Rigidity Phenomena of Surface Amalgams

by
Yandi Wu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)

at the
University of Wisconsin-Madison
2024

Date of Final Oral Exam: 3/15/2024
The dissertation is approved by the following members of the Final Oral Committee:

Tullia Dymarz, Associate Professor, Mathematics
Autumn Kent, Professor, Mathematics
Marissa Loving, Assistant Professor, Mathematics
Caglar Uyanik, Assistant Professor, Mathematics

Rigidity Phenomena of Surface Amalgams

Yandi Wu

Abstract

Mostow's famous rigidity theorem states that two closed hyperbolic n-manifolds ($n \geq 3$) are isometric if and only if they have the same fundamental groups. Two alternative forms of rigidity hold for surfaces, a setting where Mostow Rigidity does not hold: *topological* and *marked length spectrum rigidity.* Our main objects of study are *surface amalgams*, natural generalizations of surfaces with links to geometric group theory, especially in the study of Right-Angled Coxeter Groups. In this thesis, we explore topological and marked length spectrum rigidity for surface amalgams and their compact quotients.

Dedication

To all my friends who have kept me company throughout my PhD.

Declaration

I declare that the contents of this thesis reflect my own original research, except where otherwise indicated, and that all sources used have been properly cited. I further declare that I have not submitted the work contained herein as a thesis or similar document to any academic institution other than the University of Wisconsin-Madison, or in pursuit of any other degree.

Acknowledgements

This thesis (and PhD) could not have been possible without the support, encouragement, and help from my wonderful advisors, to whom I owe everything: T. Dymarz for always brightening up my day with her positive energy, and her expertise in spite of having research interests very different from my own, and C. Uyanik for being not only a valuable source of mathematical knowledge, but also a trusted confidante, cheerleader, member of my support system, and a shining example of what it means to “pay it forward.” Thanks for all the coffee, lunches, dinners, halves of veggie platters at Estacion Inka, cat photos, jokes, gossip, and fun math discussions! My advisors are everything I could ever hope to be and more; thank you for being the perfect role models.

I'd also like to thank the rest of the “dynamics” group at Madison for creating a welcoming environment in which I could thrive. I especially thank my other committee members, M. Loving for her support and many helpful math discussions and A. Kent for teaching me hyperbolic geometry during my early years. Outside of Madison, I would like to thank J.F. Lafont for introducing me to both marked length spectrum rigidity and surface amalgams, and suggesting one of the problems explored in this thesis. Additionally, I would like to thank E. Stark for inspiring my work on topological rigidity of surface amalgams, and for helpful discussions. I also thank C. Leininger for his time, expertise, and kindness, and for his feedback on some of the contents of this thesis. I am so lucky to have found mentors outside my home institution who have been so kind and generous with both their time and ideas. And many thanks to my mentors from my undergraduate years who nurtured within me a passion for math. Special mention goes to B. Strenner and C. Abbott, who inspired me to apply to the math PhD program at UW Madison, where I have experienced some of the happiest years of my life.

I am immensely grateful to my loving, supportive, inspiring, and brilliant partner P. Ford, who has been the best partner I could ever ask for during these stressful graduate school years. You inspire me to work hard with your own work ethic, give me the strength and courage to follow my dreams as you follow yours, and buy me bubble tea and yummy food in times both happy and sad. I am so thankful for the life we have built together in Madison; my experience here would not have been complete without you.

I would also like to thank my parents for providing me with the resources to get to where I am now. Although they have not always agreed with my decisions, I am grateful they have worked to accept and come to terms with them, which I acknowledge is no easy task for a wayward free spirit like myself. And of course Lilia and Michael, for being my second set of parents and always offering me life advice (even though I rarely listen), and Vesper, for always cheering me up and bringing a smile to my face.

I also gratefully acknowledge my friends, at all stages of my PhD, for making graduate school so much fun! I am so lucky to have met people who have made it so hard for me to say goodbye to such a cold, snowy place. I cannot possibly name all the people who have made an impact on me throughout the years, but I would especially like to thank A. Byars, Y. Han, A. Hof, J. Johnson, and K. Srivastava.

Thank you to all the wonderful coffee shops in Madison that have given me the caffeine to generate this thesis. Special mention goes to Crescendo on Hilldale (iced blueberry matchas and that one barista who looks like Aubrey Plaza forever), Sencha (it's a love-hate relationship), Lakeside Street Coffee House (for the views) and Cargo on Johnson and Barriques on West Washington (for both always having empty tables next to outlets).

Finally, I gratefully acknowledge support from an NSF RTG grant DMS-2230900. I am also indebted to support from my advisors' personal grants and an Elizabeth Hirschfelder Award.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

Geometric rigidity theory aims to determine a geometric object with the smallest amount of data possible. Current work in geometric rigidity is largely inspired by Mostow's famous rigidity theorem:

Theorem 1.1.1 (Mostow, [Mos68]). *Two hyperbolic (constant negative curvature Riemannian) closed manifolds of dimension* ≥ 3 *are isometric if and only if their fundamental groups are isomorphic.*

Mostow Rigidity does not hold for manifolds with variable negative curvature or surfaces, which come equipped with a Teichmüller's space worth of hyperbolic metrics (see [FM12] Chp. 10). Proving rigidity statements in these settings necessitates formulations of weaker versions of rigidity, such as *topological* and *marked length spectrum rigidity*, which are the main topics of discussion in this dissertation.

