

MARKED LENGTH SPECTRUM RIGIDITY FOR SURFACE AMALGAMS

YANDI WU

ABSTRACT. In this paper, we show that simple, thick negatively curved two-dimensional P-manifolds, a large class of surface amalgams, are marked length spectrum rigid. That is, if two piecewise negatively curved Riemannian metrics (satisfying certain smoothness conditions) on a simple, thick two-dimensional P-manifold assign the same lengths to all closed geodesics, then they differ by an isometry up to isotopy. Our main theorem is a natural generalization of Croke and Otal’s celebrated results about marked length spectrum rigidity of negatively curved surfaces.

1. INTRODUCTION

Mostow’s famous rigidity theorem states that for closed hyperbolic manifolds of dimension greater than two, the metric is completely determined by the fundamental group up to isotopy. Mostow’s rigidity theorem does not hold for hyperbolic surfaces or negatively curved Riemannian metrics without constant sectional curvature. For these cases, in order to conclude two metrics are equivalent up to isotopy, further restrictions are needed, such as requiring the surfaces to have the same *marked length spectra*, defined as follows:

Definition 1.1 (Marked length spectrum). The *marked length spectrum of a metric space* (M, g) is the class function

$$MLS : \pi_1(M) \rightarrow \mathbb{R}^+, [\alpha] \mapsto \inf_{\gamma \in [\alpha]} \ell_g(\gamma)$$

which assigns to each free homotopy class $[\alpha] \in \pi_1(M)$ the infimum of lengths in the class.

In particular, if g is negatively curved or locally CAT(-1), each homotopy class of curves has a unique geodesic representative, so the marked length spectrum assigns a length to each closed geodesic in (M, g) . We say that (M, g_0) and (M, g_1) have the same marked length spectrum if for every $[\alpha] \in \pi_1(M)$, $MLS_0([\alpha]) = MLS_1([\alpha])$. A class of metrics \mathcal{M} is *marked length spectrum rigid* if whenever (M, g_0) and $(M, g_1) \in \mathcal{M}$ have the same marked length spectrum, there exists an isometry $\varphi : (M, g_0) \rightarrow (M, g_1)$ that is isotopic to the identity.

In this paper, we will study a class of objects, defined below, that are natural generalizations of surfaces. Definition 1.2 is adapted from Definition 2.3 of [31].

Received by the editors December 22, 2023, and, in revised form, November 16, 2024, December 3, 2024, and December 4, 2024.

2020 *Mathematics Subject Classification.* Primary 57M50; Secondary 53C24, 37D40.

The author was supported by an NSF RTG grant DMS-2230900.

Definition 1.2 (Negatively curved two-dimensional P-manifold). A compact metric space X is a *negatively curved two-dimensional P-manifold* if there exists a closed subset $Y \subset X$ (the *gluing curves* of X) that satisfies the following:

- (1) Each connected component of Y is homeomorphic to S^1 ;
- (2) The closure of each connected component of $X - Y$ is homeomorphic to a compact surface with boundary endowed with a negatively curved (Riemannian) metric, and the homeomorphism takes the component of $X - Y$ to the interior of a surface with boundary. We will call each $\overline{X - Y}$ a *chamber* in X ;
- (3) There exists a negatively-curved metric on each chamber which coincides with the original metric.

If Y forms a totally geodesic subspace of X consisting of disjoint simple closed curves, we say that X is *simple*. If each connected component of Y (gluing curve) is attached to at least three distinct boundary components of chambers, then we say X is *thick* (note this definition differs slightly from the one given in [31] but follows the one given in [32] and [30]). Like Lafont in [31], we will only be considering simple, thick two-dimensional P-manifolds. Doing so ensures X is locally CAT(-1) (in other words, its universal cover, \tilde{X} , is CAT(-1)) and guarantees some useful properties, such as one pointed out in Lemma 2.9. Lafont has defined higher dimensional P-manifolds as well (see [32]); we focus on the two-dimensional ones, which consist of surfaces with boundary glued together along their boundary components. Throughout this paper, all P-manifolds will be assumed to be two-dimensional.




FIGURE 1. An example of a simple, thick (two-dimensional) P-manifold with four chambers

We will equip a simple, thick negatively-curved P-manifold X with a metric in a class we denote as \mathcal{M}_{\leq} , following the notation from [9]. For the precise definition of \mathcal{M}_{\leq} , we refer the reader to Section 2.1, but roughly speaking, metrics in \mathcal{M}_{\leq} are locally CAT(-1) piecewise Riemannian metrics with an additional condition that limits pathological behavior around the gluing curves of X . We now state the main result of the paper:

Theorem 1.3. Suppose (X, g_1) and (X, g_2) are simple, thick negatively-curved P-manifolds equipped with metrics $g_1, g_2 \in \mathcal{M}_{\leq}$. Furthermore, suppose (X, g_1) and (X, g_2) have the same marked length spectrum. Then there is an isometry $\phi : (X, g_1) \rightarrow (X, g_2)$ that is isotopic to the identity.

The marked length spectrum rigidity problem has a long history. Fricke and Klein showed the class of closed hyperbolic surfaces is marked length spectrum rigid (see [18]). Burns and Katok then conjectured that closed negatively curved manifolds of all dimensions are marked length spectrum rigid in [7]. A major breakthrough occurred when Croke and Otal independently proved the conjecture in dimension two (see [35], [14]). Croke and Otal's results led to a large quantity of generalizations of the marked length spectrum rigidity problem for surfaces, some (but far from all) of which are listed here. Croke, Fathi, and Feldman extended Croke's result to the case of non-positively curved metrics in [13]. In another direction, Hersonksy and Paulin extended Otal's result to the case of negatively curved metrics with finitely many cone singularities in [25]. Duchin, Leininger, and Rafi proved marked length spectrum rigidity for closed surfaces endowed with flat metrics with finitely many cone singularities arising from quadratic differentials in [17], a result extended by Bankovic and Leininger in [2], who proved the result for all flat metrics with finitely many cone singularities. Finally, by combining previous methods, Constantine proved marked length spectrum rigidity for non-positively curved metrics on surfaces with finitely many cone singularities in [8]. We remark that marked length spectrum rigidity is unlikely to hold for surfaces with boundary, although Guillarmou and Mazzucchelli have shown a weaker form of rigidity, *marked boundary rigidity*, for a large family of negatively curved metrics on surfaces with strictly convex boundary (see [21]).

While the marked length spectrum rigidity problem has been well-studied in the case of surfaces, for higher dimensions, the conjecture remains largely open. Hamenstädt showed marked length spectrum rigidity for rank one locally symmetric closed manifolds of dimension greater than 2 (see [24]). There is also a local marked length spectrum rigidity result for closed, compact manifolds of all dimensions with Anosov geodesic flow due to Guillarmou and Lefeuvre (see [20]). For metric spaces that are not manifolds, the question also remains largely unstudied, although there are a few notable results. Work by Culler and Morgan (see [15]) and Alperin and Bass (see [1]) leads to a marked length spectrum rigidity result for Culler-Vogtmann Outer Space, built to study the group $\text{Out}(\mathbb{F}_n)$ in analogy to the relationship between the Teichmüller space of S , $\mathcal{T}(S)$, and the mapping class group, $\text{Mod}(S)$. Generalizing the work from [15] and [1], Constantine and Lafont prove a version of marked length spectrum rigidity of compact, non-contractible one-dimensional spaces in [10]. In [9], they show that certain compact quotients of Fuchsian buildings, including those with piecewise hyperbolic metrics, are marked length spectrum rigid.

Outline of the paper. In Section 2, we define precisely the class of metrics \mathcal{M}_{\leq} and review some fundamental facts about CAT(-1) spaces. Furthermore, we detail some separability properties of hyperbolic groups developed by Haglund and Wise which we will use in Section 4. In Section 3, we patch together isometries constructed in [35] to prove the theorem for P-manifolds with the property that each chamber can be included into a closed surface. Since there is no well-defined unit tangent bundle, proving ergodicity of the geodesic flow map, a key component of Otal's original proof, requires some heavy machinery included in the Appendix. Finally, in Section 4, we prove the general case by constructing finite-sheeted covers to reduce to the base case examined in Section 3. The argument relies on separability properties of the fundamental groups of simple, negatively-curved surface

amalgams, which can be realized as non-positively curved cube complexes. In particular, we show that such groups are *QCERF*, which allows one to promote certain immersions to embeddings in finite-sheeted covers.

2. PRELIMINARIES

2.1. The class of metrics \mathcal{M}_{\leq} . We first precisely define the class of metrics \mathcal{M}_{\leq} under consideration in Theorem 1.3 and point out some important properties. We say $g \in \mathcal{M}_{\leq}$ if g satisfies the following properties:

- (1) Each chamber of $C \subset X$ is equipped with a negatively-curved Riemannian metric with sectional curvature bounded above by -1 so that C has geodesic boundary components;
- (2) The restrictions of g to the chambers of X are “compatible” in the sense that if two boundary components b_1 and b_2 of two (possibly the same) chambers C_1 and C_2 are both attached to a gluing curve $\gamma \subset X$, then the gluing maps $b_1 \hookrightarrow \gamma$ and $b_2 \hookrightarrow \gamma$ are isometries (in particular, we do not allow circle maps of degree two or greater like those explored in [26]);
- (3) For any two boundary components $b_1 \in C_1$ and $b_2 \in C_2$, the restriction of g to $N_{b_1} \bigcup_{b_1 \sim b_2} N_{b_2}$ is a negatively curved smooth Riemannian metric with sectional curvature bounded above by -1 , where N_{b_1} and N_{b_2} are ϵ -neighborhoods around b_1 and b_2 respectively for some $\epsilon > 0$.

We impose the third condition to ensure that we can exploit previous marked length spectrum rigidity results for surfaces which in particular require Riemannian negatively curved metrics with at most a finite number of cone singularities. We remark that Theorem 1.3 still holds if one introduces a finite number of cone singularities (points with cone angles greater than 2π) to each chamber; instead of using Theorem 1 from [35] in the proof of Lemma 3.8, one would instead use Theorem C from [25] or Corollary 2 from [8]. Furthermore, we suspect Condition (3) could be eliminated if one were to carefully modify Otal’s proof to allow for metrics that are Riemannian outside a singular set of gluing geodesics, but more work would need to be done.

Finally, we point out that a surface amalgam is *non-positively curved (NPC)* if each chamber admits a non-positively curved Riemannian metric. Note that the only difference between negatively curved and NPC surface amalgams is the latter may have chambers which are cylinders. For the sake of completeness, we will briefly allow NPC surface amalgams in some auxiliary lemmas in Section 4.

We now discuss some properties of \mathcal{M}_{\leq} that will be useful in the proof of Theorem 1.3.

2.1.1. Properties of \mathcal{M}_{\leq} . First, we remark that metrics in \mathcal{M}_{\leq} are locally CAT(-1):

Remark 2.1. If (X, g) is a negatively-curved P-manifold where $g \in \mathcal{M}_{\leq}$, then (X, g) is locally CAT(-1).

Indeed, suppose X is equipped with a metric $g \in \mathcal{M}_{\leq}$ and $C \subset X$ is a chamber in X . Recall a generalization of the Cartan-Hadamard Theorem which states that a smooth Riemannian manifold M has sectional curvature $\leq \kappa$ if and only if M is locally CAT(κ) (see [5, Theorem 1A.6]). As a result, the restriction of g to C is locally CAT(-1) since C is endowed with a negatively curved metric with sectional curvature bounded above by -1 . If $\kappa \in \mathbb{R}$ and X_1 and X_2 are locally CAT(κ)




FIGURE 2. The universal cover of a simple, thick P-manifold (X, g)

spaces glued isometrically along a convex, complete metric subspace $A \subset X_1 \cap X_2$, then $X_1 \sqcup_A X_2$ is locally $\text{CAT}(\kappa)$ (see Theorem 2.11.1 in [5]). As a result, since we have endowed each chamber in a P-manifold X with a locally $\text{CAT}(-1)$ metric, (X, g) will also be locally $\text{CAT}(-1)$, which shows the remark.

The locally $\text{CAT}(-1)$ property of metrics in \mathcal{M}_{\leq} will prove useful in the proof of the base case of the main result in the Section 3. Note that $\text{CAT}(-1)$ spaces are also Gromov (δ) -hyperbolic, so (X, g) where $g \in \mathcal{M}_{\leq}$ is also Gromov hyperbolic, another useful property we will exploit in Section 3.

Remark 2.2. Another way to show that (X, g) where $g \in \mathcal{M}_{\leq}$ is Gromov hyperbolic is to observe that any two-dimensional P-manifold (a metric space satisfying only Conditions (1) and (2) from Definition 1.2) has a fundamental group that is an amalgamated product of surface groups and free groups over cyclic subgroups (without any \mathbb{Z}^2 subgroups), which is Gromov hyperbolic by the Bestvina-Feighn Combination Theorem (see [3]).

2.2. The universal cover of a simple, thick P-manifold. Next, we describe the universal cover of a simple, thick P-manifold, which we will be working with extensively in Section 3. Roughly speaking, \tilde{X} will look like an infinite collection of totally geodesic subspaces of disks (homeomorphic to \mathbb{H}^2) glued together in a tree-like fashion. Following Lafont, we will call disks in \tilde{X} *apartments*. Throughout this paper, also following Lafont, we will also call polygonal lifts of chambers in the universal cover a *chamber*. Following Lafont, we will call geodesics in \tilde{X} that project to gluing curves under the covering map *branching geodesics*.

While in general, \tilde{X} is very hard to visualize, the case where each chamber can be included into a closed surface is much easier to describe. We will focus on describing \tilde{X} in this special case since our proof strategy relies on reducing to it. Each gluing geodesic γ will lift to a countably infinite collection of branching geodesics in \tilde{X} , with a countably infinite collection in each apartment containing lifts of chambers that intersect γ . Each of these lifts of γ are adjacent to n half planes, where n is the number of boundary components attached to γ . Each closed surface will lift to an infinite collection of apartments tiled by lifts of chambers that are embedded in the closed surface. See Figure 2 for an illustration of \tilde{X} .




