSL}_2(\mathbb{R})$, as the lemma below shows. Recall that a *triangle group* is a discrete subgroup of $\mathrm{Isom}(\mathbb{H})$ having a triangle Δ as its fundamental domain and such that it is generated by reflections in the sides of Δ. By a *Fuchsian triangle group* we mean the orientation preserving subgroup of a triangle group.

Lemma 2.8. *Suppose $G_1, G_2 < \mathrm{PSL}_2(\mathbb{R})$ and that there exists an orientation preserving homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ conjugating G_1 to G_2. If G_1 is a Fuchsian triangle group, then $h \in \mathrm{PSL}_2(\mathbb{R})$.*

Proof. The Douady-Earle extension [8] extends any self-homeomorphism of $\partial\mathbb{H}$ to a self-homeomorphism of the closure $\overline{\mathbb{H}}$ of the hyperbolic plane in an equivariant way: pre- or post-composing the extension with an element of $\mathrm{PSL}_2(\mathbb{R})$ is the same as first pre- or post-composing the homeomorphism with this element and then extending the resulting homeomorphism. Let $\bar{h} : \overline{\mathbb{H}} \to \overline{\mathbb{H}}$ be the Douady-Earle extension of h. The equivariance property implies that $\bar{h}\gamma = \rho(\gamma)\bar{h}$ for all $\gamma \in G_1$, where $\rho : G_1 \to G_2$ is the isomorphism $\gamma \mapsto h\gamma h^{-1}$ induced by conjugation by h. It follows that $\bar{h}$ descends to a homeomorphism $\mathcal{O}_1 \to \mathcal{O}_2$ between the orbifolds $\mathcal{O}_i = \mathbb{H}/G_i$. However, $\mathcal{O}_1$ is a triangle orbifold—which has a trivial Teichmüller space—and hence the orbifold deformation must be isotopic to an isometry. It follows that the boundary map h is the extension of an isometry and hence an element of $\mathrm{PSL}_2(\mathbb{R})$ (in fact, so is $\bar{h}$ by construction). $\quad \square$

3. Deformations

Here we prove that the construction briefly described in the introduction does indeed yield examples with the desired properties. Specifically, we prove the following.
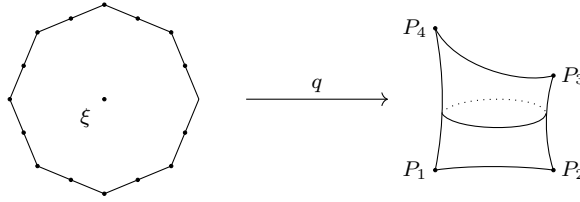
Fig. 3. A genus 2 surface obtained by identifying opposite sides admits an obvious order-eight symmetry, and ρ is its order-four square. The branched cover $q\colon S \to S/\langle\rho\rangle$ is branched over the four points P_1, P_2, P_3, P_4, and branches at all the points indicated on the octagon, which are identified in S to 6 points. We consider $S/\langle\rho\rangle$ as a hyperbolic orbifold with P_1, P_2, P_3, P_4 as orbifold points of orders $2, 2, 2, 4$, respectively.

Theorem 3.1. *Suppose $\varphi_1 \in \widetilde{\mathrm{Hyp}}_c(S)$, $q\colon (S, \varphi_1) \to \mathcal{O}_1$ is a locally isometric branched covering of an orbifold with $q(\mathrm{cone}(\varphi_1)) \subset \mathcal{E}_1$, the even order orbifold points. Then for any (orbifold) homeomorphism $F\colon \mathcal{O}_1 \to \mathcal{O}_2$, the metric φ_2 obtained by pulling back the hyperbolic metric by $F \circ q$ has $\mathcal{G}_{\tilde\varphi_1} = \mathcal{G}_{\tilde\varphi_2}$ (in fact, $\mathcal{G}_{\hat\varphi_1} = \mathcal{G}_{\hat\varphi_2}$). If F is not orbifold isotopic to an isometry, then $\varphi_1 \nsim \varphi_2$.*

Before we begin the proof, we provide an example exhibiting the behavior in this theorem.

Example 3.2. Let S be a genus 2 surface, which we view as an octagon with opposite sides identified. There is an order eight rotational symmetry, and we let $\rho\colon S \to S$ denote its square, which thus has order four. The quotient $S/\langle\rho\rangle$ is a sphere, and the quotient map $q\colon S \to S/\langle\rho\rangle$ is a branched cover. The center of the octagon, $\xi \in S$, projects to a single point $P_1 \in S/\langle\rho\rangle$, and q is locally four-to-one near that point. The midpoints and vertices of the octagon account for the rest of the branch points. The midpoints define four points of S near which q is locally two-to-one, mapping to two points we call P_2 and P_3, and the vertices define a single point of S near which q is locally four-to-one, mapping to a single point we call P_4; see Fig. 3.

Now, if we were to impose a hyperbolic orbifold structure on $S/\langle\rho\rangle$ so that P_1, P_2, P_3, P_4 were orbifold points with orders $4, 2, 2, 4$, respectively, then the pull back metric via q would be an honest hyperbolic structure and q would be an orbifold covering map. Instead, we make $S/\langle\rho\rangle$ into a hyperbolic orbifold $\mathcal{O}_1$ so that the orders of P_1, P_2, P_3, P_4 are $2, 2, 2, 4$, instead. Pulling back the metric via q now gives S a hyperbolic cone metric φ_1 on S so that the cone angle of ξ is 4π, and q is a branched cover. We may deform this hyperbolic orbifold by a homeomorphism to another hyperbolic orbifold $F\colon \mathcal{O}_1 \to \mathcal{O}_2$, and we obtain a new hyperbolic cone metric φ_2 by pulling this metric back via $F \circ q$. According to Theorem 3.1, $\mathcal{G}_{\tilde\varphi_1} = \mathcal{G}_{\tilde\varphi_2}$.

Delaunay cell decompositions provide a convenient tool for analyzing hyperbolic orbifolds (and elsewhere below), and so we recall their construction in a special case of interest for us. Let $G < \mathrm{PSL}_2(\mathbb{R})$ be a discrete, cocompact subgroup and assume that it contains an even order elliptic element. Let $\widetilde{\mathcal{E}} \subset \mathbb{H}$ denote the G–invariant set of fixed

points of elliptic elements of even order. The associated *Delaunay cell decomposition* $\tilde{\Delta}(\widetilde{\mathcal{E}})$ is the unique G–invariant cell decomposition of $\mathbb{H}$ whose 2–cells are compact, convex polygons with vertices in $\widetilde{\mathcal{E}}$ characterized by the property that their circumscribing circles contain no points of $\widetilde{\mathcal{E}}$ in their interiors and so that each 2–cell is the convex hull of its vertices which are all contained in its circumscribing circle; see e.g. [4]. We collect some observations about this in the next lemma.

Lemma 3.3. *Suppose $G < \mathrm{PSL}_2(\mathbb{R})$ is a cocompact discrete subgroup containing an even order elliptic element and $\widetilde{\mathcal{E}}$ is the set of fixed points of even order elliptic elements. Then*

1. *The 1–skeleton $\tilde{\Delta}(\widetilde{\mathcal{E}})^{(1)}$ is a G–invariant union of biinfinite geodesics through $\widetilde{\mathcal{E}}$.*
2. *If $H < G$ is a nontrivial, maximal, cyclic subgroup, and H has odd order, then H fixes a point in the interior of a 2–cell of $\tilde{\Delta}(\widetilde{\mathcal{E}})$. Moreover, each 2–cell contains at most one elliptic fixed point in its interior.*

Note in condition 2, since the even order elliptic elements fix vertices, any elliptic fixed point in the interior of a 2–cell is necessarily the fixed point of an odd order elliptic subgroup.

Proof. $\tilde{\Delta}(\widetilde{\mathcal{E}})^{(1)}$ is a G–invariant union of geodesic segments. Given any $\zeta \in \widetilde{\mathcal{E}}$ there is a unique elliptic element $\tau_\zeta \in G$ of order two fixing ζ. For any edge e of $\tilde{\Delta}(\widetilde{\mathcal{E}})^{(1)}$ having ζ as an endpoint, $\tau_\zeta(e)$ is also an edge having ζ as an endpoint, and together $e \cup \tau_\zeta(e)$ is a geodesic segment (because it makes angle π on both sides at ζ). Applying the same reasoning at the other endpoints of e and $\tau_\zeta(e)$, and repeating recursively, we see that the unique biinfinite geodesic containing e is a union of edges of $\tilde{\Delta}(\widetilde{\mathcal{E}})$, proving the first statement.

For the second statement, observe that the unique fixed point ξ cannot be in the 1–skeleton. To see this, note first that if ξ were a vertex, then H would contain an even order element, and hence $|H|$ would be even, a contradiction. On the other hand, if ξ were on the interior of an edge e, then by G–invariance of $\tilde{\Delta}(\widetilde{\mathcal{E}})$, $H \cdot e$ would consist of more than one edge containing ξ, meaning that ξ is a vertex, another contradiction. Therefore, ξ is necessarily in the interior of some 2–cell σ. Since H fixes ξ, it fixes σ, and hence the (unique) circumcircle of σ is invariant by H and thus has ξ as its center. Any other elliptic subgroup with a unique fixed point in the interior of H must also have its fixed point as the center of this circumcenter, and hence is contained in H. That is, there is a unique elliptic fixed point in the interior of σ, as required. $\quad\square$

We will use the Delaunay cell structure to construct maps from $\mathbb{H}$ to itself. To describe the construction of the maps, we first recall that given any two k–simplices $\sigma_1, \sigma_2 \in \mathbb{H}$ and a bijection between the vertices $\sigma_1^{(0)} \to \sigma_2^{(0)}$, there is a natural way of extending this to a map $\sigma_1 \to \sigma_2$ as follows. In the hyperboloid model, the vertices on the hyperboloid

represent linearly independent vectors and the bijection between the vertices of σ_1 and σ_2 determines a linear isomorphism between their respective spans. Restricting this map to the simplex σ_1 and radially projecting back to the hyperboloid defines the *canonical map $\sigma_1 \to \sigma_2$* extending the bijection on the vertices. We observe that (1) pre- and post-composing a canonical map with isometries of $\mathbb{H}$ is again a canonical map, and (2) the restriction of a canonical map to a face is again a canonical map (cf. [20, Section 11.4]).

Next suppose $\mathcal{T}_1, \mathcal{T}_2$ are *geometric triangulations* of the hyperbolic plane, a pair of hyperbolic orbifolds, or hyperbolic (cone) surfaces: that is, these are Δ–complex structures (in the sense of [14]) so that each triangle of either structure admits a locally isometric parameterization by a geodesic triangle in $\mathbb{H}$. A cellular homeomorphism is a *canonical map* if its restriction to each simplex of $\mathcal{T}_1$ is a canonical map to a simplex of $\mathcal{T}_2$; or more precisely, the lift via the locally isometric parameterizations of the restriction to each simplex is a canonical map. Every cellular homeomorphism is isotopic through cellular maps to a canonical map—indeed, it is straightforward to construct the isotopy inductively over the skeleta.

Lemma 3.4. *Let $G_1, G_2 < \mathrm{PSL}_2(\mathbb{R})$ be two discrete, cocompact groups, containing even order elliptic elements, and $\rho: G_1 \to G_2$ an isomorphism. If $\widetilde{\mathcal{E}}_1, \widetilde{\mathcal{E}}_2 \subset \mathbb{H}$ are the sets of even order elliptic fixed points, then there is a G_2–invariant cell structure $\tilde{\Delta}_2$ with $\tilde{\Delta}_2^{(0)} = \widetilde{\mathcal{E}}_2$ and a ρ–equivariant cellular homeomorphism $\tilde{F}: (\mathbb{H}, \tilde{\Delta}_1) \to (\mathbb{H}, \tilde{\Delta}_2)$ where $\tilde{\Delta}_1 = \tilde{\Delta}(\widetilde{\mathcal{E}}_1)$. Moreover, $\tilde{\Delta}_2$ satisfies the same conclusions as in Lemma 3.3.*

In fact, there are subdivisions of the cell structures to triangulations whose vertex set is the set of all elliptic fixed points so that the map is an equivariant canonical map. The descent $F: \mathbb{H}/G_1 \to \mathbb{H}/G_2$ of $\tilde{F}$ is likewise a canonical map with respect to geometric triangulations that have exactly one edge incident to each odd-order orbifold point.

Proof. Set $\tilde{\Delta}_1 = \Delta(\widetilde{\mathcal{E}}_1)$. By Lemma 3.3 part 1, there is a set Θ_1 of biinfinite geodesics so that

$$\tilde{\Delta}_1 = \bigcup_{\eta \in \Theta_1} \eta.$$

For any $\zeta \in \tilde{\Delta}_1^{(0)}$ there is a maximal elliptic subgroup $E_1(\zeta) < G_1$ of even order, fixing ζ. Let

$$\Theta_1(\zeta) = \{\eta \in \Theta_1 \mid \zeta \in \eta\}$$

and observe that $\Theta_1(\zeta)$ is $E_1(\zeta)$–invariant.

The isomorphism ρ is given as a topological conjugacy by a homeomorphism

$$h: \partial\mathbb{H} \to \partial\mathbb{H},$$

since the isomorphism is a quasi-isometry and thus extends to an equivariant homeomorphism of the boundaries at infinity. The homeomorphism h defines a straightening map

$\eta \mapsto h_*(\eta)$ that sends any biinfinite geodesic η to the unique geodesic with endpoints $h(\partial\eta)$. Since h is ρ–equivariant, so is h_*. We note that since h is a homeomorphism, two geodesics η, η' intersect in $\mathbb{H}$ if and only if the geodesics $h_*(\eta)$ and $h_*(\eta')$ intersect as well. This is because h must preserve the linking (or lack thereof) of $\partial\eta$ and $\partial\eta'$.

We will define $\tilde{F}$ inductively over the skeleta, and in the process, we define the cell structure $\tilde{\Delta}_2$. For any vertex $\zeta \in \tilde{\Delta}_1^{(0)}$, we define $\tilde{F}(\zeta)$, to be the unique fixed point of $\rho(E_1(\zeta)) < G_2$. By construction, the restriction of $\tilde{F}$ to $\tilde{\Delta}_1^{(0)}$ is ρ–equivariant. We claim that for any $\zeta \in \tilde{\Delta}_1^{(0)}$, we have

$$\tilde{F}(\zeta) \in \bigcap_{\eta \in \Theta_1(\zeta)} h_*(\eta).$$

To see this, let $\tau_\zeta \in E_1(\zeta)$ be the unique element of order 2, and note that $\tau_\zeta(\eta) = \eta$ for all $\eta \in \Theta_1(\zeta)$, or for any geodesic passing through ζ. Since ρ is induced by conjugation by $h\colon \partial\mathbb{H} \to \partial\mathbb{H}$, restricting to $\partial\mathbb{H}$ we have

$$\tau_{\tilde{F}(\zeta)} = \rho(\tau_\zeta) = h \circ \tau_\zeta \circ h^{-1}.$$

Therefore, for any η passing through ζ, we have

$$\tau_{\tilde{F}(\zeta)}(\partial h_*(\eta)) = \tau_{\tilde{F}(\zeta)}(h(\partial\eta)) = h(\tau_\zeta(h^{-1}(h(\partial\eta)))) = h(\tau_\zeta(\partial\eta)) = h(\partial\eta) = \partial h_*(\eta). \quad (1)$$

That is, $h_*(\eta)$ is invariant by $\tau_{\tilde{F}(\zeta)}$ and consequently $\tilde{F}(\zeta)$ is on $h_*(\eta)$, as required.

For every edge e of $\tilde{\Delta}_1$ connecting a pair of vertices $\zeta, \zeta' \in \tilde{\Delta}_1^{(0)}$, define $\tilde{F}$ to send e by the canonical map to the geodesic segment between $\tilde{F}(\zeta)$ and $\tilde{F}(\zeta')$. We observe that if $\eta \in \Theta_1$ is the unique geodesic containing e, then $h_*(\eta)$ contains $\tilde{F}(e)$. Moreover, $h_*(\eta)$ preserves the order of the vertices along η. To see this, orient η and suppose that $\zeta, \zeta' \in \eta$ such that ζ is encountered before ζ' along η with respect to this orientation. Let $\tau_\zeta, \tau_{\zeta'} \in G_1$ be order two elliptic elements fixing ζ, ζ' respectively. Note that τ_ζ and $\tau_{\zeta'}$ both fix η (but interchange its endpoints) and the hyperbolic element $\tau_{\zeta'} \circ \tau_\zeta$ has η as its axis and translates along it in the positive direction; see Fig. 4. It follows that $\rho(\tau_{\zeta'} \circ \tau_\zeta)$ has $h_*(\eta)$ as its axis and translates along it in the positive direction. Since $\rho(\tau_{\zeta'} \circ \tau_\zeta) = \rho(\tau_{\zeta'}) \circ \rho(\tau_\zeta)$ and $\rho(\tau_\zeta)$ and $\rho(\tau_{\zeta'})$ are elliptic elements fixing $\tilde{F}(\zeta)$ and $\tilde{F}(\zeta')$ it follows that $F(\zeta)$ is encountered before $F(\zeta')$ along $h_*(\eta)$ and so the order is preserved. In particular, $\tilde{F}$ maps each $\eta \in \Theta_1$ to $h_*(\eta)$ by a map which is the canonical map on each segment between consecutive vertices of $\tilde{\Delta}_1$. For any two geodesics η, η', we have $\eta \cap \eta' = \emptyset$ if and only if $h_*(\eta) \cap h_*(\eta') = \emptyset$ and it follows that $\tilde{F}$ is injective on the 1–skeleton $\tilde{\Delta}_1^{(1)}$. Since the map on the zero–skeleton is ρ–equivariant and the maps on the edges are canonical maps, it follows that $\tilde{F}$ is a ρ–equivariant homeomorphism onto its image. Consequently, $\tilde{F}(\tilde{\Delta}_1^{(1)})$ is the 1–skeleton of a cell structure $\tilde{\Delta}_2$ on the image.

Evidently, $\tilde{\Delta}_2$ has property (1) of Lemma 3.3. Because of this ρ–equivariance, $\tilde{F}$ preserves the cyclic ordering of edges incident to each vertex. This implies that cycles
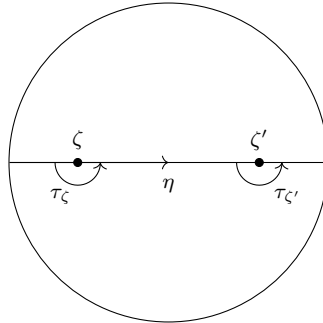
Fig. 4. The composition of τ_ζ and $\tau_{\zeta'}$ is a hyperbolic isometry with axis η.
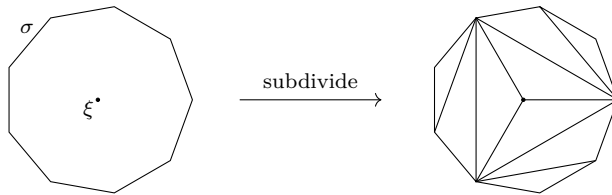

Fig. 5. Subdividing a 2–cell σ containing the fixed point ξ of a maximal cyclic subgroup $H_1 < G_1$ of order $k = 3$ in its interior. In this example, σ is the 9–gon on the left. The subdivision is as shown on the right.

that bound a 2–cell in $\tilde{\Delta}_1$ are sent to cycles that bound a 2–cell in $\tilde{\Delta}_2$. This suggests how we may extend $\tilde{F}$ to a map on all of $\mathbb{H}$. The details of this follow shortly.

To see that property (2) also holds, we note that the boundary of any 2–cell containing an elliptic fixed point in its interior is sent to the boundary of a 2–cell of $\tilde{\Delta}_2$ by ρ–equivariance, and that this defines a bijection between the 2–cells containing elliptic fixed points in their interiors. We also note that each 2–cell of $\tilde{\Delta}_2$ is a convex polygon: the 1–skeleton is a union of biinfinite geodesics and so the interior angles are all less than π.

Finally, we subdivide each $\tilde{\Delta}_i$ to triangulations $\tilde{\Delta}_i'$ as follows. For each non-triangular 2–cell of $\tilde{\Delta}_1$ not containing an elliptic fixed point in its interior, we note the interior is mapped disjoint from itself by any nontrivial element of G_1, and so the same is true for the corresponding 2–cell of $\tilde{\Delta}_2$. We subdivide the union of all such 2–cells of $\tilde{\Delta}_1$ in any G_1–invariant way, without introducing new vertices, so that each new 1–cell is a geodesic segment. Convexity of the 2–cells of $\tilde{\Delta}_2$ implies that there is a G_2–invariant subdivision of 2–cells not containing elliptic fixed points in their interiors so that $\tilde{F}$ extends to a cellular homeomorphism on the subdivision. Finally, for each 2–cell σ of $\tilde{\Delta}_1$ that contains in its interior the fixed point ξ of some maximal, odd order elliptic subgroup $H_1 < G_1$, first pick a vertex of v of σ and begin the subdivision by connecting consecutive vertices of the orbit $H_1 \cdot v$ (if they are not already connected by an edge), doing so in a G_1–invariant way. The cell σ now contains an H_1–invariant k–gon with ξ in its interior and k polygons permuted by H_1, where $k = |H_1|$, unless σ was already a k–gon. Now we cone off the k–gon to ξ and subdivide the remaining polygons in a G_1–invariant way. See Fig. 5.

We have thus subdivided $\tilde{\Delta}_1$ to a G_1–invariant triangulation $\tilde{\Delta}'_1$, and the ρ–equivariant homeomorphism $\tilde{F}|_{\tilde{\Delta}_1}$, together with the convexity of the 2–cells of $\tilde{\Delta}_2$ determines an associated subdivision to a triangulation $\tilde{\Delta}'_2$ so that $\tilde{F}|_{\tilde{\Delta}_1}$ extends to a canonical map with respect to $\tilde{\Delta}'_1$ and $\tilde{\Delta}'_2$. The properties of the descent to the orbifolds have been built into the definition, completing the proof. $\square$

We will also need the following property of isomorphic, cocompact, discrete groups G_1, G_2 as in Lemma 3.4.

Lemma 3.5. *Suppose G_1, G_2 are any two cocompact, discrete groups containing even order elliptic elements, $\rho\colon G_1 \to G_2$ is an isomorphism given by conjugation by $h\colon \partial\mathbb{H} \to \partial\mathbb{H}$, and suppose*

$$\tilde{F}\colon (\mathbb{H}, \tilde{\Delta}_1) \to (\mathbb{H}, \tilde{\Delta}_2)$$

is the cellular homeomorphism from Lemma 3.4. Then for any biinfinite geodesic η in $\mathbb{H}$, η nontrivially intersects a k–cell σ of $\tilde{\Delta}(\widetilde{\mathcal{E}}_1)$ if and only if the straightening $h_*(\eta)$ of $\tilde{F}(\eta)$ nontrivially intersects $\tilde{F}(\sigma)$.

Proof. The fact that the lemma holds for vertices $\zeta \in \tilde{\Delta}_1^{(0)} = \widetilde{\mathcal{E}}_1$ and geodesics η that intersect ζ follows immediately from (1) (the only thing being used in that argument was that ζ was a point on η).

Now we show that the lemma holds for an edge e. By the previous paragraph, it suffices to show that η intersects the interior of e if and only if $h_*(\eta)$ intersects the interior of $\tilde{F}(e)$. For this, observe that η intersects the interior of e if and only if there are geodesics η_+ and η_- through the endpoints ζ_+ and ζ_- of e, respectively, so that η is in the subset of $\mathbb{H}$ bounded by η_+ and η_-. Since this configuration of η, η_-, η_+ is determined by the endpoints on $\partial\mathbb{H}$, it follows that η is in the subset of $\mathbb{H}$ bounded by η_+ and η_- if and only if $h_*(\eta)$ is contained in the subset of $\mathbb{H}$ bounded by $h_*(\eta_+)$ and $h_*(\eta_-)$, and this happens (for some η_+, η_-) if and only if $h_*(\eta)$ intersects the interior of $\tilde{F}(e)$ (since $h_*(\eta_+)$ and $h_*(\eta_-)$ contain the endpoints $\tilde{F}(\zeta_+)$ and $\tilde{F}(\zeta_-)$).

Since a geodesic that intersects a 2–cell must also intersect its boundary, the case of 0–cells and 1–cells implies the case of 2–cells. This observation completes the proof. $\square$

We are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. Suppose $G_1, G_2 < \mathrm{PSL}_2(\mathbb{R})$ are two cocompact discrete groups with quotient orbifolds $p_i\colon \mathbb{H} \to \mathcal{O}_i = \mathbb{H}/G_i$, and suppose that $F\colon \mathcal{O}_1 \to \mathcal{O}_2$ is a homeomorphism. We also let $\mathcal{E}_i \subset \mathcal{O}_i$ be the even order orbifold points, so that $p_i(\widetilde{\mathcal{E}}_i) = \mathcal{E}_i$.

Now suppose $q_1\colon (S, \varphi_1) \to \mathcal{O}_1$ is a locally isometric branched cover of the orbifold, branched only over even order orbifold points. Let $q_2 = F \circ q_1\colon S \to \mathcal{O}_2$ be the composition, and φ_2 be the pull back of the hyperbolic metric on $\mathcal{O}_2$ (so that with respect to φ_2, q_2 is a locally isometric branched cover).

The homeomorphism $F\colon \mathcal{O}_1 \to \mathcal{O}_2$ induces an isomorphism $F_*\colon G_1 \to G_2$. Lemma 3.4 produces a homeomorphism $\mathcal{O}_1 \to \mathcal{O}_2$ inducing $F_*\colon G_1 \to G_2$ that is isotopic to F (by the Dehn-Nielsen-Baer Theorem). Therefore, by applying an isotopy if necessary, we may assume that F is the homeomorphism from Lemma 3.4 (the isotopy lifts to S, and thus changing F in this way results in a metric equivalent to φ_2). For $i = 1, 2$, we let $\tilde{\Delta}_i$ be the cell structures of $\mathbb{H}$ from that lemma, and Δ_i the corresponding cell structures on $\mathcal{O}_i$, for $i = 1, 2$. We pull these cell structures back to cell structures $\Lambda = \Lambda_1 = \Lambda_2$ on S. Each 2–cell of Λ is a(n immersed) geodesic polyon with respect to either metric φ_i, whose interior is embedded. The restriction of q_i to (the interior of) each such polygon is either a locally isometric homeomorphism, or else it is a k-to-1 orbifold quotient, where k is an odd integer, the order of the orbifold point in the image of the interior. This is perhaps easiest to see by lifting these cell structures back to cell structure $\hat{\Lambda}$ on $\hat{S}$, which we can view as lifting the cell structures $\tilde{\Delta}_i$ by the developing maps $D_i\colon (\hat{S}, \hat{\varphi}_i) \to \mathbb{H}$ (note that $D_2 = \tilde{F} \circ D_1$).