A collection of topological spaces $\mathcal{X}$ is *topologically rigid* if for any $X_1, X_2 \in \mathcal{X}$, if $\pi_1(X_1) \cong \pi_1(X_2)$, then X_1 and X_2 are homeomorphic (rather than isometric). It is classical that closed surfaces are topologically rigid since a closed surface is determined completely, up to homeomorphism, by its genus. A collection of negatively curved metric spaces $\mathcal{M}$ is *marked length spectrum rigid (MLSR)* if whenever $\pi_1(M_1) \cong \pi_1(M_2)$ *and*

$M_1, M_2 \in \mathcal{M}$ have the same length assignments to closed geodesics, M_1 and M_2 are isometric (see Chapter 4 for precise definitions). Marked length spectrum rigidity (MLSR) has attracted considerable interest in the recent decades due to a longstanding conjecture of Burns and Katok:

Conjecture 1.1.2 (MLSR Conjecture, [BK85])**.** Closed, negatively curved Riemannian manifolds are marked length spectrum rigid.

Croke and Otal independently proved in their seminal papers (see [Cro90] and [Ota90]) the MLSR conjecture for surfaces. Their results have inspired a wide range of generalizations (see the Introduction section of Chapter 3 for a comprehensive list). In this thesis, we will provide another generalization of their result by showing marked length spectrum rigidity for *2-dimensional P-manifolds.*

Roughly speaking, n-dimensional P-manifolds are constructed by gluing together a finite collection of compact n-manifolds with boundary along their boundary components. Hence, an n-dimensional P-manifold can be seen as a natural generalization of an n-manifold, which can be thought of as a P-manifold consisting of two n-manifolds with boundary glued together. In a series of papers (see [Laf04], [Laf07b], and [Laf07a]), Lafont proves topological rigidity for *simple, thick, hyperbolic* n-dimensional P-manifolds for all $n \geq 2$ (see 2 for a precise definition). Lafont also proves Mostow Rigidity for (simple, thick) 3-dimensional P-manifolds equipped with piecewise hyperbolic metrics, which suggests that in many ways such objects behave like n-manifolds despite displaying much greater topological and geometric complexity.

In this dissertation, we focus on the $n = 2$ case, and use *surface amalgam* in place of "2-dimensional P-manifold". Our terminology comes from the fact that the fundamental groups of such objects arise as cyclic amalgams of surface groups. While simple, thick hyperbolic surface amalgams are topologically rigid by [Laf07b], motivated by the open problem of the abstract commensurability classification of Right-angled Coxeter Groups, Stark shows in [Sta18] that their compact quotients are not. In particular, Stark finds a pair of nonhomeomorphic but homotopic *Davis orbicomplexes* (see Chapter chapter 3 for a precise definition). We generalize Stark's results to a broader class of Davis orbicomplexes by finding infinitely many pairs of nonhomeorphic but homotopic Davis orbicomplexes and explore the extent to which topological rigidity is preserved under quotienting. We also show that similar to the case of negatively-curved surfaces, simple, thick surface amalgams equipped with negatively-curved metrics is marked length spectrum rigid.

1.2 Main Results

In geometric group theory, *Right-angled Coxeter Groups (RACGs)* are a popular object of study. This is in part due to the fact that despite their relatively elementary definition, they exhibit a surprising variety of geometric phenomena which make them a rich source of examples for developing intuition for geometric properties and conjectures (see [Dan20]). See Chapter 3 for precise definitions of RACGs.

Recall that two finitely-generated groups G_1 and G_2 are *abstractly commensurable* if they have isomorphic finite-index subgroups. The abstract commensurability classification of RACGs is a big open problem in geometric group theory.

Dani, Stark, and Thomas demonstrate in [DST18] an interesting property of surface amalgams: they may arise as finite-sheeted covers of *Davis orbicomplexes*, which are well-behaved K(G,1) spaces of RACGs (see the Preliminaries section of Chapter 3 for details). They use this fact to provide a partial abstract commensurability classification of RACGs, exploiting the topological rigidity property of surface amalgams. There are a few obstructions to extending the results of [DST18]. First, it is unknown which Davis orbicomplexes have finite-sheeted covers that are surface amalgams. One solution to this problem is to extend commensurability results by using topologically rigid collections of Davis orbicomplexes, since every Davis orbicomplex will have finite-sheeted covers that are themselves Davis orbicomplexes. However, Stark proved Davis orbicomplexes are in general *not* topologically rigid in [Sta18] by constructing an example of two homotopic Davis orbicomplexes that are not homeomorphic. The first half of this thesis partially answers the following question posed in [Sta18]:

Question 1.2.1. *For which set W of Coxeter groups is the set of associated Davis orbicomplexes D_Γ together with their finite-sheeted covers topologically rigid?*

The results of [Sta18] suggest that topological rigidity is in fact very hard to achieve for families of Davis orbicomplexes. We further prove this in [Wu23b]. Even so, we also prove the following (see [Wu23b] and Chapter 3):

Theorem 1.2.2. *There exists an infinite class $\mathcal{C}$ of Davis orbicomplexes such that for any $D_\Gamma \in \mathcal{C}$, the collection of finite-sheeted covers of D_Γ forms a topologically rigid set.*

In the second half of this thesis, we shift our attention from topological rigidity to marked length spectrum rigidity. We focus our attention on surface amalgams with piecewise negatively-curved metrics. We prove the following, given that the metrics satisfy certain smoothness conditions (see Chapter 4 for details):

Theorem 1.2.3. *The collection of simple, thick, negatively-curved surface amalgams (see Definition 2.0.1) is marked length spectrum rigid.*

We point out that the proof of Theorem 1.2.3 differs from the well-explored case of surfaces (e.g. [Cro90], [Ota90]) because the universal cover of a surface amalgam is no longer simply a disk but an infinite collection of half-disks glued together in a tree-like fashion. This is a major obstruction to proving marked length spectrum rigidity using the classical methods of Otal from [Ota90]. The MLSR Conjecture (Conjecture 1.1.2) remains unsolved in large part because the boundary of the universal cover of a higher dimensional manifold behaves very differently from that of a surface. Thus, proving MLSR for geometric objects, such as surface amalgams, with more complex boundaries could help develop necessary tools with which to attack the conjecture. In fact, surface amalgams share similarities with 3-manifolds (see [HST20]), a setting where Conjecture 1.1.2 remains open.

1.3 Outline

We now give an outline for this dissertation.

In Section 2, we introduce surface amalgams. Other definitions will be introduced in the preliminaries sections of the chapters where they are needed.

Section 3, which is based off the work in [Wu23b], contains the proof of Theorem 1.2.2. We also identify key combinatorial obstructions to topological rigidity for Davis orbicomplexes. These obstructions are encoded in the defining graphs of the associated Right-angled Coxeter Groups.

Section 4, which details the proof of Theorem 1.2.3, is based off the work in [Wu23a]. While the proof uses many of the same tools as those used in [Ota90], it also relies on standard techniques in geometric group theory, such as exploiting the coarse geometry of the universal cover, doubling, and passing to finite-sheeted covers. We also modify tools from Otal's original paper, namely intersection numbers and Liouville currents, to better suit the fractal-like nature of the boundaries of the universal covers.

Chapter 2

Surface Amalgams

We now precisely define a particularly well-behaved family of surface amalgams, which will be the focus of this dissertation. First, though, we recall the precise definition of a surface amalgam. Roughly speaking, a surface amalgam consists of a finite collection of compact surfaces with boundary glued together along their boundary components. The following definition is adapted from Definition 2.3 of [Laf07b].

Definition 2.0.1 (Negatively curved surface amalgam)**.** A compact metric space X is a *negatively curved surface amalgam* 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

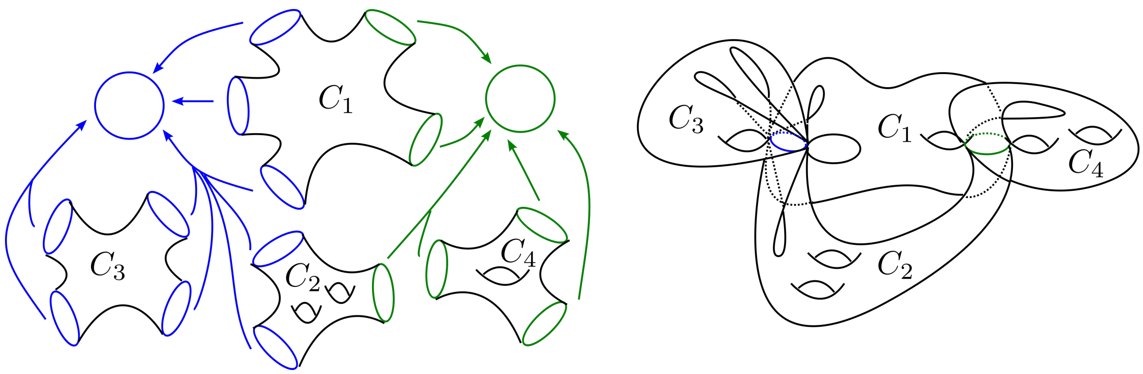

Figure 2.1: An example of a simple, thick surface amalgam with four chambers.

at least three distinct boundary components of chambers, then we say X is *thick* (note this definition differs slightly from the one given in [Laf07b] but follows the one given in [Laf07a] and [Laf06]). Like Lafont in [Laf07b], we will only be considering simple, thick surface amalgams. Doing so ensures X is CAT(-1) and guarantees some useful properties, such as the one pointed out in Lemma 4.2.9.