FIGURE 3. An illustration of the calculation of $[\xi\xi'\eta\eta']$ in \mathbb{H}^2

Remark 2.3. Some readers may notice similarities between \tilde{X} and a two-dimensional hyperbolic building. Even the terminology of P-manifolds borrows heavily from that of buildings; both have “chambers,” “apartments” and sets of branching geodesics, which are called *walls* in building terminology. This allows us to borrow some definitions, such as intersection numbers, from [9]. However, for buildings, boundaries of chambers are not totally geodesic in the space, so their walls are not totally geodesic. The proof methods from [9] are different from those seen in this paper in part due to this fact. Furthermore, all the chambers in buildings are assumed to be isometric, which is not the case in \tilde{X} . See Section 1.4.3 of [33] for a more complete list of differences between buildings and P-manifolds.

2.3. Visual boundaries. In this section, we will assume that (\tilde{X}, \tilde{g}) is a proper (closed balls are compact) metric space that is either CAT(-1) or Gromov hyperbolic; the definitions are the same for both. For more details and background, we refer the reader to [28]. We say that two geodesic rays in \tilde{X} , $\alpha_1 : \mathbb{R}_{\geq 0} \rightarrow \tilde{X}$, and $\alpha_2 : \mathbb{R}_{\geq 0} \rightarrow \tilde{X}$ are *asymptotic* if they lie within bounded distance of one another; in other words, there exists some finite $M \geq 0$ such that for all $t \in \mathbb{R}_{\geq 0}$, $d(\alpha_1(t), \alpha_2(t)) \leq M$. To define $\partial^\infty(\tilde{X}, \tilde{g})$, fix any basepoint $x_0 \in (\tilde{X}, \tilde{g})$ and consider the set of all geodesic rays originating from x_0 . Two geodesic rays based at x_0 are equivalent if they are asymptotic.

Definition 2.4 (Visual boundary). We define equivalence classes of such geodesic rays based at x_0 as the *visual boundary* of \tilde{X} , $\partial^\infty(\tilde{X}, \tilde{g})$.

When \tilde{g} is CAT(-1), there is a unique geodesic in (\tilde{X}, \tilde{g}) between any two points in $\partial^\infty(\tilde{X}, \tilde{g})$ (see Proposition 1.4 of [5]), so the space of oriented geodesics in \tilde{X} is identified with $\partial^\infty(\tilde{X}, \tilde{g}) \times \partial^\infty(\tilde{X}, \tilde{g}) \setminus \Delta$, where Δ indicates the diagonal.

We define the cross ratio as follows:

Definition 2.5 (Cross ratio and Möbius maps). If (\tilde{X}, \tilde{g}) has boundary $\partial^\infty(\tilde{X}, \tilde{g})$, one can define the *cross ratio* of a four-tuple of distinct boundary points $(\xi, \xi', \eta, \eta') \in (\partial^\infty(\tilde{X}))^4$:

$$(1) \quad [\xi\xi'\eta\eta'] = \lim_{(a, a', b, b') \rightarrow (\xi, \xi', \eta, \eta')} \frac{1}{2}(\tilde{g}(a, b) + \tilde{g}(a', b') - \tilde{g}(a, b') - \tilde{g}(a', b)).$$

See Figure 3 for an example in \mathbb{H}^2 .

A map $\partial^\infty f : \partial^\infty(\tilde{X}, \tilde{g}_1) \rightarrow \partial^\infty(\tilde{X}, \tilde{g}_2)$ is *Möbius* if it preserves the cross ratio (i.e. $[\xi\xi'\eta\eta'] = [f(\xi)f(\xi')f(\eta)f(\eta')]$).

From the definition of the cross-ratio, it follows that:

$$(2) \quad [abcd] = [cdab] \text{ and } -[abdc] = [abcd] = -[bacd].$$

We are now ready to define the Gromov Product:

Definition 2.6 (Gromov product). Let $x \in (\tilde{X}, \tilde{g})$ and $a, b \in \partial^\infty(\tilde{X})$. Then the *Gromov Product* of a and b with respect to a basepoint $x \in \tilde{X}$ is:

$$\langle a, b \rangle_x = \lim_{(a_i, b_i) \rightarrow (a, b)} \frac{1}{2} (\tilde{g}(a_i, x) + \tilde{g}(x, b_i) - \tilde{g}(a_i, b_i)).$$

The Gromov product measures how long two geodesics travel close together in (\tilde{X}, \tilde{g}) . Indeed, if $a_i, b_i, x \in \tilde{X}$ are three arbitrary points in a δ -hyperbolic metric space, then

$$\langle a_i, b_i \rangle_x = \frac{1}{2} (\tilde{g}(a_i, x) + \tilde{g}(x, b_i) - \tilde{g}(a_i, b_i))$$

approximates within 2δ the distance between x and the geodesic segment $[a_i, b_i]$. As a result, the Gromov product provides a convenient way to topologize $\partial^\infty(\tilde{X}, \tilde{g})$. Indeed, we can define a countable neighborhood base for any $p \in \partial^\infty(\tilde{X})$ and $d \in \mathbb{N}$ as follows:

$$(3) \quad B(p, d) = \{q \in \partial^\infty(\tilde{X}) \mid \langle p, q \rangle_x > d\}.$$

Using the Gromov Product, we can also endow $\partial^\infty(\tilde{X}, \tilde{g})$ with a *visual metric* that induces the aforementioned topology on $\partial^\infty(\tilde{X}, \tilde{g})$:

Definition 2.7 (Visual metric). Let (\tilde{X}, \tilde{g}) be a proper CAT(-1) space. Suppose $a, b \in \partial^\infty(\tilde{X}, \tilde{g})$. Then for some fixed basepoint $x \in (\tilde{X}, \tilde{g})$, we can assign $\partial^\infty(\tilde{X}, \tilde{g})$ a *visual metric*:

$$(4) \quad \tilde{g}_{\infty, x}(a, b) = \begin{cases} e^{-\langle a, b \rangle_x} & \text{if } a \neq b, \\ 0 & \text{otherwise.} \end{cases}$$

Note that if $x' \in (\tilde{X}, \tilde{g})$ is a different basepoint, then we have that for $A = e^{2\delta} e^{\tilde{g}(x, x')} > 1$:

$$A^{-1} \tilde{g}_{\infty, x}(a, b) \leq \tilde{g}_{\infty, x'}(a, b) \leq A \tilde{g}_{\infty, x}(a, b).$$

2.3.1. The visual boundary of a P -manifold. We now present some useful properties of the visual boundary of a simple, thick P -manifold (\tilde{X}, \tilde{g}) . As before, we will assume all the chambers in X are negatively curved.

Bass-Serre theory provides a useful characterization of points on $\partial^\infty(\tilde{X})$. Consider a bipartite *graph of groups decomposition* \mathcal{G} of $G = \pi_1(X)$, where there is a vertex group $\pi_1(C)$ for each chamber $C \subset X$ and $\langle \gamma \rangle \cong \mathbb{Z}$ for each gluing curve $\gamma \subset X$. Furthermore, there is a cyclic edge group between $\langle \gamma \rangle$ and $\pi_1(C)$ for each boundary component of C that is attached to γ . Thus, $\mathcal{G} = V_0 \sqcup V_1$ can be partitioned into collection of vertices labeled with $\langle \gamma \rangle$, which we will call V_0 , and those labeled with $\pi_1(C)$, which we will call V_1 (see [16] for more details on and examples of this construction). The *Bass-Serre tree* of G is a tree T on which G acts minimally (i.e. there is no proper invariant subtree of T) and without inversions on edges with quotient $\mathcal{G} = T/G$.

To better describe the Bass-Serre tree T , we briefly summarize the main points from Section 4.1 of [34], which uses Bowditch's construction to characterize boundary points for *geometric amalgams of free groups* (e.g. $\pi_1(X)$). Given a point in

$x \in \partial^\infty(\tilde{X})$, Bowditch characterizes x by $\text{val}(x)$, the number of topological ends of $\partial^\infty(\tilde{X}) \setminus \{x\}$. In particular, for a simple, thick P-manifold with negatively-curved chambers, $\text{val}(x) = n \geq 3$ if x is a branching geodesic attached to n distinct boundary components and $\text{val}(x) = 2$ otherwise. Bowditch defines separate equivalence classes for $M(2) = \{x \in \partial^\infty(\tilde{X}) : \text{val}(x) = 2\}$ and $M(3+) = \{x \in \partial^\infty(\tilde{X}) : \text{val}(x) \geq 3\}$:

- (1) *For points in $M(3+)$:* We say $x \approx y$ if either $x = y$ or the number of connected components of $\partial^\infty(\tilde{X}) \setminus \{x, y\}$ is equal to $\text{val}(x)$ and $\text{val}(y)$.
- (2) *For points in $M(2)$:* We say $x \sim y$ if $x = y$ or the number of connected components of $\partial^\infty(\tilde{X}) \setminus \{x, y\}$ is equal to 2.

In both cases, equivalent pairs of points form a cut pair of $\partial^\infty(\tilde{X})$. The Bass-Serre tree is a bipartite infinite-valence tree consisting of vertices labeled by equivalence classes in $M(3+)$ and $M(2)$, which represent conjugacy classes of vertices in V_0 and V_1 respectively. Two vertices Λ_e and $\Lambda_v \in T$ corresponding to a \approx -class in $M(3+)$ and a \sim -class in $M(2)$ respectively are connected by an edge if any two distinct points $x, y \in \Lambda_v$ lie on the same connected component of $\partial^\infty(\tilde{X}) \setminus \{a, b\}$, where a, b are the two distinct points in Λ_e corresponding to endpoints of some branching geodesic. Bowditch thus formulates a trichotomy of points in $\partial^\infty(\tilde{X})$:

Proposition 2.8 (Proposition 1.3 of [4] and Lemma 10 of [29]). *If $\pi_1(X)$ acts on Bass-Serre tree T , then each point $x \in \partial^\infty(\tilde{X})$ corresponds to exactly one of the following equivalence classes of points in $\partial^\infty(\tilde{X})$:*

- (1) *A point in Λ_e corresponding to some conjugate of $V_0 \in V(\mathcal{G})$;*
- (2) *A point in Λ_v corresponding to some conjugate of $V_1 \in V(\mathcal{G})$;*
- (3) *A point of $\partial^\infty(T)$, with a unique x for each point in $\partial^\infty(T)$.*

In other words, \approx -classes in $M(3+)$ correspond to vertices in category (1) in Proposition 2.8 while vertices in category (2) correspond to \sim -classes of $M(2)$.

Using Proposition 2.8, we can topologically characterize points in $\partial^\infty(\tilde{X})$. As mentioned earlier, points in category (1) correspond to endpoints of branching geodesics, lifts of gluing curves in $\partial^\infty(\tilde{X})$. Equivalence classes of points corresponding to a vertex in category (2) consist of points x and y such that the unique geodesic (x, y) between x and y does not intersect any branching geodesics. Indeed, if (x, y) were to intersect a branching geodesic, $\partial^\infty(\tilde{X}) \setminus \{x, y\}$ would still be connected. Topologically, one can check that this means x and y in fact lie on the boundary of the same connected component of $p^{-1}(C) \setminus \{\mathcal{BG}\}$, where C is some chamber in X , $p : \tilde{X} \rightarrow X$ is the universal covering map, and \mathcal{BG} is the full set of branching geodesics in \tilde{X} . A vertex Λ_e is connected to a vertex Λ_v if the branching geodesic corresponding to Λ_e is a boundary component of the connected component of $p^{-1}(C) \setminus \{\mathcal{BG}\}$ associated with Λ_v . From this information, we see that traveling along an edge in T is equivalent to crossing a branching geodesic in \tilde{X} . Thus, the unique points in category (3) correspond to geodesic rays that intersect infinitely many branching geodesics in \tilde{X} .

We now point out that points on $\partial^\infty(\tilde{X})$ lie on boundaries of embedded disks, a fact stated in the proof of Lemma 3.1 in [31]. For the convenience of the reader, we provide a proof as well.

Lemma 2.9. *Given a geodesic $\gamma \in \tilde{X}$ with endpoints $p, q \in \partial^\infty(\tilde{X})$, there exists an isometrically embedded apartment containing γ . That is, p and q lie in an embedded $S^1 \subset \partial^\infty(\tilde{X})$. As a consequence, the boundary $\partial^\infty(\tilde{X})$ is path-connected.*

Proof. First, we consider the case where γ is a gluing geodesic. If $p(\gamma)$, the gluing geodesic lifting to γ in \tilde{X} , is a closed geodesic in a closed surface, then there is automatically an apartment containing γ . Otherwise, consider two distinct boundary components $b_1 \subset C_1$ and $b_2 \subset C_2$ attached to $p(\gamma)$; note that b_1 and b_2 exist due to the thickness hypothesis and C_1 and C_2 are not required to be distinct chambers. In \tilde{X} , there is a set of lifts of C_1 , which we will call K_1 , that forms a totally geodesic subset of a half-disk containing γ . Similarly, there is a set of lifts of C_2 forming a totally geodesic subset K_2 (disjoint from K_1) of a half-disk containing γ . Note that even though $K_1 \cup K_2$ is only a subset of a disk containing γ , we can “fill out” $K_1 \cup K_2$ to obtain an apartment A embedded in \tilde{X} . Indeed; if there is part of a disk missing from $K_1 \cup K_2$, its boundary must necessarily be a branching geodesic γ' in \tilde{X} . By the thickness hypothesis, there is some collection of polygons \mathcal{P} disjoint from K_1 and adjacent to γ' that project to a chamber C adjacent to a gluing geodesic $p(\gamma')$ in X . We can “continue” A by attaching a subset of a half-plane disjoint from K_1 and tiled by copies the polygons in \mathcal{P} . We then iterate, attaching collections of lifts of chambers and eventually limiting to a half-plane $H_1 \supset K_1$ embedded in $\partial^\infty(\tilde{X})$. Apply the same procedure to obtain a half-disk H_2 containing K_2 . Note that this construction may be counterintuitive because $A = H_1 \cup H_2$ does not project under the covering map to any surfaces in X (see Figure 4).

The case where γ is not a gluing geodesic is similar. Note that in this case, γ will pass through a sequence of branching geodesics $\{\gamma_n\}_{n \in \mathbb{N}}$. For each γ_n that γ passes through, γ will also pass through a collection \mathcal{P}_n of lifts of C_n , a chamber attached to $p(\gamma_n)$. Note \mathcal{P}_n can be chosen to be a totally geodesic subspace of a disk, K_n , containing both γ_n and γ_{n+1} ; similar to the procedure from before, we can “fill in” K_n to a connected section of a disk bordered by γ_n and γ_{n+1} , which we will call H_n . Iterate the procedure to obtain an apartment $A = \bigcup_{n \in \mathbb{N}} H_n$. For a proof of why this implies $\partial^\infty(\tilde{X})$ is path-connected, we refer the reader to [31]. \square

Lemma 2.9 allows us to define intervals in $\partial^\infty(\tilde{X})$:

Definition 2.10 (Interval on $\partial^\infty(\tilde{X})$). Suppose X is a simple, thick negatively-curved P-manifold. An *interval* on $\partial^\infty(\tilde{X})$ with endpoints $a, b \in \partial^\infty(\tilde{X})$ is an interval on the boundary of an apartment containing the unique geodesic in \tilde{X} with endpoints a and b .