Next, consider the commutative diagram, in which all maps are cellular (by construction).
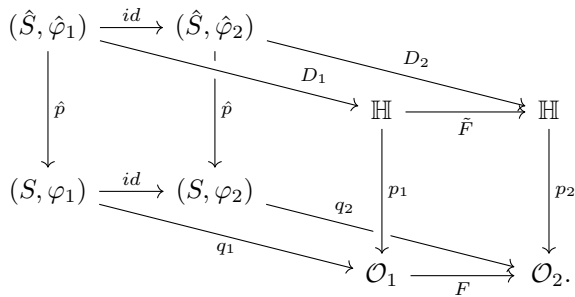
By Lemma 2.4, if we can show that $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$, then we will have $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, as required. For this, assume we have any nonsingular geodesic δ in $(\hat{S}, \hat{\varphi}_1)$. Then $D_1(\delta_1)$ is a biinfinite geodesic. For any cell σ of $\tilde{\Delta}_1$, Lemma 3.5 tells us that the straightened image geodesic $\tilde{F}_*(D_1(\delta))$ intersects $\tilde{F}(\sigma)$ if and only if $D_1(\delta)$ intersects σ. Since $\tilde{F}$ is cellular, this means that $\tilde{F}_*(D_1(\delta))$ meets the cell $\tilde{F}(\sigma)$ if and only if $\tilde{F}(D_1(\delta))$ does. Therefore, there is a homotopy from $\tilde{F}(D_1(\delta))$ to $\tilde{F}_*(D_1(\delta))$ with bounded tracks that respects the cell structure, and thus lifts to a homotopy from δ to the $\hat{\varphi}_2$–geodesic δ', also preserving the cell structures (and having bounded tracks). Note that δ meets a cell σ of Λ if and only if δ' does, and hence δ' is also nonsingular. Therefore, $\mathcal{G}_{\hat{\varphi}_1} \subset \mathcal{G}_{\hat{\varphi}_2}$. By a symmetric argument, they are equal. Lemma 2.4 completes the proof. $\square$

4. Geodesics and triangulations

Recall that $\mathcal{G}(\tilde{\varphi}_i)$ denotes the set of all basic $\tilde{\varphi}_i$–geodesics and that the maps $\partial_{\tilde{\varphi}_i}\colon \mathcal{G}(\tilde{\varphi}_i) \to \mathcal{G}_{\tilde{\varphi}_i}$, for $i = 1, 2$, which maps each geodesic to its endpoints are homeomorphisms. Hence, under the assumption that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, this gives rise to a homeomorphism

$$g : \mathcal{G}(\tilde{\varphi}_1) \to \mathcal{G}(\tilde{\varphi}_2)$$

which maps each basic $\tilde{\varphi}_1$–geodesic η to the basic $\tilde{\varphi}_2$–geodesic having the same endpoints as η. In other words, $g(\eta)$ is the $\tilde{\varphi}_2$-*straightening* of the $\tilde{\varphi}_1$–geodesic η. Note that orienting η naturally induces an orientation on $g(\eta)$.

In this section, we will show that the straightening map preserves many properties of basic geodesics. Along the way we will begin adjusting the metric φ_2 to equivalent metrics (under the assumption $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$). A useful tool in the analysis that we also develop are triangulations of S adapted to the two metrics.

4.1. Concurrency points and partitions

Recall by Lemma 2.1 that only countably many basic $\tilde{\varphi}_i$–geodesics pass through more than one cone point and that we denote the set of the corresponding pairs of endpoints by $\mathcal{G}^2_{\tilde{\varphi}_1}$. Also, given $\zeta \in \widetilde{S}$, we will write

$$\mathcal{G}(\tilde{\varphi}_i, \zeta) = \{\gamma \in \mathcal{G}(\tilde{\varphi}_i) \mid \zeta \in \gamma\} \subset \mathcal{G}(\tilde{\varphi}_i), \quad \text{and} \quad \mathcal{G}_{\tilde{\varphi}_i}(\zeta) = \partial_{\tilde{\varphi}_i}(\mathcal{G}(\tilde{\varphi}_i, \zeta)) \subset \mathcal{G}_{\tilde{\varphi}_i}$$

for each $i = 1, 2$. That is, $\mathcal{G}(\tilde{\varphi}_i, \zeta)$ is the set of basic $\tilde{\varphi}_i$–geodesics that contain ζ and $\mathcal{G}_{\tilde{\varphi}_i}(\zeta)$ the set of all pairs of endpoints at infinity of such geodesics. For $\zeta \in \hat{S}$, will use the corresponding notation $\mathcal{G}(\hat{\varphi}_i, \zeta)$ and $\mathcal{G}_{\hat{\varphi}_i}(\zeta)$ for the analogously defined subsets of $\mathcal{G}(\hat{\varphi}_i)$ and $\mathcal{G}_{\hat{\varphi}_i}$, respectively, for $i = 1, 2$.

Using the notion of *chains*, the following is proved in [3] for Euclidean cone metrics, but the proof is essentially identical in the case of hyperbolic cone metrics.

Proposition 4.1. *Suppose* $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ *with* $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. *There is a countable,* $\pi_1 S$–*invariant set* $\Omega \subset \mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ *containing* $\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2}$ *with the following property. For any* $\zeta_1 \in \mathrm{cone}(\tilde{\varphi}_1)$ *there exists* $\zeta_2 \in \mathrm{cone}(\tilde{\varphi}_2)$ *so that*

$$\mathcal{G}_{\tilde{\varphi}_1}(\zeta_1) - \Omega = \mathcal{G}_{\tilde{\varphi}_2}(\zeta_2) - \Omega.$$

Moreover, sending ζ_1 *to* ζ_2 *determines a* $\pi_1 S$–*equivariant bijection* $\mathrm{cone}(\tilde{\varphi}_1) \to \mathrm{cone}(\tilde{\varphi}_2)$.

Idea of the proof. A *chain* associated to any cone point $\zeta_i \in \mathrm{cone}(\tilde{\varphi}_i)$, for $i = 1, 2$, is a bi-infinite sequence $\mathbf{x} = \{x_k\}_{k \in \mathbb{Z}} \subset S^1_\infty$ so that for any $k \in \mathbb{Z}$, x_{k-1}, x_k, x_{k+1} is a counter-clockwise ordered triple in S^1_∞, and so that

$$\{x_k, x_{k+1}\} \in \mathcal{G}_{\tilde{\varphi}_1}(\zeta) \setminus (\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2}).$$

From the fact that the cone angles are greater than 2π at ζ_i, it is straightforward to see that for any $x \in S^1_\infty$ with x_{k+1}, x, x_{k-1} counterclockwise ordered, $\{x, x_k\}$ is *not* in $\mathcal{G}_{\tilde{\varphi}_i}$. Any biinfinite sequence $\{x_k\} \subset S^1_\infty$ with the property that $\{x_k, x_{k+1}\} \subset \mathcal{G}_{\tilde{\varphi}_i} \setminus (\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2})$

and the property for triples x_{k+1}, x, x_{k-1} just described is in fact a chain associated to a unique cone point of $\tilde{\varphi}_i$. See [3, Proposition 4.1].

Because chains depend only on $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and $\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2}$, we get functions from the set of chains to both $\mathrm{cone}(\tilde{\varphi}_1)$ and $\mathrm{cone}(\tilde{\varphi}_2)$. In fact, there is an equivalence relation on chains, describable in terms of $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and $\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2}$, so that two chains are equivalent if and only if they are sent to the same cone point of $\mathrm{cone}(\tilde{\varphi}_i)$, for each $i = 1, 2$ (see [3, Lemma 4.4]). Therefore, the set of equivalence classes of chains is in a bijective correspondence with $\mathrm{cone}(\tilde{\varphi}_1)$ and $\mathrm{cone}(\tilde{\varphi}_2)$. This in turn determines a bijection $\mathrm{cone}(\tilde{\varphi}_1) \to \mathrm{cone}(\tilde{\varphi}_2)$. The chain condition is $\pi_1 S$–invariant, and the map to cone points is equivariant, from which we deduce that the bijection $\mathrm{cone}(\tilde{\varphi}_1) \to \mathrm{cone}(\tilde{\varphi}_2)$ to $\pi_1 S$–equivariant. For any ζ_i, $i = 1, 2$, the set of pairs of endpoints of geodesics in $\mathcal{G}_{\tilde{\varphi}_i}(\zeta_i)$ which are not consecutive terms in some chain is countable, and thus taking Ω to be the union of all such pairs produces the countable set Ω containing $\mathcal{G}^2_{\tilde{\varphi}_1} \cup \mathcal{G}^2_{\tilde{\varphi}_2}$. See [3, Section 4] for more details. $\quad\square$

Going forward we let $C_0 = \mathrm{cone}(\varphi_1)$ and $\widetilde{C}_0 = p^{-1}(C_0) = \mathrm{cone}(\tilde{\varphi}_1)$. For the remainder of this subsection, we assume $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, and let Ω be the set from Proposition 4.1. The $\pi_1 S$–equivariant bijection from that proposition can be extended to a $\pi_1 S$–equivariant homeomorphism, which is thus the lift of a homeomorphism $S \to S$, isotopic to the identity. We replace φ_2 with its pull back by this homeomorphism, which does not change the equivalence class, but allows us to say that $\mathrm{cone}(\tilde{\varphi}_1) = \widetilde{C}_0 = \mathrm{cone}(\tilde{\varphi}_2)$. More precisely, we can promote the conclusion of the proposition to say that for all $\zeta \in \widetilde{C}_0$, we have

$$\mathcal{G}_{\tilde{\varphi}_1}(\zeta) - \Omega = \mathcal{G}_{\tilde{\varphi}_2}(\zeta) - \Omega.$$

The conclusion of Proposition 4.1 may hold for *other* points ζ, and adjusting φ_2 to take this into account will be useful later. Say that $\zeta_1 \in \widetilde{S}$ is a $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$–*concurrence point* (or briefly *concurrence point*) if there exists $\zeta_2 \in \widetilde{S}$ so that

$$\mathcal{G}_{\tilde{\varphi}_1}(\zeta_1) - \Omega = \mathcal{G}_{\tilde{\varphi}_2}(\zeta_2) - \Omega$$

where Ω is as in Proposition 4.1. In particular, the proposition says that every point in $\widetilde{C}_0$ is a concurrence point. We note that since Ω and $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ are $\pi_1 S$–invariant, if ζ is a concurrence point, then so is $\gamma \cdot \zeta$, for any $\gamma \in \pi_1 S$. We say $\widetilde{C} \subset \widetilde{S}$ is a *concurrency set* (for $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$), if $\widetilde{C}$:

1. consists of $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$–concurrence points,
2. contains $\widetilde{C}_0$,
3. is $\pi_1 S$–invariant, and
4. contains only finitely many $\pi_1 S$–orbits of points.

If $\widetilde{C}$ is a concurrency set, then we let $C = p(\widetilde{C})$, and also call $C \subset S$ a concurrency set (by $\pi_1 S$–invariance, note that $\widetilde{C} = p^{-1}(C)$).

As we did with the cone points, note that for any concurrency set $\widetilde{C}$ we can adjust φ_2 by a homeomorphism isotopic to the identity such that for all $\zeta \in \widetilde{C}$ we have

$$\mathcal{G}_{\tilde{\varphi}_1}(\zeta) - \Omega = \mathcal{G}_{\tilde{\varphi}_2}(\zeta) - \Omega.$$

We will say that φ_2 has been *C–normalized* with respect to φ_1. When the reference to φ_1 is obvious, we will simply say that φ_2 has *C–normalized*. We summarize the above discussion for later reference:

Lemma 4.2. *Assume $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and let $\widetilde{C}$ be any concurrency set for $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$. Then $\tilde{\varphi}_2 \sim \tilde{\varphi}_2'$ for a C–normalized $\tilde{\varphi}_2' \in \widetilde{\mathrm{Hyp}}_c(S)$; that is, $\tilde{\varphi}_2'$ satisfies*

$$\mathcal{G}_{\tilde{\varphi}_1}(\zeta) - \Omega = \mathcal{G}_{\tilde{\varphi}_2'}(\zeta) - \Omega$$

for all $\zeta \in \widetilde{C}$. □

The upshot of this lemma is that, in order to prove the Current Support Theorem, we can assume that we have adjusted φ_2 such that it is *C–normalized*, for some concurrency set C. In fact, for most of the proof of the Current Support Theorem, we can assume our concurrency set is just the set of cone points, $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2)$. However, later we will find additional concurrency points, and rather than returning to all of the setup and preliminaries, we will proceed with an arbitrary concurrency set.

Now, let $\widetilde{C} \subset \widetilde{S}$ be a concurrency set for $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$, assume φ_2 is *C–normalized*, and let η be a $\tilde{\varphi}_i$–geodesic. As before, we say that η is singular if it contains a cone point and nonsingular otherwise. We say that η is *C–singular* if it contains a point of $\widetilde{C}$, and *non-C–singular* otherwise. Let

$$\mathcal{G}(\tilde{\varphi}_i, \Omega) = \partial_{\varphi_i}^{-1}(\mathcal{G}_{\tilde{\varphi}_i} - \Omega) \subset \mathcal{G}(\tilde{\varphi}_i)$$

for $i = 1, 2$. The homeomorphism $g : \mathcal{G}(\tilde{\varphi}_1) \to \mathcal{G}(\tilde{\varphi}_2)$ restricts to a homeomorphism

$$g \colon \mathcal{G}(\tilde{\varphi}_1, \Omega) \to \mathcal{G}(\tilde{\varphi}_2, \Omega)$$

Note that by Proposition 4.1, $\eta \in \mathcal{G}(\tilde{\varphi}_1, \Omega)$ is nonsingular if and only if $g(\eta)$ is nonsingular. Moreover, η passes through $\zeta \in \widetilde{C}$ if and only if $g(\eta)$ does (assuming, as we have, that φ_2 is *C–normalized*) by Lemma 4.2. It is worth pointing out that, a priori, $\mathcal{G}(\tilde{\varphi}_1, \Omega)$ might not contain all nonsingular $\tilde{\varphi}_1$–geodesics because of information loss in Ω: there is the potential for (at most countably many) nonsingular $\tilde{\varphi}_1$–geodesics to be missing from $\mathcal{G}(\tilde{\varphi}_1, \Omega)$ (though this ultimately turns out not to be the case). Below we will need to approximate basic geodesics by sequences of nonsingular geodesics in $\mathcal{G}(\tilde{\varphi}_1, \Omega)$ and we note that we can always do this: any basic geodesic is, by definition, a limit of nonsingular geodesics and note that the proof of Lemma 2.2 actually shows

that there are uncountably many sequences of nonsingular geodesics approximating each basic geodesic.

Let $\eta \in \mathcal{G}(\tilde{\varphi}_i)$ be a basic $\tilde{\varphi}$–geodesic and $\zeta \in \widetilde{C}$. Recall that η divides $\widetilde{S}$ into two half-planes, the positive (left) half plane $\mathcal{H}^+(\eta)$ and the negative (right) half plane $\mathcal{H}^-(\eta)$. If $\zeta \in \widetilde{C}$ is disjoint from η, we say that ζ lies *to the left* of η if it belongs to $\mathcal{H}^+(\eta)$ and that it lies *to the right* if it belongs to $\mathcal{H}^-(\eta)$. If $\zeta \in \eta$, then we say that ζ is *on* η. If $\zeta \in \widetilde{C}_0$ and is on η, then η must make angle π on exactly one side at ζ: if η makes angle π on the left of ζ (respectively, on the right of ζ), and we say that ζ is *on and to the right of η* (respectively, *on and to the left of η*). If $\zeta \in \widetilde{C} \smallsetminus \widetilde{C}_0$ and its on η, then η makes angle π on both sides at ζ, and we say that ζ is *on and splits* η. Therefore, $\eta \in \mathcal{G}(\tilde{\varphi}_i)$ defines a partition of $\widetilde{C}$ into five sets depending on whether a point is to the left of η, to the right of η, on and to the left of η, on and to the right of η, or on and splits η (where any of the last three could determine empty sets). Note that while the labels left, right, $+$, $-$ are determined by arbitrarily choosing an orientation on η, the partition itself is defined by the unoriented geodesic. We will see that g preserves this partition (cf. [9, Lemma 12]). We start with the following special case.

Lemma 4.3. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, Ω is as in Proposition 4.1, $\widetilde{C} \subset \widetilde{S}$ is a concurrency set, and φ_2 is C–normalized. Each nonsingular $\eta \in \mathcal{G}(\tilde{\varphi}_1, \Omega)$ defines the same partition of $\widetilde{C}$ as $g(\eta)$ does.*

Proof. Let $\eta \in \mathcal{G}(\tilde{\varphi}_1, \Omega)$ be nonsingular and fix an orientation on η and recall that this induces an orientation on $g(\eta)$. Let $\zeta \in \widetilde{C}$. Note that ζ is on η if and only if it is on $g(\eta)$ by the definition of concurrency points and the assumption that φ_2 is C–normalized. Moreover, it cannot be a cone point since η is nonsingular, and hence ζ splits both η and ζ.

So suppose, without loss of generality, that ζ is to the left of η, i.e. $\zeta \in \mathcal{H}^+(\eta)$. We claim that there exists a geodesic $\delta \in \mathcal{G}(\tilde{\varphi}_1, \Omega)$ passing through ζ and contained in $\mathcal{H}^+(\eta)$.

To see that such δ exists, let μ be a minimal length geodesic path from ζ to η meeting η at a point x (thus μ meets η orthogonally if x is not a cone point, and with angle at least $\pi/2$ on both sides if it is a cone point, cf. Fig. 6). Construct δ by concatenating two basic geodesic rays emanating from ζ and orthogonal to μ such that it makes angle π at ζ on the same side as x is on (see the image at left in Fig. 6). Then δ and η do not intersect, since if they did, this would create a geodesic triangle with angle sum greater than π which is not possible. Furthermore, δ cannot be asymptotic to η since this would create a partially ideal triangle with angle sum π. Therefore, δ and η have distinct unlinked end points in S_∞^1. If necessary we perturb the angle at which δ meets μ while maintaining unlinked distinct endpoints from η so that δ is in $\mathcal{G}(\tilde{\varphi}_1, \Omega)$.

Now, note that $g(\delta)$ also passes through ζ and that δ and $g(\delta)$ have the same endpoints as η and $g(\eta)$. If ζ is to the right of $g(\eta)$ there would be a bigon bounded by subsegments of $g(\eta)$ and $g(\delta)$ as seen in Fig. 6, a contradiction. Hence ζ lies to the left of $g(\eta)$. $\quad\square$
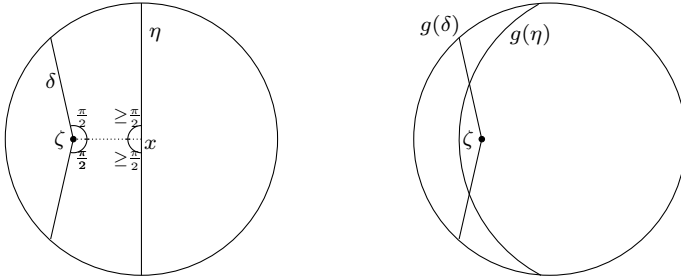
Fig. 6. Left: The point $\zeta \in \widetilde{C}$ and basic geodesic δ lie in $\mathcal{H}^+(\eta)$. **Right:** The bigon created if $g(\eta)$ did not preserve the partition.

The next lemma proves a stronger property for concurrence points and extends Lemma 4.3.

Lemma 4.4. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, Ω is as in Proposition 4.1, $\widetilde{C} \subset \widetilde{S}$ is a concurrency set, and φ_2 is C–normalized. For any $\zeta \in \widetilde{C}$ we have $\mathcal{G}_{\tilde{\varphi}_1}(\zeta) = \mathcal{G}_{\tilde{\varphi}_2}(\zeta)$. Moreover, any $\eta \in \mathcal{G}(\tilde{\varphi}_1)$ defines the same partition of $\widetilde{C}$ (into five sets) as $g(\eta)$ does.*

Proof. Let $\zeta \in \widetilde{C}$ and let $\eta \in \mathcal{G}(\tilde{\varphi}_1, \zeta)$ be an oriented basic $\tilde{\varphi}_1$–geodesic passing through ζ. Choose $\{\eta_n\} \subset \mathcal{G}(\tilde{\varphi}_1, \Omega)$ to be a sequence of nonsingular oriented $\tilde{\varphi}_1$–geodesics limiting to η. If ζ is a cone point which is, say, on and to the left of η, we have that ζ lies to the left of η_n for all large n. If ζ is not a cone point we can choose a sequence of nonsingular geodesics η_n that converge to η such that ζ lies to the left of η_n for all n (see proof of Lemma 2.2 for a construction of such a sequence).

Now, $\{g(\eta_n)\}$ is a sequence of nonsingular geodesics which limits to $g(\eta)$ and ζ lies to the left of each $g(\eta_n)$ by Lemma 4.3, and hence ζ lies to the left of $g(\eta)$ or is on $g(\eta)$. Note that η and $g(\eta)$ have the same endpoints. Suppose ζ is *not* on $g(\eta)$. Then there exists (an oriented) $\delta \in \mathcal{G}(\tilde{\varphi}_2, \Omega)$ such that ζ is to the left of δ (it is in $\mathcal{H}^+(\delta)$) and $g(\eta)$ is to the right of δ (it is contained in $\mathcal{H}^-(\delta)$). Such a geodesic can be constructed by a similar method as was used in the proof of Lemma 4.3: take the minimal length geodesic path μ from ζ to $g(\eta)$ meeting $g(\eta)$ at a point x (orthogonally if x is not a cone point). For any point y in the interior of μ that is not a cone point, let δ_y be any basic $\tilde{\varphi}_2$–geodesic intersecting μ orthogonally at y, and if necessary, perturb y slightly to avoid the countably many geodesics not in $\mathcal{G}(\tilde{\varphi}_2, \Omega)$; the resulting geodesic $\delta = \delta_y$ cannot intersect $g(\eta)$ since if it did it would create a geodesic triangle with angle sum at least π. Again by Lemma 4.3, $g^{-1}(\delta) \in \mathcal{G}(\tilde{\varphi}_1, \Omega)$ also has ζ on the left (and has the same endpoints as δ). However, since η passes through ζ we have then that η and $g^{-1}(\delta)$ contain segments bounding a geodesic bigon, a contradiction. Hence we must have that ζ is on $g(\eta)$ as desired. Therefore, $\mathcal{G}_{\tilde{\varphi}_1}(\zeta) \subset \mathcal{G}_{\tilde{\varphi}_2}(\zeta)$. The other containment is proved by a symmetric argument, and thus $\mathcal{G}_{\tilde{\varphi}_1}(\zeta) = \mathcal{G}_{\tilde{\varphi}_2}(\zeta)$.

The second claim is proved similarly, appealing partially to the first part of the lemma. Let $\eta \in \mathcal{G}(\tilde{\varphi}_1)$ be an oriented geodesic and $\zeta \in \widetilde{C}$. If ζ lies *on* η then by the first argument,

ζ also lies *on* $g(\zeta)$. If $\zeta \in \widetilde{C} \smallsetminus \widetilde{C}_0$, then ζ splits η if and only if $g(\eta)$ splits η. So suppose ζ is either to the left of η or on and to the left of η. As above, we can approximate η with nonsingular geodesics $\{\eta_n\} \subset \mathcal{G}(\tilde{\varphi}_1, \Omega)$ such that ζ lies to the left of each η_n. Then ζ lies to the left of each $g(\eta_n)$ by Lemma 4.3, and since $g(\eta_n)$ converges to $g(\eta)$, it must be that either ζ is to the left of $g(\eta)$ or on and to the left of $g(\eta)$. Since ζ is on η if and only if its on $g(\eta)$ (by the first part of the lemma), it follows that ζ is on and to the left of η if and only if it is on and to the left of $g(\eta)$. Consequently, ζ is to the left of η if and only if ζ is to the left of $g(\eta)$. The case of ζ to the right of η or on and to the right of η is handled in exactly the same way. This completes the proof. $\square$

Remark. According to Lemma 4.4 a $(\tilde{\varphi}_1, \tilde{\varphi}_2, \Omega)$–concurrency point ζ has the property that (after normalizing φ_2) all basic $\tilde{\varphi}_1$–geodesics which contain it straighten to $\tilde{\varphi}_2$–geodesics that contain it. Thus we are justified in simply referring to ζ as a $(\tilde{\varphi}_1, \tilde{\varphi}_2)$–concurrency point (and likewise referring to concurrency sets for $(\tilde{\varphi}_1, \tilde{\varphi}_2)$) without reference to Ω. Indeed, going forward Ω will play no further role and so we discard it from the discussion.

4.2. Saddle connections, rays, and triangles

For the discussion in the remainder of this subsection, assume $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}_c}(S)$ with $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, we fix a concurrency set C for (φ_1, φ_2), and assume φ_2 has been C–normalized.

Given a C–singular geodesic η, choosing an orientation of η induces an order on the points of $\widetilde{C}$ it encounters. We say that the (oriented) geodesics $\eta \in \mathcal{G}(\tilde{\varphi}_1)$ and $g(\eta) \in \mathcal{G}(\tilde{\varphi}_2)$ have the *same combinatorics* (with respect to $\widetilde{C}$) if they pass through exactly the same set of points of $\widetilde{C}$ and in the same order. The goal of this subsection is to show that η and $g(\eta)$ have the same combinatorics for all $\eta \in \mathcal{G}(\tilde{\varphi}_1)$. Observe that by Lemma 4.4, η and $g(\eta)$ encounter the same set of points of $\widetilde{C}$, so we will need only show that the order in which they encounter the concurrence points is preserved by g. First we introduce some more terminology.

We say a $\tilde{\varphi}_i$–geodesic ray is a $(C, \tilde{\varphi}_i)$–*ray* (or just a C-*ray* if the metric is understood) if its initial point is in $\widetilde{C}$ but its interior is disjoint from $\widetilde{C}$. We say that $\zeta \in \widetilde{C}$ and $x \in S^1_\infty$ determine a $(C, \tilde{\varphi}_i)$–ray if the $\tilde{\varphi}_i$–geodesic ray from ζ to x is a $(C, \tilde{\varphi}_i)$–ray.

Let $\zeta, \xi \in \widetilde{C}$ and δ be a $\tilde{\varphi}_i$–geodesic segment connecting ζ and ξ. If δ contains no point of $\widetilde{C}$ in its interior, we say that δ is a $(C, \tilde{\varphi}_i)$–*saddle connection*, or just a C–*saddle connection* when the metric is clear. We say that ζ, η determine a $(C, \tilde{\varphi}_i)$–saddle connection when the $\tilde{\varphi}_i$–geodesic segment between them is a $(C, \tilde{\varphi}_i)$–saddle connection. A (C, φ_i)–saddle connection in S is the image of a $(C, \tilde{\varphi}_i)$–saddle connection in $\widetilde{S}$.