Throughout this thesis, we make heavy use of the following topological rigidity result by Lafont, which we mentioned in the introduction:

Theorem 2.0.2 ([Laf07b], Theorem 1.2)**.** *Let X_1, X_2 be a pair of simple, thick surface amalgams, and assume that $\phi : \pi_1(X_1) \to \pi_1(X_2)$ is an isomorphism. Then there exists a homeomorphism $\phi : X_1 \to X_2$ that induces ϕ on the level of the fundamental groups.*

Chapter 3

Topological Rigidity

3.1 Introduction

Often, determining whether two topological objects are homeomorphic is a significantly harder problem than determining whether their fundamental groups are isomorphic. In some cases, however, if we impose enough conditions on the topological spaces we are studying, the weaker equivalence relation (isomorphism between fundamental groups) implies the stronger and often more useful equivalence relation (homeomorphism between the topological objects). We can often exploit the topological rigidity of such sets of spaces to derive useful results (recall that a collection of topological objects $\mathcal{X}$ is *topologically rigid* if for any $X_1, X_2 \in \mathcal{X}$, if $\pi_1(X_1) \cong \pi_1(X_2)$, then X_1 and X_2 are homeomorphic). For example, to determine if two groups G_1 and G_2 are *abstractly commensurable* (i.e. have isomorphic finite-index subgroups), we often construct two finite-sheeted homeomorphic covers of X_1 and X_2, where $\pi_1(X_i) = G_i$ for $i = 1, 2$. These homeomorphic covers can be hard to construct. However, if the finite-sheeted covers $\tilde{X}_1$ and $\tilde{X}_2$ belong to a topologically rigid class of spaces and $\pi_1(\tilde{X}_1) \cong \pi_1(\tilde{X}_2)$, then we know $\tilde{X}_1$ and $\tilde{X}_2$ are homeomorphic.

There are several well-established examples of topologically rigid classes. For example, the set of closed orientable surfaces is topologically rigid. The Poincare Conjecture implies the set of simply-connected, closed 3-manifolds is topologically rigid. Recall that

Lafont proves the set of of simple, thick n-dimensional hyperbolic P-manifolds, a subclass of piecewise CAT(-1) spaces, is topologically rigid for $n \geq 2$ (see [Laf04], [Laf07b], and [Laf07a]). In this chapter, we consider certain *orbicomplexes*, unions of collections of orbifolds identified along homeomorphic suborbifolds, associated with Right-Angled Coxeter Groups (RACGs), defined below.

Definition 3.1.1 (Right-Angled Coxeter Group)**.** Given a finite simplicial graph Γ with edge set E and vertex set V, the *Right-Angled Coxeter Group (RACG)* W_Γ with defining graph Γ is the group with presentation $\langle v_i \in V : v_i^2 = 1, [v_i, v_j] = 1 \text{ if } [v_i, v_j] \in E \rangle$.

A RACG W_Γ acts properly discontinuously by isometeries on a space called the *Davis Complex* Σ_Γ. The quotient $\mathcal{D}_\Gamma = \Sigma_\Gamma / W_\Gamma$, which we call a *Davis orbicomplex*, is one of the aforementioned orbicomplexes and comes equipped with cell stabilizer data defined by the action of W_Γ on Σ_Γ. To clarify, recall that if an amalgamated free product or HNN extension G acts on a Basse Serre tree T, the resulting quotient T/G is a graph of groups whose vertices and edges are labelled by subgroups of G isomorphic to vertex and edge stabilizers of T. Similarly, each edge and vertex of a Davis orbicomplex $\mathcal{D}_\Gamma$ may be labeled by a subgroup of W_Γ that stabilizes a lift of the edge or vertex in Σ_Γ (in this paper, we do not specify such labels as they are not crucial for our proofs). For further background on the Davis complex and Coxeter groups, refer to [Dav08]. The Davis orbicomplex has been studied extensively by Stark, who posed Question 1.2.1 in [Sta18], which we will give a partial answer to in this chapter.

Despite the simplicity of the problem statement, the answer to Question 1.2.1 is very nuanced. In this paper, we focus our attention on RACGs that are one-ended (Γ has no separating edges or vertices and is connected) and hyperbolic (Γ is square-free, or has no cycles of length four). One example of a class of defining graphs that gives rise to W_Γ satisfying these conditions is a subclass of *generalized* Θ *graphs*, defined as follows:

Definition 3.1.2 (Generalized Θ-graph)**.** For $k \geq 1$, $0 \leq n_1 \leq ... \leq n_k$, let $\Theta = \Theta(n_1, n_2, ..., n_k)$ be the graph with two vertices a and b, each of valence k, and k edges

$e_1, e_2, ..., e_k$ connecting them, which we will call the branches of Θ. Furthermore, for $1 \leq i \leq k$, e_i is subdivided into $n_i + 1$ edges by inserting n_i new vertices.