2.4. Subgroup separability, QCERF, and NPC cube complexes. We now briefly introduce key tools for the arguments in Section 4 and give a preview of some claims we will prove.

Definition 2.11 (Separability of a subgroup). Let G be a group and let $H \leq G$ be a subgroup. We say that H is *separable* in G if for any $g \in G \setminus H$, there exists a finite-index subgroup $K < G$ which contains H but $g \notin K$.

If the trivial subgroup $H = \{1\}$ is separable in G , we say that G is *residually finite*. Separability properties have attracted considerable interest in recent decades




FIGURE 4. An example of how to choose an apartment containing γ . In this example, $p(\gamma)$ is completely contained in the torus with one boundary component C_1 , which lifts to a totally geodesic subspace of a half-disk. We can fill in the missing portions of the half-disk (shaded in gray) with portions of disks that are copies of the universal cover of C_2 . While this is an especially straightforward case of the filling in procedure, the construction works in general.

due to an associated powerful topological property which allows one to promote certain immersions to embeddings:

Theorem 2.12 (Scott [38, Lemma 1.3]). *Let G be a finitely generated group and X a CW complex such that $\pi_1(X) = G$. Let $H \leq G$ and $Y \rightarrow X$ a cover corresponding to H . Then H is separable if and only if for every finite subcomplex $K \subseteq Y$, there exists an intermediate finite-sheeted cover $Y \rightarrow Z \rightarrow X$ such that K embeds as a subcomplex of Z .*

A group G is *subgroup separable*, or, for historical reasons, LERF, if every finitely-generated subgroup of G is separable. In general, it is unknown which fundamental groups of simple, thick, negatively curved surface amalgams are subgroup separable, although results due to Gitik in [19] provide partial results. For example, the fundamental groups of the surface amalgams from Figure 2 and Figure 4 are subgroup separable.

When subgroup separability of a group G is unknown, one can sometimes show that well-behaved subgroups of G are separable. A subspace Y of a geodesic metric space X is *quasiconvex* if there exists a $K \geq 0$ such that, for all $y_1, y_2 \in Y$ and all $x \in [y_1, y_2]$, $d(x, Y) \leq K$. Note that if $K = 0$, Y is convex. A group H acting on a geodesic metric space X is *quasiconvex* if the orbit Hx is a quasiconvex subspace of X for some (any) $x \in X$. We point out a stronger notion: if H acts cocompactly on and stabilizes C , a convex subset of X , then H is *convex*. A group G is *quasiconvex subgroup separable* (QCERF) if every quasiconvex subgroup of G is separable.

A particularly well-behaved class of groups are those which act properly and cocompactly on *non-positively curved (NPC) cube complexes*, also known as CAT(0) cube complexes. Recall that a *cube complex* is built from gluing unit cubes along their faces by isometries. The *link* of a vertex v , $lk(v)$, in a cube complex is a simplicial complex whose vertices lie on the edges adjacent to v . A collection of vertices in $lk(v)$ spans a $(d-1)$ -dimensional complex if and only if they lie on edges of a common d -dimensional complex. A link $lk(v)$ is *flag* if for every complete graph in $lk(v)$, there is a simplex in $lk(v)$ whose 1-skeleton is the complete graph. A cube

complex is NPC if it satisfies a combinatorial link condition: all vertex links of the cube complex are flag complexes. We will later show that simple, non-positively curved surface amalgams are NPC cube complexes.

In their seminal paper [22], Haglund and Wise define *special cube complexes*. We refer the reader to Section 3 of their paper for precise definitions, but special cube complexes satisfy three conditions: (1) every hyperplane embeds; (2) no hyperplane directly self-oscillates; and (3) no two hyperplanes inter-oscillate. A group G is *virtually special* if it has a finite-index subgroup which is the fundamental group of a special cube complex. One main result in [22] is that Gromov hyperbolic virtually special groups are QCERF:

Theorem 2.13 ([22, Corollary 7.4]). *Let X be a compact cube complex with Gromov hyperbolic fundamental group. If X is virtually special, then every quasiconvex subgroup of $\pi_1(X)$ is separable (i.e., $\pi_1(X)$ is QCERF).*

Finally, we mention a convenient criterion developed by Wise for determining whether a Gromov hyperbolic group is virtually special:

Definition 2.14 ([40, Definition 11.5]). Let \mathcal{QVH} denote the smallest collection of hyperbolic groups closed under the following four operations:

- (1) $1 \in \mathcal{QVH}$;
- (2) If $G = A *_C B$, $A, B \in \mathcal{QVH}$, and C is finitely generated and quasiconvex in G , then $G \in \mathcal{QVH}$;
- (3) If $G = A *_C$, $A \in \mathcal{QVH}$, and C is finitely-generated and quasiconvex in G , then $G \in \mathcal{QVH}$;
- (4) Let $H \subset G$ be a finite-index subgroup and let $H \in \mathcal{QVH}$. Then $G \in \mathcal{QVH}$.

The following deep theorem relates the hierarchy of groups \mathcal{QVH} to virtual specialness:

Theorem 2.15 ([40, Theorem 13.3]). *A torsion-free, Gromov hyperbolic group G is virtually special if and only if $G \in \mathcal{QVH}$.*

Again, it is unknown which fundamental groups of surface amalgams are special; however, using Theorem 2.15, we will prove that fundamental groups of simple negatively curved surface amalgams are virtually special, and therefore QCERF.

3. THE BASE CASE

Recall that for any proper metric space X , the identity map $(X, g_1) \rightarrow (X, g_2)$ lifts to a quasi-isometry $f : (\tilde{X}, \tilde{g}_1) \rightarrow (\tilde{X}, \tilde{g}_2)$ by the Švarc-Milnor Lemma. By the Morse Lemma, f in turn induces a boundary homeomorphism $\partial^\infty f : \partial^\infty(\tilde{X}, \tilde{g}_1) \rightarrow \partial^\infty(\tilde{X}, \tilde{g}_2)$. The following section is devoted to proving $\partial^\infty f$ induces an isometry, a fact summarized in Proposition 3.1:

Proposition 3.1. *Suppose X is a simple P -manifold endowed with a pair of metrics g_1 and $g_2 \in \mathcal{M}_\leq$. Furthermore, suppose that there is an inclusion of every chamber C into a closed surface $S \subset X$. If (X, g_1) and (X, g_2) have the same marked length spectrum, then there exists an isometry $\phi : (X, g_1) \rightarrow (X, g_2)$ that is induced by the boundary homeomorphism $\partial^\infty f : \partial^\infty(\tilde{X}, \tilde{g}_1) \rightarrow \partial^\infty(\tilde{X}, \tilde{g}_2)$ discussed above.*

We first provide an outline of the proof. First, in the proof of Proposition 3.2 of Section 3.1, we show that the boundary homeomorphism $\partial^\infty f : \partial^\infty(\tilde{X}, \tilde{g}_1) \rightarrow$

$\partial^\infty(\tilde{X}, \tilde{g}_2)$ induced by the identity map is Möbius (see Definition 2.5). The proof of Proposition 3.2 requires the ergodicity of the geodesic flow map, which can be deduced from general theory developed by Kaimanovich in [27] for Gromov hyperbolic spaces (see the Appendix). Using ergodicity of the geodesic flow map, the proof of Lemma 3.5 shows the cross ratio of any distinct $(a, b, c, d) \in (\partial^\infty(\tilde{X}))^4$ can be approximated arbitrarily well by lengths of closed geodesics in (X, g) . We can thus conclude that $\partial^\infty f$ is Möbius since (X, g_1) and (X, g_2) have the same marked length spectra, which determine lengths of closed geodesics and thus cross ratios. Next, in Section 3.2, we patch together isometries $\phi_S : (S, g_1|_S) \rightarrow (S, g_2|_S)$ constructed by Otal in [35] on every closed subsurface $S \subset X$ to construct a global isometry $\phi : (X, g_1) \rightarrow (X, g_2)$. In particular, we show that given two surfaces S and S' that intersect in X , $\phi_S|_{S \cap S'} = \phi_{S'}|_{S \cap S'}$ using the fact that $\partial^\infty f$ is Möbius.

3.1. The boundary map is Möbius. In this section, we prove Proposition 3.2 by adapting Otal's proof of Theorem 1 in [35]. Note that in this section, the metrics are only assumed to be locally CAT(-1).

Proposition 3.2. *Suppose (X, g_1) and (X, g_2) are two simple, thick locally CAT(-1) P-manifolds. If (X, g_1) and (X, g_2) have the same marked length spectrum, then the boundary map is Möbius.*

A key ingredient in Otal's proof is ergodicity of the geodesic flow map on surfaces, a well-known and classical result. Recall that in the setting of surfaces, geodesic flow is defined on the unit tangent bundle. For P-manifolds, however, the unit tangent bundle is undefined on the gluing curves. One, however, has the notion of a widely-studied *generalized unit tangent bundle*, which can be identified with the usual unit tangent bundle in the setting of Riemannian manifolds.

Recall that for a metric space (X, g) , a geodesic $\gamma : \mathbb{R} \rightarrow X$ has *speed* $s \geq 0$ if for every $t \in \mathbb{R}$, there exists a neighborhood $U \subset \mathbb{R}$ such that for all $t_1, t_2 \in U$, $g(\gamma(t_1), \gamma(t_2)) = s|t_1 - t_2|$. In particular, if $s = 1$, γ is a *unit-speed geodesic*.

Definition 3.3 (Generalized unit tangent bundle). Given a geodesically complete metric space X , the *generalized unit tangent bundle* of X , SX , is the space of unit-speed geodesics in X .

Suppose two unit-speed geodesics $\gamma \sim \gamma'$ if $\gamma(t) = \gamma'(-t)$ for all $t \in \mathbb{R}$. Then there is a natural identification of SX/\sim with $\mathcal{G}(X) \times \mathbb{R}$, where as before, $\mathcal{G}(X)$ denotes the space of unoriented, unparametrized geodesics in X . We say a point $x \in X$ is a *basepoint* of $\xi \in SX$ if $x = \xi(0)$. We now recall the definition of the geodesic flow map on SX :

Definition 3.4 (Geodesic flow). The *geodesic flow* on SX is the map $\phi_t : SX \rightarrow SX$, $\phi_t(\xi)(s) = \xi(t + s)$, where $s, t \in \mathbb{R}$.

For a closed negatively curved Riemannian surface, there is a unique direction in which to continue a geodesic via the exponential map. On the other hand, for a P-manifold, there are $n - 1$ directions in which to continue a geodesic hitting a point on a gluing curve γ , where n is the number of boundary components glued to γ . Due to this key difference between surfaces and P-manifolds, the usual proof of the ergodicity of the geodesic flow map on surfaces does not generalize to the setting of P-manifolds and different techniques are needed. We refer the reader to the Appendix, which summarizes a general result by Kaimanovich about the ergodicity

of the geodesic flow map on proper Gromov hyperbolic spaces. Kaimanovich's results apply in our setting since simple, thick P-manifolds equipped with locally CAT(-1) metrics are proper Gromov hyperbolic spaces.

Recall from [28] that a sequence of points $(x_n)_{n \in \mathbb{N}}$ in a Gromov hyperbolic space X converges to infinity if for any choice of $x \in X$,

$$\lim_{i,j \rightarrow \infty} \inf \langle x_i, x_j \rangle_x = \infty,$$

where $\langle \cdot \rangle_x$ denotes the Gromov Product from Definition 2.6; this definition is independent of choice of basepoint x . Furthermore, two sequences $(x_n)_{n \in \mathbb{N}}$ and $(y_n)_{n \in \mathbb{N}}$ are equivalent if

$$\lim_{i,j \rightarrow \infty} \inf \langle x_i, y_j \rangle_x = \infty.$$

Endowing $\overline{X} = X \cup \partial^\infty(X)$ with the topology described in Definition 2.13 of [28], we can think of points in ∂^∞ as limit points of equivalence classes of sequences of points in X that converge to infinity.

With the above paragraph in mind, we introduce a lemma whose proof is adapted from Otal's proof of Theorem 1 in [36]:

Lemma 3.5. *Suppose X is simple, thick P-manifold endowed with a locally CAT(-1) metric. Then given a 4-tuple of distinct points $(a, b, c, d) \in (\partial^\infty(\tilde{X}))^4$, $[abcd]$ can be approximated arbitrarily well by lengths of closed geodesics.*

Proof. Fix a 4-tuple of distinct points $(a, b, c, d) \in (\partial^\infty(\tilde{X}, \tilde{g}))^4$. By the Birkhoff Ergodic Theorem (see Section 5.3.1), we can conclude that there exists $v \in SX$ with dense orbit under the geodesic flow map ϕ_t on SX . Thus, there exist sequences of vectors $(\phi_n(\tilde{v}_1))_{n \in \mathbb{N}}$, $(\phi_{-n}(\tilde{v}_1))_{n \in \mathbb{N}}$, $(\phi_n(\tilde{v}_2))_{n \in \mathbb{N}}$, and $(\phi_{-n}(\tilde{v}_2))_{n \in \mathbb{N}}$ whose basepoints approach a , b , c , and d respectively (here, \tilde{v}_1 and \tilde{v}_2 are lifts of v in $S\tilde{X}$). As a result, Then there exist $n_i \in \mathbb{N}$ such that

$$[\phi_{n_1}(\tilde{v}_1)(0)\phi_{-n_2}(\tilde{v}_1)(0)\phi_{n_3}(\tilde{v}_2)(0)\phi_{-n_4}(\tilde{v}_2)(0)]$$

approximates $[abcd]$ arbitrarily well.