Lemma 4.5. *If $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, $\widetilde{C}$ is a concurrency set, φ_2 is C–normalized, and $\zeta, \xi \in \widetilde{C}$, $x \in S^1_\infty$, then:*

1. ζ, ξ *determine a* $(C, \tilde{\varphi}_1)$*–saddle connection if and only if they determine a* $(C, \tilde{\varphi}_2)$*–saddle connection.*
2. ζ, x *determine a* $(C, \tilde{\varphi}_1)$*–ray if and only if they determine a* $(C, \tilde{\varphi}_2)$*–ray.*

Proof. Let δ_1 be a $(C, \tilde{\varphi}_1)$–saddle connection between ζ and ξ and δ_2 the unique φ_2–geodesic segment between ζ and ξ. We need to show that δ_2 contains no points of $\widetilde{C}$ in its interior. Suppose to the contrary that there were some $\omega \in \widetilde{C}$, distinct from ζ, ξ on δ_2. Since ω is not on δ_1, we can find an oriented, non-C–singular $\tilde{\varphi}_1$–geodesic $\eta \in \mathcal{G}(\tilde{\varphi}_1)$ so that ζ, ξ are to the left of δ_1 and ω is to the right. By Lemma 4.4, $g(\eta)$ is non-C–singular and ζ, ξ are to the left and ω is to the right. But then subsegments of δ_2 and $g(\eta)$ bound a bigon, which is a contradiction. Therefore, δ_2 is a $(C, \tilde{\varphi}_2)$–saddle connection. A symmetric argument proves the other implication.

The proof of the second claim is similar. Suppose r_i is the $\tilde{\varphi}_i$–ray from ζ to x and assume r_1 is a C–ray while r_2 contains some point $\omega \in \widetilde{C}$ different than ζ. There is an oriented non-C–singular $\tilde{\varphi}_1$–geodesic $\eta \in \mathcal{G}(\tilde{\varphi}_1)$ so that r_1 is to the left of η *and does not share an endpoint with it* while ω is to the right. By Lemma 4.4 $g(\eta)$ is non-C–singular, ζ is to the left, and ω is to the right, but r_2 must therefore travel from the left of $g(\eta)$ starting at ζ, to the right passing through ω, and back to the left again to limit to x. Therefore r_2 and $g(\eta)$ contain segments that bound a bigon, which is a contradiction. So r_2 is a C–ray. Again a symmetric argument proves the other implication. $\square$

If r is a $(C, \tilde{\varphi}_1)$–ray from $\zeta \in \widetilde{C}$ to $x \in S^1_\infty$, then Lemma 4.5 implies there is a $(C, \tilde{\varphi}_2)$–ray from ζ to x, which we suggestively denote $g(r)$. Similarly, the lemma implies that for each $(C, \tilde{\varphi}_1)$–saddle connection δ between a pair $\zeta, \xi \in \widetilde{C}$ there is a $(C, \tilde{\varphi}_2)$–saddle connection we denote $g(\delta)$ between ζ, ξ as well.

Corollary 4.6. *Suppose* $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, $\widetilde{C}$ *is a concurrency set, and* φ_2 *is* C*–normalized. Then for every* $\eta \in \mathcal{G}(\tilde{\varphi}_1)$, *we have that* η *and* $g(\eta)$ *have the same combinatorics. Equivalently,* η *is non-*C*–singular if and only if* $g(\eta)$ *is, and if* η *is a (finite, infinite, bi-infinite) concatenation*

$$\eta = \dots \delta_i \cdot \delta_{i+1} \cdot \delta_{i+2} \dots$$

of $(C, \tilde{\varphi}_1)$*–saddle connections and/or* $(C, \tilde{\varphi}_1)$*–rays, then*

$$g(\eta) = \dots g(\delta_i) \cdot g(\delta_{i+1}) \cdot g(\delta_{i+2}) \dots.$$

Proof. Given an oriented $\eta \in \mathcal{G}(\tilde{\varphi}_1)$, if η is non-C–singular, then by Lemma 4.4, so is $g(\eta)$. If η is a nontrivial concatenation $\eta = \dots \delta_i \cdot \delta_{i+1} \dots$, then all the points of $\widetilde{C}$ that appear as endpoints of the saddle connections/rays of this concatenation also appear in $g(\eta)$ by Lemma 4.4. In particular, each $g(\delta_i)$ is a $(C, \tilde{\varphi}_2)$–saddle connection or $(C, \tilde{\varphi}_2)$–ray contained in $g(\eta)$. If the points of $\widetilde{C}$ did not appear in the same order along $g(\eta)$ as they
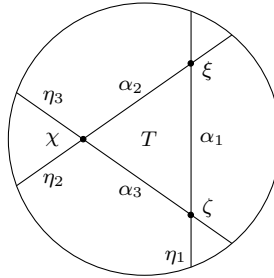
Fig. 7. The points ζ, ξ are on and to the right of η_1, while χ is to the left of η_1. The intersection of the positive half-planes is the triangle: $\mathcal{H}^+(\eta_1) \cap \mathcal{H}^+(\eta_2) \cap \mathcal{H}^+(\eta_3) = T$.

did in η, then $g(\delta_i)$ would overlap with some $g(\delta_j)$ for some $i \neq j$, contradicting the fact that these are $(C, \tilde{\varphi}_2)$–saddle connections and/or $(C, \tilde{\varphi}_2)$–rays, by Lemma 4.5. $\quad\square$

We say that T is a $(C, \tilde{\varphi}_i)$-*triangle* if it is a $\tilde{\varphi}_i$–geodesic triangle such that the vertices lie in $\widetilde{C}$, its edges are C–saddle connections, and there are no points of $\widetilde{C}$ in its interior. Note that, since $\widetilde{C}_0 \subset \widetilde{C}$, any $(C, \tilde{\varphi}_i)$–triangle is isometric to a hyperbolic triangle in $\mathbb{H}$. If the covering map p restricted to a $(C, \tilde{\varphi}_i)$–triangle is injective on the interior, then we call the image in S a (C, φ_i)–triangle (in fact, in this case, the negative curvature of φ_i implies that the restriction of p can only fail to be injective on the vertices). We say that three points of $\widetilde{C}$ *determine a* $(C, \tilde{\varphi}_i)$–*triangle* if each pair determines a $(C, \tilde{\varphi}_i)$–saddle connection, and the region bounded by these C–saddle connections is such a triangle.

Lemma 4.7. *Suppose* $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, $\widetilde{C}$ *is a concurrency set, and* φ_2 *is* C–*normalized. If* $\zeta, \xi, \chi \in \widetilde{C}$ *determine a* $(C, \tilde{\varphi}_1)$–*triangle then they also determine a* $(C, \tilde{\varphi}_2)$–*triangle, and the two triangles have the same orientation.*

Proof. Let T be the $(C, \tilde{\varphi}_1)$–triangle determined by $\zeta, \xi, \chi \in \widetilde{C}$, and orient it positively (its interior is to the left of its boundary) and suppose this induces the cyclic orientation $\zeta < \xi < \chi < \zeta$ on its vertices. Let α_1, α_2, and α_3 be the $(C, \tilde{\varphi}_1)$–saddle connection from ζ and ξ, from ξ to χ, and from χ to ζ, respectively. By Lemma 4.5, these determine $(C, \tilde{\varphi}_2)$–saddle connections $\alpha'_1 = g(\alpha_1)$, $\alpha'_2 = g(\alpha_2)$, and $\alpha'_3 = g(\alpha_3)$, and these bound a triangular region T'. We need to show that T' has no points of $\widetilde{C}$ in its interior, and that it has the same orientation as T.

Extend $\alpha_1, \alpha_2, \alpha_3$ to oriented basic $\tilde{\varphi}_1$–geodesics η_1, η_2, η_3 such that each of the three points ζ, ξ, χ lie to the left or on and to the right of each of the geodesics; see Fig. 7. This choice ensures that for each $i \neq j$, η_i and η_j cross and intersect in a single point (namely one of ζ, ξ, χ). By Corollary 4.6 $g(\eta_i)$ is a basic geodesic which has $g(\alpha_i)$ as a subsegment for each $i = 1, 2, 3$.

Note that the intersection of the three left half-planes $\mathcal{H}^+(\eta_1), \mathcal{H}^+(\eta_2)$, and $\mathcal{H}^+(\eta_3)$ is T and hence contains no other points of $\widetilde{C}$ besides the vertices ζ, ξ, χ. By Lemma 4.4 (preservation of partitions), the same is true for the intersection of $\mathcal{H}^+(g(\eta_1)), \mathcal{H}^+(g(\eta_2))$,

and $\mathcal{H}^+(g(\eta_3))$. However, by the above, this intersection is exactly T'. Hence T' is indeed a $(C, \tilde{\varphi}_2)$–triangle. By Corollary 4.6, ζ precedes ξ along $g(\eta_1)$ and since χ is to the left of $g(\eta_1)$, T' has the same orientation as T. $\square$

We end this subsection with one more fact we will need for future reference.

Lemma 4.8. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, $\widetilde{C}$ is a concurrency set and assume φ_2 is C–normalized. Let η be a basic $\tilde{\varphi}_1$–geodesic or a $(C, \tilde{\varphi}_1)$–saddle connection. If η transversely crosses the interior of a $(C, \tilde{\varphi})$–saddle connection δ, then $g(\eta)$ transversely crosses the interior of $g(\delta)$.*

Proof. Extend η and δ to basic $\tilde{\varphi}_1$–geodesics η' and δ' (where $\eta' = \eta$ in the case η is a basic geodesic to start). Since η' and δ' cross, their endpoints link, and since $g(\eta')$ and $g(\delta')$ have the same endpoints as η' and δ', respectively, $g(\eta')$ and $g(\delta')$ also cross.

Suppose $\zeta, \xi \in \widetilde{C}$ are the endpoints of δ. Note that ζ and ξ are disjoint from and lie on opposite sides of η'. Hence, since g preserves partitions and concurrency, ζ and ξ are disjoint from and lie on opposite sides of $g(\eta')$. It follows that $g(\eta')$ transversely crosses $g(\delta')$ at a point on the interior of the segment connecting ζ and ξ, which is exactly $g(\eta)$ by Corollary 4.6.

Now, if $\eta = \eta'$ is a geodesic we are done. If η is a saddle connection, say with endpoints ζ', ξ' we get by the same argument as above that a basic extension of δ partition ζ' and ξ' on different sides and, as above, we have that the intersection must happen at a point belonging to the saddle connection η. $\square$

4.3. Triangulations

In this section we continue with the assumption that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ for some $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}_c}(S)$, and that φ_2 has been C–normalized for a concurrency set C. We will explain how to adjust φ_2 to an equivalent metric φ_2' and find a *common triangulation* with vertices in C and edges which are (C, φ_i)–saddle connections for $i = 1$ *and* 2 (see below for precise definitions). We note that the isotopy to the identity arising in the adjustment of φ_2 may nontrivially braid C. Once this adjustment is made, however, we will show that for any other triangulation, the analogous adjustment of φ_2' can be done *without* additional braiding of C.

Now suppose $\mathcal{T}$ is a (C, φ_1)–*triangulation* (or just φ_1–triangulation, if C is understood). This is a Δ–complex structure $\mathcal{T}$ for S (in the sense of [14]) so that the 0–skeleton is precisely C, and so that each edge is a (C, φ_1)–saddle connection. In particular, each 2–cell is a (C, φ_1)–triangle as defined above.

We define a map $f_{\mathcal{T}} \colon S \to S$ to be the descent to S of a $\pi_1 S$–equivariant map $\tilde{f}_{\mathcal{T}} \colon \widetilde{S} \to \widetilde{S}$ which is the identity on $\widetilde{C}$ and which sends each $\tilde{\varphi}_1$–triangle T_1 to the straightened $\tilde{\varphi}_2$–triangle T_2. By Lemma 4.7, T_2 is also isometric to a geodesic triangle in the hyperbolic plane, and $\tilde{f}_{\mathcal{T}}$ is defined to be the canonical triangle map from T_1

to T_2 (conjugated by the isometries to the hyperbolic triangles). Equivariance of $\tilde{f}_{\mathcal{T}}$ is immediate.

Proposition 4.9. *The map $f_{\mathcal{T}}$ is a homeomorphism, isotopic to the identity on S.*

Proof. Since $\tilde{f}_{\mathcal{T}}$ is the identity on $\widetilde{C}$ and is equivariant, $(f_{\mathcal{T}})_*$ is the identity, and so $f_{\mathcal{T}}$ has degree 1 and is homotopic to the identity. Since it maps each oriented triangle to a similarly oriented triangle, it is in fact a homeomorphism. Furthermore, since isotopy and homotopy are the same equivalence relation for homeomorphisms for closed surfaces, $f_{\mathcal{T}}$ is isotopic to the identity, as required. $\square$

Given a (C, φ_1)–triangulation $\mathcal{T}$, we may adjust the metric φ_2 by pulling back via $f_{\mathcal{T}}$ to produce an equivalent metric. Having done so, and replacing φ_2 with the equivalent metric we say that φ_2 has been $(C, \mathcal{T})$–*normalized* and we note that $\mathcal{T}$ is simultaneously a (C, φ_1)–triangulation and a (C, φ_2)–triangulation, and the map $f_{\mathcal{T}}$ is the identity. In this case we say that $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–*triangulation*. In other words, we have the following result:

Lemma 4.10. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_1}$. If $\widetilde{C}$ is any concurrency set and $\mathcal{T}$ a (C, φ_1)–triangulation, then $\varphi_2 \sim \varphi_2'$ for some $(C, \mathcal{T})$–normalized φ_2'. That is, up to equivalence, we can always assume that $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–triangulation.* $\square$

Having adjusted φ_2 so that $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–triangulation, the next lemma says that this adjustment "anchors" the cone points in relation to the geodesics.

Lemma 4.11. *Suppose $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–triangulation. For any $\eta \in \mathcal{G}(\tilde{\varphi}_1)$, $g(\eta)$ is isotopic to η by an isotopy $h_t \colon \widetilde{S} \to \widetilde{S}$ for $t \in [0, 1]$ that preserves the triangles of $\mathcal{T}$. The same is true for any $(C, \tilde{\varphi}_1)$–saddle connection η.*

Proof. Let η be any basic $\tilde{\varphi}_1$–geodesic or $(C, \tilde{\varphi}_i)$–saddle connection. For any triangle T_1 of $\widetilde{\mathcal{T}}$, we have that $\eta \cap T_1$ is either empty, a vertex of T_1, a side of T_1, or an arc cutting through the interior of T_1, by convexity. In the last case, the arc may run between a vertex and the interior of a side or from the interior of one side to the interior of another. If $\eta \cap T_1$ is not an arc, then by Lemmas 4.4 and 4.5 $g(\eta) \cap T_1 = \eta \cap T_1$. On the other hand, if $\eta \cap T_1$ is an arc cutting through the interior of T_1, then by Lemma 4.8, $g(\eta) \cap T_1$ is also an arc connecting the same vertex to opposite sides, or the same pair of sides.

Now we produce the required isotopy h_t in two steps. The first step takes place during times $t \in [0, \frac{1}{2}]$ and the second step takes place during times $t \in [\frac{1}{2}, 1]$. We define h_t for $t \in [0, \frac{1}{2}]$ so that it is the identity outside a small neighborhood of the 1–skeleton $\widetilde{\mathcal{T}}^{(1)}$, has $h_t(\widetilde{\mathcal{T}}^{(1)}) = \widetilde{\mathcal{T}}^{(1)}$, and so that $h_{\frac{1}{2}}(\eta \cap \widetilde{\mathcal{T}}^{(1)}) = g(\eta) \cap \widetilde{\mathcal{T}}^{(1)}$. Now for any triangle T_1 of $\widetilde{\mathcal{T}}$, $h_{\frac{1}{2}}(\eta) \cap T_1 = g(\eta) \cap T_1$, or else $h_{\frac{1}{2}}(\eta) \cap T_1$ and $g(\eta) \cap T_1$ are arcs with the same endpoints. It is easy to complete the isotopy so that for $t \in [\frac{1}{2}, 1]$, h_t is the identity on $\widetilde{\mathcal{T}}^{(1)}$ and so that $h_1(\eta) \cap T_1 = g(\eta) \cap T_1$ for every triangle T_1. $\square$

If an isotopy is the identity on a set C throughout the isotopy, then we say that it is an *isotopy relative to C*, and that homeomorphism differing by such an isotopy is *isotopic relative to C*.

Corollary 4.12. *Suppose $\mathcal{T}$ is a (C, φ_1)–triangulation and that φ_2 has been $(C, \mathcal{T})$–normalized. If $\mathcal{T}'$ is another (C, φ_1)–triangulation, then the associated homeomorphism $f_{\mathcal{T}'} \colon S \to S$ is isotopic to the identity relative to C.*

Proof. Each (C, φ_1)–saddle connection making up the edges of $\mathcal{T}'$ is homotopic *in the complement of C* to its (C, φ_2)–straightening (since their lifts are isotopic in $\widetilde{S}$ fixing $\widetilde{C}$). Homotopy of arcs in the complement of the cone points implies isotopy of said arcs (fixing C throughout the isotopy). Since the saddle connections making up the edges of the triangulation have pairwise disjoint interiors, there is an isotopy of the identity, relative to C, that sends each (C, φ_1)–saddle connection of $\mathcal{T}'$ to a (C, φ_2)–saddle connection. After a further isotopy preserving the 1–skeleton of $\mathcal{T}'$, we obtain a map which is the canonical triangle map on each triangle of $\mathcal{T}'$, and is thus $f_{\mathcal{T}'}$. $\square$

We say that φ_2 is *C–uber-normalized* (with respect to φ_1) if for any (C, φ_1)–triangulation $\mathcal{T}$ of S, the isotopy to the identity involved in adjusting φ_2 so $\mathcal{T}$ is a (C, φ_2)–triangulation fixes C. By the corollary, if φ_2 is $(C, \mathcal{T})$–normalized, then it is C–uber-normalized. We record the following fact about the examples from §3 which will be useful later.

Lemma 4.13. *Suppose $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ are obtained from branched covers $q \colon (S, \varphi_1) \to \mathcal{O}_1$ and $F \circ q \colon (S, \varphi_2) \to \mathcal{O}_2$, for some homeomorphism $F \colon \mathcal{O}_1 \to \mathcal{O}_2$ as in Theorem 3.1, so that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. Then the preimage of the even order orbifold points, $C = q^{-1}(\mathcal{E}_1)$, is a concurrency set and φ_2 is C–uber-normalized with respect to φ_1.*

Proof. As in the proof of Theorem 3.1, we adjust F by a homeomorphism isotopic to the identity so that $\mathcal{O}_1 \to \mathcal{O}_2$ is the homeomorphism of Lemma 3.4. Since an isotopy on an orbifold preserves the orbifold points, the lift of the isotopy is an isotopy of the identity $S \to S$ relative to C. It follows that the adjustment to F changes φ_2 within its equivalence class by a homeomorphism isotopic to the identity relative to C. Therefore, it suffices to prove the theorem under this assumption. From the proof of Theorem 3.1, we have a cell structure on S with vertices in C for which the 2–cells are convex hyperbolic polygons with respect to both φ_1 and φ_2. As in the proof of Theorem 3.1, pulling back these cell structures to $\widetilde{S}$ with vertex set $\widetilde{C} = p^{-1}(C)$ we found that for every nonsingular $\tilde{\varphi}_1$–geodesic η, the $\tilde{\varphi}_2$–straightening $g(\eta)$ is nonsingular and it intersects the exact same set of cells as η (this is proved in $\hat{S}$, but this property pushes down to $\widetilde{S}$). Consequently, η and $g(\eta)$ define the same partition of the set $\widetilde{C}$. Now we may make a nearly identical argument to the one in the proof of Lemma 4.4 to see that $\mathcal{G}_{\tilde{\varphi}_1}(\zeta) = \mathcal{G}_{\tilde{\varphi}_2}(\zeta)$ for all $\zeta \in \widetilde{C}$, proving that $\widetilde{C}$ is a concurrency set (by assumption, $\widetilde{C}$ contains all cone points) and that φ_2 is C–normalized with respect to φ_1.

Now subdivide each 2–cell into (C, φ_1)–triangles. We can adjust φ_2 by a homeomorphism isotopic to the identity relative to C so that the triangles are also (C, φ_2)–triangles, and the identity is the canonical map. Since the isotopy is relative to C, the new metric is C–uber-normalized if and only if the original one is. However, the new metric is C–uber-normalized by Corollary 4.12. This completes the proof. $\square$

5. Holonomy

A key ingredient in the proof of the Current Support Theorem is Proposition 5.4 below, which states that if $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ (and φ_2 is suitable normalized), then there is an orientation preserving homeomorphism $h \colon \partial\mathbb{H} \to \partial\mathbb{H}$ topologically conjugating the associated holonomy homomorphisms ρ_1 to ρ_2 (see §2.5). The majority of this section is devoted to proving this proposition, and ends by showing that if the conjugating homeomorphism h is in $\mathrm{PSL}_2(\mathbb{R})$, then $\varphi_1 \sim \varphi_2$.

For the remainder of this section we shall assume $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ for $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$, that C is a concurrency set for (φ_1, φ_2), and that φ_2 has been C–uber-normalized. In particular we have $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2) \subseteq C$.

Recall from §2.2 that $\hat{S}$ is the completion of the universal cover of $\dot{S} = S \setminus C_0$ with respect to (either of) the pulled back metrics $\hat{\varphi}_i$ of φ_i, for $i = 1, 2$, and that the universal covering map extends to the completion which we denote by $\hat{p} \colon \hat{S} \to S$. Because φ_2 is C–uber-normalized, for any (C, φ_1)–triangulation $\mathcal{T}$ we may adjust φ_2 by a homeomorphism isotopic to the identity by an isotopy which fixes C, so that $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–triangulation. Since $C_0 \subseteq C$, the C–uber-normalization implies that the isotopy lifts to $\hat{S}$, and thus $\mathcal{G}_{\hat{\varphi}_2}$ does not change under such an adjustment. The developing map for $\hat{\varphi}_i$ extends to the completion $D_i \colon \hat{S} \to \mathbb{H}$, and is equivariant with respect to the holonomy homomorphism $\rho_i \colon \pi_1 \dot{S} \to \mathrm{PSL}_2(\mathbb{R})$, for $i = 1, 2$, where $\dot{S} = S \setminus C_0$; see §2.4.

5.1. Combinatorics of basic geodesics in $\hat{S}$

Analogous to the definition for $\tilde{\varphi}_i$–geodesics in $\widetilde{S}$, we say that a $\hat{\varphi}_i$–geodesic is *non-C–singular* if it is disjoint from $\hat{C} = \hat{p}^{-1}(C)$ (and simply nonsingular if it is disjoint from the completion points $\hat{C}_0 = \hat{p}^{-1}(C_0)$). A $(C, \hat{\varphi}_i)$–saddle connection is a $\hat{\varphi}_i$–geodesic segment between two points of $\hat{C}$ such that its interior is disjoint from $\hat{C}$. Also, recall that $\mathcal{G}(\hat{\varphi}_i)$ denotes the set of all basic $\hat{\varphi}_i$–geodesics, i.e. the closure of the set of nonsingular $\hat{\varphi}_i$–geodesics. Moreover, as we did in $\widetilde{S}$, we say that two basic geodesics $\eta_1 \in \mathcal{G}(\hat{\varphi}_1)$ and $\eta_2 \in \mathcal{G}(\hat{\varphi}_2)$ have the *same combinatorics* if they have the same endpoints in $\hat{S}^1_\infty$ and pass through exactly the same set of points of $\hat{C}$, in the same order.

Lemma 5.1. *Let $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ with $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, C any concurrency set, and suppose $\tilde{\varphi}_2$ is C–uber-normalized. Then $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ and hence we have a well-defined straightening map $\hat{g} \colon \mathcal{G}(\hat{\varphi}_1) \to \mathcal{G}(\hat{\varphi}_2)$. Moreover, for all $\eta \in \mathcal{G}(\hat{\varphi}_1)$, η and $\hat{g}(\eta)$ have the same combinatorics.*

As a consequence of Lemma 5.1, the $\hat{\varphi}_2$–straightening map $\hat{g}$ may be extended to a map from all basic $\hat{\varphi}_1$–geodesic segments and rays (as well as lines), with endpoints (possibly at infinity) in the set $\hat{C} \cup \hat{S}^1_\infty$, to basic $\hat{\varphi}_2$–geodesic segments and rays having the same combinatorics. This is obtained by extending any basic $\hat{\varphi}_1$–geodesic segment or ray to a (biinfinite) basic $\hat{\varphi}_1$–geodesic (see Lemma 2.2), applying $\hat{g}$, then appealing to the lemma to find an appropriate subsegment or ray of the resulting basic $\hat{\varphi}_2$–geodesic.

Proof. Fix a $(C, \varphi_1, \varphi_2)$–triangulation $\mathcal{T}$, adjusting φ_2 to an equivalent metric if necessary. Let $\widetilde{\mathcal{T}}$ be the lifted triangulation to $\widetilde{S}$ and further lift $\widetilde{\mathcal{T}}$ to a triangulation $\widehat{\mathcal{T}}$ in $\hat{S}$. As noted above, because φ_2 was C–uber-normalized, the isotopy to the identity of the homeomorphism used to adjust φ_2 fixes C (and hence C_0) and so lifts to $\hat{S}$. In particular, proving the conclusion of the lemma for the adjusted metric implies it for the original metric.

As in Lemma 2.2, every non-C–singular $\hat{\varphi}_i$–geodesic projects to a non-C–singular geodesic in $\widetilde{S}$ by $\tilde{p}$ since this map is a local isometry (away from the cone points). Let η be a non-C–singular $\hat{\varphi}_1$–geodesic. The $\tilde{\varphi}_2$–straightening $g(\tilde{p}(\eta))$ is isotopic to $\tilde{p}(\eta)$ in the complement of $\widetilde{C}$, by an isotopy that preserves each triangle of $\widetilde{\mathcal{T}}$ by Lemma 4.11. We can lift this isotopy, triangle-by-triangle, to an isotopy from η to a non-C–singular $\hat{\varphi}_2$–geodesic η' passing through the same triangles of $\widehat{\mathcal{T}}$, which therefore has the same endpoints in $\hat{S}^1_\infty$. Any basic $\hat{\varphi}_1$–geodesic η is a limit of nonsingular $\hat{\varphi}_1$–geodesics, and by analyzing the limits in each triangle (and appealing to Corollary 4.6), it follows that there is a basic $\hat{\varphi}_2$–geodesic η' running through precisely the same set of triangles as η (and through the same set of vertices, in the same order). It follows that $\mathcal{G}_{\hat{\varphi}_1} \subset \mathcal{G}_{\hat{\varphi}_2}$ and the straightening map $\hat{g}$ preserves combinatorics. A symmetric argument proves the other inclusion, hence $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$, and we are done. $\quad\square$

The following is an immediate consequence of Lemma 5.1. Recall that $\mathcal{G}(\hat{\varphi}_i, \zeta)$ is the set of all basic $\hat{\varphi}_i$–geodesics through a point $\zeta \in \hat{C}$.