For the purposes of this paper, we will require that $n_i > 0$ for all $1 \leq i \leq k$ and $n_2 > 1$ in order to ensure W_Γ is hyperbolic.

Associated to each generalized Θ-graph is an Euler characteristic vector, which captures the Euler characteristics of the orbifolds in the Davis orbicomplex $\mathcal{D}_\Gamma$. The Euler characteristic vector is often used to classify Davis orbicomplexes; in [DST18], the Euler characteristic vector is used to list abstract commensurability criteria. In this paper, we will use Euler characteristic vectors to list criteria for topological rigidity.

Definition 3.1.3 (Euler characteristic vectors of generalized Θ-graphs)**.** Let $\Theta = \Theta(n_1, n_2, ..., n_k)$ be a generalized Θ graph. Then the *Euler characteristic vector* of Θ is the vector $v = (x_1, x_2, ..., x_n)$, where $x_i = \frac{1-n_i}{4}$. Two Euler characteristic vectors v_1 and v_2 are said to be *commensurable* if there exist $K, L \in \mathbb{Z}_{\neq 0}$ such that $Kv = Lw$.

Dani, Stark, and Thomas show in Theorem 5.2 of [DST18] that finite covers of Davis orbicomplexes with $\Gamma = \Theta(n_1, n_2, ..., n_k)$ are topologically rigid. In this paper, we focus on cycles of generalized Θ graphs introduced in [DST18], which consist of generalized Θ graphs identified along their essential vertices.

Definition 3.1.4 (Cycle of generalized Θ-graphs)**.** Let $N \geq 3$ and let $b_1, b_2, ..., b_N$ be positive integers so that for each i, $1 \leq i \leq N$, at most one of b_i and b_{i+1} where i is taken mod N is equal to 1. Let Θ_i be a generalized Θ graph with b_i edges between two vertices a_i and c_i. We can construct a cycle of N generalized Θ-graphs Γ by identifying c_i with a_{i+1}.

We call the vertex of a cycle of generalized Θ graphs with valence greater than two an *essential vertex*. For the rest of the paper, we will use $\{v_i\}_{i=1}^N$ to denote the set of essential vertices of all the graphs involved. The indices of all v_i's will also taken mod N, where N is the number of essential vertices (or equivalently generalized Θ graphs) in a cycle of generalized Θ-graphs Γ.

Let Γ be a cycle of generalized Θ graphs with Davis complex Σ_Γ, and G a finite index, torsion-free subgroup of W_Γ. Stark proves in [Sta18] that the set of quotients Σ_Γ/G, which correspond to finite-sheeted covers of the Davis orbicomplexes $\mathcal{D}_\Gamma$, is not topologically rigid by constructing $X_1 = \Sigma_\Gamma/G_1$ and $X_2 = \Sigma_\Gamma/G_2$ that are homotopic but not homeomorphic. Theorem 3.1.6 in Section 3.2 generalizes the construction from [Sta18] to create a class of orbicomplexes where topological rigidity fails. Our construction of homotopic but not homeomorphic covers relies on the fact that one set of orbifolds in the orbicomplex is a finite cover of another set of orbifolds.

Definition 3.1.5. Suppose Γ is a cycle of generalized Θ graphs with essential vertices $\{v_i\}_{i=1}^N$, and there exist two essential vertices v_j, v_k such that the generalized Θ graphs Θ_j and Θ_k between v_j and $v_{j+1} \mod N$ and v_k and $v_{k+1} \mod N$ (where $k \neq j$) respectively have commensurable Euler characteristic vectors u and w ($Ku = Lw$ for some $K, L \in \mathbb{Z}_{\neq 0}$). Then we say Γ is *repetitive.* If K or $L = 1$, then we say Γ is *strongly repetitive.*

Theorem 3.1.6. *Suppose a class of finite-sheeted covers of Davis orbicomplexes $\mathcal{X}$ contains all the finite-sheeted covers of some Davis orbicomplex $\mathcal{D}_\Gamma$ where Γ is strongly repetitive. Then $\mathcal{X}$ is not topologically rigid.*

It is not known whether Theorem 3.1.6 is true if we only assume Γ is repetitive.

Theorem 3.1.6 as well as Stark's proof in [Sta18] rely on constructions of finite-sheeted covers of the same Davis orbicomplex $\mathcal{D}_\Gamma$. One can also prove, however, that two finite-sheeted covers of nonhomeomorphic Davis orbicomplexes can also violate topological rigidity.

Definition 3.1.7 (Permuted pairs)**.** Two cycles of generalized Θ graphs Γ_1 and Γ_2 form a *permuted pair* if they are obtained from identifying the essential vertices of the same set of generalized Θ graphs. Equivalently, the set of Euler characteristic vectors of Γ_1 is some permutation of the set of Euler characteristic vectors of Γ_2.

Remark 3.1.8. Note that if we use the definition above, it is possible for a permuted pair Γ_1 and Γ_2 to be isomorphic. For example, if Γ_1 and Γ_2 each consist of three generalized

Θ graphs glued together, they are isomorphic (see the proof of Lemma 3.3.4 for details). We do not consider such pairs in Theorem 3.1.9, stated below.