Fix any $u \in S\tilde{X}$. Then since the orbit of v is dense in $SX = S\tilde{X}/\Gamma$, for every sufficiently large $m_1 \in \mathbb{N}$ there exists some $\gamma_1 \in \Gamma$ such that $\gamma_1 \cdot \phi_{m_1}(\tilde{v}_1)$ lies arbitrarily close to $\gamma_1 \cdot u$ in $S\tilde{X}/\Gamma$. Similarly, there exists some $\gamma_3 \in \Gamma$ and sufficiently large $m_3 \in \mathbb{N}$ such that $\gamma_3 \cdot \phi_{m_3}(\tilde{v}_2)$ and $\gamma_3 \cdot u$ are arbitrarily close. From this, by choosing $\gamma_l, \gamma'_m \in \Gamma$ accordingly, we can construct sequences of vectors $(\gamma_m \cdot u)_{m \in \mathbb{N}}$ and $(\gamma'_l \cdot u)_{l \in \mathbb{N}}$ whose basepoints limit to a and c respectively (see Figure 5). Note that the translation lengths of γ_l and γ'_m will become arbitrarily large.

Let $\ell_{\gamma_m^{-1}\gamma'_l}$ be the length of the fundamental domain of the action of $\gamma_m^{-1} \circ \gamma'_l$ on its translation axis. Note that $\ell_{\gamma_m^{-1}\gamma'_l}$, the length of a closed geodesic corresponding to $\gamma_m^{-1} \circ \gamma'_l$, is approximated arbitrarily well by $\tilde{g}((\gamma_m \cdot u)(0), (\gamma'_l \cdot u)(0))$ since by construction, the basepoints of $(\gamma_m \cdot u)_{m \in \mathbb{N}}$ and $(\gamma'_l \cdot u)_{l \in \mathbb{N}}$ approach the endpoints of a geodesic $(a, c) \in (\tilde{X}, \tilde{g})$ and $\gamma_1 \cdot u = (\gamma_1 \circ \gamma_3^{-1})(\gamma_3 \cdot u)$. We apply a similar argument to construct sequences $(\gamma''_k \cdot u)_{k \in \mathbb{N}}$ and $(\gamma'''_j \cdot u)_{j \in \mathbb{N}}$ that tend to b and d respectively. Furthermore, we can construct the sequences such that \tilde{g} -distances between basepoints of terms of the sequences are arbitrarily close to lengths of long closed geodesics.

Since $(\gamma_m \cdot u)_{m \in \mathbb{N}}$, $(\gamma''_k \cdot u)_{k \in \mathbb{N}}$, $(\gamma'_l \cdot u)_{l \in \mathbb{N}}$, and $(\gamma'''_j \cdot u)_{j \in \mathbb{N}}$ are sequences whose basepoints tend to a , b , c , and d respectively, we have that $[abcd]$ is approximated




FIGURE 5. An illustration of the setup of Lemma 3.5

arbitrarily well by $[(\gamma_m \cdot u)(0)(\gamma_k'' \cdot u)(0)(\gamma_l' \cdot u)(0)(\gamma_j''' \cdot u)(0)]$. Since we can approximate $[(\gamma_m \cdot u)(0)(\gamma_k'' \cdot u)(0)(\gamma_l' \cdot u)(0)(\gamma_j''' \cdot u)(0)]$ by adding and subtracting lengths of closed geodesics, the result then follows. \square

We can now prove Proposition 3.2. By Lemma 3.5, for all $i \in \mathbb{N}$, there is some quadruple of geodesic arcs $\alpha_i = (x_i, y_i)$, $\beta_i = (x'_i, y'_i)$, $\gamma_i = (x_i, y'_i)$, and $\delta_i = (x'_i, y_i)$ projecting to closed geodesics in X such that:

$$|[abcd] - (\tilde{g}_1(x_i, y_i) + \tilde{g}_1(x'_i, y'_i) - \tilde{g}_1(x_i, y'_i) - \tilde{g}_1(x'_i, y_i))| < \frac{1}{i}.$$

By construction, (x_i, y_i, x'_i, y'_i) converges to (a, b, c, d) in (\tilde{X}, \tilde{g}_1) so $(f(x_i), f(y_i), f(x'_i), f(y'_i))$ will converge to $(\partial^\infty f(a), \partial^\infty f(b), \partial^\infty f(c), \partial^\infty f(d))$ in (\tilde{X}, \tilde{g}_2) since f is continuous. Thus, by definition of cross ratio, we have that:

$$[abcd] = \lim_{i \rightarrow \infty} \tilde{g}_1(x_i, y_i) + \tilde{g}_1(x'_i, y'_i) - \tilde{g}_1(x_i, y'_i) - \tilde{g}_1(x'_i, y_i)$$

and

$$\begin{aligned} \lim_{i \rightarrow \infty} \tilde{g}_2(f(x_i, y_i)) + \tilde{g}_2(f(x'_i, y'_i)) - \tilde{g}_2(f(x_i, y'_i)) - \tilde{g}_2(f(x'_i, y_i)) \\ = [\partial^\infty f(a) \partial^\infty f(b) \partial^\infty f(c) \partial^\infty f(d)]. \end{aligned}$$

Recall that since (x_i, y_i) , (x'_i, y'_i) , (x'_i, y_i) and (x_i, y'_i) each project to closed geodesics in X , the distances between their endpoints are all lengths of elements in $\pi_1(X)$. Thus, since (X, g_1) and (X, g_2) have the same marked length spectrum, $\tilde{g}_1(x_i, y_i) = \ell_{g_1}(\alpha_i) = \ell_{g_2}(\alpha_i) = \tilde{g}_2(f(x_i, y_i))$, and a similar statement holds for the other distances as well. It then follows that $[abcd] = [\partial^\infty f(a) \partial^\infty f(b) \partial^\infty f(c) \partial^\infty f(d)]$.

3.2. Proof of Proposition 3.1. In order to prove Proposition 3.1, we need a few auxiliary lemmas. The first lemma is a basic fact about patching isometries together.

Recall a metric space (X, d) is *convex* if for any two points $x, y \in X$, there exists $z \in X$ distinct from x and y such that $d(x, z) + d(z, y) = d(x, y)$.

Lemma 3.6. *Suppose U_i and V_i are complete, convex, locally compact metric spaces, where $i = 1, 2$. Suppose $\phi_i : U_i \rightarrow V_i$ are (invertible) isometries, and $\phi_1|_{U_1 \cap U_2} = \phi_2|_{U_1 \cap U_2}$. Then there exists an isometry $\phi : U_1 \cup U_2 \rightarrow V_1 \cup V_2$ such that $\phi|_{U_i} = \phi_i|_{U_i}$ for $i = 1, 2$.*




FIGURE 6. An illustration of the setup of Lemma 3.6

Proof. In the following proof, for a metric space X , d_X will denote the distance function of X . We define ϕ in the natural way: $\phi(x) = \phi_i(x)$ if $x \in U_i$. It suffices to show the theorem is true when p and q are not both in U_1 or U_2 . Suppose without loss of generality $p \in U_1$ and $q \in U_2$. By Hopf-Rinow, since $U_1 \cup U_2$ is complete, convex, and locally compact, then there exists a minimizing geodesic between p and q , which we will call $[p, q]$. Let $r \in (U_1 \cap U_2) \cap [p, q]$, so $d_{U_1 \cup U_2}(p, r) + d_{U_1 \cup U_2}(r, q) = d_{U_1 \cup U_2}(p, q)$ (see Figure 6).

Since ϕ_i is an isometry for $i = 1, 2$, it follows that:

$$\begin{aligned} d_{V_1 \cup V_2}(\phi(p), \phi(r)) + d_{V_1 \cup V_2}(\phi(r), \phi(q)) &= d_{V_1}(\phi_1(p), \phi_1(r)) + d_{V_2}(\phi_2(r), \phi_2(q)) \\ &= d_{U_1}(p, r) + d_{U_2}(r, q) = d_{U_1 \cup U_2}(p, q). \end{aligned}$$

By triangle inequality, $d_{V_1 \cup V_2}(\phi(p), \phi(q)) \leq d_{U_1 \cup U_2}(p, q)$. For the reverse direction, apply a symmetric argument: since $V_1 \cup V_2$ is complete, convex, and locally compact, there exists a minimizing geodesic between $\phi(p)$ and $\phi(q)$, $[\phi(p), \phi(q)]$. Choose $r' \in [\phi(p), \phi(q)] \cap (V_1 \cap V_2)$ such that $d_{V_1}(\phi(p), r') + d_{V_2}(r', \phi(q)) = d_{V_1 \cup V_2}(\phi(p), \phi(q))$. Since ϕ_i is invertible for $i = 1, 2$,

$$\begin{aligned} d_{V_1 \cup V_2}(\phi(p), \phi(q)) &= d_{V_1 \cup V_2}(\phi_1(p), \phi_2(q)) = d_{V_1}(\phi_1(p), r') + d_{V_2}(r', \phi_2(q)) \\ &= d_{U_1}(p, \phi_1^{-1}(r')) + d_{U_2}(\phi_2^{-1}(r'), q) = d_{U_1 \cup U_2}(p, \phi^{-1}(r')) \\ &\quad + d_{U_1 \cup U_2}(\phi^{-1}(r'), q) \\ &\geq d_{U_1 \cup U_2}(p, q). \end{aligned}$$

Then it follows $d_{V_1 \cup V_2}(\phi(p), \phi(q)) = d_{U_1 \cup U_2}(p, q)$. \square

The following is a technical lemma about the cross-ratio defined in Equation (1).

Lemma 3.7. *Suppose $(\tilde{S}, \tilde{g}|_{\tilde{S}})$ and $(\tilde{S}', \tilde{g}|_{\tilde{S}'})$ are two lifts of $(S, g|_S)$ and $(S', g|_{S'})$ respectively that meet at a branching geodesic $\tilde{\gamma} \subset (\tilde{X}, \tilde{g})$. Then for every $p \in \tilde{\gamma}$, there exist four points a, b, c, d such that the geodesics (a, c) , (b, d) , (a, d) , and (b, c) all meet at p and $[abcd] = 0$.*

Proof. Choose some $p \in \tilde{\gamma}$. We will find a set of four geodesics (a, c) , (b, d) , (a, d) , and (b, c) that all meet at p . Choose an arbitrary $(a, c) \subset (\tilde{S}, \tilde{g}|_{\tilde{S}})$ that transversely intersects $\tilde{\gamma}$ at p in \tilde{S} . Suppose $\tilde{\gamma}$ divides \tilde{S}' into two half planes, which we will call H_1' and H_2' . Consider a map $h : \partial^\infty(H_1') \rightarrow \tilde{\gamma}$ that maps a point $x \notin \partial^\infty(H_1')$ to the point $h(x) = \tilde{\gamma} \cap (x, c)$ where (x, c) denotes the geodesic arc with endpoints




FIGURE 7. By an intermediate value theorem argument, there exists some $x \in \partial^\infty H'_1$ such that $\tilde{h}(x) = p$ for a given point $p \in \tilde{\gamma}$

x and c . Note that h extends continuously to a map $\tilde{h} : \partial^\infty(H'_1) \rightarrow \tilde{\gamma} \cup \{q, r\}$, where $q, r \in \partial^\infty(\tilde{S}) \cap \partial^\infty(\tilde{S}')$ are endpoints of $\tilde{\gamma}$ in the following way: $\tilde{h}(q) = q$ and $\tilde{h}(r) = r$. As we vary x along $\partial^\infty(H'_1)$, $h(x)$ varies continuously along $\tilde{\gamma}$, so by an Intermediate Value Theorem argument, since $\tilde{\gamma}$ and $\partial^\infty(H'_1)$ are both connected, it follows that $\tilde{h}(x) = p$ for some $x \in \partial^\infty(H'_1)$ (see Figure 7). We set b equal to this x .

Note that (b, c) is a geodesic that passes through p . Continue the geodesic ray (b, p) into $H'_2 \subset \tilde{S}'$ to obtain a bi-infinite geodesic $(b, d) \subset \tilde{S}'$. Note that (a, d) will also be a geodesic that passes through p . Thus, we have found four geodesics (a, c) , (b, d) , (a, d) , and (b, c) that all intersect at p . It is then possible to find a sequence (a_i, b_i, c_i, d_i) where $a_i \in (p, a)$, $b_i \in (p, b)$, $c_i \in (p, c)$ and $d_i \in (p, d)$ such that:

$$\begin{aligned} [abcd] &= \lim_{(a_i, b_i, c_i, d_i) \rightarrow (a, b, c, d)} \tilde{g}_1(a_i, c_i) + \tilde{g}_1(b_i, d_i) - \tilde{g}_1(a_i, d_i) - \tilde{g}_1(b_i, c_i) \\ &= \lim_{(a_i, b_i, c_i, d_i) \rightarrow (a, b, c, d)} \tilde{g}_1(a_i, p) + \tilde{g}_1(p, c_i) + \tilde{g}_1(b_i, p) + \tilde{g}_1(p, d_i) \\ &\quad - \tilde{g}_1(a_i, p) - \tilde{g}_1(p, d_i) - \tilde{g}_1(b_i, p) - \tilde{g}_1(p, c_i) \\ &= 0. \end{aligned}$$

□

Let \mathcal{S} be a collection of closed surfaces that covers X . We know such a covering exists; for each chamber $C_i \subset X$, note that $C_i \subset S_i$ by assumption, where S_i is a closed surface. Then $\{S_i\}_{i=1}^n$ is a collection of closed surfaces that covers X .

Lemma 3.8. *Suppose $S, S' \in \mathcal{S}$ are two closed surfaces in X that are identified along some set of gluing curves $\{\gamma_1, \dots, \gamma_n\}$. Then there exist isometries $\phi_S : (S, g_1|_S) \rightarrow (S, g_2|_S)$ and $\phi_{S'} : (S', g_1|_{S'}) \rightarrow (S', g_2|_{S'})$ where $\phi_S|_{\gamma_i} = \phi_{S'}|_{\gamma_i}$ for all $1 \leq i \leq n$.*

Proof. Suppose (X, g_1) and (X, g_2) have the same marked length spectra. As before, $\partial^\infty f$ will denote the boundary homeomorphism between $\partial^\infty(\tilde{X}, \tilde{g}_1)$ and $\partial^\infty(\tilde{X}, \tilde{g}_2)$.

Recall \tilde{S} and \tilde{S}' are apartments that are arbitrary lifts of S and S' in \tilde{X} such that $\tilde{S} \cap \tilde{S}' \neq \emptyset$. We can then use the restrictions of the boundary homeomorphism




FIGURE 8. Given a sequence of points $\{x_i\}_{i=1}^n$ converging radially to x , one can find a sequence of geodesics ξ_i and η_i with endpoints a_i, b_i, c_i, d_i converging to x . The images of these endpoints under the boundary homeomorphism $\partial^\infty f$ will also converge to $\partial^\infty f(x) = \lim_{i \rightarrow \infty} \widetilde{\phi}_S(x_i)$.