Corollary 5.2. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, C is a concurrency set, and φ_2 is C–uber-normalized. Then we have $\hat{g}(\mathcal{G}(\hat{\varphi}_1, \zeta)) = \mathcal{G}(\hat{\varphi}_2, \zeta)$ for all $\zeta \in \hat{C}$.* $\quad\square$

For this reason, we will refer to the points of $\hat{C}$ as $(\hat{\varphi}_1, \hat{\varphi}_2)$–*concurrence points* and $\hat{C}$ a concurrency set for $(\hat{\varphi}_1, \hat{\varphi}_2)$. Another useful fact is the following, which relates angles between intersecting basic geodesic rays.

Lemma 5.3. *Given oriented basic $\hat{\varphi}_1$–geodesic rays η, η' emanating from a cone point $\zeta \in \hat{C}_0$ making angle $\theta_1 > 0$, let $\theta_2 > 0$ be the angle between $\hat{g}(\eta)$ and $\hat{g}(\eta')$. Then $\frac{\theta_1}{\pi} \in \mathbb{Z}$ if and only if $\frac{\theta_2}{\pi} \in \mathbb{Z}$, and in this case the angles are equal. In general, we have $\left\lfloor \frac{\theta_1}{\pi} \right\rfloor = \left\lfloor \frac{\theta_2}{\pi} \right\rfloor$.*

As a consequence, if two oriented $\hat{\varphi}_1$–geodesics η, η' intersect at a cone point, and $D_1(\eta) = D_1(\eta')$, then $D_2(\hat{g}(\eta)) = D_2(\hat{g}(\eta'))$.

Proof. First note that the angles θ_1 and θ_2 depend only on the initial segments of η and η' (the case of θ_2 requires an application of Lemma 5.1). In particular, θ_1 remains unchanged by modifying the side on which η and η' make angle π at various cone points they meet, since their initial segments are preserved by this modification.

First suppose the angle between η and η' is π. Without loss of generality, we may assume that the angle from η to η' is π (i.e. counterclockwise) and that η makes angle π on the left and η' makes angle π on the right at every cone point they encounter. Then reversing the orientation on η and concatenating with η' we get $\eta \cdot \eta'$ is a basic $\hat{\varphi}_1$–geodesic making angle π on the right at every cone point encountered. Straightening, we see that $\hat{g}(\eta \cdot \eta') = \hat{g}(\eta) \cdot \hat{g}(\eta')$ is also a basic geodesic, implying that $\hat{g}(\eta)$ and $\hat{g}(\eta')$ also make angle π (in fact, it's counter clockwise from $\hat{g}(\eta)$ to $\hat{g}(\eta')$ using the cyclic ordering on $\hat{S}^1_\infty$).

Now given any two rays η and η', without loss of generality assume the angle from η to η' is positive (counterclockwise) so that the angle from $\hat{g}(\eta)$ to $\hat{g}(\eta')$ is also positive. Let $\eta = \eta_0, \eta_1, \ldots, \eta_k, \eta_{k+1}$ be a finite sequence of geodesic rays emanating from ζ so that the angle from η_i to η_{i+1} is π (counterclockwise) and the angle from η_k to η' is $\theta_1^0 \in [0, \pi)$. If $\eta_k = \eta'$ then $\theta_1^0 = 0$, but otherwise, let η_{k+1} be such that the angle from η' to η_{k+1} is $\pi - \theta_0^1 > 0$. Then $\theta_1 = k\pi + \theta_1^0$ and $\theta_1^0 \in [0, \pi)$. By the first part of the proof, the angle from $\hat{g}(\eta_i)$ to $\hat{g}(\eta_{i+1})$ is π. If $\theta_1^0 = 0$, then $\hat{g}(\eta_k) = \hat{g}(\eta')$ and thus $\theta_2 = k\pi = \theta_1$. In this case $\frac{\theta_1}{\pi} = \frac{\theta_2}{\pi}$ and these are integers. On the other hand, if $\theta_1^0 \in (0, \pi)$, then θ_2 must be between $k\pi$ (the measure of the angle from $\hat{g}(\eta_0)$ to $\hat{g}(\eta_k)$) and $(k+1)\pi$ (the measure of the angle from $\hat{g}(\eta_0)$ to $\hat{g}(\eta_{k+1})$). Therefore $\lfloor \frac{\theta_1}{\pi} \rfloor = k\pi = \lfloor \frac{\theta_2}{\pi} \rfloor$.

Finally note that if $\eta, \eta' \in \mathcal{G}(\hat{\varphi}_1, \zeta)$ and $D_1(\eta) = D_1(\eta')$, then η, η' must make angle at ζ which is an integral multiple of π. By the first part of the lemma the same is true for $\hat{g}(\eta), \hat{g}(\eta')$, and thus $D_2(\hat{g}(\eta)) = D_2(\hat{g}(\eta'))$. $\quad\square$

5.2. Conjugating circle actions

The fundamental group, $\pi_1 \dot{S}$, acts on $\hat{S}$ by isometries with respect to $\hat{\varphi}_1$ and $\hat{\varphi}_2$. More generally we consider any group G acting by homeomorphisms on $\hat{S}$ that are isometries with respect to both $\hat{\varphi}_1$ and $\hat{\varphi}_2$. In this situation, there are holonomy homomorphisms $\rho_i : G \to \mathrm{PSL}_2(\mathbb{R})$, for $i = 1, 2$.

The goal of this subsection is to use the developing maps and behavior of basic geodesics to prove the following, which is a key ingredient in the proof of the Current Support Theorem.

Proposition 5.4. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, C is a concurrency set and φ_2 is C–uber-normalized. Suppose a group G acts on $\hat{S}$ by homeomorphisms that are isometries with respect to $\hat{\varphi}_1$ and $\hat{\varphi}_2$. Then there is an orientation preserving homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ topologically conjugating ρ_1 to ρ_2. That is, for all $x \in \partial\mathbb{H}$ and $\gamma \in G$, we have*

$$h(\rho_1(\gamma) \cdot x) = \rho_2(\gamma) \cdot h(x). \tag{2}$$

Moreover, for all $\eta \in \mathcal{G}(\hat{\varphi}_1)$, $h(\partial D_1(\eta)) = \partial D_2(\hat{g}(\eta))$.

We will define the map $h : \partial\mathbb{H} \to \partial\mathbb{H}$ as follows. Suppose $x \in \partial\mathbb{H}$, and take any oriented geodesic α in $\mathbb{H}$ with forward end point x. Let η be any basic $\hat{\varphi}_1$–geodesic with $D_1(\eta) = \alpha$ (see Lemma 2.5) and define $h(x)$ to be the forward endpoint of $D_2(\hat{g}(\eta))$ in $\partial\mathbb{H}$. An immediate concern is whether h is well-defined, and §5.2.1 is dedicated to showing that it is.

5.2.1. Developed asymptotic geodesics

Given oriented basic $\hat{\varphi}_i$–geodesics η_1, η_2, we say that they are *developed forward asymptotic* if $D_i(\eta_1)$ and $D_i(\eta_2)$ are forward asymptotic in $\mathbb{H}$ (note we allow the possibility that $D_i(\eta_1) = D_i(\eta_2)$).

In this section, our goal is to show that if η and η' are basic $\hat{\varphi}_1$–geodesics that are developed forward asymptotic, then $\hat{g}(\eta)$ and $\hat{g}(\eta')$ are developed forward asymptotic as well, which is an important ingredient in verifying that h is well-defined. The proof of this preservation of developed forward asymptoticity by $\hat{g}$ is somewhat complicated. In fact, it turns out to be easier to prove that $\hat{g}$ preserves the property that developed geodesics *intersect* in the forward direction, and then observe that being developed forward asymptotic occurs as the limiting case of developed geodesics intersecting. The intersecting behavior we prove is described by the next proposition.

Proposition 5.5. *Suppose η and η' are oriented basic $\hat{\varphi}_1$–geodesics and let σ be a $\hat{\varphi}_1$–saddle connection with one end point on η and the other on η'. Suppose further that*

(i) *the angle between σ and the forward rays of each of η and η' measures strictly between 0 and π, and*
(ii) *$D_1(\eta)$ and $D_1(\eta')$ intersect forward of the end points of $D_1(\sigma)$.*

Then $\hat{g}(\eta)$, $\hat{g}(\eta')$, and $\hat{g}(\sigma)$ satisfy (i) and (ii) with respect to D_2 as well.

Proof. Observe that (i) holds for $\hat{g}(\eta)$, $\hat{g}(\eta')$, and $\hat{g}(\sigma)$ by Lemma 5.3, so we need only prove (ii). Let ζ and ξ be the cone points at which σ intersects η and η' respectively, and we orient σ in the direction from ζ to ξ so that η and η' cross σ from left to right; see Fig. 8. Without loss of generality we can replace η and η' by geodesics that agree on their initial segments forward of ζ and ξ, respectively, so that they make angles π on the left and right, respectively, at every cone point they encounter (by Lemma 5.3 if the conclusion holds for these replacements, it must hold for the original geodesics).

Creating a foliated triangle in $\mathbb{H}$. The assumptions of the proposition imply that $D_1(\sigma)$, $D_1(\eta)$ and $D_1(\eta')$ form a triangle T. Let a be the point of intersection $D_1(\eta) \cap D_1(\eta')$.
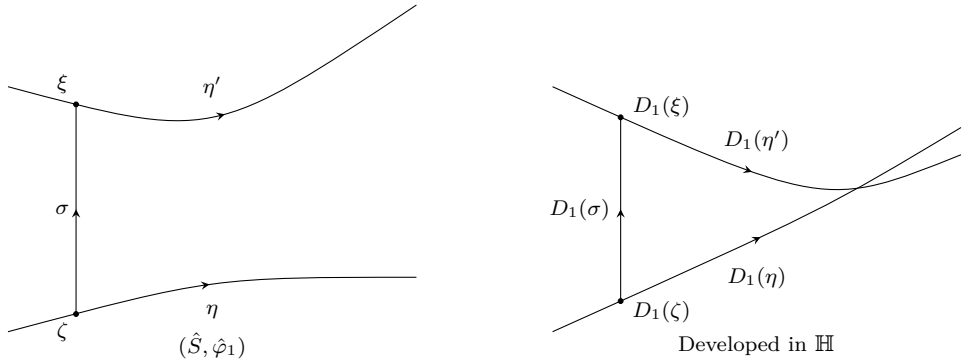
Fig. 8. A pair of developed forward intersecting geodesics connected by a saddle connection and their developed image.
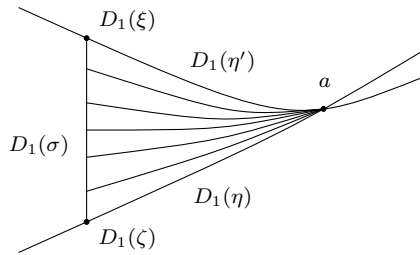

Fig. 9. Foliated triangle T in $\mathbb{H}$.

Foliate $T \subseteq \mathbb{H}$ by geodesic segments from $D_1(\sigma)$ to a, see Fig. 9. Orient each segment from $D_1(\sigma)$ to a. The set of leaves can be parametrized by their intersection with $D_1(\sigma)$.

The goal for what follows is to "pull this foliation back" to a foliation of a region containing σ whose leaves are basic $\hat{\varphi}_1$–geodesic segments. We will then argue that a sufficiently robust subset of these basic geodesics are configured in a similar way after $\hat{\varphi}_2$–straightening so that we may conclude that $\hat{g}(\eta)$ and $\hat{g}(\eta')$ are developed forward intersecting.

Pulling back the foliation. Now we describe a foliated region in $\hat{S}$ that maps to the foliated triangle T by D_1. Consider a point $z \in \sigma$. If there is a unique basic $\hat{\varphi}_1$–geodesic arc starting at z whose D_1–image is a leaf of the foliation of T through $D_1(z)$, then denote this arc f_z. If not, then all such arcs must agree on some initial subarc, δ, from z to a cone point. In this case we let f_z^l and f_z^r be the arcs of the unique basic $\hat{\varphi}_1$–geodesics through z, containing δ, making angle π on the left and right, respectively, at every cone point they encounter and that each map to the leaf through $D_1(z)$. Now let $\hat{T}$ be the union of the arcs f_z (or f_z^l and f_z^r) over all $z \in \sigma$, which is a region in $\hat{S}$ (singularly) foliated by these arcs.

Finitely many cone points. Next we want to show that $\hat{T}$ contains only finitely many cone points, and hence only finitely many leaves of the foliation of $\hat{T}$ encounter cone
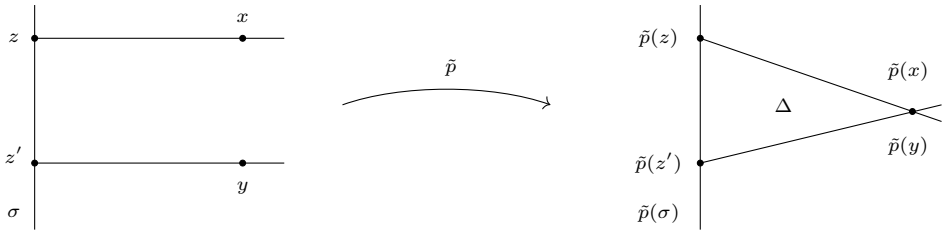
Fig. 10. Failure of injectivity produces a triangle Δ in $\widetilde{S}$.

points. To see this, we first note that $\hat{T}$ has finite diameter: if d is the maximal length of a leaf of T, then every point of $\hat{T}$ is within d of the compact segment σ. Next, we claim that $\hat{T}$ projects injectively to $\widetilde{S}$ by $\tilde{p}$. Assuming this claim, observe that since there are only finitely many cone points in the bounded image of $\hat{T}$ in $\widetilde{S}$, there are only finitely many cone points in $\hat{T}$. We are left to prove the claim.

To prove $\hat{T}$ projects injectively to $\widetilde{S}$, we argue by contradiction. If it does not project injectively, there must be two distinct points x and y in $\hat{T}$ that project to a single point in $\widetilde{S}$. The $\hat{\varphi}_1$–geodesic leaves through x and y must be distinct since basic geodesics project injectively from $\hat{S}$ to $\widetilde{S}$. In the case where the two leaves are f_z^l and f_z^r for some z, x and y projecting to the same point of $\widetilde{S}$ would contradict the fact that the total cone angle at each cone point of $\widetilde{S}$ exceeds 2π. In the case where the two leaves through x and y intersect σ at distinct points z and z', we call the leaves f_z and $f_{z'}$, suppressing the r, l superscripts that may decorate these since they do not make a difference in the argument that follows.

The subsegment of σ from z to z', together with the segments of f_z and $f_{z'}$ from z to x and z' to y, respectively, together project to a geodesic triangle Δ in $\widetilde{S}$; see Fig. 10. Note that this triangle must either have a cone point in its interior, or a cone point on the boundary which makes an angle greater than π toward the inside: otherwise, the entire triangle would lift to $\hat{S}$, and hence x and y would be the same point of $\hat{T}$, a contradiction. We will say that such cone points are *fully inside* Δ.

Because Δ is compact, there are only finitely many cone points fully inside Δ. Observe that there must be some w in σ between z and z' (possibly equal to one of z or z') so that f_w^r and f_w^l pass through one of these cone points fully inside Δ. One or both of these must project to a geodesic segment in $\widetilde{S}$ that cuts off a "sub-triangle", coming from two more distinct points x' and y' in $\hat{T}$ that project to the same point; see Fig. 11. This new triangle Δ' contains fewer cone points. We may repeat this procedure, replacing x, y with x', y' and eventually arrive at a triangle containing no cone points fully inside it, which is a contradiction. Therefore, $\hat{T}$ must inject into $\widetilde{S}$, and consequently $\hat{T}$ contains only finitely many cone points.

$\hat{T}$ is a union of hyperbolic triangles. We note that $\hat{T}$ has piecewise $\hat{\varphi}_1$–geodesic boundary (consisting of subarcs of the arcs f_z^l and f_z^r, together with arcs of η, η', and σ), does not contain cone points in the interior, and the angle between consecutive boundary segments is less than or equal to 2π; see Fig. 12. Therefore, D_1 maps $\hat{T}$ to T, and is injective except
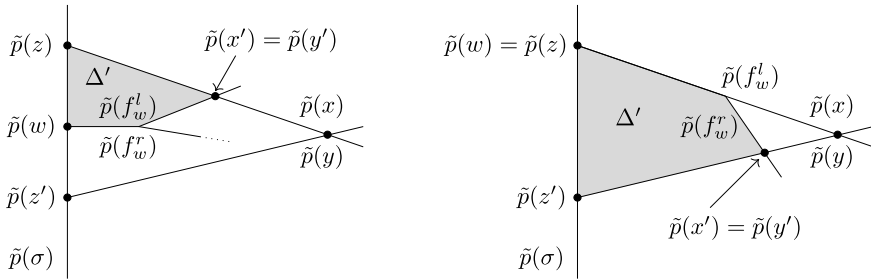
Fig. 11. Two possibilities for a subtriangle Δ' of Δ in $\widetilde{S}$ (shaded) with fewer cone points fully inside.
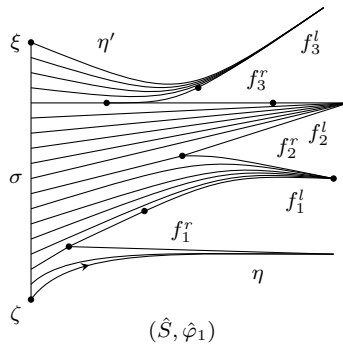

Fig. 12. The "polygon" $\hat{T}$ created in $(\hat{S}, \hat{\varphi}_1)$.

along consecutive boundary arcs that make angle 2π: these are arcs of f_z^l and f_z^r which are "zipped together" by D_1. The leaves of $\hat{T}$ are sent to leaves of T by construction.

The orientation of σ in the direction from ζ to ξ orders the leaves of the foliation of $\hat{T}$, except pairs f_z^l and f_z^r that meet cone points, which we order so that $f_z^r < f_z^l$. Now write $f_1^r, f_1^l, f_2^r, f_2^l, \ldots, f_n^r, f_n^l$ for the (possibly empty) set of all leaves in $\hat{T}$, in order, that encounter cone points away from σ, and for convenience, write f_0^l and f_{n+1}^r for the leaves contained in η and η', respectively. For each $0 \leq i \leq n$, let $\hat{T}_i \subset \hat{T}$ be the region bounded by f_i^l, f_{i+1}^r, and σ. Note that this region contains no cone points in its interior, and maps isometrically by D_1 to a sub-triangle of T in $\mathbb{H}$ which is a union of leaves. We view $\hat{T}$ as a "polygon" decomposed into a union of hyperbolic triangles in this way. See Fig. 12. The special case that $n = 0$ is possible (then $\hat{T}$ is itself a hyperbolic triangle) and our argument is valid in that case as well, though one could also argue it directly.

Intersections and orientations preserved. Going forward, we abuse notation and let

$$\sigma, f_0^l, f_1^r, f_1^l, \ldots, f_n^r, f_n^l, f_{n+1}^r$$

denote oriented basic $\hat{\varphi}_1$–geodesic containing the arc of the same name. For f_0^r and f_{n+1}^l, this basic $\hat{\varphi}_1$–geodesic is η and η', respectively. For $1 \leq i \leq n$, we assume the basic $\hat{\varphi}_1$–geodesics f_i^l and f_i^r are as described above, thus making angles π on the left and
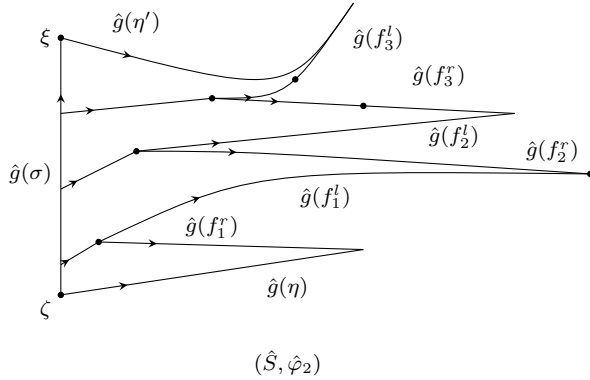
Fig. 13. The $\hat{\varphi}_2$–straightenings of the important $\hat{\varphi}_1$–geodesics.

right, respectively, at every cone point they meet. Finally, we choose the extension for σ so that it crosses η and η' at ζ and ξ, respectively.

The pairs (f_i^l, f_i^r) have the same D_1–image in $\mathbb{H}$ because they share a common sub-arc from their point of intersection with σ to their first cone point after that in the forward direction. Therefore the basic $\hat{\varphi}_2$–geodesics $\hat{g}(f_i^l)$ and $\hat{g}(f_i^r)$ must also have the same D_2–images since they also share a common arc (since f_i^l and $\hat{g}(f_i^l)$ have the same combinatorics, as do f_i^r and $\hat{g}(f_i^r)$, by Lemma 5.1).

We also observe that f_i^r and f_i^l cross σ from left to right and f_i^l crosses f_{i+1}^r from right to left, for all indices i for which the corresponding geodesics are defined. By construction, these geodesics clearly intersect in a single point (possibly a cone point) and cross because the intersection is transverse or because of the choices of the sides on which geodesics make angle π if at a cone point. Since the crossing of geodesics and orientations of such (e.g. left to right or right to left) is encoded by the endpoints at infinity, it follows that $\hat{g}(f_i^r)$ and $\hat{g}(f_i^l)$ cross $\hat{g}(\sigma)$ from left to right and $\hat{g}(f_i^l)$ crosses $\hat{g}(f_{i+1}^r)$ from right to left, for all i as above. Moreover, because δ and $\hat{g}(\delta)$ have the same combinatorics for every $\delta \in \mathcal{G}(\hat{\varphi}_1)$ (Lemma 5.1), these intersections occur at a single point (see Fig. 13).

When developed, $\hat{g}(\eta)$ and $\hat{g}(\eta')$ intersect. As noted above, $D_2(\hat{g}(f_i^l)) = D_2(\hat{g}(f_i^r))$ since they share an arc. We denote this image simply as $D_2(\hat{g}(f_i))$, noting that $D_2(\hat{g}(f_0)) = D_2(\hat{g}(\eta))$ and $D_2(\hat{g}(f_{n+1})) = D_2(\hat{g}(\eta'))$. Since $\hat{g}(f_i^l)$ crosses $\hat{g}(f_{i+1}^r)$ from right to left for each $0 \leq i \leq n$ it follows that

(1) $D_2(\hat{g}(f_i))$ crosses $D_2(\hat{g}(f_{i+1}))$ from right to the left, for all $0 \leq i \leq n$.

Similarly, since $\hat{g}(f_i^r)$ and $\hat{g}(f_i^l)$ cross $\hat{g}(\sigma)$ from left to right for all $1 \leq i \leq n$, as do $\hat{g}(\eta) = \hat{g}(f_0^l)$ and $\hat{g}(\eta') = \hat{g}(f_{n+1}^r)$, we also have

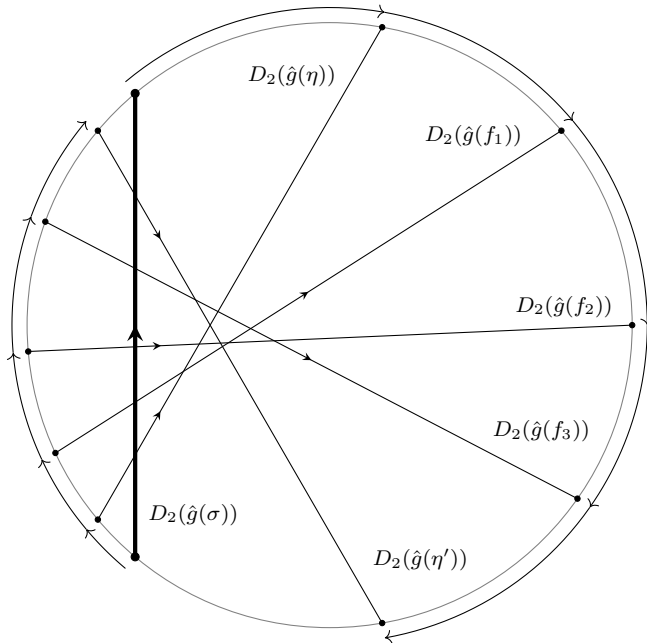(2) $D_2(\hat{g}(f_i))$ crosses $D_2(\hat{g}(\sigma))$ from left to right for all $0 \leq i \leq n+1$.

Fig. 14. Linking and ordering of end points of D_2–images of the important $\hat{\varphi}_2$–straightened geodesics (shown here in the Klein model of $\mathbb{H}$).

We now claim that condition (ii) in Proposition 5.5 holds for $\hat{g}(\eta)$, $\hat{g}(\eta')$, and $\hat{g}(\sigma)$. For this, note that the end points of $D_2(\hat{g}(\sigma))$ divide $\partial\mathbb{H}$ into two intervals: the left interval (left of $D_2(\hat{g}(\sigma))$) and the right interval (right of $D_2(\hat{g}(\sigma))$). Statement (2) implies that $D_2(\hat{g}(f_i))$ must have initial endpoint on the left interval and final endpoint on the right interval, for all $0 \leq i \leq n+1$. Statement (1) implies that for all $0 \leq i \leq n$, the initial (final) endpoint of $D_2(\hat{g}(f_i))$ must occur clockwise of the initial (final) endpoint of $D_2(\hat{g}(f_{i+1}))$ within the left (right) intervals. Since clockwise ordering is a total order on each of the left and right intervals, transitivity implies that the initial (final) endpoint of $D_2(\hat{g}(\eta')) = D_2(\hat{g}(f_{n+1}))$ must occur clockwise of the initial (final) endpoint of $D_2(\hat{g}(\eta)) = D_2(\hat{g}(f_0))$ within the left (right) intervals. See Fig. 14. This implies not only that $D_2(\hat{g}(\eta))$ and $D_2(\hat{g}(\eta'))$ intersect (because their endpoints are linked), but that the intersection point must be forward of the intersections of $D_2(\hat{g}(\eta))$ and $D_2(\hat{g}(\eta'))$ with $D_2(\hat{g}(\sigma))$. The latter is because $D_2(\xi)$ is forwards of $D_2(\zeta)$ along $D_2(\hat{g}(\sigma))$. This proves condition (ii), and thus completes the proof. $\square$

We now prove an asymptotic version of Proposition 5.5.

Proposition 5.6. *Suppose η and η' are oriented basic $\hat{\varphi}_1$–geodesics that are developed forward asymptotic. Suppose there is a saddle connection σ with one end point on η and*

one end point on η', such that the angles between σ and the forward rays of η and η' lie in the interval $[0, \pi]$. Then $\hat{g}(\eta)$ and $\hat{g}(\eta')$ are developed forward asymptotic as well.