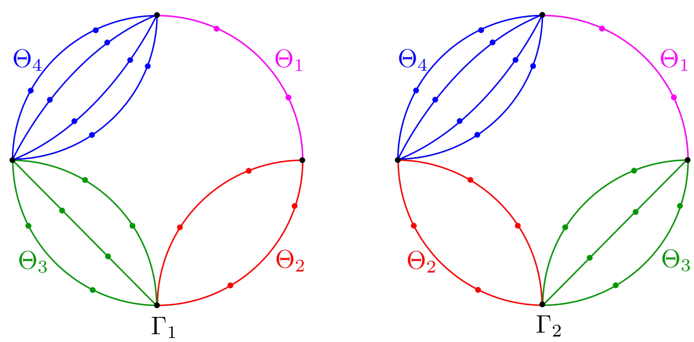

Figure 3.1: An example of a permuted pair. Note that Γ_1 and Γ_2 both consist of Θ_i (where $1 \leq i \leq 4$) glued along essential vertices.

Theorem 3.1.9. *Suppose a class of finite-sheeted covers of Davis orbicomplexes $\mathcal{X}'$ contains all finite sheeted covers of two Davis orbicomplexes $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$, where Γ_1 and Γ_2 form a permuted pair and Γ_1 and Γ_2 are not isomorphic. Then $\mathcal{X}'$ is not topologically rigid.*

In the proof of Theorem 3.1.9, we find two homotopic finite-sheeted covers of $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$ that are not homemomorphic. As a side note, this means that W_{Γ_1} and W_{Γ_2} are commensurable, so having two defining graphs that form a permuted pair is a sufficient condition for commensurability. Recall that in Theorem 1.12 of [DST18], Dani, Stark, and Thomas provide two necessary and sufficient conditions for commensurability of RACGs with defining graphs that are cycles of generalized Θ graphs.

In [Sta18], Stark constructs X_1 and X_2, two homotopic finite covers of a Davis orbicomplex $\mathcal{D}_\Gamma$ with non-homeomorphic *singular sets* (e.g. sets along which the orbifolds are identified). In her example, Γ is a cycle of generalized Θ graphs, proving that the set of finite-sheeted covers of Davis complexes with defining graphs that are cycles of generalized Θ graphs is not topologically rigid. In light of these results, in Section 3 of [DST18], Dani, Stark, and Thomas construct a different set of orbicomplexes that is topologically rigid, which they use to prove abstract commensurability results. Nevertheless, in section

4 (see Theorem 3.4.6), we are able to find a topologically rigid subclass of finite-sheeted covers of Davis orbicomplexes $\mathcal{D}_\Gamma$ where Γ is a cycle of generalized Θ graphs. The subclass also takes Theorems 3.1.6 and 3.1.9 into account to exclude finite-sheeted covers of $\mathcal{D}_\Gamma$ that violate topological rigidity. Although the exact statement of the theorem is fairly technical, we refer the reader to Theorem 1.2.2 for a simplified version.

3.2 Preliminaries

We now introduce a construction of the Davis orbicomplex specific to the setting where the defining graph Γ is a cycle of generalized Θ graphs consisting of $\Theta_i = \Theta(n_{i,1}, n_{i,2}, ..., n_{i,k})$ for $1 \leq i \leq N$. For a more detailed construction of W_Γ and verification that $\pi_1(\mathcal{D}_\Gamma)$ is indeed W_Γ, we refer the reader to Section 2 of [Sta18] and Section 3 of [DST18].

First, we describe how to construct an orbifold $\mathcal{P}_{i,j}$ for a branch (edge) $b_{i,j}$ of a generalized Θ graph Θ_i. For each $b_{i,j}$, construct a $(n_{i,j}+2)$-gon with an edge of order 1, which we call a *nonreflection edge*, and $n_{i,j}+1$ reflection edges of order 2. All the vertices are order 4 vertices, with the exception of the two order 2 vertices adjacent to the non-reflection edge.

Construction 3.2.1 (Davis Orbicomplex $\mathcal{D}_\Gamma$ of a cycle of generalized Θ graphs)**.** First, we will construct an orbifold graph S. The underlying graph of S is a star with one central vertex v_0 adjacent to N valence one vertices. The valence one vertices are orbifold points of order 2. Cyclically label the orbifold points with v_l where $1 \leq l \leq N$, and use e_l to denote the edge $[v_0, v_l] \in E(S)$. Then attach the set of branch orbifolds $\mathcal{P}_{i,j}$ along their non-reflection edges to e_i and e_{i+1}, where the labels are taken mod N. An example of Construction 3.2.1 is shown in Figure 3.2. o

Note that for the cycle of generalized Θ graphs Ψ shown in Figure 3.2, the Euler characteristic vectors of Θ_1 and Θ_3 are $(-\frac{1}{4}, -\frac{1}{4})$ and $(-\frac{3}{4}, -\frac{3}{4})$, so $3w_1 = w_3$, which means Ψ is strongly repetitive. Theorem 1.7 then implies any class $\mathcal{X}$ that contains all finite-sheeted covers of $\mathcal{D}_\Psi$ is not topologically rigid.