$\partial^\infty f|_{\partial^\infty(\widetilde{S})}$ and $\partial^\infty f|_{\partial^\infty(\widetilde{S}')}$ to construct isometries $\phi_S : (S, g_1|_S) \rightarrow (S, g_2|_S)$ and $\phi_{S'} : (S', g_1|_{S'}) \rightarrow (S', g_2|_{S'})$ via the methods of Otal. In other words, given $x \in (\widetilde{S}, \widetilde{g}_1|_{\widetilde{S}})$, consider the set of all geodesics in $(\widetilde{S}, \widetilde{g}_1|_{\widetilde{S}})$ that intersect at x . By [35], the set of geodesics mapped to $(\widetilde{S}, \widetilde{g}_2|_{\widetilde{S}})$ via the boundary map $\partial^\infty f|_{\partial^\infty(\widetilde{S})}$ intersect at a single point $\widetilde{\phi}_S(x)$; $\widetilde{\phi}_{S'}$ is constructed similarly. Otal shows that $\widetilde{\phi}_S$ and $\widetilde{\phi}_{S'}$ are π_1 -equivariant.

We first check that $\widetilde{\phi}_S$ and $\widetilde{\phi}_{S'}$ extend to $\partial^\infty f|_{\partial^\infty(\widetilde{S})}$ and $\partial^\infty f|_{\partial^\infty(\widetilde{S}')}$ respectively. Indeed, consider a sequence of points $\{x_i\}_{i \in \mathbb{N}}$ that converge radially towards a point $x \in \partial^\infty \widetilde{S}$, so $x = \lim_{i \rightarrow \infty} x_i$. It is possible to find a sequence of pairs of geodesics (ξ_i, η_i) , such that $\xi_i \cap \eta_i = x_i$ and $a_i, b_i, c_i, d_i \rightarrow x$, where a_i and b_i are the endpoints of ξ_i and c_i and d_i are the endpoints of η_i (see Figure 8). Then since $\partial^\infty f$ is continuous, $\lim_{i \rightarrow \infty} \partial^\infty f(a_i) = \partial^\infty f(x)$, and the same is true for b_i, c_i , and d_i . Then:

$$\begin{aligned} \lim_{i \rightarrow \infty} \widetilde{\phi}_S(x_i) &= \lim_{i \rightarrow \infty} \widetilde{\phi}_S((a_i, b_i) \cap (c_i, d_i)) \\ &= \lim_{i \rightarrow \infty} (\partial^\infty f(a_i), \partial^\infty f(b_i)) \cap (\partial^\infty f(c_i), \partial^\infty f(d_i)) = \partial^\infty f(x). \end{aligned}$$

As before, $(\partial^\infty f(a_i), \partial^\infty f(b_i))$ denotes a geodesic with endpoints $\partial^\infty f(a_i)$ and $\partial^\infty f(b_i)$.

As a result, we conclude that $\partial^\infty f|_{\partial^\infty(\widetilde{S})} = \partial^\infty \widetilde{\phi}_S$, $\partial^\infty f|_{\partial^\infty(\widetilde{S}')} = \partial^\infty \widetilde{\phi}_{S'}$, and $\partial^\infty \widetilde{\phi}_S|_{\widetilde{S} \cap \widetilde{S}'} = \partial^\infty f|_{\partial^\infty(\widetilde{S} \cap \widetilde{S}')} = \partial^\infty \widetilde{\phi}_{S'}|_{\widetilde{S} \cap \widetilde{S}'}$. In particular, if $p \in \partial^\infty(\widetilde{S}) \cap \partial^\infty(\widetilde{S}')$, then:

$$(5) \quad \partial^\infty \widetilde{\phi}_S(p) = \partial^\infty f|_{\widetilde{S} \cap \widetilde{S}'}(p) = \partial^\infty \widetilde{\phi}_{S'}(p).$$

Since isometries between Riemannian manifolds necessarily preserve geodesics, $\widetilde{\phi}_S$ and $\widetilde{\phi}_{S'}$ must map the branching geodesic $\widetilde{\gamma} = (p, q) \subset (\widetilde{S} \cap \widetilde{S}', \widetilde{g}_1)$ to geodesics in $(\widetilde{S}, \widetilde{g}_2|_{\widetilde{S}})$ and $(\widetilde{S}', \widetilde{g}_2|_{\widetilde{S}'})$ respectively. By Equation (5), $\widetilde{\phi}_S(\widetilde{\gamma})$ and $\widetilde{\phi}_{S'}(\widetilde{\gamma})$ must share endpoints. Since $\widetilde{S}, \widetilde{S}' \subset (\widetilde{X}, \widetilde{g}_2)$ and geodesics between two given boundary




FIGURE 9. If $\widetilde{\phi}_S(x) \neq \widetilde{\phi}_{S'}(x)$, then $[\partial^\infty f(a)\partial^\infty f(a')\partial^\infty f(b)\partial^\infty f(b')] \neq 0$, but a, a', b , and b' were chosen so that $[aa'bb'] = 0$, a contradiction since $\partial^\infty f$ is Möbius. As before, the pink geodesics are branching.

points in the CAT(-1) space $(\widetilde{X}, \widetilde{g}_2)$ are unique, $\widetilde{\phi}_S(\widetilde{\gamma}) = \widetilde{\phi}_{S'}(\widetilde{\gamma})$ necessarily. Thus, for every branching geodesic $\widetilde{\gamma} \subset (\widetilde{S}, \widetilde{g}_1) \cap (\widetilde{S}', \widetilde{g}_1)$, $\widetilde{\phi}_S(\widetilde{\gamma}) = \widetilde{\phi}_{S'}(\widetilde{\gamma})$.

It then suffices to show that for $x \in \widetilde{\gamma}$, $\widetilde{\phi}_S(x) = \widetilde{\phi}_{S'}(x)$. Suppose there exists some $x \in \widetilde{\gamma}$ where $\widetilde{\phi}_S(x) \neq \widetilde{\phi}_{S'}(x)$. By Proposition 3.2, if an isomorphism of fundamental groups preserves the marked length spectrum of two CAT(-1) spaces, then the induced map at infinity is Möbius. As a consequence, $\partial^\infty f$ is Möbius. By Lemma 3.7, for any $x \in \widetilde{\gamma}$ there exists a pair of bi-infinite geodesics $\xi = (a, b) \in (\partial^\infty \widetilde{S} \times \partial^\infty \widetilde{S}) \setminus \Delta$ and $\xi' = (a', b') \in (\partial^\infty \widetilde{S}' \times \partial^\infty \widetilde{S}') \setminus \Delta$ such that $[aa'bb'] = 0$ (see Figure 9). Note that (a, b) and (a', b') map to a pair of non-intersecting geodesics since we assumed that $\widetilde{\phi}_S(x) \neq \widetilde{\phi}_{S'}(x)$.

Let $a_i, b_i \in \xi$, $a'_i, b'_i \in \xi'$, $(a_i, b_i) \rightarrow (a, b)$ and $(a'_i, b'_i) \rightarrow (a', b')$. The Parallelogram Law for CAT(κ) spaces (see Exercise 1.16 of [5]) gives us the following inequality:

$$(6) \quad \widetilde{g}_2(\widetilde{\phi}_S(a'_i), \widetilde{\phi}_S(b_i)) + \widetilde{g}_2(\widetilde{\phi}_S(a_i), \widetilde{\phi}_S(b'_i)) \neq \widetilde{g}_2(\widetilde{\phi}_S(a_i), \widetilde{\phi}_S(b_i)) + \widetilde{g}_2(\widetilde{\phi}_{S'}(a'_i), \widetilde{\phi}_{S'}(b'_i)).$$

We therefore conclude that:

$$\begin{aligned} [\partial^\infty f(a)\partial^\infty f(a')\partial^\infty f(b)\partial^\infty f(b')] &= \lim_{i \rightarrow \infty} [\widetilde{\phi}_S(a_i)\widetilde{\phi}_{S'}(a'_i)\widetilde{\phi}_S(b_i)\widetilde{\phi}_{S'}(b'_i)] \\ &= \lim_{i \rightarrow \infty} \widetilde{g}_2(\widetilde{\phi}_S(a_i), \widetilde{\phi}_S(b_i)) + \widetilde{g}_2(\widetilde{\phi}_{S'}(a'_i), \widetilde{\phi}_{S'}(b'_i)) \\ &\quad - \widetilde{g}_2(\widetilde{\phi}_S(a_i), \widetilde{\phi}_{S'}(b'_i)) - \widetilde{g}_2(\widetilde{\phi}_S(b_i), \widetilde{\phi}_{S'}(a'_i)) \\ &\stackrel{(6)}{\neq} 0 = [aa'bb']. \end{aligned}$$

We then conclude that $\partial^\infty f$ does not preserve the cross ratio, a contradiction. It then follows that $\widetilde{\phi}_S$ and $\widetilde{\phi}_{S'}$ must pointwise agree on $\widetilde{\gamma}$, as claimed. Finally, π_1 -invariance of Otal's maps gives us our desired result. \square

Lemma 3.9. *Suppose that S and S' are two closed surfaces in X , and suppose $S \cap S' = \{C_i\}_{i=1}^n$, where each C_i is a chamber. Then for all $C_i \in S \cap S'$, $\phi_S(x) = \phi_{S'}(x)$ where as before, $\phi_S : (S, g_1|_S) \rightarrow (S, g_2|_S)$ and $\phi_{S'} : (S', g_1|_{S'}) \rightarrow (S', g_2|_{S'})$ are the isometries constructed by Otal.*

Proof. By Lemma 3.8, it suffices to prove $\phi_S(x) = \phi_{S'}(x)$ for $x \in \text{Int}(C_i)$, where $C_i \subset S \cap S'$. Suppose \tilde{x} is a lift of x inside the interior of $\widetilde{C_i}$, some polygonal lift of C_i in \widetilde{X} . Furthermore, suppose $\widetilde{C_i}$ is adjacent to some $\tilde{\gamma}_1$, a lift of a gluing curves $\gamma \in S \cap S'$.

We first find $\xi = (w_1, z_1)$ and $\eta = (w_2, z_2)$ on \tilde{S} , some lift of S in \widetilde{X} , such that ξ and η intersect at \tilde{x} . Recall that $\widetilde{\phi_S(\tilde{x})} = (\partial^\infty f(w_1), \partial^\infty f(z_1)) \cap (\partial^\infty f(w_2), \partial^\infty f(z_2))$, where $\partial^\infty f$ is the boundary homeomorphism induced by the identity on X . Similarly, we choose $\xi' = (w'_1, z'_1)$ and $\eta' = (w'_2, z'_2)$ in a copy of \tilde{S}' such that $\xi' \cap \eta' = \tilde{x}$ so $\widetilde{\phi_{S'}(\tilde{x})} = (\partial^\infty f(w'_1), \partial^\infty f(z'_1)) \cap (\partial^\infty f(w'_2), \partial^\infty f(z'_2))$. We want to use our choice of ξ, ξ', η , and η' to show that $\widetilde{\phi_S(\tilde{x})} = \widetilde{\phi_{S'}(\tilde{x})}$.

In order to construct ξ , consider a geodesic arc $\alpha = [\tilde{x}, \tilde{\gamma}_1(a)]$ joining \tilde{x} to some point $\tilde{\gamma}_1(a)$ on $\tilde{\gamma}_1 \subset \tilde{X}$. Furthermore, we require that $\tilde{\gamma}_1(a)$ is chosen so that α lies entirely inside $\widetilde{C_i}$. To find ξ , extend α in \tilde{S} via the exponential map to a geodesic lying entirely inside a single lift of S , \tilde{S} . Similarly, construct $\xi' \subset \tilde{S}'$ by extending α to a geodesic lying completely inside \tilde{S}' . Along the same vein, construct η by extending $\beta = [\tilde{x}, \tilde{\gamma}_1(b)]$ on \tilde{S} , where $\tilde{\gamma}_1(b)$ is another point on the image of $\tilde{\gamma}_1$ and β also lies entirely inside $\widetilde{C_i}$. Similarly, we can extend β in \tilde{S}' to obtain η' . See Figure 10 for an example of a choice of α and β given some \tilde{x} .

We thus have two pairs of geodesics $\xi, \eta \subset \tilde{S}$, and $\xi', \eta' \subset \tilde{S}'$. Let $p^{-1}(\{C_i\})$ denote the collection of lifts of chambers in $\{C_i\}_{i=1}^n = S \cap S'$ in \widetilde{X} . Note that ξ and ξ' either agree on a geodesic arc $[\tilde{\gamma}_2(c), \tilde{\gamma}_1(a)]$ (where $\tilde{\gamma}_2$ is a lift of a gluing curve in $S \cap S'$) or a geodesic ray with endpoint $\tilde{\gamma}_1(a)$ and limit point $y \in \partial^\infty(p^{-1}(\{C_i\}))$. Similarly, η and η' agree either on a geodesic arc $[\tilde{\gamma}_3(d), \tilde{\gamma}_1(b)]$ or a geodesic ray with endpoint $\tilde{\gamma}_1(b)$ and limit point $z \in \partial^\infty(p^{-1}(\{C_i\}))$. For example, in Figure 10, $\xi \subset \tilde{S}$ and $\xi' \subset \tilde{S}'$ agree on a geodesic arc $[\tilde{\gamma}_1(a), \tilde{\gamma}_2(c)]$ while $\eta \subset \tilde{S}$ and $\eta' \subset \tilde{S}'$ agree on a geodesic ray $[\tilde{\gamma}_1(b), y_2)$.