Proof. Let $\theta_1, \theta_1' \in [0, \pi]$ be the angles between σ and the forward rays of η and η', respectively, and θ_2, θ_2' be the corresponding angles between $\hat{g}(\sigma)$ and the forward rays of $\hat{g}(\eta)$ and $\hat{g}(\eta')$, respectively. If $\theta_1 = 0$ or π, then $\theta_1' = \pi$ or 0, respectively, and then the D_1–image of σ, η, and η' are all contained in a single geodesic. Appealing to Lemma 5.3 we see that the same is true of $\hat{g}(\sigma)$, $\hat{g}(\eta)$, and $\hat{g}(\sigma)$, completing the proof in these cases.

Now assume that $\theta_1, \theta_1' \in (0, \pi)$, so by Lemma 5.3, $\theta_2, \theta_2' \in (0, \pi)$. Let ζ and ξ be the end points of σ where η and η' cross it, respectively. Observe that for all $\theta \in (0, \theta_1)$, any geodesic η_θ through ζ whose forward ray makes angle θ with σ must have $D_1(\eta_\theta)$ intersecting $D_1(\eta')$ forward of σ. By Proposition 5.5, $D_2(\hat{g}(\eta_\theta))$ must intersect $D_2(\hat{g}(\eta'))$ forward of $D_2(\hat{g}(\sigma))$.

We may assume that as $\theta \to \theta_1$, $\eta_\theta \to \eta$, after replacing η and each η_θ by geodesics through ζ with the same D_1–images, if necessary. Therefore as $\theta \to \theta_1$, we have $\hat{g}(\eta_\theta) \to \hat{g}(\eta)$, and so also $D_2(\hat{g}(\eta_\theta)) \to D_2(\hat{g}(\eta))$. Since $D_2(\hat{g}(\eta_\theta))$ intersects $D_2(\hat{g}(\eta'))$ forward of $D_2(\hat{g}(\sigma))$ for all $\theta \in (0, \theta_0)$, we see that either $D_2(\hat{g}(\eta))$ intersects $D_2(\hat{g}(\eta'))$ forward of $D_2(\hat{g}(\sigma))$ or else $D_2(\hat{g}(\eta))$ and $D_2(\hat{g}(\eta'))$ are forward asymptotic. We cannot be in the former situation since reversing the roles of $\hat{\varphi}_1$ and $\hat{\varphi}_2$ and applying Proposition 5.6 to $\hat{g}(\sigma)$, $\hat{g}(\eta)$, and $\hat{g}(\eta')$, would imply $D_1(\eta)$ and $D_1(\eta')$ intersect forward of $D_1(\sigma)$, a contradiction. We are thus in the latter situation, and so $\hat{g}(\eta)$ and $\hat{g}(\eta')$ are developed forward asymptotic. This completes the proof. $\square$

Now, given a $\hat{\varphi}_i$–saddle connection σ and $x \in \partial \mathbb{H}$, we consider the family $\Lambda(\sigma, x) \subset \mathcal{G}(\hat{\varphi}_i)$ which consists of the set of basic $\hat{\varphi}_i$–geodesics η intersecting σ that may be oriented so that $D_i(\eta)$ limits to x in the forward direction. We view $\Lambda(\sigma, x)$ as a subset of $\mathcal{G}(\hat{\varphi}_i)$, together with a choice of orientation for each geodesic.

If $D_i(\sigma)$ can be extended to an oriented geodesic η with x as its forward end point, then $\Lambda(\sigma, x)$ is the set of all oriented basic $\hat{\varphi}_i$–geodesics meeting σ with developed image equal to η. Note that in this case $\Lambda(\sigma, x)$ is countable.

If not, then note that $\Lambda(\sigma, x)$ contains the pullback (as in the proof of Proposition 5.5) of the family of geodesics in $\mathbb{H}$ that limit to x in the forward direction and intersect $D_i(\sigma)$ transversely. In particular, in this case, $\Lambda(\sigma, x)$ is uncountable.

Proposition 5.7. *For any $\hat{\varphi}_1$–saddle connection σ and $x \in \partial \mathbb{H}$, there exists a point $h(\sigma, x) \in \partial \mathbb{H}$ so that*

$$\hat{g}(\Lambda(\sigma, x)) = \Lambda(\hat{g}(\sigma), h(\sigma, x)).$$

We clarify that $\hat{g}(\Lambda(\sigma, x)) \subset \mathcal{G}(\hat{\varphi}_2)$ and each geodesic in the $\hat{g}$–image is assigned the orientation coming from its orientation in $\Lambda(\sigma, x)$.

Proof. We begin by addressing the case where $\Lambda(\sigma, x)$ is countable. In this situation, $D_1(\Lambda(\sigma, x))$ is a single oriented geodesic in $\mathbb{H}$. By Lemma 5.3, $\hat{g}(\Lambda(\sigma, x))$ consists of oriented $\hat{\varphi}_2$–geodesics that develop to a single geodesic η in $\mathbb{H}$ which contains $D_2(\hat{g}(\sigma))$. Letting the forward end point of this geodesic be $h(\sigma, x)$ gives us that $\hat{g}(\Lambda(\sigma, x)) \subset \Lambda(\hat{g}(\sigma), h(\sigma, x))$. Reversing the roles of $\hat{\varphi}_1$ and $\hat{\varphi}_2$ proves the other inclusion.

In the case where $\Lambda(\sigma, x)$ is uncountable, $D_1(\Lambda(\sigma, x))$ consists of all oriented geodesics in $\mathbb{H}$ that intersect $D_1(\sigma)$ and have x as their forward end point. As in Proposition 5.5, we systematically pull back each of these geodesics to one or two $\hat{\varphi}_1$–geodesics to get a family $\mathcal{R} \subseteq \Lambda(\sigma, x)$.

More precisely, we begin by pulling back short segments of forward asymptotic geodesics that intersect the interior of $D_1(\sigma)$ and extending them to basic geodesics in $(\hat{S}, \hat{\varphi}_1)$. Again, as in Proposition 5.5, if in the extension process, a cone point is encountered, we continue extending in two ways: by always making angle π on the right, and by always making angle π on the left. At most countably many of these pullback geodesics, $\{f_i^l, f_i^r\}_{i \in A}$, where A is a countable indexing set, are such extensions (because there are only countably many cone points in $(\hat{S}, \hat{\varphi}_1)$). The superscript tells us which side of the oriented basic geodesic makes angle π. Again as in Proposition 5.5, this construction always has $D_1(f_i^l) = D_1(f_i^r)$. For the two geodesics in $\mathbb{H}$ intersecting the end points of D_σ, we also include in $\mathcal{R}$ the two geodesics η_0 and η_1 in $\Lambda(\sigma, x)$ that intersect the end points of σ and whose forward rays make an angle less than π with σ.

We show that $\hat{g}(\mathcal{R})$ consists of basic $\hat{\varphi}_2$–geodesics that are developed forward asymptotic, and use this to prove the desired result. The geodesics of $\hat{g}(\mathcal{R})$ in $(\hat{S}, \hat{\varphi}_2)$ must also be a family of pairwise non-intersecting geodesics since the end points of any two are not linked, and because the cone points visited by each geodesic are preserved by the straightening map $\hat{g}$. Betweenness of geodesics is also preserved by $\hat{g}$. This implies that the geodesics of $\hat{g}(\mathcal{R})$ intersect $\hat{g}(\sigma)$ in the same order as in $(\hat{S}, \hat{\varphi}_1)$. In particular, $\hat{g}(\mathcal{R})$ is a family of geodesics between $\hat{g}(\eta_0)$ and $\hat{g}(\eta_1)$. Finally, $D_2(\hat{g}(f_i^l)) = D_2(\hat{g}(f_i^r))$, by Lemma 5.3, and we denote this geodesic by $D_2(\hat{g}(f_i))$.

Consider the region $\mathcal{S}$ of $(\hat{S}, \hat{\varphi}_2)$ given by the union of geodesics in $\hat{g}(\mathcal{R})$. Using an argument similar to the proof of Proposition 5.5 (finitely many cone points), we can show that $\mathcal{S}$ projects injectively to $\widetilde{S}$ and therefore has the following property. Given any distance d, there are only finitely many cone points in $\mathcal{S}$ distance d away from σ. This implies that $\mathcal{S}$ is a hyperbolic surface with boundary. $\mathcal{S}$ does not contain any cone points in its interior, but it does have at most countably many cone points on its boundary. Through the ith cone point, there are two geodesics of $\hat{g}(\mathcal{R})$: $\hat{g}(f_i^l)$ and $\hat{g}(f_i^r)$. Extending the overlap of $\hat{g}(f_i^l)$ and $\hat{g}(f_i^r)$ to an isometric identification of the two geodesics ("zipping up" the slits of $\mathcal{S}$), we obtain a new surface $\mathcal{S}_{\text{zip}}$. Note that the metric on $\mathcal{S}_{\text{zip}}$ is nonsingular, since after the identification, the total angle at each cone point is 2π. Therefore, the restriction of D_2 to $\mathcal{S}$ descends to an isometry $\mathcal{S}_{\text{zip}} \to D_2(\mathcal{S})$. Since $\hat{g}(\mathcal{R})$ is a set of non-intersecting geodesics between these two, the image of $\hat{g}(\mathcal{R})$ in $\mathcal{S}_{\text{zip}}$ is a set of parallel geodesics between the images of $\hat{g}(\eta_0)$ and $\hat{g}(\eta_1)$. Further since $\hat{g}(\eta_0)$ and $\hat{g}(\eta_1)$ are developed forward asymptotic, the image of $\hat{g}(\mathcal{R})$ must be forward

asymptotic (a family of parallel geodesics between two forward asymptotic geodesics in the hyperbolic plane must all be forward asymptotic to each other).

Note that every $\hat{\varphi}_1$-geodesic $\eta \in \Lambda(\sigma, x)$ is either contained in $\mathcal{R}$ or there exists $\eta' \in \mathcal{R}$ such that η and η' intersect at a cone point and $D_1(\eta) = D_1(\eta')$. Using Lemma 5.3, we then have $D_2(\hat{g}(\Lambda(\sigma, x))) = D_2(\hat{g}(\mathcal{R}))$. Therefore, $\hat{g}(\Lambda(\sigma, x))$ is a developed forward asymptotic family, and we let $h(\sigma, x)$ be the forward end point of $D_2(\hat{g}(\Lambda(\sigma, x)))$. From here, since every geodesic of $\hat{g}(\Lambda(\sigma, x))$ must instersect $\hat{g}(\sigma)$, the inclusion $\hat{g}(\Lambda(\sigma, x)) \subseteq \Lambda(\hat{g}(\sigma), h(\sigma, x))$ follows by definition, and the reverse inclusion follows by a symmetric argument with the roles of $\hat{\varphi}_1$ and $\hat{\varphi}_2$ exchanged. $\quad\square$

Finally, we are able to show the desired result that developed forward asymptoticity is preserved by the straightening map.

Proposition 5.8. *Suppose η_1 and η_2 are developed forward asymptotic $\hat{\varphi}_1$–geodesics. Then $\hat{g}(\eta_1)$ and $\hat{g}(\eta_2)$ are developed forward asymptotic as well.*

Proof. If η_1 and η_2 meet in a segment or cone point, then $\hat{g}(\eta_1)$ and $\hat{g}(\eta_2)$ are developed forward asymptotic by Lemma 5.3. For the rest of the proof, suppose that this is not the case. The geodesics η_1 and η_2 must not intersect transversely in $(\hat{S}, \hat{\varphi}_1)$ (otherwise they would intersect transversely when developed). Therefore, they partition $\hat{S}$ into three regions: one bounded by η_1, one bounded by η_2 and one bounded by both. Pick a cone point ζ in the region bounded by η_1 alone and another cone point ξ in the region bounded by η_2 alone. Take the geodesic segment between ζ and ξ. It must be a concatenation of saddle connections. After forgetting some initial and final saddle connections, we may assume the remaining concatenation $\sigma_1, \sigma_2, \ldots \sigma_k$ has only σ_1 intersecting η_1 and only σ_k intersecting η_2.

Let z be the forward end point of $D_1(\eta_1)$ and $D_1(\eta_2)$ in $\partial\mathbb{H}$. Then consider the developed forward asymptotic families $\Lambda(\sigma_i, z)$. We know that $\eta_1 \in \Lambda(\sigma_1, z)$ and $\eta_2 \in \Lambda(\sigma_k, z)$, and from Proposition 5.7 we know that $\hat{g}(\Lambda(\sigma_i, z)) = \Lambda(\hat{g}(\sigma_i), h(\sigma_i, z))$. To complete our proof, we will show that $\cup_i \Lambda(\hat{g}(\sigma_i), h(\sigma_i, z))$ is a developed forward asymptotic family by showing that $h(\sigma_1, z) = h(\sigma_2, z) = \ldots = h(\sigma_n, z)$.

Let c_i be the cone point $\sigma_i \cap \sigma_{i+1}$, the intersection of two consecutive saddle connections. The intersection of any two consecutive families $\Lambda(\sigma_i, z) \cap \Lambda(\sigma_{i+1}, z)$ consists of the basic $\hat{\varphi}_1$–geodesics through c_i that develop to have z as their forward end point. This set is countably infinite, so in particular, it is not empty. For any $\eta \in \Lambda(\sigma_i, z) \cap \Lambda(\sigma_{i+1}, z)$, the $\hat{\varphi}_2$–straightening $\hat{g}(\eta)$ is contained in $\Lambda(\hat{g}(\sigma_i), h(\sigma_i z)) \cap \Lambda(\hat{g}(\sigma_{i+1}), h(\sigma_{i+1}, z))$. The geodesic $\hat{g}(\eta)$ cannot have two distinct forward end points, and so $h(\sigma_i, z) = h(\sigma_{i+1}, z)$. This argument holds for all i, and so $\cup_i \Lambda(\hat{g}(\sigma_i), h(\sigma_i, z))$ is a developed forward asymptotic family. In particular, this implies that $\hat{g}(\eta_1)$ and $\hat{g}(\eta_2)$ are developed forward asymptotic. $\quad\square$

As proposed earlier, we may now define $h : \partial\mathbb{H} \to \partial\mathbb{H}$ by requiring $h(x)$ to be the forward endpoint of $D_2(\hat{g}(\eta))$ where $\eta \in \mathcal{G}(\hat{\varphi}_1)$ is a basic geodesic oriented so that x is the

forward endpoint of $D_1(\eta)$, and Proposition 5.8 implies that this is indeed well-defined. Moreover, by the proposition, for any $\hat{\varphi}_1$–saddle connection σ, we have $h(x) = h(\sigma, x)$ where $h(\sigma, x)$ is as in Proposition 5.7.

At long last, we are able to prove Proposition 5.4.

Proof of Proposition 5.4. The definition of h implies that for any biinfinite basic $\hat{\varphi}_1$–geodesic η, $h(\partial D_1(\eta)) = \partial D_2(\hat{g}(\eta))$. It remains to show that h is an orientation preserving homeomorphism, and that it topologically conjugates the holonomies.

Recall that the developing maps D_i can be extended to $\bar{D}_i\colon \hat{S} \cup \mathcal{K}_i \to \overline{\mathbb{H}}$, where $\mathcal{K}_i$ is the set of end points of basic $\hat{\varphi}_i$–geodesics (see discussion after Lemma 2.5). Note that $\mathcal{K}_1 = \mathcal{K}_2 =: \mathcal{K}$ since $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$. Using this extension, we may write $h(\bar{D}_1(k)) = \bar{D}_2(k)$, for any $k \in \mathcal{K} \subseteq \hat{S}^1_\infty$.

If h were not injective, there would be two developed forward asymptotic families of oriented basic $\hat{\varphi}_1$–geodesics that are not developed forward asymptotic to each other, however their images by $\hat{g}$ are developed forward asymptotic to each other. This contradicts Proposition 5.8, and thus h is injective. The map h is surjective because $\bar{D}_2$ maps $\mathcal{K}$ onto $\partial\mathbb{H}$ by Corollary 2.6. So, for every $y \in \partial\mathbb{H}$, there exists $k \in \mathcal{K}$ such that $\bar{D}_2(k) = y$. Then $h(\bar{D}_1(k)) = \bar{D}_2(k)$, so we have $h(\bar{D}_1(k)) = y$, as desired.

In order to show h is a homeomorphism, we show that h and h^{-1} are both open maps. Let U be an open interval in $\partial\mathbb{H}$. Choose a point $\zeta_1 \in \mathbb{H}$ such that $\zeta_1 = D_1(\zeta)$, where ζ is a cone point in $(\hat{S}, \hat{\varphi}_1)$. Consider the set of oriented geodesics through ζ_1 that have forward end point in U. These geodesics sweep out some interval of angles (α, β) measured at ζ_1. From the description of neighborhoods of cone points in $\hat{S}$ in §2.2, we may choose a family of basic $\hat{\varphi}_1$-geodesics Λ through ζ that sweep out an interval of angles of the same measure as (α, β) which develop to the above mentioned set of geodesics in $\mathbb{H}$. The $\hat{\varphi}_2$–straightening of this family, $\hat{g}(\Lambda)$, is a family of basic geodesics in $(\hat{S}, \hat{\varphi}_2)$ that are concurrent at ζ and sweep out an open interval of angles at ζ, though the measure of this interval is not necessarily $\beta - \alpha$. The image $D_2(\hat{g}(\Lambda))$ must be a family of concurrent geodesics, passing through $\zeta_2 = D_2(\zeta)$, that sweep out an open interval of angles at ζ_2. Therefore the forward end points of the geodesics in $D_2(\hat{g}(\Lambda))$ must form an open interval in $\partial\mathbb{H}$. This implies that h is an open map. The same argument with the subscripts 1 and 2 interchanged shows that h^{-1} is an open map. Therefore h is a homeomorphism. Moreover, sweeping through angles counterclockwise with respect to $\hat{\varphi}_1$ is the same as sweeping through them counterclockwise with respect to $\hat{\varphi}_2$, and thus h is orientation preserving.

We now show that h topologically conjugates ρ_1 to ρ_2. That is, for $\gamma \in G$ and $x \in \partial\mathbb{H}$, we prove Equation (2). Recall that since γ acts on $\hat{S}$ by isometries with respect to both $\hat{\varphi}_1$ and $\hat{\varphi}_2$, it extends to $\hat{S}^1_\infty$, and thus $\mathcal{K}$. Since D_i also extends to $\mathcal{K}$, we have:

$$
\begin{aligned}
h(\rho_1(\gamma) \cdot x) &= h(\rho_1(\gamma) \cdot \bar{D}_1(k)), &&\text{for some } k \in \mathcal{K}. \\
&= h(\bar{D}_1(\gamma \cdot k)) &&\text{by definition of the holonomy homomorphism.} \\
&= \bar{D}_2(\gamma \cdot k) &&\text{by definition of } h.
\end{aligned}
$$

$$= \quad \rho_2(\gamma) \cdot \bar{D}_2(k) \qquad \text{by definition of the holonomy homomorphism.}$$

$$= \quad \rho_2(\gamma) \cdot h(\bar{D}_1(k)) \qquad \text{by definition of } h \text{ and extensions of } D_i.$$

$$= \quad \rho_2(\gamma) \cdot h(x) \qquad \text{by definition of } k. \quad \square$$

5.3. Rigidity from $h \in \mathrm{PSL}_2(\mathbb{R})$

We end this section by noting that for the rigidity statement in the Current Support Theorem, we want the homeomorphism h from Proposition 5.4 to be in $\mathrm{PSL}_2(\mathbb{R})$.

Proposition 5.9. *If $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and $h \colon \partial\mathbb{H} \to \partial\mathbb{H}$ from Proposition 5.4 lies in $\mathrm{PSL}_2(\mathbb{R})$, then $\varphi_1 \sim \varphi_2$.*

This proposition will follow easily from the next lemma.

Lemma 5.10. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and $h \colon \partial\mathbb{H} \to \partial\mathbb{H}$ from Proposition 5.4 lies in $\mathrm{PSL}_2(\mathbb{R})$. If φ_2 is C_0–uber-normalized, with $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2)$, then for any two basic $\hat{\varphi}_1$–geodesics $\eta, \eta' \in \mathcal{G}(\hat{\varphi}_1)$ that intersect making angle $\theta < \pi$, $\hat{g}(\eta)$ and $\hat{g}(\eta')$ also intersect making angle θ.*

Proof. From the hypotheses and Lemma 5.1, we have $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$, $\mathrm{cone}(\hat{\varphi}_1) = \mathrm{cone}(\hat{\varphi}_2) = \hat{C}_0$, and since η and η' intersect in a single point making angle less than π, the same is true of $\hat{g}(\eta)$ and $\hat{g}(\eta')$ (by Lemma 5.3 if η intersects η' in a point of $\hat{C}_0$). Moreover the angle between η and η' (respectively, $\hat{g}(\eta)$ and $\hat{g}(\eta')$) are the same as their images by D_1 (respectively, D_2). By Proposition 5.4, $h(\partial D_1(\eta)) = \partial D_2(\hat{g}(\eta))$ and $h(\partial D_1(\eta')) = \partial D_2(\hat{g}(\eta'))$. Since $h \in \mathrm{PSL}_2(\mathbb{R})$, we extend it to a map $\bar{h} \colon \overline{\mathbb{H}} \to \overline{\mathbb{H}}$ which is an isometry in $\mathbb{H}$ and

$$\bar{h}(D_1(\eta)) = D_2(\hat{g}(\eta)) \quad \text{and} \quad \bar{h}(D_1(\eta')) = D_2(\hat{g}(\eta')).$$

Since $\bar{h}$ is an isometry in $\mathbb{H}$, the angle between $D_1(\eta)$ and $D_1(\eta')$ is equal to the angle between their $\bar{h}$–images, $D_2(\hat{g}(\eta))$ and $D_2(\hat{g}(\eta'))$. $\quad \square$

Proof of Proposition 5.9. After adjusting φ_2 if necessary, we may assume it has been C_0–uber-normalized with respect to φ_1, so that $\mathrm{cone}(\tilde{\varphi}_1) = \mathrm{cone}(\tilde{\varphi}_2) = \widetilde{C}_0$, and thus $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ and $\mathrm{cone}(\hat{\varphi}_1) = \mathrm{cone}(\hat{\varphi}_2) = \hat{C}_0$ by Lemma 5.1. Let $\mathcal{T}$ be a (C_0, φ_1)–triangulation of S and further adjust φ_2 by an isotopy if necessary so that it is a $(C_0, \varphi_1, \varphi_2)$–triangulation. Let $\widehat{\mathcal{T}}$ be the lift of $\mathcal{T}$ to $\hat{S}$.

Since $\widehat{\mathcal{T}}$ is simultaneously a $(C_0, \tilde{\varphi}_1)$–triangulation and a $(C_0, \tilde{\varphi}_2)$–triangulation, the image of each triangle $\widehat{T}$ in $\widehat{\mathcal{T}}$ under both developing maps is a geodesic triangle in $\mathbb{H}$. Further $D_1(\widehat{T})$ is isometric to $(\widehat{T}, \hat{\varphi}_1)$ and $D_2(\widehat{T})$ is isometric to $(\widehat{T}, \hat{\varphi}_2)$. This is because the developing maps are local isometries on sets that do not contain cone points in

their interiors, and since all interior angles of triangles measure less than π, this can be promoted to an isometry in the case of a triangle.

For each vertex of each triangle of $\widehat{\mathcal{T}}$, extending the two sides adjacent to that vertex to basic $\hat{\varphi}_1$–geodesics η and η', Lemma 5.10 implies the angle between η and η' agrees with the angle between $\hat{g}(\eta)$ and $\hat{g}(\eta')$. Since $\hat{g}(\eta)$ and $\hat{g}(\eta')$ are extensions of the sides of the triangle to basic $\hat{\varphi}_2$–geodesics, it follows that the interior angle of the triangle is the same when measured in both metrics. Since the identity on $\hat{S}$ is a canonical map on each triangle with respect to $\hat{\varphi}_1$ on the domain and $\hat{\varphi}_2$ on the range, and since the interior angles of the triangles are equal, this map is an isometry. Therefore, the identity $(\hat{S}, \hat{\varphi}_1) \to (\hat{S}, \hat{\varphi}_2)$ is an isometry, and hence so is the identity $(S, \varphi_1) \to (S, \varphi_2)$, proving $\varphi_1 \sim \varphi_2$. $\square$

6. Proof of the Current Support Theorem

The conclusion of the Current Support Theorem has two parts: a rigidity statement and a flexibility statement. Proposition 5.9 says that to prove the rigidity statement— that is, if $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ then $\varphi_1 \sim \varphi_2$—it suffices to show that the homeomorphism h from Proposition 5.4 lies in $\mathrm{PSL}_2(\mathbb{R})$. Proposition 2.7 easily implies that if $\Gamma_1 = \rho_1(\pi_1\dot{S})$ is indiscrete, then h lies in $\mathrm{PSL}_2(\mathbb{R})$. In this section we first describe a weaker condition that still suffices to guarantee that h lies in $\mathrm{PSL}_2(\mathbb{R})$. We then show that the only way this condition can fail is if we are in the situation described in Theorem 3.1, proving the flexibility statement.

Throughout this section, we will assume $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ with $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and that $C_0 = \mathrm{cone}(\varphi_1) = \mathrm{cone}(\varphi_2)$ with φ_2 being C_0–uber-normalized (so that $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ by Lemma 5.1). For each $i = 1, 2$, recall that $D_i : \hat{S} \to S$ is the developing map for $\hat{\varphi}_i$, ρ_i is the holonomy homomorphism defined on $\pi_1\dot{S}$ (or any group acting by homeomorphisms, isometrically with respect to both $\hat{\varphi}_1$ and $\hat{\varphi}_2$), and $\Gamma_i = \rho_i(\pi_1\dot{S})$ with $q_i : (S, \varphi_i) \to \mathcal{O}_i = \mathbb{H}/\Gamma_i$, the quotient space by Γ_i (which is an orbifold if Γ_i is discrete).

6.1. Rigidity

The right condition to ensure rigidity turns out to involve a certain enlargement of Γ_1 which we describe below. First, we need the following fact. Given $\xi \in \mathbb{H}$, let τ_ξ denote the involution given by a rotation of order 2 about ξ.

Proposition 6.1. *Suppose* $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, φ_2 *is* C_0–*uber-normalized, and* $h : \partial\mathbb{H} \to \partial\mathbb{H}$ *is the homeomorphism from Proposition 5.4. Given* $\zeta \in \hat{C}_0$, *set* $\zeta_i = D_i(\zeta)$, *for* $i = 1, 2$. *Then* $h\tau_{\zeta_1}h^{-1} = \tau_{\zeta_2}$.