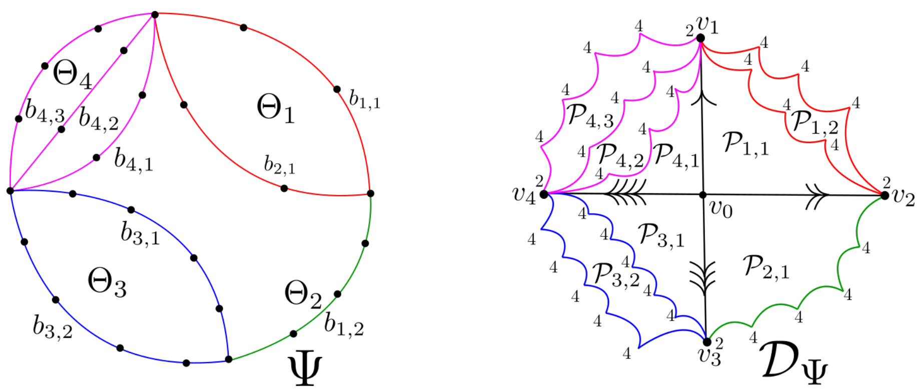

Figure 3.2: A cycle of generalized Θ graphs Ψ along with its Davis orbicomplex $\mathcal{D}_\Psi$. We label edges e_l in the singular star S with l arrows. Note that each branch $b_{i,j} \in \Psi$ determines an orbifold $\mathcal{P}_{i,j} \in \mathcal{D}_\psi$.

All finite-sheeted covers of Davis orbicomplexes that we construct will contain *jester hats*, a specific kind of orbifold defined below:

Definition 3.2.2 (Jester hats)**.** Suppose $\mathcal{O} = D^2\underbrace{(2, 2, ..., 2)}_{n}$, i.e. a disk with n order 2 points. Then we will call $\mathcal{O}$ a jester hat with n (order two) cone points.

3.3 Examples of topologically non-rigid sets

We first introduce the construction of jester hats that cover orbifolds in a Davis orbicomplex $\mathcal{D}_\Gamma$.

Lemma 3.3.1. *Suppose d is a positive even integer. Each orbifold $\mathcal{P}$ with r reflection edges is covered by the jester hat $D^2\underbrace{(2, 2, ..., 2)}_{c}$, where $c = \frac{d}{2}(r-3)+2$ and d is the degree of the cover.*

Proof. This construction is based on Stark's construction in Lemma 3.1 from [Sta18] and Crisp and Paoluzzi's construction from Section 3.1 of [CP08]. Using Crisp and Paoluzzi's construction, we observe that for any even integer $d > 0$, an orbifold $\widehat{\mathcal{O}}$ with $\frac{d}{2}(r-3)+3$ reflection edges is tiled by $\frac{d}{2}$ copies of an orbifold with r reflection edges, so $\widehat{\mathcal{O}}$ is a $\frac{d}{2}$-sheeted orbifold cover of $\mathcal{O}$. For example, in Figure 3.3, an orbifold $\widehat{\mathcal{O}}$ with 6 reflection edges is tiled by 3 copies of $\mathcal{O}$, an orbifold with 4 reflection edges. Thus, $\widehat{\mathcal{O}}$ is a 3-sheeted orbifold cover of $\mathcal{O}$. Next, if we unfold along the reflection edges of $\widehat{\mathcal{O}}$, we obtain a closed disk with $\frac{d}{2}(r-3)+2$ order two cone points, as desired. The construction is illustrated in Figure 3.3.

□

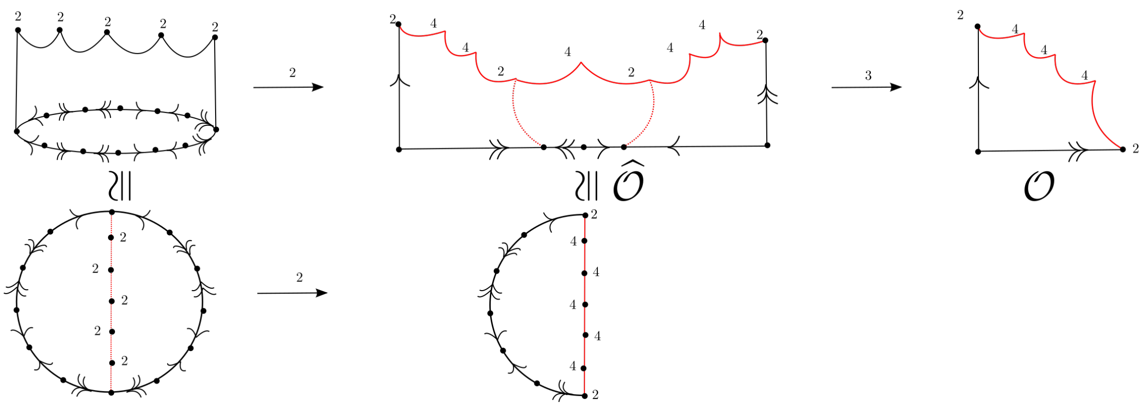

Figure 3.3: A tower of covers illustrating the lemma. Here, $D^2(2,2,2,2,2)$ is a six-fold cover of an orbifold with four reflection edges.