Suppose that ξ and ξ' agree on a geodesic arc $[\tilde{\gamma}_2(c), \tilde{\gamma}_1(a)] \subset p^{-1}(\{C_i\})$. By Lemma 3.8, we know that we can define points $x_1 := \phi_S(\tilde{\gamma}_1(a)) = \phi_{S'}(\tilde{\gamma}_1(a))$ and $x_2 := \phi_S(\tilde{\gamma}_2(c)) = \phi_{S'}(\tilde{\gamma}_2(c))$. Then it follows that $\widetilde{\phi_S(\xi)}$ and $\widetilde{\phi_{S'}(\xi')}$ agree on the unique geodesic arc between x_1 and x_2 . In other words,

$$\widetilde{\phi_S(\xi \cap \xi')} = \widetilde{\phi_S([\tilde{\gamma}_1(a), \tilde{\gamma}_2(c)])} = [x_1, x_2] = \widetilde{\phi_{S'}([\tilde{\gamma}_1(a), \tilde{\gamma}_2(c)])} = \widetilde{\phi_{S'}(\xi \cap \xi')}.$$

Suppose on the other hand that ξ and ξ' agree on a geodesic ray $[\tilde{\gamma}_1(a), y)$. Again, we can set $x_1 := \phi_S(\tilde{\gamma}_1(a)) = \phi_{S'}(\tilde{\gamma}_1(a))$. Note that we can define $y_1 := \partial^\infty \widetilde{\phi_S(y)} = \partial^\infty \widetilde{\phi_{S'}(y)}$ since in the proof of Lemma 3.8, we determined that $\partial^\infty \widetilde{\phi_S}|_{\widetilde{S \cap S'}} = \partial^\infty f|_{\widetilde{S \cap S'}} = \partial^\infty \widetilde{\phi_{S'}}|_{\widetilde{S \cap S'}}$. Since the geodesic between x_1 and y_1 is unique in $\widetilde{X} \cup \partial^\infty(\widetilde{X})$, it then follows that $\widetilde{\phi_S(\xi)}$ and $\widetilde{\phi_{S'}(\xi')}$ necessarily agree on a geodesic ray with endpoints x_1 and y_1 . Thus,

$$\widetilde{\phi_S(\xi \cap \xi')} = \widetilde{\phi_S([\tilde{\gamma}_1(a), y))} = [x_1, y_1] = \widetilde{\phi_{S'}([\tilde{\gamma}_1(a), y))} = \widetilde{\phi_{S'}(\xi \cap \xi')}.$$

It then follows that $\widetilde{\phi_S(\xi \cap \xi')}$ and $\widetilde{\phi_{S'}(\xi \cap \xi')}$ agree, and the same argument holds for $\widetilde{\phi_S(\eta \cap \eta')}$ and $\widetilde{\phi_{S'}(\eta \cap \eta')}$. We now show that $\widetilde{\phi(\xi \cap \xi')} := \widetilde{\phi_S(\xi \cap \xi')} = \widetilde{\phi_{S'}(\xi \cap \xi')}$ and $\widetilde{\phi(\eta \cap \eta')} := \widetilde{\phi_S(\eta \cap \eta')} = \widetilde{\phi_{S'}(\eta \cap \eta')}$ intersect at a point $\tilde{x} \in \widetilde{\phi_S(\tilde{S} \cap \tilde{S}')} = \widetilde{\phi_{S'}(\tilde{S} \cap \tilde{S}'})$ by doing some case work.

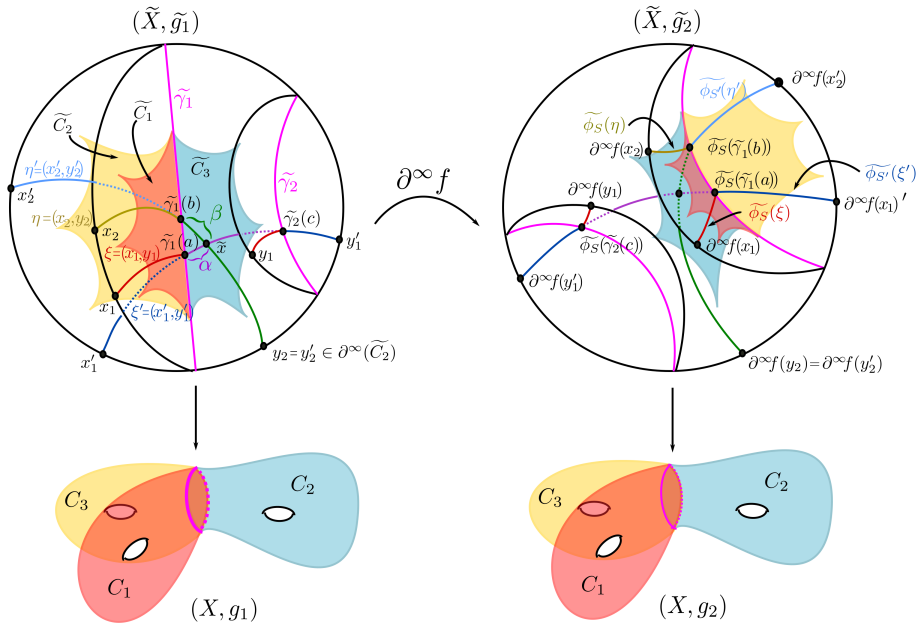


FIGURE 10. We extend $\alpha = (\tilde{x}, \tilde{\gamma}_1(a))$ in the picture on the left to $\xi \subset \tilde{S}$ and $\xi' \subset \tilde{S}'$. Similarly, we extend $\beta = (\tilde{x}, \tilde{\gamma}_1(b))$ where $b \neq a$ to $\eta \subset \tilde{S}$ and $\eta' \subset \tilde{S}'$. While ξ and ξ' agree on a geodesic arc $[\tilde{\gamma}_2(c), \tilde{\gamma}_1(a)]$, η and η' agree on a geodesic ray $[\tilde{\gamma}_1(b), y)$.




FIGURE 11. An illustration of Case 1; $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic rays




FIGURE 12. An example of Case 2; $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic arcs




FIGURE 13. An example of Case 3; $\xi \cap \eta$ is a geodesic ray and $\xi' \cap \eta'$ is a geodesic arc

- (1) **Case One:** $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic rays. Suppose $\xi \cap \xi' = [p_1, r)$ and $\eta \cap \eta' = [p_2, s)$, where $p_1 = \tilde{\gamma}_1(a)$ and $p_2 = \tilde{\gamma}_1(b)$ lie on a branching geodesic $\tilde{\gamma}_1$ and $r, s \in \partial^\infty(p^{-1}(\{C_i\}))$. Additionally, suppose $\tilde{\gamma}_1$ has endpoints p and q , $p = \lim_{t \rightarrow -\infty} \tilde{\gamma}_1(t)$, and $q = \lim_{t \rightarrow \infty} \tilde{\gamma}_1(t)$. Consider an apartment A in \tilde{X} that contains $\tilde{\gamma}_1$, r , and s (e.g. \tilde{S} or \tilde{S}'). By construction, if without loss of generality $a < b$ (so that p_1 is “closer” to p and p_2 is “closer” to q), the order of points, according to some fixed orientation on $\partial^\infty(A)$, will be p, q, r, s , and $\partial^\infty f|_A$ will preserve that order. Furthermore, since the orientation of $\tilde{\gamma}_1$ will be preserved under $\partial^\infty f|_A$, it follows that the geodesic rays $[\tilde{\phi}(p_2), \partial^\infty f(s)) = [\tilde{\phi}(p_2), \partial^\infty \tilde{\phi}(s))$ and $[\tilde{\phi}(p_1), \partial^\infty f(r)) = [\tilde{\phi}(p_1), \partial^\infty \tilde{\phi}(r))$ will necessarily intersect (see Figure 11).
- (2) **Case Two:** $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic arcs. Suppose $\xi \cap \xi' = [p_1, p_3]$ and $\eta \cap \eta' = [p_2, p_4]$. Let p and q be defined as before in case one. As before, suppose $\tilde{\gamma}_1(a) = p_1 \in \xi \cap \xi'$ and $\tilde{\gamma}_1(b) = p_2 \in \eta \cap \eta'$ where $a < b$. Additionally, suppose that $\tilde{\gamma}_2(c) = p_3$ and $\tilde{\gamma}_2(d) = p_4$ where $d < c$, and $\lim_{t \rightarrow -\infty} \tilde{\gamma}_2(t) = s$ and $\lim_{t \rightarrow \infty} \tilde{\gamma}_2(t) = r$. Again, one can find some apartment A

(e.g. \tilde{S} or \tilde{S}') that contains γ_1 and γ_2 , and as before, we have that $\partial^\infty f|_A$ is a circle homeomorphism which preserves the cyclic order of the four-tuple $[p, q, r, s]$. Since the orientation of $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ is, as before, preserved, it follows that $[\tilde{\phi}(p_1), \tilde{\phi}(p_3)]$ and $[\tilde{\phi}(p_2), \tilde{\phi}(p_4)]$ intersect (see Figure 12).

- (3) **Case Three: $\xi \cap \xi'$ or $\eta \cap \eta'$ is a geodesic ray, and the other is a geodesic arc.** For the last case, suppose without loss of generality that $\xi \cap \xi'$ is a geodesic ray $[p_1, r)$ and $\eta \cap \eta'$ is a geodesic arc $[p_2, p_3]$. Again, suppose $p_1 = \tilde{\gamma}_1(a)$ and $p_2 = \tilde{\gamma}_1(b)$ for some $a < b$ where $\tilde{\gamma}_1$ has endpoints $p = \lim_{t \rightarrow -\infty} \tilde{\gamma}_1(t)$ and $q = \lim_{t \rightarrow \infty} \tilde{\gamma}_1(t)$. Suppose $p_3 = \tilde{\gamma}_2(c)$ where $c \in \mathbb{R}$ and the endpoints of $\tilde{\gamma}_2$ are s and t . Again, choose some apartment A (e.g. \tilde{S} or \tilde{S}') that contains all five points $(p, q, r, s, \text{ and } t)$. Then, after fixing an orientation, the points will be ordered p, q, r, s, t , and their order will be preserved under $\partial^\infty f|_A$. Again, since the orientation of $\tilde{\gamma}_1$ is fixed, it follows that $[\tilde{\phi}(p_1), \partial^\infty f(r)) = [\tilde{\phi}(p_1), \partial^\infty \tilde{\phi}(r))$ and $[\tilde{\phi}(p_2), \tilde{\phi}(p_3)]$ will intersect. (See Figure 13.)

In conclusion, $\widetilde{\phi_S}(\xi \cap \xi') = \widetilde{\phi_{S'}}(\xi \cap \xi')$ and $\widetilde{\phi_S}(\eta \cap \eta') = \widetilde{\phi_{S'}}(\eta \cap \eta')$ intersect at some point in $\widetilde{\phi_S}(p^{-1}(\{C_i\})) = \widetilde{\phi_{S'}}(\{p^{-1}(\{C_i\})\})$. Since $\widetilde{\phi_S}(\xi \cap \eta) = \widetilde{\phi_S}(\tilde{x})$, $\widetilde{\phi_{S'}}(\xi' \cap \eta') = \widetilde{\phi_{S'}}(\tilde{x})$ and $\widetilde{\phi_S}$ and $\widetilde{\phi_{S'}}$ are injective, it follows that:

$$\begin{aligned} \widetilde{\phi_{S'}}(\tilde{x}) &= \widetilde{\phi_{S'}}(\xi \cap \xi' \cap \eta \cap \eta') = \widetilde{\phi_{S'}}(\xi \cap \xi') \cap \widetilde{\phi_{S'}}(\eta \cap \eta') = \widetilde{\phi_S}(\xi \cap \xi') \cap \widetilde{\phi_S}(\eta \cap \eta') \\ &= \widetilde{\phi_S}(\xi \cap \xi' \cap \eta \cap \eta') = \widetilde{\phi_S}(\tilde{x}). \end{aligned}$$

By the $\pi_1(S)$ and $\pi_1(S')$ -equivariance of ϕ_S and $\phi_{S'}$ respectively, it follows that $\phi_S(x) = \phi_{S'}(x)$. \square

The proof of Proposition 3.1 then follows easily:

Proof of Proposition 3.1. As before, consider a minimal collection of closed surfaces \mathcal{S} that covers X . Consider any arbitrary $S, S' \in \mathcal{S}$. By Lemmas 3.8 and 3.9, ϕ_S and $\phi_{S'}$ pointwise agree on $S \cap S'$. Since X is complete, convex, and compact, by Lemma 3.6, one can thus patch the collection of isometries $\{\phi_S\}_{S \in \mathcal{S}}$ together to obtain a global isometry $\phi : (X, g_1) \rightarrow (X, g_2)$. \square

4. THE GENERAL CASE

We now tackle the general case of Theorem 1.3. We will begin by showing that simple (not necessarily thick), NPC surface amalgams are NPC cube complexes. Then, using some powerful machinery mentioned in Section 2.4, we deduce that simple, negatively curved surface amalgams are QCERF. Next, we show that simple, thick, negatively curved surface amalgams can be covered by finitely many immersed closed surfaces. Finally, since surface subgroups of $\pi_1(X)$ are quasi-convex, we use Scott's theorem (Theorem 2.12) to promote the immersed closed surfaces to embedded closed surfaces in finite-sheeted covers, allowing us to reduce to the base case.

As promised, we first prove the following:

Lemma 4.1. *Let X be a simple, NPC surface amalgam. Then X is a NPC cube complex.*

Proof. We will realize X as a cube complex and then apply the combinatorial link condition to show X is NPC. First, we assume X is negatively curved. Since the

gluing geodesics of X are simple and disjoint, there exists a pants decomposition \mathcal{P} of X such that the gluing curves are contained in the set of pants curves in \mathcal{P} . We first realize each pair of pants $P \in \mathcal{P}$ as cube complex:




FIGURE 14. A pair of pants as a cube complex. The boundary components are red. The link $lk(v)$ is shown in blue, $lk(u)$ and $lk(w)$ are shown in purple, and $lk(x)$ is in orange.

Note that each boundary component of P contains four vertices. Given a collection of boundary components $\{b_i\}$ identified in X , one can subdivide each b_i so that the vertices are aligned when the b_i are glued together. The result is a cube complex structure for X .

We claim that for every vertex v in the cube complex structure of X , $lk(v)$ is triangle-free. Indeed, we list all the possibilities:

- (1) If v is one of two vertices in each pair of pants adjacent to six edges, $lk(v)$ will be a hexagon (see blue link in Figure 14);
- (2) If v is one of three vertices in each pair of pants adjacent to four edges, then $lk(v)$ will be a square (see orange link in Figure 14);
- (3) If v lies on a pants curve not attached to any other pants curve, then its link is a graph with three vertices and two edges (see purple links in Figure 14);
- (4) Finally, if v lies on a pants curve glued to at least one other pants curve, then $lk(v)$ is a union of squares all sharing two vertices lying on boundary components of the pants decompositions. For instance, if u and w are identified in Figure 14, then the union of the purple links would form a square.