Proof. Fix any $\zeta \in \hat{C}_0$. By Corollary 5.2, $\hat{g}(\mathcal{G}(\hat{\varphi}_1, \zeta)) = \mathcal{G}(\hat{\varphi}_2, \zeta)$. Moreover, for any $\eta \in \mathcal{G}(\hat{\varphi}_1, \zeta)$ we have $h(\partial D_1(\eta)) = \partial D_2(\hat{g}(\eta))$ by Proposition 5.4. Next observe that τ_{ζ_1} interchanges the endpoints of $D_1(\eta)$, while τ_{ζ_2} interchanges the endpoints of $D_2(\hat{g}(\eta))$. Therefore, $\partial D_1(\eta)$ has the form $\{x, \tau_{\zeta_1}(x)\}$ for some $x \in \partial\mathbb{H}$, and $\partial D_2(\hat{g}(\eta)) = \{h(x), \tau_{\zeta_2}(h(x))\}$. Combining these facts, we have $\tau_{\zeta_2}(h(x)) = h(\tau_{\zeta_1}(x))$. Every $x \in \partial\mathbb{H}$ is in $\partial D_1(\eta)$ for some $\eta \in \mathcal{G}(\hat{\varphi}_1, \zeta)$ by Lemma 2.5, and thus $\tau_{\zeta_2}h = h\tau_{\zeta_1}$, as required. $\square$

For each $\zeta \in \hat{C}_0 \subset \hat{S}$ and $i = 1, 2$, let $\zeta_i = D_i(\zeta)$, let $\tau_{\zeta_i} \in \mathrm{PSL}_2(\mathbb{R})$ be the involution as above, and set

$$\Gamma_i^0 = \langle \Gamma_i, \{\tau_{\zeta_i}\}_{\zeta \in \hat{C}_0} \rangle.$$

Corollary 6.2. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and φ_2 is C_0–uber-normalized. The homeomorphism $h\colon \partial \mathbb{H} \to \partial \mathbb{H}$ from Proposition 5.4 topologically conjugates Γ_1^0 to Γ_2^0 sending τ_{ζ_1} to τ_{ζ_2} for all $\zeta \in \hat{C}_0$.*

Proof. As a consequence of Proposition 5.4, h topologically conjugates Γ_1 to Γ_2, and by Proposition 6.1, $h\tau_{\zeta_1}h^{-1} = \tau_{\zeta_2}$, for all $\zeta \in \hat{C}_0$. Therefore h topologically conjugates Γ_1^0 to Γ_2^0. $\square$

For each $i = 1, 2$, we write $p_i\colon \mathbb{H} \to \mathcal{O}_i^0 = \mathbb{H}/\Gamma_i^0$ for the quotient space. If Γ_i^0 is discrete, then $\mathcal{O}_i^0$ is a hyperbolic orbifold. Otherwise, we view it simply as the quotient topological space. Since $\rho_i(\pi_1 \dot{S}) = \Gamma_i < \Gamma_i^0$, and the map $D_i\colon \hat{S} \to \mathbb{H}$ is ρ_i–equivariant, it follows that D_i descends to a continuous map $q_i^0\colon (S, \varphi_i) \to \mathcal{O}_i^0$ fitting into the following diagram

$$
\begin{array}{ccc}
\hat{S} & \xrightarrow{\ D_i\ } & \mathbb{H} \\
{\scriptstyle \hat{p}} \downarrow & & \downarrow {\scriptstyle p_i} \\
S & \xrightarrow{\ q_i^0\ } & \mathcal{O}_i^0.
\end{array}
$$

As with q_i, the map q_i^0 is a locally isometric branched cover with respect to φ_i if Γ_i^0 is discrete, but is only a continuous map in general.

The next proposition provides a sufficient condition for rigidity.

Proposition 6.3. *If $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and Γ_1^0 is indiscrete, then $\varphi_1 \sim \varphi_2$.*

Proof. Observe that $\mathcal{O}_1 = \mathbb{H}/\Gamma_1^0$ is compact, being the image of S by q_1^0. Adjusting φ_2 by a homeomorphism isotopic to the identity if necessary, we may assume that φ_2 is C_0–uber-normalized. By Corollary 6.2, Γ_1^0 and Γ_2^0 are topologically conjugate by h, and since Γ_1^0 is indiscrete, Proposition 2.7 implies $h \in \mathrm{PSL}_2(\mathbb{R})$. Therefore, $\varphi_1 \sim \varphi_2$ by Proposition 5.9. $\square$

The next corollary describes one situation where Γ_1^0 is discrete, and yet the homeomorphism h topologically conjugating Γ_1^0 to Γ_2^0 is necessarily in $\mathrm{PSL}_2(\mathbb{R})$, giving us another case where we have rigidity.

Corollary 6.4. *If $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and Γ_1^0 is a Fuchsian triangle group, then $\varphi_1 \sim \varphi_2$.*

Proof. Since h conjugates Γ_1^0 to Γ_2^0, it follows from Lemma 2.8 that $h \in \mathrm{PSL}_2(\mathbb{R})$, and Proposition 5.9 again implies $\varphi_1 \sim \varphi_2$. $\square$

6.2. Flexibility

We now turn to the situation that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, but $\varphi_1 \nsim \varphi_2$. We continue with the assumptions on φ_1, φ_2, C_0, etc., and with the definitions of Γ_i^0, $\mathcal{O}_i^0$, etc., as above. According to Corollary 6.2, h topologically conjugates Γ_1^0 to Γ_2^0, and we let

$$\psi^0 : \Gamma_1^0 \to \Gamma_2^0$$

denote the isomorphism; that is, in terms of the action on $\partial \mathbb{H}$, $\psi^0(\gamma) = h\gamma h^{-1}$, for all $\gamma \in \Gamma_1^0$.

To prove the Current Support Theorem, we need to show that φ_1 comes from a branched cover of an orbifold and that φ_2 is (up to equivalence) obtained by deforming the orbifold as in Theorem 3.1. By Proposition 6.3, we need only consider the case that Γ_1^0 (and hence also Γ_2^0) is discrete. The goal is to show that for the branched covers of orbifolds $q_i^0 : S \to \mathcal{O}_i^0$, for $i = 1, 2$, there is a homeomorphism $F : \mathcal{O}_1^0 \to \mathcal{O}_2^0$ such that q_2^0 and $F \circ q_1^0$ differ by a homeomorphism of S isotopic to the identity. Let $\mathcal{E}_i^0$ be the set of even order orbifold points in $\mathcal{O}_i^0$. We begin with the following.

Lemma 6.5. *Suppose $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and Γ_1^0 is discrete. The set $C = (q_1^0)^{-1}(\mathcal{E}_1^0)$ is a concurrency set for (φ_1, φ_2).*

Proof. Note that by the definition of Γ_1^0, C is a finite set and $\mathrm{cone}(\varphi_1) = C_0 \subset C$. Now fix any $\zeta \in C$. If $\zeta \in C_0$ we know that it is a concurrency point and its lifts $\hat{p}^{-1}(\zeta) \subset \hat{S}$ are concurrency points by Corollary 5.2. So assume $\zeta \in (q_1^0)^{-1}(\mathcal{E}_1^0) \setminus C_0$ and consider any $\hat{\zeta} \in \hat{p}^{-1}(\zeta) \subset \hat{S}$. Recall that $\mathcal{G}(\hat{\varphi}_1, \hat{\zeta})$ consists of all basic $\hat{\varphi}_1$–geodesics passing through $\hat{\zeta}$. Since $\hat{\zeta}$ is not a cone point, any two geodesics in $\mathcal{G}(\hat{\varphi}_1, \hat{\zeta})$ intersect in exactly the one point $\hat{\zeta}$.

Claim. *The set of basic $\hat{\varphi}_2$–geodesics $\{\hat{g}(\eta) \mid \eta \in \mathcal{G}(\hat{\varphi}_1, \hat{\zeta})\}$ are concurrent.*

Proof. First, for any two geodesics $\eta_1, \eta_2 \in \mathcal{G}(\hat{\varphi}_1, \hat{\zeta})$, their $\hat{\varphi}_2$–straightenings $\hat{g}(\eta_1), \hat{g}(\eta_2)$ must intersect in exactly one point, and that point is not a cone point. To see this, note that the endpoints of $\hat{g}(\eta_i)$ are the same as the endpoints of η_i, for $i = 1, 2$, and the endpoints of η_1 and η_2 link. The intersection is a single non-cone point because η_i and $\hat{g}(\eta_i)$ have the same combinatorics, for $i = 1, 2$ by Lemma 5.1. Now, suppose $\eta_0, \eta_1, \eta_2 \in \mathcal{G}(\hat{\varphi}_1, \hat{\zeta})$ are any three distinct geodesics, let $\hat{g}(\eta_0) \cap \hat{g}(\eta_1) = \hat{\xi}_1$ and $\hat{g}(\eta_0) \cap \hat{g}(\eta_2) = \hat{\xi}_2$. We must show that $\hat{\xi}_1 = \hat{\xi}_2$. To do this, we first analyze the point $\hat{\zeta}$ in more detail.

Since $p_1(D_1(\hat{\zeta})) \in \mathcal{O}_1^0$ is an even order orbifold point, there exists an involution $\tau \in \Gamma_1^0$, rotating about the fixed point $D_1(\hat{\zeta})$ through angle π. Since ψ^0 is induced by conjugation by h, $\psi^0(\tau)$ is also a rotation through angle π about some point $\xi \in \mathbb{H}$.

For each $i = 0, 1, 2$, $D_1(\eta_i)$ is a geodesic through $D_1(\hat{\zeta})$, and we let x_i and $y_i = \tau \cdot x_i$ denote its endpoints in $\partial \mathbb{H}$. By Proposition 5.4, the endpoints of $D_2(\hat{g}(\eta_i))$ are $h(x_i)$ and $h(y_i)$. On the other hand, by Proposition 6.1, we have

$$h(y_i) = h(\tau \cdot x_i) = (h\tau h^{-1}) \cdot h(x_i) = \psi^0(\tau) \cdot h(x_i).$$

Therefore, $D_2(\hat{g}(\eta_i))$ is also invariant by $\psi^0(\tau)$, and thus passes through ξ, for $i = 0, 1, 2$. Since ξ is the unique point of $D_2(\hat{g}(\eta_0)) \cap D_2(\hat{g}(\eta_i))$, for $i = 1, 2$, we must have $D_2(\hat{\xi}_1) = \xi = D_2(\hat{\xi}_2)$. But $\hat{\xi}_1, \hat{\xi}_2 \in \hat{g}(\eta_0)$, and D_2 is injective on $\hat{g}(\eta_0)$, and therefore $\hat{\xi}_1 = \hat{\xi}_2$, as required. $\square$

Let $\hat{\xi}$ be the concurrent point of $\{\hat{g}(\eta) \mid \eta \in \mathcal{G}(\hat{\varphi}_1, \hat{\zeta})\}$. We have thus shown that

$$\mathcal{G}_{\hat{\varphi}_1}(\hat{\zeta}) \subset \mathcal{G}_{\hat{\varphi}_2}(\hat{\xi}).$$

We may reverse the roles of $\hat{\varphi}_1$ and $\hat{\varphi}_2$ since $D_2(\hat{\xi})$ is the fixed point of an even order elliptic element of Γ_2^0, and is not a cone point. Doing so proves the reverse inclusion, and shows that ζ is a concurrency point for (φ_1, φ_2), as required. $\square$

The next proposition is the final ingredient for the proof of the Current Support Theorem.

Proposition 6.6. *Suppose* $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, Γ_1^0 *is discrete,* $C = (q_1^0)^{-1}(\mathcal{E}_1^0)$ *is the concurrency set from Lemma 6.5, and that* φ_2 *is* C*–uber-normalized. Then there is a homeomorphism* $f: S \to S$, *isotopic to the identity relative to* C, *and a homeomorphism* $F: \mathcal{O}_1^0 \to \mathcal{O}_2^0$ *so that* $q_2^0 \circ f = F \circ q_1^0$.

Observe that the conclusion of this proposition implies that φ_1 and φ_2 arise as in Theorem 3.1, and thus proving this proposition will essentially complete the proof of the Current Support Theorem.

Proof. Let $\tilde{\Delta}_i$ be the cell structures on $\mathbb{H}$ from Lemma 3.4 applied to $G_i = \Gamma_i^0$ and the isomorphism ψ^0, let $\tilde{\Delta}'_i$ be the triangulation of $\mathbb{H}$ obtained by subdividing $\tilde{\Delta}_i$, for $i = 1, 2$, and $\tilde{F}: (\mathbb{H}, \tilde{\Delta}'_1) \to (\mathbb{H}, \tilde{\Delta}'_2)$ the canonical map given by the second half of Lemma 3.4. We write $F: \mathcal{O}_1^0 \to \mathcal{O}_2^0$ to denote the descent of $\tilde{F}$, by the universal (orbifold) coverings $p_i: \mathbb{H} \to \mathcal{O}_i^0$, for $i = 1, 2$. The map F is the canonical map with respect to the descended triangulations, Δ'_1 and Δ'_2 on $\mathcal{O}_1^0$ and $\mathcal{O}_2^0$, respectively.

We will adjust φ_2 by a homeomorphism $f: S \to S$ isotopic to the identity relative to C (without changing its name, as usual). Thus to prove the proposition, it suffices to show that after this adjustment, the following diagram commutes.
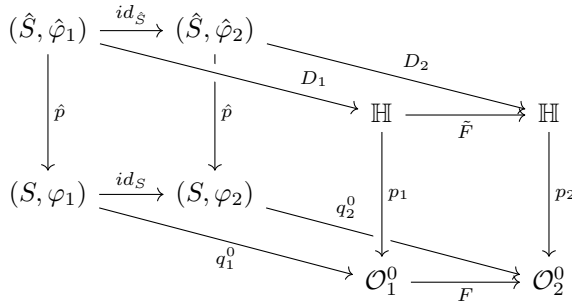
The "back" square obviously commutes ($\hat{p} = \hat{p}$), and since F is the descent of the equivariant map $\tilde{F}$ this implies that the "front" square commutes ($p_2 \circ \tilde{F} = F \circ p_1$). Similarly, the two "sides" commute ($p_i \circ D_i = q_1^0 \circ \hat{p}$, for $i = 1, 2$) by equivariance and descent. Since the "bottom" square is what we want to prove, and is the descent of the "top" square, it suffices to prove the top commutes; that is, we must show that $\tilde{F} \circ D_1 = D_2$.

Let $\hat{\Delta}$ be the cell structure on $\hat{S}$ obtained by pulling back $\tilde{\Delta}_1$ with respect to $D_1 : \hat{S} \to \mathbb{H}$. The vertex set of $\hat{\Delta}$ is the concurrency set, that is, $\hat{\Delta}^{(0)} = \hat{C}$. Since the 1–skeleton of $\tilde{\Delta}_1$ is a union of biinfinite geodesics, the 1–skeleton of $\hat{\Delta}$ is a union of basic $\hat{\varphi}_1$–geodesics. Specifically, there is a maximal set of basic geodesics $\Upsilon_1 \subset \mathcal{G}(\hat{\varphi}_1)$ so that

$$\hat{\Delta}^{(1)} = \bigcup_{\eta \in \Upsilon_1} \eta.$$

Since D_1 is ρ_1–equivariant, and $\rho(\pi_1 \dot{S}) = \Gamma_1 < \Gamma_1^0$, it follows that $\hat{\Delta}$ is $\pi_1 \dot{S}$–invariant, and hence by maximality, Υ_1 is $\pi_1 \dot{S}$–invariant.

The 2–cells of $\hat{\Delta}$ are convex hyperbolic polygons that map isometrically by D_1 to their image 2–cells in $\mathbb{H}$. Since the action of $\pi_1 \dot{S}$ on $\hat{S}$ is free except at the points of $\hat{C}_0$, it follows that the stabilizer of each 2–cell (a compact polygon) is trivial: otherwise it would have a fixed point in the interior. It follows that $\hat{p}$ maps the interior of every 2–cell homeomorphically to S, and thus $\hat{\Delta}$ descends to a cell structure Δ on S, all of whose 2–cells are hyperbolic polygons with respect to φ_1 (possibly with some edges identified). The zero-skeleton is $\Delta^{(0)} = \hat{p}(\hat{C}) = C$.

Subdivide Δ to a triangulation $\mathcal{T}$, then adjust φ_2 by a homeomorphism isotopic to the identity so that $\mathcal{T}$ is a $(C, \varphi_1, \varphi_2)$–triangulation. Since φ_2 is C–uber-normalized, the adjustment is by a homeomorphism isotopic to the identity relative to C. This pulls back to a $\pi_1 \dot{S}$–invariant triangulation $\widehat{\mathcal{T}}$. Any geodesic $\eta \in \Upsilon_1$ is a concatenation of edges of $\hat{\Delta}$, and hence edges in $\widehat{\mathcal{T}}$, and thus it follows that $\hat{g}(\eta) = \eta$. Moreover, the identity restricts to the canonical map on each C–saddle connection of η, with respect to $\hat{\varphi}_1$ on the domain and $\hat{\varphi}_2$ on the range.

Now we see how $\tilde{F}$ is related to D_1 and D_2. For any vertex $\zeta \in \hat{\Delta}^{(0)} = \hat{C}$, if $\zeta_1 = D_1(\zeta)$ and $\zeta_2 = D_2(\zeta)$, then because

$$\psi^0(\tau_{\zeta_1}) = h\tau_{\zeta_1}h^{-1} = \tau_{\zeta_2},$$

we have $\tilde{F}(\zeta_1) = \zeta_2$. That is

$$\tilde{F} \circ D_1(\zeta) = D_2(\zeta).$$

Furthermore, for any $\eta \in \Upsilon_1$, since η is a basic geodesic for both $\hat{\varphi}_1$ and $\hat{\varphi}_2$, $D_1(\eta)$ and $D_2(\eta)$ are both geodesics in $\mathbb{H}$. The geodesic η passes through infinitely many points of $\hat{C}$. Let $\zeta, \zeta' \in \eta \cap \hat{C}$ be two distinct such points, and observe that by the equation above we have

$$\tilde{F} \circ D_1(\zeta) = D_2(\zeta) \quad \text{and} \quad \tilde{F} \circ D_1(\zeta') = D_2(\zeta').$$

Since $D_1(\eta)$ and $D_2(\eta)$ are determined by two points on them, it follows that $\tilde{F} \circ D_1(\eta) = D_2(\eta)$. Moreover, for every edge e of $\hat{\Delta}^{(1)}$ contained in η, the identity in $\hat{S}$ is the canonical geodesic segment map when restricted to e, and since $\tilde{F}$ is also the canonical map from $D_1(e)$ to $D_2(e)$, naturality of canonical maps implies $\tilde{F} \circ D_1$ and D_2 agree on e. So, $\tilde{F} \circ D_1$ and D_2 agree on $\hat{\Delta}^{(1)}$.

For any 2–cell σ of $\hat{\Delta}$, the restrictions of D_1 and D_2 to σ are isometries with respect to $\hat{\varphi}_1$ and $\hat{\varphi}_2$, respectively, and the composition

$$D_2|_\sigma^{-1} \circ \tilde{F} \circ D_1|_\sigma$$

is the identity on $\partial\sigma$. We may therefore adjust φ_2 by a homeomorphism $S \to S$ isotopic to the identity by an isotopy which is the identity outside $\hat{p}(\sigma)$, so that $D_2|_\sigma^{-1} \circ \tilde{F} \circ D_1|_\sigma$ is the identity. Equivalently, $\tilde{F} \circ D_1|_\sigma = D_2|_\sigma$. Performing this adjustment on each of the finitely many 2–cell of Δ we have $\tilde{F} \circ D_1 = D_2$, as required. This completes the proof. $\square$

Proof of the Current Support Theorem. Suppose $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ and $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. If Γ_1^0 is indiscrete, then Proposition 6.3 implies $\varphi_1 \sim \varphi_2$. We may therefore assume Γ_1^0 is discrete, in which case Proposition 6.6 guarantees that φ_1 and φ_2 come from branched covers of orbifolds, as in Theorem 3.1 completing the proof. $\square$

Remark. Given $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}}_c(S)$ with $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and $\varphi_1 \not\sim \varphi_2$, we know that the φ_2 is obtained by deforming $\mathcal{O}_1^0$ and pulling back via q_1^0. It follows that the space of equivalence classes of all such metrics in $\mathrm{Hyp}_c(S)$,

$$\{\varphi \in \widetilde{\mathrm{Hyp}}_c(S) \mid \mathcal{G}_{\tilde{\varphi}} = \mathcal{G}_{\tilde{\varphi}_1}\}/\sim$$

is parameterized by the Teichmüller space of the orbifold $\mathcal{O}_1^0$. In particular, the dimension of this space can be arbitrarily large (depending on $\mathcal{O}_1^0$).

6.3. Too many cone points implies rigidity

As another consequence of the Current Support Theorem, we show that when there are too many cone points, the metric is forced to be rigid.

Corollary 1.1. *If $\varphi \in \mathrm{Hyp}_c(S)$ has at least $32(g-1)$ cone points (where g is the genus of S) then φ is rigid.*

Proof. If φ is flexible then by the Current Support Theorem there is a locally isometric branched covering $(S, \varphi) \to \mathcal{O}$ over a hyperbolic orbifold $\mathcal{O}$ sending each of the cone points to an even order orbifold point. We will thus show that if there is such a branched covering, then the number of cone points is less than $32(g-1)$.

The proof requires a repeated application of the Gauss-Bonnet Theorem, and is similar in spirit to the proof of the "$84(g-1)$ Theorem"; see e.g. [11, Theorem 7.4]. The Gauss-Bonnet Theorem implies in our case that the total area of a hyperbolic cone surface or orbifold of genus g_0 with n cone/orbifold points having cone angles $\theta_1, \ldots, \theta_n$ is given by

$$\mathrm{Area} = 4\pi(g_0 - 1) + 2\pi n - \sum_{i=1}^n \theta_i. \tag{3}$$

Fix any $\varphi \in \mathrm{Hyp}_c(S)$ and observe that by (3), we have $\mathrm{Area}(\varphi) < 4\pi(g-1)$. Suppose that $p \colon (S, \varphi) \to \mathcal{O}$ is a branched cover as above and let d be the degree. Then we have

$$d = \frac{\mathrm{Area}(\varphi)}{\mathrm{Area}(\mathcal{O})} < \frac{4\pi(g-1)}{\mathrm{Area}(\mathcal{O})}. \tag{4}$$

Next, let k be the number of cone points of φ and let r be the number of even-order orbifold points of $\mathcal{O}$. We write $b_1, \ldots, b_r \in 2\mathbb{Z}_+$ to denote the orders of the orbifold points, so that the i^{th} orbifold point has cone angle $\frac{2\pi}{b_i}$. Let k_i be the number of cone points that map to the i^{th} orbifold point so that $k = \sum_i k_i$. Now observe that since the cone angle of any cone point of φ is *strictly greater* than 2π, it follows that the local degree of p near any cone point is at least 3. This gives us

$$d \geq 3 \max_i k_i \geq 3\frac{k}{r}.$$

Combining this with (4) we get

$$k < \frac{4r\pi(g-1)}{3\mathrm{Area}(\mathcal{O})}. \tag{5}$$

All that remains is to show that the right-hand side of (5) is at most $32(g-1)$. For this, we consider a case-by-case analysis depending on $r \geq 1$: for any fixed r, the right-hand side is maximized when $\mathrm{Area}(\mathcal{O})$ is minimized. For $r \geq 5$, we note that (3) implies that the area of $\mathcal{O}$ is at least $\pi(r-4)$, which occurs when $\mathcal{O}$ is a sphere with r orbifold points, all of order 2. In this case we have

$$k < \tfrac{4\pi r(g-1)}{3\pi(r-4)} = \tfrac{4(g-1)}{3}\left(\tfrac{r}{r-4}\right) = \tfrac{4(g-1)}{3}\left(1 + \tfrac{4}{r-4}\right) \le \tfrac{20(g-1)}{3} < 32(g-1).$$

Again appealing to (3), for the remaining values of r one can check that the minimum area of an orbifold with at least r even-order orbifold points is realized by a sphere with n orbifold points of orders $(b_1, \ldots, b_n)$ described by the following table:

r	Area	$(b_1, \ldots, b_n)$	$\tfrac{4r\pi(g-1)}{3\mathrm{Area}(\mathcal{O})}$
1	$\tfrac{\pi}{21}$	$(2,3,7)$	$28(g-1)$
2	$\tfrac{\pi}{12}$	$(2,3,8)$	$32(g-1)$
3	$\tfrac{\pi}{6}$	$(2,4,6)$	$24(g-1)$
4	$\tfrac{\pi}{2}$	$(2,2,2,4)$	$\tfrac{32}{3}(g-1)$

The case $r = 1$ is classically known to be the minimal area of any (orientable) hyperbolic orbifold (see e.g. [11, Theorem 7.10]). The cases $r = 2, 3, 4$ follow a similar argument to this classical result by enumerating the possibilities, appealing to the following two facts: (1) the genus must be zero (since positive genus and at least 2 orbifold points give area at least 2π) and (2) we must have

$$n - \sum_{i=1}^{n} \tfrac{1}{b_i} > 2$$

if $(b_1, \ldots, b_n)$ are the orders of orbifold points of a hyperbolic orbifold of genus zero.

Since each quantity in the last column is at most $32(g-1)$, this completes the proof. □

7. Billiards

Consider a finite sided (simply connected, but not necessarily convex) polygon P in the hyperbolic plane and trajectories of "billiard balls", modeled by continuous piecewise geodesic paths $\beta : \mathbb{R} \to P$. The billiard trajectories are geodesic segments in the interior of P and reflect off the sides of the polygon with angle of incidence equal to angle of reflection. Note that we do not consider trajectories that encounter vertices of P. We call these dynamical systems polygonal billiard tables in the hyperbolic plane.

As in the introduction, we assume an n–gon P comes equipped with a labeling of the sides by elements of the set $\mathcal{A} = \{1, 2, \ldots, n\}$, in counterclockwise cyclical order. Given a billiard trajectory $\beta : \mathbb{R} \to P$ its corresponding *bounce sequence* is the biinfinite sequence $\mathbf{b}(\beta) = (\ldots, b_{-1}, b_0, b_1, \ldots) \in \mathcal{A}^{\mathbb{Z}}$ corresponding to the ordered sequence of labels of the sides of P that β encounters, where b_0 is the first side encountered by $\beta([0, \infty))$. The *bounce spectrum* of P, denoted $\mathcal{B}(P)$, is the set of all bounce sequences realized by billiard trajectories on P. We say that two labeled n–gons P_1, P_2 have the *same bounce spectrum* if $\mathcal{B}(P_1) = \mathcal{B}(P_2)$.

We recall the main theorem about billiards from the introduction.

Billiard Rigidity Theorem. *Given hyperbolic polygons P_1, P_2, we have $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ if and only if*

1. *P_1 is isometric to P_2 by a label preserving isometry, or*
2. *P_1, P_2 are reflectively tiled and there exists a label-preserving homeomorphism $H \colon P_1 \to P_2$ that maps tiles to tiles, preserving their interior angles.*

We note that the two cases are not mutually exclusive as the homeomorphism H may be an isometry. Before launching into the proof of the Billiard Rigidity Theorem, we make the notion of reflective tilings precise and introduce the tools we need in the sections below.

7.1. Reflective tilings

A *tiling* of $\mathbb{H}$ is a cell decomposition so that the closed 2–cells, called the *tiles*, are compact convex hyperbolic polygons. For each tile, we require the intersection with the 0–skeleton to be the vertices (thus having interior angles less than π). The 1–cells are called the *edges* of the tiling; an edge contained in a tile will also be called a *side* of the tile. Given a tiling T of $\mathbb{H}$ let $R(T)$ be the group generated by reflections in the (geodesic lines containing the) edges.

We say that a polygon in $\mathbb{H}$ is *good* if it is a finite sided compact polygon with each of its interior angles being an integral submultiple of π, that is, of the form $\frac{\pi}{k}$ for some $k \in \mathbb{Z}$. The Poincaré polygon theorem implies that the group R generated by reflections in the sides of a good polygon $\mathbf{t}$ is discrete and the polygon is a fundamental domain (see, e.g. [20, Theorem 7.1.3]). The translates of $\mathbf{t}$ by the group R determine a tiling T with $R = R(T)$. We call such a tiling of $\mathbb{H}$ a *reflective tiling*.

Lemma 7.1. *A tiling T of $\mathbb{H}$ is reflective if and only if the fixed point set of any reflection in $R(T)$ is contained in the 1–skeleton.*

Proof. If T is a reflective tiling, then all tiles sharing a vertex differ by an element of $R(T)$ fixing that vertex. Since the interior angles of tiles are integral submultiples of π, we see that the biinfinite geodesic containing any edge is contained in the 1–skeleton. Therefore, the reflections in sides of any tile are contained in the 1–skeleton. If there were any other reflection in $R(T)$ whose fixed point set was not contained in the 1–skeleton, then it would necessarily pass through the interior of a tile. Since any tile is a fundamental domain, this is impossible, and hence the fixed point set of any reflection is in the 1–skeleton, as required.

Now we assume that the fixed point set of any reflection in $R(T)$ is contained in the 1–skeleton, and prove that T is reflective.

First observe that the 1–skeleton is a union of biinfinite geodesics: since the reflection in any edge is in $R(T)$, it must have its entire fixed line in the 1–skeleton by assumption,

and thus it follows that for every edge of T, the unique geodesic containing it is in the 1–skeleton.

Next, we claim that $R(T)$ preserves the tiling. To prove this, it suffices to show that it preserves the collection of biiinfinite geodesics whose union is the 1–skeleton (since $R(T)$ preserves the tiling if and only if it preserves the 1–skeleton). Thus let $g \in R(T)$ be any element and $\eta \subset \mathbb{H}$ a biinfinite geodesic in the 1–skeleton. The reflection r_η in η is a generator of $R(T)$, and hence $g r_\eta g^{-1} = r_{g(\eta)}$, the reflection in $g(\eta)$, is in $R(T)$. By assumption $g(\eta)$ is contained in the 1–skeleton, as required.

Since $R(T)$ preserves the titling, it is discrete and acts on the set of tiles. In fact, this action on the tiles is transitive. To see this, take any two tiles $\mathbf{t_1}, \mathbf{t_2}$ and consider a geodesic segment from a point in the interior of $\mathbf{t_1}$ to a point in the interior of $\mathbf{t_2}$ that misses the vertex set. Then the composition of the reflections in the sides of tiles encountered by the geodesic (in order) maps $\mathbf{t_1}$ to $\mathbf{t_2}$.

From the proof of transitivity, we see that $R(T)$ is generated by reflections in the sides of any fixed tile $\mathbf{t}$. All that remains to prove is that the interior angle at each vertex of a tile is an integral submultiple of π. This follows because the stabilizer of any vertex ζ is a dihedral group whose lines of reflection are precisely the biinfinite geodesic extensions of the edges adjacent to ζ. Therefore, $\mathbf{t}$ is a good polygon, and the tiling is reflective. $\square$

Corollary 7.2. *Given a discrete group R generated by reflections, including reflections in the sides of a compact polygon P, then the union of the fixed point sets of all reflections in R is the 1–skeleton of a reflective tiling.*

Proof. Since R is discrete, the set of fixed point sets of the reflections is a locally finite collection of lines: otherwise there would be a sequence of such lines that meet a fixed compact set, and from this it is easy to construct a nonconstant sequence of elements of R as products of pairs of reflections that converge to the identity. For any point in the interior of P which is disjoint from the union of the lines, the intersection of half planes containing the point and bounded by the lines in our set is contained in P, and is thus a compact, convex polygon $\mathbf{t}$. That is, the closure of one component of the complement of the lines in our set is a compact polygon $\mathbf{t}$. Given any other point ζ in the complement of the union of the lines, we can construct a geodesic from a point in $\mathbf{t}$ to ζ, and as in the proof above, from this construct a sequence of reflections whose composition takes $\mathbf{t}$ to the closure of the complementary component containing ζ. It follows that the lines form the 1–skeleton of a tiling, and by construction, the fixed point set of any reflection of R is contained in it. Thus, by Lemma 7.1, the tiling is reflective. $\square$

We now define a *reflective tiling of a polygon P* similarly to be a cell decomposition of P such that:

- the closed 2-cells (the *tiles*) are isometric copies of a good polygon, and any tile meets the 0–skeleton precisely in its vertices,

- if two tiles share a *side* (or *edge*, i.e. 1–cell) then they differ by a reflection in that side, and
- if a tile **t** shares a vertex with P, then the interior angle of **t** at that vertex is an *even* integer submultiple of π, i.e. is of the form $\frac{\pi}{2k}$ for some $k \in \mathbb{N}$.

The third condition comes up naturally in the billiards theorem, though it is not clear what relevance it has outside our context. We will need the following lemma in our proof of the Billiard Rigidity Theorem.

Lemma 7.3. *If P is a reflectively tiled polygon, then we can extend this tiling to a reflective tiling of $\mathbb{H}$. That is, there is a reflective tiling T of $\mathbb{H}$ such that the reflective tiling of P is a subcomplex of T.*

Proof. Because our polygons are assumed simply connected (hence connected), any two tiles differ by a sequence of reflections in tiles. The entire group generated by reflections in any given tile defines a reflective tiling since the tile is a good polygon. $\square$

7.2. Unfoldings

Next we define the concept of unfolding a polygonal billiard table. This connects polygonal billiard tables in the hyperbolic plane to the focus of the rest of this paper, hyperbolic cone metrics on closed orientable surfaces. This a generalization (as in [9]) of a classical construction (see e.g. [24]).

Given a polygon P in $\mathbb{H}$, consider its *double*

$$\mathcal{D}P = P \times \{0, 1\}/\{(x, 0) \sim (x, 1) \;\forall x \in \partial P\}$$

which admits a hyperbolic cone metric for which the cone points are the vertices (or more precisely, the identified vertices of the two copies of P). Observe that all cone angles are less than 2π, however, so these metrics are *not* negatively curved. Note that there is a canonical map $\mathcal{D}P \to P$ sending $(x, t) \mapsto x$ for $t = 0, 1$, which is a local isometry away from the identified boundaries. An *unfolding* of P is any locally isometric branched cover $(S, \varphi) \to \mathcal{D}P$, branched over the cone points, so that the total angle around any point in the pre-image of a cone point of P is greater than 2π. In particular $\varphi \in \widetilde{\mathrm{Hyp}}_c(S)$ and all cone points map to cone points.

Given a hyperbolic polygon P, observe that every (biinfinite) nonsingular geodesic on $\mathcal{D}P$ projects to a (biinfinite) billiard trajectory. Moreover, every billiard trajectory arises in this way; indeed, this sets up a 2-to-1 correspondence between nonsingular geodesics on $\mathcal{D}P$ and billiard trajectories on P. If $(S, \varphi) \to \mathcal{D}P$ is an unfolding (with degree k, say) then the nonsingular φ–geodesics on S project to nonsingular geodesics on $\mathcal{D}P$, and all nonsingular geodesics arise in this way. Combined with the above, this sets up a $2k$-to-1 correspondence between nonsingular φ–geodesics on S and billiard trajectories on P.

For any two labeled n–gons P_1 and P_2 in $\mathbb{H}$ we may construct a common unfolding in the following sense. First, fix a labeled *topological n–gon P*: this is a CW complex with n vertices and n edges attached to form a circle, with the edges labeled by $\mathcal{A}$ in order cyclically, and a single 2–cell (a disk) whose attaching map is a homeomorphism to the 1–skeleton. The cell structure on P lifts to one on its double, $\mathcal{D}P$. For each $i = 1, 2$ fix a homeomorphism $P \to P_i$ respecting the labeling, which induces a homeomorphism of the doubles $\mathcal{D}P \to \mathcal{D}P_i$. A *common unfolding* of P_1 and P_2 is a branched covering $S \to \mathcal{D}P$ branched over the vertices together with a pair of metrics $\varphi_1, \varphi_2 \in \widetilde{\mathrm{Hyp}_c}(S)$ so that the composition $(S, \varphi_i) \to \mathcal{D}P \to \mathcal{D}P_i$ is an unfolding for each $i = 1, 2$. Observe that there is a natural cell structure on S obtained by lifting the one on $\mathcal{D}P$ and that with respect to φ_i, the closed 2–cells are isometric copies of P_i: indeed, the map $(S, \varphi_i) \to \mathcal{D}P \to \mathcal{D}P_i \to P_i$ restricts to an isometry on each closed 2–cell. Moreover, the edges have labels in $\mathcal{A}$, induced from those on the edges of $\mathcal{D}P$.

Now, let S be a common unfolding of two n-gons P_1 and P_2 with corresponding metrics φ_1 and φ_2. Let $p : \tilde{S} \to S$ denote the universal cover of S, and as usual let $\tilde{\varphi}_1, \tilde{\varphi}_2$ be the pull back metrics by p. Through the correspondence between billiard trajectories and nonsingular φ_i–geodesics on S we get the following:

Lemma 7.4. *Suppose P_1 and P_2 are two n–gons (with labeled sides) and let S be a common unfolding with metrics φ_1 and φ_2, respectively. If $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ then $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and φ_2 is C_0–uber-normalized with respect to φ_1 (where $C_0 = \mathrm{cone}(\varphi_1)$, as usual).*

Proof. First observe that the cell structure on S with labeled edges lifts to a cell structure on $\tilde{S}$ with labeled edges. With respect to $\tilde{\varphi}_i$, each closed 2–cell is isometric to P_i, and as with S, the map $(\tilde{S}, \tilde{\varphi}_i) \to (S, \varphi_i) \to P_i$ restricts to an isometry on each 2–cell.

Now suppose $\mathcal{B}(P_1) = \mathcal{B}(P_2)$. Let $\{x, y\} \in \mathcal{G}_{\tilde{\varphi}_1}$ be the endpoints of a nonsingular $\tilde{\varphi}_1$–geodesic η. Reading off the labels crossed gives a sequence $\mathbf{b} \in \mathcal{A}^{\mathbb{Z}}$. Since this nonsingular $\tilde{\varphi}_1$–geodesic projects to a nonsingular φ_1–geodesic in S, and then down to billiard trajectory in P_1, it follows that $\mathbf{b} \in \mathcal{B}(P_1)$. Since $\mathcal{B}(P_1) = \mathcal{B}(P_2)$, $\mathbf{b} \in \mathcal{B}(P_2)$ too, and hence there is a nonsingular $\tilde{\varphi}_2$–geodesic η' crossing the exact same set of edges of the 1–skeleton as η which projects to the billiard trajectory in P_2 with bounce sequence $\mathbf{b}$. Since η and η' cross the exact same set of edges in $\tilde{S}$, they remain a bounded distance apart, and hence the endpoints of η' are $\{x, y\} \in \mathcal{G}_{\tilde{\varphi}_2}$. Since endpoints of nonsingular $\tilde{\varphi}_i$–geodesics are dense in $\mathcal{G}_{\tilde{\varphi}_i}$, it follows that $\mathcal{G}_{\tilde{\varphi}_1} \subset \mathcal{G}_{\tilde{\varphi}_2}$. The other inclusion follows by a symmetric argument and we have $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$.

To see that φ_2 is C_0–uber-normalized, first observe that $\tilde{C}_0 = \mathrm{cone}(\tilde{\varphi}_1) = \mathrm{cone}(\tilde{\varphi}_2)$ (since the metrics come from a common unfolding of P_1 and P_2). Moreover, the $\pi_1 S$–equivariant bijection $\mathrm{cone}(\tilde{\varphi}_1) \to \mathrm{cone}(\tilde{\varphi}_2)$ from Proposition 4.1 is the *identity*. This follows directly from the fact that $g : \mathcal{G}(\tilde{\varphi}_1) \to \mathcal{G}(\tilde{\varphi}_2)$ preserves partitions (Lemma 4.4) and the proof above showing that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. Next we subdivide the cell structure on S to a (C_0, φ_1)–triangulation $\mathcal{T}$. We may adjust φ_2 by a homeomorphism isotopic to the identity so that each of the triangles are also (C_0, φ_2)–triangles. Since the 1–cells are

already saddle connections with respect to both φ_1 and φ_2, the isotopy can be assumed to preserve the 1–skeleton of the cell structure, and hence is relative to C_0. After further isotopy relative to C_0 if necessary we can assume that the identity is the canonical map for $\mathcal{T}$, and thus $\mathcal{T}$ is a $(C_0, \varphi_1, \varphi_2)$–triangulation. Since all adjustments to φ_2 were by homeomorphism isotopic to the identity relative to C_0 and the result is a metric that is C_0–uber-normalized (by Corollary 4.12), the original φ_2 is C_0–uber-normalized. $\square$

7.3. Reflection groups

Let P_1, P_2 be two n–gons and $S \to \mathcal{D}P \to \mathcal{D}P_i$, $i = 1, 2$, a common unfolding with corresponding metrics φ_1, φ_2. If $\mathcal{B}(P_1) = \mathcal{B}(P_2)$, then Lemma 7.4 implies $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ and φ_2 is C_0–uber-normalized. Let $\hat{p} \colon \hat{S} \to S$ be the extension to the completion of the universal cover of $\dot{S} = S \setminus C_0$, as defined in §2.2, and $\hat{\varphi}_1, \hat{\varphi}_2$ the pull backs of φ_1, φ_2 by $\hat{p}$. By Lemma 5.1 we also have $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$. We assume that S and $\widetilde{S}$ are given the cell structures with labeled edges as described in the previous section, as is $\hat{S}$ in a similar fashion.

Viewing $\dot{P} = P \setminus \{\text{vertices of } P\}$ as an orbifold with reflectors as sides, let $R = \pi_1^{orb}(\dot{P})$. Then $\pi_1 \mathcal{D}\dot{P} < R$ is an index 2 subgroup, and the maps $\dot{S} \to \mathcal{D}\dot{P} \to \dot{P}$ induce inclusions

$$\pi_1 \dot{S} < \pi_1 \mathcal{D}\dot{P} < R$$

and all three groups act on $\hat{S}$ by isometries with respect to both $\hat{\varphi}_1$ and $\hat{\varphi}_2$. For $i = 1, 2$ let $D_i : \hat{S} \to \mathbb{H}$ be the developing maps and $\rho_i : R \to \text{Isom}(\mathbb{H})$ the corresponding holonomy homomorphisms. We fix some 2–cell B in $\hat{S}$ and assume that D_i maps this 2–cell to P_i, for $i = 1, 2$. Then, ρ_i sends the reflection in the j^{th} side of B to reflection in the j^{th} side of P_i.

For each $i = 1, 2$, let $R_{P_i} = \rho_i(R) < \text{Isom}(\mathbb{H})$ and note that R_{P_i} is the group generated by reflections in the sides of P_i. For any $\zeta \in \mathbb{H}$ let τ_ζ denote a rotation about ζ of order 2 and let

$$R_{P_i}^0 = \langle R_{P_i}, \{\tau_v\}_{v \in V(P_i)} \rangle$$

where $V(P_i)$ denotes the vertex set of P_i. Note that $R_{P_i}^0$ is also generated by reflections because each τ_v is in $\text{Stab}_{R_{P_i}^0}(v)$, which is a dihedral group and is thus generated by its reflections.

Lemma 7.5. *Suppose $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ and that $S \to \mathcal{D}P \to \mathcal{D}P_i$, for $i = 1, 2$, is a common unfolding with associated metrics φ_1, φ_2. Further suppose that we have chosen developing maps $D_1, D_2 \colon \hat{S} \to \mathbb{H}$ as above with associated holonomy homomorphisms ρ_1, ρ_2, respectively. Then there exists an orientation preserving homeomorphism $h : \partial\mathbb{H} \to \partial\mathbb{H}$ topologically conjugating ρ_1 to ρ_2. In particular, R_{P_1} to R_{P_2} are topologically conjugate by h, with h conjugating the reflection in the j^{th} sides of P_1 to the reflection in the j^{th} side of P_2.*

Proof. Since $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ we have that $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ with φ_2 being C_0–uber-normalized by Lemma 7.4, and $\mathcal{G}_{\hat{\varphi}_1} = \mathcal{G}_{\hat{\varphi}_2}$ by Lemma 5.1. Since R acts on $\hat{S}$ by isometries with respect to $\hat{\varphi}_1$ and $\hat{\varphi}_2$, the existence of the conjugating homeomorphism h is guaranteed by Proposition 5.4. From the same proposition we have that

$$h(\rho_1(\gamma) \cdot x) = \rho_2(\gamma) \cdot h(x)$$

for all $\gamma \in R$ and $x \in \partial \mathbb{H}$. Hence if $\gamma \in R$ corresponds to reflection in the j^{th} side of P then $\rho_i(\gamma)$ is the reflection across the j^{th} side of P_i. Therefore, the equation above expresses precisely that h conjugates reflection in the j^{th} sides of P_1 and P_2. $\square$

The following is useful in the proof of the Billiard Rigidity Theorem.

Lemma 7.6. If $R^0_{P_1}$ is either indiscrete or contains reflections in all three sides of a triangle, then P_1 is billiard rigid. That is, if P_2 is another polygon with $\mathcal{B}(P_1) = \mathcal{B}(P_2)$, then there exists a label preserving isometry from P_1 to P_2.

Proof. Suppose $\mathcal{B}(P_1) = \mathcal{B}(P_2)$. By Lemma 7.5 R_{P_1} and R_{P_2} are topologically conjugate by a homeomorphism $h \colon \partial \mathbb{H} \to \partial \mathbb{H}$. It follows from Proposition 6.1 (after choosing a common unfolding and developing maps D_1, D_2 as above) that $R^0_{P_1}$ and $R^0_{P_2}$ are also topologically conjugate by h since the j^{th} vertex of P_1 and P_2 are given by $D_1(\zeta^j)$ and $D_2(\zeta^j)$ for some $\zeta^j \in \hat{C}_0$. If $R^0_{P_1}$ is indiscrete then it follows by Proposition 2.7 that h is in $\mathrm{PSL}_2(\mathbb{R})$.

Next suppose there is a hyperbolic triangle such that $R^0_{P_1}$ contains reflections in all three of its sides; call these reflections r, s, t and let R^Δ be the subgroup of $R^0_{P_1}$ they generate. If R^Δ is indiscrete then so is $R^0_{P_1}$ and h is in $\mathrm{PSL}_2(\mathbb{R})$ by the above, so assume R^Δ is discrete. Let $\Gamma^\Delta = \langle sr, tr \rangle$ be the subgroup of R^Δ generated by the two elliptic elements sr and tr (which must be of finite order since Γ^Δ is discrete). Note that Γ^Δ is orientation preserving since its generators are. Next, observe that $R^\Delta = \langle \Gamma^\Delta, r \rangle$ and that $r(sr)r = (sr)^{-1} \in \Gamma^\Delta$ and $r(tr)r = (tr)^{-1} \in \Gamma^\Delta$ and hence $r\Gamma^\Delta = \Gamma^\Delta r$. It follows that Γ^Δ is an index 2 subgroup of R^Δ and in particular it is cocompact since R^Δ is. We have shown that Γ^Δ is a discrete cocompact subgroup of $\mathrm{PSL}_2(\mathbb{R})$ generated by two finite order elements; the only such groups are Fuchsian triangle groups (see [19]). It follows by Lemma 2.8 that h must lie in $\mathrm{PSL}_2(\mathbb{R})$.

Since h is in $\mathrm{PSL}_2(\mathbb{R})$ we can extend it to an isometry of $\mathbb{H}$ which we also denote by h. Since h conjugates the reflection in the j^{th} side of P_1 to the reflection in the j^{th} side of P_2, it sends the geodesic extension of the j^{th} side of P_1 to the geodesic extension of the j^{th} side of P_2. This implies that h sends sides of P_1 to sides of P_2, so defines a label preserving isometry $P_1 \to P_2$. $\square$

Before proving the full billiard rigidity theorem in the next section, we already now see as a direct consequence of the above lemma that billiard rigidity is generic:

Corollary 7.7. *If any interior angle of a polygon P is an irrational multiple of π, then P is billiard rigid.*

Proof. The assumption on P implies R_P is indiscrete, and hence Lemma 7.6 implies P is billiard rigid. $\square$

7.4. The Billiard Rigidity Theorem and examples

We are now ready to prove the theorem.

Proof of the Billiard Rigidity Theorem. Suppose $\mathcal{B}(P_1) = \mathcal{B}(P_2)$. If $R_{P_1}^0$ is indiscrete we are done by Lemma 7.6 so we will assume that it is discrete. Then, for each $i = 1, 2$ the union of the lines of reflection for $R_{P_i}^0$ defines a reflective tiling T_i by Corollary 7.2, and P_i is a union of tiles.

The homeomorphism $h : \partial \mathbb{H} \to \partial \mathbb{H}$ from Lemma 7.5 topologically conjugates $R_{P_1}^0$ to $R_{P_2}^0$, and so defines a bijection, h_*, between the set of oriented lines of reflection of $R_{P_1}^0$ to those of $R_{P_2}^0$ via the action on the endpoints of such lines. In particular, the lemma guarantees that h_* sends the line containing the j^{th} side of P_1 to the line containing the j^{th} side of P_2. Note that h_* also preserves concurrence of any pair of lines of reflection because the point of intersection is the unique fixed point of a finite subgroup and h conjugates this subgroup of R_1^0 to the corresponding subgroup of R_2^0. Hence, h determines an equivariant bijection H' from the vertices of T_1 to the vertices of T_2.

We would like to extend H' to a map $H \colon (\mathbb{H}, T_1) \to (\mathbb{H}, T_2)$ that sends P_1 to P_2 by first extending over the 1–skeleton, $T_1^{(1)}$, and then over the tiles. For each edge $[v_1, v_2]$ of the 1–skeleton of T_1, we define $H|_{[v_1, v_2]}$ to send $[v_1, v_2]$ to the geodesic between $H'(v_1)$ and $H'(v_2)$ by the canonical map. We have thus defined an equivariant map $H|_{T_1^{(1)}}$, and for any line ℓ of reflection for $R_{P_1}^0$, $H|_{T_1^{(1)}}$ maps all vertices in ℓ to vertices in $h_*(\ell)$, and hence $H|_{T_1^{(1)}}(\ell) = h_*(\ell)$. To show that $H|_{T_1^{(1)}}$ is a homeomorphism, it suffices to show that for any line of reflection, ℓ, the ordering of the vertices along ℓ agrees with the H'–image ordering of vertices along $h_*(\ell)$.

In order to get a contradiction, suppose ℓ is an oriented line of reflection of $R_{P_1}^0$ and let v_1 and v_2 be vertices along ℓ such that $H|_{T_1^{(1)}}(v_1)$ and $H|_{T_1^{(1)}}(v_2)$ occur in the opposite order along $h_*(\ell)$. Take two lines of reflection ℓ_1 and ℓ_2 of $R_{P_1}^0$ that are transverse to ℓ at v_1 and v_2 respectively. If ℓ_1 and ℓ_2 do not intersect, then $H|_{T_1^{(1)}}$ must not reorder v_1 and v_2, since h preserves the cyclic order of the end points of ℓ, ℓ_1, ℓ_2. If they do intersect, then there is a subgroup of $R_{P_1}^0$ generated by reflections in sides of a triangle, and P_1 and P_2 are isometric by Lemma 7.6, which is impossible if the order of v_1 and v_2 are reversed along ℓ (for, in this case, $H|_{T_1^{(1)}}$ is the restriction of an isometry conjugating $R_{P_1}^0$ to $R_{P_2}^0$). Therefore, the order of v_1 and v_2 along ℓ must be preserved by $H|_{T_1^{(1)}}$, and hence this map is an equivariant homeomorphism from $T_1^{(1)}$ to $T_2^{(1)}$. Note that the cyclic order of edges incident to a given vertex given by the embedding of the 1–skeleton of T_1

into $\mathbb{H}$ is preserved by $H|_{T_1^{(1)}}$, since h preserves the cyclic order of the endpoints of the geodesics containing these edges. Therefore, we may extend $H|_{T_1^{(1)}}$ further to the faces of the tiling T_1, giving a equivariant homeomorphism $H: (\mathbb{H}, T_1) \to (\mathbb{H}, T_2)$. Since h_* sends the line of reflection in the j^{th} side of P_1 to the line of reflection in the j^{th} side of P_2, it follows that H sends P_1 to P_2 and preserves labels.

For every v a vertex of P_1, we have $\tau_v \in \operatorname{Stab}_{R_{P_1}^0}(v)$, so this stabilizer is the dihedral group of an even sided polygon, and hence the angle of a tile at v must be of the form $\pi/2k$ for some $k \in \mathbb{N}$. Since the stabilizer of v in $R_{P_1}^0$ and $H(v)$ in $R_{P_2}^0$ is isomorphic dihedral groups, H preserves the angles of the tiles. This proves one direction of the theorem.

For the converse, assuming (1), we clearly get $\mathcal{B}(P_1) = \mathcal{B}(P_2)$, so assume (2). The reflective tiling T_i of P_i for $i = 1, 2$, determines reflective tilings of $\mathbb{H}$ of the same names. Changing H by an isotopy if necessary, we may assume that it is the restriction of a homeomorphism $\mathbb{H} \to \mathbb{H}$ that conjugates $R(T_1)$ to $R(T_2)$. The homeomorphism H induces a correspondence between the edges of P_1 and P_2 that preserves labelings, and we construct a common unfolding $S \to \mathcal{D}P$ with associated metric φ_1 and φ_2 with respect to label preserving homeomorphisms $P \to P_i$. The restricted map $H|_{P_1}$ lifts to a map $H_{\mathcal{D}P}: \mathcal{D}P_1 \to \mathcal{D}P_2$, and we assume that our homeomorphism $\mathcal{D}P \to \mathcal{D}P_2$ factors as the composition of $H_{\mathcal{D}P}$ with the homeomorphism $\mathcal{D}P \to \mathcal{D}P_1$.