In all four cases, $lk(v)$ is triangle-free, and thus would not span a simplex of dimension greater than 1. Then all the links are flag complexes, as desired.

The generalization to the NPC setting follows easily; one can naturally represent each cylindrical chamber with a square with one pair of opposite sides identified. Subdivide the unidentified sides so that each side has four vertices, including the pair of identified vertices on the corners of the original square. As before, attach boundary components together so that vertices are identified with other vertices, giving the surface amalgam a cube complex structure. Now, the links of each vertex on the cylinders can only fall under case (3) or (4) from the aforementioned list, thus satisfying the combinatorial link condition. \square

Next, we use Lemma 4.1 to show that the fundamental group of a simple, negatively curved surface amalgam is QCERF:

Lemma 4.2. *Let X be a simple, negatively curved surface amalgam. Then $\pi_1(X)$ is QCERF.*

Proof. We collect a few useful facts and observations:

- (1) $\mathbb{F}_n \in \mathcal{QVH}$ for all $n \in \mathbb{N}$. Recall that by Theorem 2.15, a torsion-free Gromov hyperbolic group is in \mathcal{QVH} if and only if it is virtually special. From this, we can naturally deduce that any finite rank free group, which is subgroup separable by a classic theorem by Hall [23], is QCERF and therefore virtually special, and in \mathcal{QVH} .
- (2) *Infinite cyclic subgroups of $\pi_1(X)$ are quasiconvex.* This follows from the fact that if a group G (such as $\pi_1(X)$) is Gromov hyperbolic, every infinite cyclic subgroup of G is quasiconvex.
- (3) *The fundamental group of a surface amalgam is torsion-free.* Due to Proposition 2 of [39], every finite order element of an amalgamated product $A *_C B$ (resp. HNN extension $A *_C$) is conjugate to an element of A or B (resp. A). Since free groups are torsion free, one can inductively show that the fundamental group of any surface amalgam is torsion-free.

By definition of \mathcal{QVH} , (1) and (2), $\pi_1(X) \in \mathcal{QVH}$. By (3), Gromov hyperbolicity of $\pi_1(X)$, and Theorem 2.15, $\pi_1(X)$ is virtually special. Thus, by Lemma 4.1 and Theorem 2.13, it follows that $\pi_1(X)$ is QCERF as well. \square

We briefly mention Corollary 4.2.1, which could be of independent interest:

Corollary 4.2.1. *Let (X, g) be a simple, negatively curved surface amalgam. Then $\pi_1(X)$ is residually finite.*

Finally, we show that every chamber is embedded in an immersed closed surface in X :

Lemma 4.3. *Let (X, g) be a simple, thick, negatively curved surface amalgam. Then each chamber in (X, g) can be included in an embedded closed surface in some finite-sheeted cover $(\widehat{X}, \widehat{g})$ of (X, g) .*

Proof. We claim that each surface amalgam can be covered by a disjoint union of immersed closed surfaces. We first construct this collection of closed surfaces.

Take two copies of each chamber in $C_i \subset X$ to obtain two identical collections of chambers $\{C_i^1\} \cup \{C_i^2\}$. Consider a gluing geodesic $\gamma \subset X$, and suppose that $\{b_k\} \subset X$ is the collection of boundary components identified with γ in X indexed by some collection of natural numbers $k = 1, 2, \dots, N$. We will label the two copies of each b_k with b_k^1 and b_k^2 . For $1 \leq k \leq N$, identify b_k^1 with b_{k+1}^2 , where $k+1$ is taken modulo N . Repeat the same process for every gluing curve in X so that each boundary component of each chamber is glued to exactly one other boundary component. Due to the thickness assumption, the result is a (possibly disconnected) collection of finitely many closed surfaces $\{N_s\}$. An example of the process is illustrated in Figure 15.




FIGURE 16. Given two points $x \in \gamma_1$ and $y \in \gamma_2$, we illustrate examples of hemispheres of unit tangent vectors based at x and y . Here, $S'_x = S'_x C_1 \cup S'_x C_2 \cup S'_x C_4$ and $S'_y = S'_y C_2 \cup S'_y C_3 \cup S'_y C_5$.




FIGURE 15. An immersion of a closed surface into a surface amalgam X using the construction from the proof of Lemma 4.3

We now show that there is an immersion $f : \{N_s\} \looparrowright X$, given by the projection of $\{N_s\}$ back onto X . More specifically, if $y \in \{N_s\} \setminus (\{b_k^1\} \cup \{b_k^2\})$, then it projects to the interior of a chamber in X . Otherwise, y projects to a point on a gluing geodesic in X .

Usually, an immersion between manifolds is defined as a map with a locally injective derivative. An analogous definition holds for surface amalgams, but we need to define the local unit tangent bundle carefully. Let $S'_x C$ be the open hemisphere of unit tangent vectors based at x pointing inside a chamber C containing or adjacent to x . Then define $S'_x := \bigcup_{i=1}^m S'_x C_i$, the collection of the m open hemispheres of unit tangent vectors based at x and pointing into the m chambers adjacent to or containing x (see Figure 16). Note that if $x \in X$ lies on the interior of a chamber, $S'_x(X)$ resembles the usual circle of unit tangent vectors based at a point on a surface.

With this definition in mind, note that for all $y \in \{N_s\}$, $Df_y : T_y(\{N_s\}) \rightarrow S'_{f(y)}(X)$ is injective. We point out that injectivity comes from the fact that there do not exist b_k^1 and b_k^2 which are glued together in $\{N_s\}$. In conclusion, each chamber C sits inside an immersed closed surface in X , which we will call N_C .

The preimage of N_C in the universal cover \tilde{X} under the covering map is a convex set, as it is a union of polygons with geodesic boundary. As a consequence, $\pi_1(N_C)$ is a quasiconvex subgroup of $\pi_1(X)$. Since $\pi_1(X)$ is QCERF, there is some finite-sheeted cover \hat{X} of X such that $N_C \subset \widehat{X_C}$ embeds as a closed surface due to Theorem 2.12. Given a collection of finite-sheeted covers $\{\widehat{X_C}\}$, there is a common finite-sheeted cover \hat{X} such that every chamber C embeds in a closed surface, as the intersection of a finite-index subgroup is still a finite-index subgroup. This proves the lemma. \square

4.1. Proof of Theorem 1.3. We now have all the ingredients needed for proving Theorem 1.3.

Proof. Suppose (X, g_1) and (X, g_2) have the same marked length spectra. Suppose every chamber in X can be included into a closed surface. Then by Proposition 3.1, (X, g_1) and (X, g_2) are isometric via a map isotopic to identity, $\phi : (X, g_1) \rightarrow (X, g_2)$.

Otherwise, by Lemma 4.3, there exists some finite-sheeted cover \hat{X} of X such that every chamber is embedded in \hat{X} . Then by Proposition 3.1, there exists an isometry isotopic to the identity $\hat{\phi} : (\hat{X}, \hat{g}_1) \rightarrow (\hat{X}, \hat{g}_2)$. Recall that $\hat{\phi}$ is constructed by projecting the $\pi_1(\hat{X})$ -equivariant isometry $\tilde{\phi} : (\tilde{X}, \tilde{g}_1) \rightarrow (\tilde{X}, \tilde{g}_2)$ between copies of the universal cover \tilde{X} of both X and \hat{X} . By construction, $\tilde{\phi}$ is also $\pi_1(X)$ -equivariant, so there exists some isometry $\phi : (X, g_1) \rightarrow (X, g_2)$ isotopic to the identity as well, as desired. \square

5. APPENDIX

The main purpose of this appendix is to prove the following statement, which is essential to the proof of Lemma 3.5:

Proposition 5.1. *The geodesic flow map on (Gromov) hyperbolic P -manifolds is ergodic with respect to the Bowen-Margulis measure.*

Notice that for the purposes of the proof of Lemma 3.5, it does not really matter what measure the geodesic flow map is ergodic with respect to. We remark that in the case of symmetric spaces, the Bowen-Margulis measure coincides with the Liouville measure.

The ideas and terminology of this section are mostly taken from [27], which we give a summary of for the convenience of the reader. We remark that we can also deduce ergodicity of the geodesic flow map from general theory in the setting of CAT(0) spaces with rank-one axes developed in [37], but for simplicity, we stick to the ideas in [27].

For the rest of this subsection, let (X, d) be a proper, connected Gromov hyperbolic space under a proper (preimages of compact sets are compact), non-elementary, isometric action by a group Γ and $\partial^\infty(X, d)$ be its visual boundary, the set of equivalence classes of asymptotic geodesic rays (see Section 2.3 for details). Then the (Gromov) hyperbolic compactification \overline{X} of X may be defined as $\overline{X} = X \cup \partial^\infty(X)$.

Following [27], we impose two additional assumptions:

Assumption 5.2 (Uniqueness of geodesics). For all $x_1, x_2 \in \overline{X}$, there exists a unique geodesic $[x_1, x_2]$ joining them.

Assumption 5.3 (Existence of convergent geodesic rays). If α_1 and α_2 are asymptotic (e.g. they lie within bounded distance of each other) geodesic rays, then there exists some $c > 0$ such that:

$$\lim_{t \rightarrow \infty} d(\alpha_1(t), \alpha_2(t + c)) = 0.$$

We remark that both assumptions hold if (X, d) is locally CAT(-1).

5.1. Patterson-Sullivan measures. We now define a family of measures $\{\mu_p\}_{p \in X}$ on $\partial^\infty(X)$ which are defined for every point $p \in X$. Intuitively, such measures, called *conformal densities*, measure the proportion of elements of $\Gamma x = \{\gamma x \mid \gamma \in \Gamma\}$ that land within a specified subset of $\partial^\infty(X)$.

Given a point $a \in \partial^\infty(X)$ and $x, y \in X$, we set:

$$B_a(x, y) = \lim_{t \rightarrow \infty} d(x, r(t)) - d(y, r(t)),$$

where $r : \mathbb{R} \rightarrow X$ is a geodesic ray with endpoint a . Sometimes in the literature, $B_a(x, y)$ is called the *horospherical distance between x and y* , and B_a , which does not depend on the choice of $r(t)$, is sometimes known as the Busmann cocycle function.

Definition 5.4. A family $\{\mu_p\}_{p \in X}$ of finite Borel (Radon) measures on $\partial^\infty(X)$ is called a *conformal density of dimension δ* if:

- (1) For all $\gamma \in \Gamma$, $\gamma_*\mu_p = \mu_{\gamma \cdot p}$ (μ_p is Γ -invariant);
- (2) For all $p, q \in X$ and $a \in \partial^\infty(X)$, μ_p and μ_q are equivalent with Radon-Nikodym derivative

$$\frac{d\mu_q}{d\mu_p}(a) = e^{-\delta B_a(x, y)}.$$

One can define a *quasiconformal density of dimension δ* by relaxing the second condition in Definition 5.4 to say for all $p, q \in X$ and $a \in \partial^\infty(X)$, there exists some $C \geq 1$ such that:

$$\frac{1}{C} e^{-\delta B_a(x, y)} \leq \frac{d\mu_p}{d\mu_q}(a) \leq C e^{-\delta B_a(x, y)}.$$

The *Poincaré series* associated to Γ is the series $P(x, s) = \sum_{\gamma \in \Gamma} e^{s d(x, \gamma \cdot x)}$, where $x \in X$ and $s \in \mathbb{R}$. There is some $\delta_\Gamma \in \mathbb{R}$ aptly named the *critical exponent of Γ* . The series converges if $s < \delta_\Gamma$ and diverges when $s > \delta_\Gamma$; if $s = \delta_\Gamma$, the series could either converge or diverge (see Proposition 5.3 of [11]). Note that if X is a proper geodesic space and X/Γ is compact, then δ_Γ is finite (see Proposition 1.7 of [6]). Furthermore, if the action of Γ on X is non-elementary, then the critical exponent is non-zero. Thus, in the case where X is a simple, thick P-manifold and $\Gamma = \pi_1(X)$, δ_Γ is non-zero and finite.

While the existence of conformal densities is not guaranteed for a Gromov hyperbolic space (X, d) , one can guarantee the existence of quasiconformal densities given that the critical exponent δ_Γ is a finite positive number (as is the case when X is a simple, thick P-manifold), summarized in Theorem 5.5. The *limit set of Γ* ,

the set of accumulation points in $\partial^\infty(X)$ of Γx for some (any) $x \in X$, is denoted by Λ_Γ .

Theorem 5.5 ([11, Theorem 5.4]). *Suppose a group Γ acts properly discontinuously via isometries on (X, d) , a proper (Gromov) hyperbolic metric space (and therefore $\delta_\Gamma \in (0, \infty)$). Then there exists a quasiconformal density of dimension δ_Γ supported on Λ_Γ .*

For every point $x \in X$, Cooernart constructs a *Patterson-Sullivan measure* μ_x from the quasiconformal density of dimension δ_Γ . Its construction depends on whether $P(x, s)$ converges or diverges at $s = \delta_\Gamma$; we refer the reader to the proof of Theorem 5.4 in [11] for details. In particular, if Γ acts convex cocompactly on X , then $P(x, s)$ diverges when $s = \delta_\Gamma$ (see Corollary 7.3 of [11]) and the Patterson-Sullivan measure is defined as the weak limit of probability measures:

$$\mu_x = \lim_{n \rightarrow \infty} \frac{1}{\sum_{\gamma \in B_n} e^{-\delta_\Gamma d(x, \gamma \cdot x)}} \sum_{\gamma \in B_n} e^{-\delta_\Gamma d(x, \gamma \cdot x)} \text{Dirac}_{\gamma \cdot x},$$

where $B_n = \{\gamma \in \Gamma \mid \gamma \cdot o \in B(o, n)\}$ for some fixed $o \in X$. Furthermore, he proves that $\text{supp}(\mu_x) = \Lambda_\Gamma \subseteq \partial^\infty(X)$.

Next, one can use the Patterson-Sullivan measures to define measures on $\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta$.

5.2. Bowen-Margulis measures. Given μ_x with density δ , Kaimanovich defines a Γ -invariant measure on $\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta$ (see Section 2.4.1 in [27]):

$$d\nu_x(a, b) = d\mu_x(a)d\mu_x(b)(e^{\langle a, b \rangle_x})^{2\delta} = \frac{d\mu_x(a)d\mu_x(b)}{(g_{\infty, x}(a, b))^{2\delta}},$$

where $a, b \in \partial^\infty(X)$, $x \in X$ and $g_{\infty, x}$ is the visual metric from Definition 2.7. One can check that due to the scale factor $(e^{\langle a, b \rangle_x})^{2\delta}$, the construction of ν_x is independent of choice of basepoint. We can then denote $\nu = \nu_x$ where x is any arbitrary point in X . We can thus define a Γ -invariant Radon measure on $\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta$, which we will also denote as ν .

While ν is defined on $\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta$, there is a natural extension of ν to a measure m on $SX = (\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta) \times \mathbb{R}$:

$$m = \nu \times dt,$$

where dt is the usual Lebesgue measure on \mathbb{R} . Note that m descends to a measure, m_Γ , on SX/Γ .

Suppose that \tilde{X} is a (Gromov) hyperbolic covering space with Γ its deck group. Unsurprisingly, for a Gromov hyperbolic metric space $X = \tilde{X}/\Gamma$, there is a one-to-one correspondence between *geodesic currents*, Γ -invariant Radon measures on $\partial^\infty(X) \times \partial^\infty(X) \setminus \Delta$, and Radon measures on $SX = S\tilde{X}/\Gamma$ that are invariant under the geodesic flow (see Theorem 2.2 of [27]). Thus, since ν is Γ -invariant, m_Γ is invariant under geodesic flow.

Definition 5.6 (Bowen-Margulis measures). The geodesic flow-invariant measure m_Γ defined above is a *Bowen-Margulis measure* on SX/Γ .

Remark 5.7. Note that $d\nu(a, b) = d\nu(b, a)$, so it follows that ν is invariant under the action of $\mathbb{Z}/2\mathbb{Z}$. The Bowen-Margulis measure thus provides a “natural” example of a geodesic current on X . Its relationship with another naturally-arising current, the

Liouville current, is an interesting open question for spaces that are not symmetric. While we will not define Liouville currents in this paper, we refer the reader to Example 5.4 of [9] for a definition suited to compact quotients of Fuchsian buildings which extends to the surface amalgam setting.

5.3. Proof of Proposition 5.1. We are now ready to state a key theorem from [27]:

Theorem 5.8 ([27, Theorem 2.6]). *Let \tilde{X} be a (Gromov) hyperbolic covering space satisfying Assumptions 5.2 and 5.3. Let μ_x be a Patterson-Sullivan measure on $\partial^\infty(X)$ used to construct the geodesic current ν from before and m_Γ the corresponding Bowen-Margulis measure that is invariant under geodesic flow on SX . Then either:*

- (1) $\mu_x(\Lambda_\Gamma) = 1$ and the geodesic flow on $SX = S\tilde{X}/\Gamma$ is ergodic with respect to m_Γ or
- (2) $\mu_x(\Lambda_\Gamma) = 0$ and the geodesic flow on SX is completely dissipative with respect to m_Γ .

In the case where X is a P-manifold with a locally CAT(-1) metric, recall from Theorem 5.5 that $\text{supp}(\mu_x) = \Lambda_\Gamma$; therefore, the geodesic flow on SX is ergodic with respect to m_Γ , as desired.

5.3.1. An application of Proposition 5.1. One application of the ergodicity of geodesic flow is there exists $v \in SX$ with dense orbit under the geodesic flow map. Indeed, recall the Birkhoff Ergodic Theorem:

Theorem 5.9 (Birkhoff Ergodic Theorem). *Let (Y, \mathcal{B}, μ) be a probability space, and let $T : Y \rightarrow Y$ be an ergodic measure-preserving transformation. Let $A \in \mathcal{B}$ be a measurable set of positive measure $\mu(A) > 0$. Then for all $f \in \mathcal{L}^1(Y, \mathcal{B}, \mu)$ and μ -almost everywhere $y \in Y$:*

$$(7) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(y)) = \int f d\mu.$$

If Γ is a non-elementary discrete group acting properly and cocompactly by isometries on a Gromov hyperbolic space X , then m_Γ is finite (see Corollary 4.17 and Section 3.2 of [12]) and in particular can be scaled to a probability measure, so we can apply the Birkhoff Ergodic Theorem. For example, if X is a locally CAT(-1) P-manifold and $\Gamma = \pi_1(X)$, then we can apply the Birkhoff Ergodic Theorem. By the Hopf-Tsuji-Sullivan Theorem (Theorem 4.2) and Proposition 4.4 of [12], since the geodesic flow map is ergodic with respect to m_Γ , the flow map is conservative, so m_Γ has full support on SX/Γ .

In particular, if f is the indicator function χ_A , $\int f d\mu = \mu(A)$ while the left hand side of Equation (7) denotes the frequency that the orbit of T visits A . In our scenario, T^k is the geodesic flow map ϕ_k on SX , which is invariant with respect to the scaled probability measure of m_Γ with full support on SX . We thus conclude for every simple, thick locally CAT(-1) P-manifold, there exists a geodesic in SX that intersects any open neighborhood of any $u \in SX$ (where X satisfies the conditions of Theorem 5.9), as every open set of SX has non-zero measure.

ACKNOWLEDGMENTS

First and foremost, I am deeply indebted to Caglar Uyanik for many helpful and illuminating discussions concerning the project, and for his edits and comments on previous drafts of this paper. Furthermore, I would like to thank Chris Leininger for his sharp and invaluable insight, and for contributing ideas towards the general case of the argument. I would also like to thank Tullia Dymarz for many useful conversations and insights, and help with proof checking. Many thanks also to Jean-Francois Lafont for generously suggesting the problem as a sequel to his work on P-manifolds, and for very helpful discussions during a visit to the Ohio State University. Thanks also goes to David Constantine for discussing the project with me in depth. Thanks goes to Alex Hof for suggesting the idea behind the proof of Lemma 3.7. Additionally, I sincerely thank the anonymous referee for carefully reading and proofreading the paper and pointing out a gap in my original argument.

REFERENCES

- [1] Roger Alperin and Hyman Bass, *Length functions of group actions on Λ -trees*, Combinatorial group theory and topology (Alta, Utah, 1984), Ann. of Math. Stud., vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 265–378. MR895622
- [2] Anja Bankovic and Christopher J. Leininger, *Marked-length-spectral rigidity for flat metrics*, Trans. Amer. Math. Soc. **370** (2018), no. 3, 1867–1884, DOI 10.1090/tran/7005. MR3739194
- [3] M. Bestvina and M. Feighn, *A combination theorem for negatively curved groups*, J. Differential Geom. **35** (1992), no. 1, 85–101. MR1152226
- [4] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186, DOI 10.1007/BF02392898. MR1638764
- [5] Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 319, Springer-Verlag, Berlin, 1999, DOI 10.1007/978-3-662-12494-9. MR1744486
- [6] M. Burger and S. Mozes, *CAT(-1)-spaces, divergence groups and their commensurators*, J. Amer. Math. Soc. **9** (1996), no. 1, 57–93, DOI 10.1090/S0894-0347-96-00196-8. MR1325797
- [7] K. Burns and A. Katok, *Manifolds with nonpositive curvature*, Ergodic Theory Dynam. Systems **5** (1985), no. 2, 307–317, DOI 10.1017/S0143385700002935. MR796758
- [8] David Constantine, *Marked length spectrum rigidity in non-positive curvature with singularities*, Indiana Univ. Math. J. **67** (2018), no. 6, 2337–2361, DOI 10.1512/iumj.2018.67.7545. MR3900371
- [9] David Constantine and Jean-François Lafont, *Marked length rigidity for Fuchsian buildings*, Ergodic Theory Dynam. Systems **39** (2019), no. 12, 3262–3291, DOI 10.1017/etds.2018.12. MR4027549
- [10] David Constantine and Jean-François Lafont, *Marked length rigidity for one-dimensional spaces*, J. Topol. Anal. **11** (2019), no. 3, 585–621, DOI 10.1142/S1793525319500250. MR3999514
- [11] Michel Coornaert, *Mesures de Patterson-Sullivan sur le bord d’un espace hyperbolique au sens de Gromov* (French, with French summary), Pacific J. Math. **159** (1993), no. 2, 241–270. MR1214072
- [12] Rémi Coulon, Rhiannon Dougall, Barbara Schapira, and Samuel Tapie, *Twisted Patterson-Sullivan measures and applications to amenability and coverings*, E-prints, arXiv:1809.10881, 2018.
- [13] C. Croke, A. Fathi, and J. Feldman, *The marked length-spectrum of a surface of nonpositive curvature*, Topology **31** (1992), no. 4, 847–855, DOI 10.1016/0040-9383(92)90013-8. MR1191384
- [14] Christopher B. Croke, *Rigidity for surfaces of nonpositive curvature*, Comment. Math. Helv. **65** (1990), no. 1, 150–169, DOI 10.1007/BF02566599. MR1036134
- [15] Marc Culler and John W. Morgan, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc. (3) **55** (1987), no. 3, 571–604, DOI 10.1112/plms/s3-55.3.571. MR907233

- [16] Pallavi Dani, Emily Stark, and Anne Thomas, *Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups*, Groups Geom. Dyn. **12** (2018), no. 4, 1273–1341, DOI 10.4171/GGD/469. MR3874643
- [17] Moon Duchin, Christopher J. Leininger, and Kasra Rafi, *Length spectra and degeneration of flat metrics*, Invent. Math. **182** (2010), no. 2, 231–277, DOI 10.1007/s00222-010-0262-y. MR2729268
- [18] Robert Fricke and Felix Klein, *Vorlesungen über die Theorie der automorphen Funktionen. Band 1: Die gruppentheoretischen Grundlagen. Band II: Die funktionentheoretischen Ausführungen und die Anwendungen* (German), Bibliotheca Mathematica Teubneriana, Bande 3, vol. 4, Johnson Reprint Corp., New York; B. G. Teubner Verlagsgesellschaft, Stuttgart, 1965. MR183872
- [19] Rita Gitik, *Graphs and separability properties of groups*, J. Algebra **188** (1997), no. 1, 125–143, DOI 10.1006/jabr.1996.6847. MR1432351
- [20] Colin Guillarmou and Thibault Lefeuvre, *The marked length spectrum of Anosov manifolds*, Ann. of Math. (2) **190** (2019), no. 1, 321–344, DOI 10.4007/annals.2019.190.1.6. MR3990606
- [21] Colin Guillarmou and Marco Mazzucchelli, *Marked boundary rigidity for surfaces*, Ergodic Theory Dynam. Systems **38** (2018), no. 4, 1459–1478, DOI 10.1017/etds.2016.94. MR3789172
- [22] Frédéric Haglund and Daniel T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), no. 5, 1551–1620, DOI 10.1007/s00039-007-0629-4. MR2377497
- [23] Marshall Hall Jr., *A topology for free groups and related groups*, Ann. of Math. (2) **52** (1950), 127–139, DOI 10.2307/1969513. MR36767
- [24] U. Hamenstädt, *Cocycles, symplectic structures and intersection*, Geom. Funct. Anal. **9** (1999), no. 1, 90–140, DOI 10.1007/s000390050082. MR1675892
- [25] Sa’ar Hersonsky and Frédéric Paulin, *On the rigidity of discrete isometry groups of negatively curved spaces*, Comment. Math. Helv. **72** (1997), no. 3, 349–388, DOI 10.1007/s000140050022. MR1476054
- [26] G. Christopher Hruska, Emily Stark, and Hung Cong Tran, *Surface group amalgams that (don’t) act on 3-manifolds*, Amer. J. Math. **142** (2020), no. 3, 885–921, DOI 10.1353/ajm.2020.0020. MR4101334
- [27] Vadim A. Kaimanovich, *Ergodicity of harmonic invariant measures for the geodesic flow on hyperbolic spaces*, J. Reine Angew. Math. **455** (1994), 57–103, DOI 10.1515/crll.1994.455.57. MR1293874
- [28] Ilya Kapovich and Nadia Benakli, *Boundaries of hyperbolic groups*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 39–93, DOI 10.1090/conm/296/05068. MR1921706
- [29] Michael Kapovich and Bruce Kleiner, *Hyperbolic groups with low-dimensional boundary* (English, with English and French summaries), Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 5, 647–669, DOI 10.1016/S0012-9593(00)01049-1. MR1834498
- [30] Jean-François Lafont, *Strong Jordan separation and applications to rigidity*, J. London Math. Soc. (2) **73** (2006), no. 3, 681–700, DOI 10.1112/S0024610706022745. MR2241974
- [31] Jean-François Lafont, *Diagram rigidity for geometric amalgamations of free groups*, J. Pure Appl. Algebra **209** (2007), no. 3, 771–780, DOI 10.1016/j.jpaa.2006.07.004. MR2298855
- [32] Jean-François Lafont, *Rigidity of hyperbolic P-manifolds: a survey*, Geom. Dedicata **124** (2007), 143–152, DOI 10.1007/s10711-007-9136-x. MR2318542
- [33] Jean-François Rajesh Lafont, *Rigidity results for singular spaces*, ProQuest LLC, Ann Arbor, MI, 2002. Thesis (Ph.D.)–University of Michigan. MR2703583
- [34] William Malone, *Topics in geometric group theory*, ProQuest LLC, Ann Arbor, MI, 2010. Thesis (Ph.D.)–The University of Utah. MR2941275
- [35] Jean-Pierre Otal, *Le spectre marqué des longueurs des surfaces à courbure négative* (French), Ann. of Math. (2) **131** (1990), no. 1, 151–162, DOI 10.2307/1971511. MR1038361
- [36] Jean-Pierre Otal, *Sur la géométrie symplectique de l’espace des géodésiques d’une variété à courbure négative* (French), Rev. Mat. Iberoamericana **8** (1992), no. 3, 441–456, DOI 10.4171/RMI/130. MR1202417
- [37] Russell Ricks, *Flat strips, Bowen-Margulis measures, and mixing of the geodesic flow for rank one CAT(0) spaces*, Ergodic Theory Dynam. Systems **37** (2017), no. 3, 939–970, DOI 10.1017/etds.2015.78. MR3628926
- [38] Peter Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565, DOI 10.1112/jlms/s2-17.3.555. MR494062

- [39] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell; Corrected 2nd printing of the 1980 English translation. MR1954121
- [40] Daniel T. Wise, *The structure of groups with a quasiconvex hierarchy*, Annals of Mathematics Studies, vol. 209, Princeton University Press, Princeton, NJ, [2021] ©2021. MR4298722

DEPARTMENT OF MATHEMATICS, RICE UNIVERSITY, HOUSTON, TEXAS
Email address: `yandi.wu@rice.edu`