The doubles $\mathcal{D}P_1$ and $\mathcal{D}P_2$ must branch-cover orbifolds $\mathcal{O}_1$ and $\mathcal{O}_2$, respectively, given by taking the quotient of $\mathbb{H}$ by the orientation preserving subgroup $\Gamma_i = R(T_i) \cap \operatorname{PSL}_2(\mathbb{R})$. This can be seen as follows. Let P_i and $\bar{P}_i$ be the two copies of the polygon used to construct $\mathcal{D}P_i$. Define the map $f_i: \mathcal{D}P_i \to \mathbb{H}/\Gamma_i$ on each of these copies. Take $f_i|_{P_i}: P_i \to \mathbb{H}/\Gamma_i$ to be the inclusion of P_i into the hyperbolic plane, followed by the quotient by the action of Γ_i. Take $f_i|_{\bar{P}_i}: \bar{P}_i \to \mathbb{H}/\Gamma_i$ to be the embedding of the polygon into the hyperbolic plane (same as before) *followed by reflecting it across one of its edges*, followed by the quotient by the action of Γ_i. Note that it does not matter what edge we reflect across in this construction, since any two choices differ by the action of an element of Γ_i. This is a branched covering because it is possible to find charts in $\mathbb{H}$ covering $\mathcal{D}P_i$ on which f_i is a local branched covering.

Since the homeomorphism H conjugates $R(T_1)$ to $R(T_2)$, it induces an orbifold homeomorphism $H_{\mathcal{O}}$ fitting into the following commutative diagram.

$$
\begin{array}{ccc}
(S, \varphi_1) & \xrightarrow{\ id_S\ } & (S, \varphi_2) \\
\downarrow & & \downarrow \\
\mathcal{D}P_1 & \xrightarrow{\ H_{\mathcal{D}P}\ } & \mathcal{D}P_2 \\
\downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} \\
\mathcal{O}_1 & \xrightarrow{\ H_{\mathcal{O}}\ } & \mathcal{O}_2
\end{array}
$$

Note that the cone points of S are sent to the cone points/vertices of $\mathcal{D}P_i$, which are then sent to even order orbifold points in $\mathcal{O}_i$. This is because at each polygon vertex, the

incident tiles of T_i have angle $\frac{\pi}{2k}$, so Γ_i must contain a rotation through angle $\frac{\pi}{k}$, which has even order.

Theorem 3.1 implies $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$, while by Lemma 4.13, the preimage C of the even order orbifold points of $\mathcal{O}_1$ is a concurrency set for (φ_1, φ_2) with φ_2 being C–uber-normalized. The cell structure on S has 2–cells that map isometrically to P_i with respect to φ_i via the common unfolding, and sides labeled according to the labels of P (and hence P_i). As usual, we lift this to a cell structure on $\widetilde{S}$ with induced labels on the edges. Since every edge is simultaneously a $\tilde{\varphi}_1$–saddle connection and $\tilde{\varphi}_2$–saddle connection, Lemma 4.8 implies that for every nonsingular $\tilde{\varphi}_1$–geodesic η, $g(\eta)$ crosses the same set of saddle connections.

Now given any $\mathbf{b} \in \mathcal{B}(P_1)$, let η be a nonsingular $\tilde{\varphi}_1$–geodesic that projects to the billiard trajectory in P_1 with bounce sequence $\mathbf{b}$. Recording the labels on the edges crossed by η we get the sequence $\mathbf{b}$. Since $g(\eta)$ is a $\tilde{\varphi}_2$–geodesic that crosses the same set of edges as η, it projects to a billiard trajectory in P_2 with the same bounce sequence $\mathbf{b}$, and thus $\mathcal{B}(P_1) \subset \mathcal{B}(P_2)$. A symmetric argument proves the other containment, and hence $\mathcal{B}(P_1) = \mathcal{B}(P_2)$. $\quad\square$

The next proposition is an application to the Billiard Rigidity Theorem, giving concrete examples of both rigid and flexible polygons.

Proposition 7.8. *Let $n \geq 4$. Suppose P, P' are n–gons (with labeled sides) and write α_i and α'_i to denote the interior angle between the sides of P and P', respectively, labeled i and $i + 1$ (indices taken modulo n). Then we have the following:*

(1) *If for all i, α_i is an even submultiple of π, then $\mathcal{B}(P) = \mathcal{B}(P')$ if and only if $\alpha_i = \alpha'_i$, for all i.*
(2) *If for all i, α_i is not an even submultiple of π, then $\mathcal{B}(P) = \mathcal{B}(P')$ if and only if P and P' are isometric by a label preserving isometry.*

Proof. First assume that we are in case (1), and hence α_i is an even submultiple of π for all i. If $\mathcal{B}(P) = \mathcal{B}(P')$ it follows immediately from the Billiard Rigidity Theorem that $\alpha_i = \alpha'_i$ for all i. For the converse, assume $\alpha_i = \alpha'_i$ for all i. In this case P (respectively, P') is reflectively tiled by the single tile P (respectively, P'). Then any label preserving homeomorphism $H \colon P \to P'$ trivially maps tiles to tiles and preserves interior angles, since $\alpha_i = \alpha'_i$, and so by the Billiard Rigidity Theorem we have $\mathcal{B}(P) = \mathcal{B}(P')$.

Now suppose none of the angles α_j are even submultiples of π. If P and P' are isometric by a label preserving isometry, then we must have $\mathcal{B}(P) = \mathcal{B}(P')$. For the converse, assume $\mathcal{B}(P) = \mathcal{B}(P')$. First note that if there exists i with $\alpha_i \notin \mathbb{Q}\pi$ then P and P' are isometric by a label preserving isometry by Corollary 7.7. Hence we will assume that each α_i is of the form $\frac{p_i \pi}{q_i}$ where $p_i, q_i \in \mathbb{N}$ with either $p_i > 1$ or $p_i = 1$ and q_i odd. We will show that R_P^0 contains reflections in all three sides of a triangle, and hence by Lemma 7.6 we must have that P and P' are again isometric by a label preserving isometry. To obtain such a triangle, we first show the following claim. Let v_i

be the vertex of P corresponding to angle α_i and let l_i be the geodesic ray emanating from v_i, traveling into P, and bisecting the angle α_i.

Claim. *There exists an i such that l_i and l_{i+1} intersect (where indices are taken modulo n).*

Proof of Claim. Pick any vertex v_j and cut P along l_j until the polygon has been divided into two. Note that at least one of the resulting polygons is an n_1–gon with $3 \leq n_1 < n$, call it P_1. Let v_k be the vertex adjacent to v_j in P_1 and not incident to l_j. Note that v_k is either v_{j+1} or v_{j-1}. If l_k intersects l_j we are done. If not, cut P_1 along l_k, resulting in two new polygons, and let P_2 be the one not having a subsegment of l_j as a side. Then P_2 is an n_2–gon with $3 \leq n_2 < n_1$. We repeat the process by considering the vertex v_l adjacent to v_k in P_2 and not incident to l_k. Note that either $(k,l) = (j+1, j+2)$ or $(k,l) = (j-1, j-2)$. This process must terminate since each step reduces the number of sides of the resulting polygon and if we arrive at a triangle, having a subsegment of a bisector l as one of its sides, the vertex v_i not incident to l must have l_i intersecting l. $\quad\square$

Now, using the claim we choose i such that l_i and l_{i+1} intersect. Let s denote the side of P labeled $i+1$, that is, the side between v_i and v_{i+1}. For $j = 1, 2$ write $\alpha_j = \frac{p_j \pi}{q_j}$ as above. If $p_j = 1$ set $\hat{l}_j$ to be the bisecting ray l_j defined above and if $p_j > 1$ let $\hat{l}_j$ be the ray emanating from v_j and making angle $\frac{\pi}{q_j}$ with s inside P. Note that in all cases, $\hat{l}_i$ and $\hat{l}_{i+1}$ must intersect and so s together with subsegments of $\hat{l}_i$ and $\hat{l}_{i+1}$ bound a triangle Δ. We claim that R_P^0 contains reflections in all three of its sides.

To see this, first note that R_P^0 contains the reflection r in the side s by definition and it contains a rotation about v_j of angle $2\alpha_j$ for $j = 1, 2$. If $\alpha_j = \frac{\pi}{q_j}$ for some odd q_j then, since R_P^0 contains rotations about v_j both of angle $\frac{2\pi}{q_j}$ and of angle π, it also contains the rotation σ_j about v_j of angle $\frac{\pi}{q_j}$. Reflection in $\hat{l}_j = l_j$ is then given by $\sigma_i \circ r$ for $j = i$ or by $\sigma_{i+1}^{-1} \circ r$ for $j = i+1$. If $\alpha_j = \frac{p_j \pi}{q_j}$, for $p_j > 1$, then R_P^0 contains a rotation about v_j of angle $\frac{2 p_j \pi}{q_j}$ and hence also the rotation σ_j about v_j of angle $\frac{2\pi}{q_i}$ and reflection in $\hat{l}_j$ is then given by $\sigma_i \circ r$ or $\sigma_{i+1}^{-1} \circ r$.

Hence R_P^0 contains reflections in all three sides of the triangle Δ as claimed and by Lemma 7.6 we are done. $\quad\square$

In case (1) of the proposition we can also compute the dimension of the space of polygons having the same bounce spectrum. To state this precisely, let

$$\mathfrak{P}(P) = \{P' \subset \mathbb{H} \mid \mathcal{B}(P') = \mathcal{B}(P)\}/\sim,$$

where $P' \sim P''$ if there is a label preserving isometry from P' to P''.

Corollary 7.9. *If P is an n–gon and all interior angles are even submultiples of π, then $\mathfrak{P}(P) \cong \mathbb{R}^{n-3}$.*
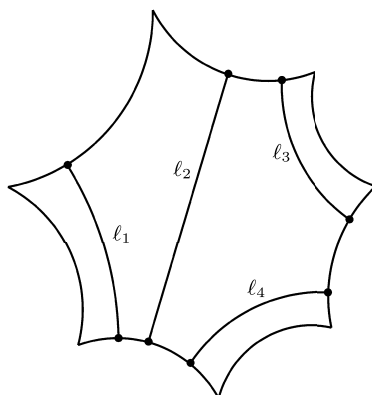
Fig. 15. Arcs subdividing a 7–gon with all acute interior angles into one pentagon, one hexagon, and three quadrilaterals. (The dots indicate orthogonal intersections.) The deformations of the polygon, preserving the interior angles are parameterized by the lengths $\ell_1, \ell_2, \ell_3, \ell_4$.

Proof. Let α_j be the interior angle between side j and $j+1$ of P for all j (indices taking modulo n). According to Proposition 7.8, $\mathfrak{P}(P)$ is precisely the set of polygons P' (up to label preserving isometry) for which the interior angle between the side j and j' is α_j for all j. This space is well-known to be homeomorphic to $\mathbb{R}_+^{n-3}$, but we briefly explain this for completeness. One approach to this is to construct a polygonal analogue of a pants decomposition and Fenchel-Nielsen coordinates (see e.g. [22, Section 4.6]) which we sketch now.

Take a maximal collection of pairwise disjoint arcs in P connecting distinct, non-adjacent sides, so that no two arcs connect the same pair of sides. The number of such arcs is precisely $n - 3$. Replacing each arc by a minimal length geodesic connecting the same sides, these geodesic segments produce a decomposition of P into right-angled hexagons, pentagons with at most one non-right angle (which occurs at a vertex of P) and quadrilaterals with at most two non-right angles (each of which occur at a vertex of P). We note that if angles α_j and α_{j+1} are both $\frac{\pi}{2}$ and if one of the arcs has endpoints on side j and $j+2$ (indices modulo n), then the geodesic replacement arc **is** the side labeled $j + 1$.

The lengths of the geodesic arcs $(\ell_1, \ldots, \ell_{n-3}) \in \mathbb{R}_+^{n-3}$ serve as parameters since the space of deformations of each of the complementary polygons, preserving all angles, is parameterized by these lengths (see e.g. [22, Section 4.6] for the case of a right-angled hexagon–the other cases are simpler; cf. [1, §7.17-7.19]). Fig. 15 illustrates an example. $\square$

By the Billiard Rigidity Theorem, a polygon P_1 can only be billiard flexible if it admits a reflective tiling for which the tile **t** is a good polygon. Moreover, we can deform P_1 to a non-isometric polygon P_2 with $\mathcal{B}(P_1) = \mathcal{B}(P_2)$ precisely by deforming the tile **t**. The deformation space of **t**, preserving the interior angles, has dimension $(n - 3)$ (and

is a point if $n = 3$) by the same argument as the one given in the proof of Corollary 7.9. This gives us the following.

Corollary 1.2 *A hyperbolic polygon P is billiard flexible if and only if it is reflectively tiled with a non-triangular tile.* □

7.5. Generalized diagonals

As discussed in the introduction, a *generalized diagonal* in a polygon P is finite billiard trajectory $\gamma \colon [a, b] \to P$ that starts and ends at a vertex of the polygon. Such a segment has a *finite* bounce sequence $\mathbf{b}(\gamma) = (b_1, \ldots, b_n)$ and we denote the set of all such finite bounce sequences $\mathcal{B}_\Delta(P)$.

Theorem 1.3 *Given polygons $P_1, P_2 \subset \mathbb{H}$, we have $\mathcal{B}_\Delta(P_1) = \mathcal{B}_\Delta(P_2)$ if and only if one of the two conclusions of the Billiard Rigidity Theorem holds.*

Proof. First, suppose $\mathcal{B}_\Delta(P_1) = \mathcal{B}_\Delta(P_2)$ and let S be a common unfolding with metrics φ_1 and φ_2, respectively. The same idea as in the proof of Lemma 7.4 implies that every $\tilde{\varphi}_1$–saddle connection is uniformly close to some $\tilde{\varphi}_2$–saddle connection (possibly having different initial/terminal endpoints) by considering the associated bounce sequences. Now observe that any non-singular $\tilde{\varphi}_i$–geodesic is a limit of $\tilde{\varphi}_i$–saddle connections, for $i = 1, 2$. Given a $\tilde{\varphi}_1$–geodesic η which is a limit of $\tilde{\varphi}_1$–saddle connections, η straightens to a basic $\tilde{\varphi}_2$–geodesic, and hence $\mathcal{G}_{\tilde{\varphi}_1} \subset \mathcal{G}_{\tilde{\varphi}_2}$. Interchanging the roles of $\tilde{\varphi}_1$ and $\tilde{\varphi}_2$, we get equality $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$. At this point, the rest of the proof of the Billiard Rigidity Theorem can be carried out verbatim since $\mathcal{G}_{\tilde{\varphi}_1} = \mathcal{G}_{\tilde{\varphi}_2}$ is the only fact that is used after choosing a common unfolding.

For the reverse implication, we assume either (1) or (2) of the Billiard Rigidity Theorem holds. As in the proof of the Billiard Rigidity Theorem we choose a common unfolding S with associated metrics φ_1 and φ_2 and deduce that $\mathcal{G}_{\varphi_1} = \mathcal{G}_{\varphi_2}$. As in that proof, appealing to Lemma 4.13 we see that the preimage of the even order orbifold points is a concurrency set $C \subset S$ for (φ_1, φ_2), with φ_2 C–uber-normalized with respect to φ_1. By Lemma 4.11 any $\tilde{\varphi}_1$–saddle connection straightens to a $\tilde{\varphi}_2$–saddle connection in the complement of C. As above we have a cell structure on $\widetilde{S}$ whose 2–cells are copies of the topological polygon P with labeled 1–cells, and it follows that a $\tilde{\varphi}_1$–saddle connection crosses exactly the same sequence of labeled 1–cells as the $\tilde{\varphi}_2$–straightening (since the 1–cells are geodesics with respect to both metrics). Therefore, every sequence in $\mathcal{B}_\Delta(P_1)$ is a sequence in $\mathcal{B}_\Delta(P_2)$, and vice-versa, proving $\mathcal{B}_\Delta(P_1) = \mathcal{B}_\Delta(P_2)$, as required. □

7.6. Connection to prior results and questions

In [23], Ullmo and Giannoni consider compact hyperbolic polygons $P \subset \mathbb{H}$ for which the interior angles are all acute. The authors provide a nearly complete description of

the bounce spectra for such a polygon. The ideal objective is to describe finite words in the label set $\mathcal{A}$ that are forbidden in a bounce sequence of a billiard trajectory in P, and then characterize bounce sequences precisely as those that do not contain such subwords. The authors do not quite obtain such a set, but do provide approximations to this. Their analysis divides up into two cases: billiards in good polygons (which they call *tiling billiards*) and the rest.

In the case that the polygon P is not good, for every $N > 0$, Ullmo-Giannoni produce two sets of words of length at most N; forbidden words and *authorized* words. The authorized words have the property that any sequence for which all subwords of length at most N are authorized are necessarily bounce sequences of a billiard trajectory in P (thus, forbidden words provide necessary conditions and authorized words provide sufficient conditions). They also show that as N tends to infinity, the ratio of the number of authorized to non-forbidden words tends to 1. It would be interesting to see if, at least for acute polygons, one could push this analysis to give another proof of rigidity for such polygons.

To describe the situation for a good polygon P, let $\frac{\pi}{k_j}$ denote the interior angle at the vertex of P between sides j and $j + 1$ (modulo n), where $k_j > 2$ is an integer for all j. In this case, Ullmo-Giannoni show that length 2 subwords which simply repeat a label are forbidden as are sequences of the form $j, j + 1, j, j + 1, \ldots$ or $j + 1, j, j + 1, j, \ldots$ of length greater than k_j and prove that this *almost* characterizes bounce sequences of billiard trajectories. More precisely, they show that given a sequence that contains no forbidden words as above, after *possibly* replacing some of the length k_j subwords of the form $j, j + 1, j, j + 1, \ldots$ with an equal length word of the form $j + 1, j, j + 1, j, \ldots$ (or vice versa), for each j, then the resulting sequence is the bounce sequence of a billiard trajectory in P. See also Nagar–Singh [17].

In the case that P_1 and P_2 are nonisometric, good, acute n–gons for which the corresponding interior angles (with respect to fixed labelings) are equal, their results provide an interesting connection to our work. For concreteness, say the angle between sides j and $j + 1$ is $\frac{\pi}{k_j}$, as above. If P_1 is billiard rigid (for example, P_1 could be as in case (2) of 7.8) then $\mathcal{B}(P_1) \neq \mathcal{B}(P_2)$. On the other hand, the results of [23] described in the previous paragraph imply that given any bounce sequence $\mathbf{b} \in \mathcal{B}(P_2)$, after possibly replacing some of the subwords of the form $j, j + 1, j, j + 1, \ldots$ of length k_j by the length k_j word of the form $j + 1, j, j + 1, j, \ldots$ (or vice versa) for all j, one obtains a bounce sequence $\mathbf{b}' \in \mathcal{B}(P_1)$. In particular, the difference between the difference between $\mathcal{B}(P_1)$ and $\mathcal{B}(P_2)$ is quite subtle.

Questions. We end with some questions and possible future directions. To start, it would be interesting to investigate further to what extent Ullmo-Giannoni's methods and ours are complementary. In particular:

(Q1) Can one prove the Billiard Rigidity Theorem using the techniques of [23] (perhaps restricted to polygons with only acute interior angles)?

In an earlier paper Ullmo-Giannoni [13] characterized the bounce spectra for ideal hyperbolic polygons; that is, polygons all of whose vertices are on $\partial \mathbb{H}$ (see also Nagar–Singh [17]). In this setting, the authors find that bounce sequences of an ideal n–gon are precisely the biinfinite sequences never repeating a label and not containing an infinite repeating string of the form

$$\overline{i, i+1} = i, i+1, i, i+1, i, i+1, \ldots$$

either forward or backward, with indices taken modulo n. In particular, any two ideal n–gons, P_1 and P_2 have the same bounce spectra, $\mathcal{B}(P_1) = \mathcal{B}(P_2)$. This fits into the Billiard Rigidity Theorem if one extends the notion of reflective tiling to allow vertices at infinity (where there are no longer any conditions on interior angles at such points). Indeed, then P_1 and P_2 are reflectively tiled with a single tile (themselves). Observe that reflections in the sides of an ideal polygon P generate a reflective tiling of $\mathbb{H}$ by copies of P.

Since ideal polygons are only a special case of noncompact polygons, one could further ask what happens for noncompact polygons in general (some special cases are considered in Nagar–Singh [17]).

(Q2) Is the Billiard Rigidity Theorem true if one allows for finite area noncompact hyperbolic polygons? How about infinite area polygons?

Note that noncompact polygons could also have infinitely many sides (even in the finite area, convex case), and so some care should be taken in clearly defining what is meant by a cyclic labeling of the sides. This question also makes sense in the case of Euclidean polygons of infinite area.

(Q3) Is there some version of the Billiard Rigidity Theorem for infinite area Euclidean polygons?

Another natural question is whether a version of Billiard Rigidity Theorem holds in higher dimensions. Coding billiards in three-dimensional hyperbolic polyhedra was considered in [21].

(Q4) Is there a version of the Billiard Rigidity Theorem for billiards in Euclidean or hyperbolic n–space, $n \geq 3$?

Returning to dimension 2, we note that in order to answer the special case of (Q2) using our methods, one would first need a more general version of the Current Support Theorem. Of course, we could also formulate this question for infinite area hyperbolic or Euclidean cone surfaces (cf. [15]), but we stick to the following concrete instance.

(Q5) Is the Current Support Theorem true for noncompact, finite area, complete, negatively curved, hyperbolic cone surfaces?

The universal cover has a circle at infinity in this case, though now the action has parabolic fixed points on the boundary associated to the "cusps" of S. We note that even under the finite area assumption, one could still have infinitely many cone points, so the cusps here are not necessarily standard, and there are various technical issues one must consider. For such a metric, we would take $\mathcal{G}_{\tilde{\varphi}}$ to be the closure of the space of endpoints of nonsingular geodesics, which should coincide with the support of a Liouville current, though again some care must be taken in defining such an object in this setting.

There are other ways in which one could try to generalize the Current Support Theorem. As we mentioned in the introduction, not only are hyperbolic and Euclidean cone metrics determined by their Liouville currents, but so are any (variable) negative or non-positively curved cone metrics, by a result of Constantine [6]. Hence it is natural to ask to what extent they are determined by only their support:

(Q6) Is there a version of the Current Support Theorem for negative or nonpositively curved cone metrics?

We note that the methods used in this paper to prove the Current Support Theorem used the fact that $\hat{S}$ is locally isometric to $\mathbb{H}$ in a crucial way (and in [9] the Euclidean setting was essential). In particular, it seems highly unlikely that (Q6) would follow by direct generalizations of the techniques used in this paper. The statement itself is likely to be quite different: for example, the second conclusion of the Current Support Theorem would certainly have to be relaxed to allow lifting deformations via branched covers of negatively curved orbifolds, but an even more general alternative seems likely. There is also nothing stopping one from asking similar questions in the case of spherical surfaces/polygons, but it seems likely the techniques will stray even further from those considered here.

References

[1] A.F. Beardon, The Geometry of Discrete Groups, Graduate Texts in Mathematics, vol. 91, Springer-Verlag, New York, 1983.
[2] M.R. Bridson, A. Haefliger, Metric Spaces of Non-positive Curvature, Grundlehren der Mathematischen Wissenschaften (Fundamental Principles of Mathematical Sciences), vol. 319, Springer-Verlag, Berlin, 1999.
[3] A. Bankovic, C.J. Leininger, Marked-length-spectral rigidity for flat metrics, Trans. Am. Math. Soc. 370 (3) (2018) 1867–1884.
[4] M. Bogdanov, M. Teillaud, G. Vegter, Delaunay triangulations on orientable surfaces of low genus, in: 32nd International Symposium on Computational Geometry, in: LIPIcs. Leibniz Int. Proc. Inform., vol. 51, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016, 20.
[5] C. Croke, A. Fathi, J. Feldman, The marked length-spectrum of a surface of nonpositive curvature, Topology 31 (4) (1992) 847–855.
[6] D. Constantine, Marked length spectrum rigidity in non-positive curvature with singularities, Indiana Univ. Math. J. 67 (6) (2018) 2337–2361.
[7] C.B. Croke, Rigidity for surfaces of nonpositive curvature, Comment. Math. Helv. 65 (1) (1990) 150–169.
[8] A. Douady, C.J. Earle, Conformally natural extension of homeomorphisms of the circle, Acta Math. 157 (1–2) (1986) 23–48.
[9] M. Duchin, V. Erlandsson, C.J. Leininger, C. Sadanand, You can hear the shape of a billiard table: symbolic dynamics and rigidity for flat surfaces, Comment. Math. Helv. 96 (3) (2021) 421–463.

[10] M. Duchin, C.J. Leininger, K. Rafi, Length spectra and degeneration of flat metrics, Invent. Math. 182 (2) (2010) 231–277.

[11] B. Farb, D. Margalit, A Primer on Mapping Class Groups, Princeton Univ. Press, Princeton, N.J., 2010.

[12] J. Frazier, Length Spectral Rigidity of Non-positively Curved Surfaces, Thesis (Ph.D.), University of Maryland, College Park, ProQuest LLC, Ann Arbor, MI, 2012.

[13] M.-J. Giannoni, D. Ullmo, Coding chaotic billiards. I. Noncompact billiards on a negative curvature manifold, Physica D 41 (3) (1990) 371–390.

[14] A. Hatcher, Algebraic Topology, Cambridge University Press, Cambridge, 2002.

[15] W.P. Hooper, An infinite surface with the lattice property I: Veech groups and coding geodesics, Trans. Am. Math. Soc. 366 (5) (2014) 2625–2649.

[16] S. Hersonsky, F. Paulin, On the rigidity of discrete isometry groups of negatively curved spaces, Comment. Math. Helv. 72 (3) (1997) 349–388.

[17] A. Nagar, P. Singh, Finiteness in polygonal billiards on hyperbolic plane, Topol. Methods Nonlinear Anal. 58 (2) (2021) 481–520.

[18] J.-P. Otal, Le spectre marqué des longueurs des surfaces à courbure négative, Ann. Math. (2) 131 (1) (1990) 151–162.

[19] N. Purzitsky, All two-generator Fuchsian groups, Math. Z. 147 (1) (1976) 87–92.

[20] J.G. Ratcliffe, Foundations of Hyperbolic Manifolds, Graduate Texts in Mathematics, vol. 149, Springer-Verlag, New York, 1994.

[21] P. Singh, Coding of billiards in hyperbolic 3-space, preprint, arXiv:2009.14427, 2020.

[22] W.P. Thurston, Three-Dimensional Geometry and Topology, vol. 1, in: S. Levy (Ed.), Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997.

[23] D. Ullmo, M.-J. Giannoni, Coding chaotic billiards. II. Compact billiards defined on the pseudo-sphere, Physica D 84 (3–4) (1995) 329–356.

[24] A.N. Zemljakov, A.B. Katok, Topological transitivity of billiards in polygons, Mat. Zametki 18 (2) (1975) 291–300.