Construction 3.3.2 (The double of a singular set)**.** The Davis orbicomplex $\mathcal{D}_\Gamma$ has a singular subset S consisting of an orbifold star graph with N order two points, where N is the number of essential vertices in Γ. We can construct a double cover of S, which we call $\widehat{S}$, by unfolding along order two points to obtain a subdivided generalized Θ-graph with N branches and one vertex on each branch between the essential vertices. For an illustration, refer to Figure 3.4. All the singular sets constructed in this section will be a finite-sheeted cover of $\widehat{S}$.

For all of the covers described in this paper, all the edges will be subdivided by a copy of the lift of an order two point $\tilde{v}_i$ from the Davis orbicomplex. Both subdivisions will be oriented towards $\tilde{v}_i$ and labeled with the label of the edge, which we will specify in the construction. For simplicity, we will count a subdivided edge as one edge when calculating cycle lengths.

Proposition 3.3.3. *Let Γ be strongly repetitive. Then there exist homotopic but non-homeomorphic finite-sheeted covers of $\mathcal{D}_\Gamma$.*

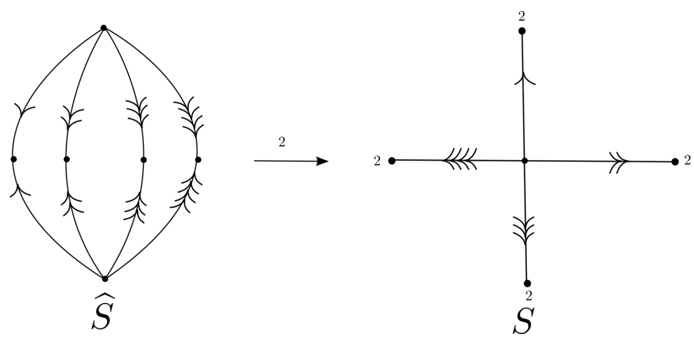

Figure 3.4: A generalized Θ graph with four branches two-fold covers the singular subset of the orbicomplex from Figure 3.2. Here, $N = 4$.

Proof. Suppose u and w are Euler characteristic vectors of Θ_i and Θ_k respectively where $Ku = w$ for some $K \in \mathbb{Z}_+$. Without loss of generality, assume $i = 1$ since we can rotate the labels of the essential vertices otherwise. Suppose Θ_1 and Θ_k have l branches. Let $n_{s,b}$ denote the number of vertices on the bth branch of Θ_s. Then if $u = (\frac{1-n_{1,1}}{4}, \frac{1-n_{1,2}}{4}, ..., \frac{1-n_{1,l}}{4})$ and $w = (\frac{1-n_{k,1}}{4}, \frac{1-n_{k,2}}{4}, ..., \frac{1-n_{k,l}}{4})$ then for $1 \leq b \leq l$, $K(\frac{1-n_{1,b}}{4}) = (\frac{1-n_{k,b}}{4})$ and $K(1 - n_{1,b}) = 1 - n_{k,b}$. If $r_{s,b}$ denotes the number of reflection edges on the orbifold in the Davis orbicomplex constructed from the bth branch of Θ_s, then $r_{s,b} = n_{s,b} + 2$, so we have $K(r_{1,b} - 3) + 3 = r_{k,b}$.

Case 1 (K = 1): We will first consider some general cases before addressing the edge case where $N \neq 3$. First, suppose $K = 1$ and $k < N - 1$. Then $r_{1,b} = r_{k,b}$ for $1 \leq b \leq l$. We will construct two non-homeomorphic but homotopic four-sheeted covers of $\mathcal{D}_\Gamma$, which we will call $\widetilde{X_1}$ and $\widetilde{X_2}$. First, we construct their singular sets $\widetilde{S_1}$ and $\widetilde{S_2}$ with four essential vertices, $\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4$ and $\tilde{v}_1', \tilde{v}_2', \tilde{v}_3', \tilde{v}_4'$. To construct $\widetilde{S_1}$, add k edges between two pairs of vertices: $\tilde{v}_1$ and $\tilde{v}_4$, as well as $\tilde{v}_2$ and $\tilde{v}_3$. The edges will be labeled with all integers between 1 and k and subdivided as described earlier in the section. Between two other pairs of vertices, $\tilde{v}_1$ and $\tilde{v}_2$, as well as $\tilde{v}_3$ and $\tilde{v}_4$, construct $N - k$ subdivided edges labeled with all integers between $k + 1$ and N. Thus, in total, there are two (subdivided) four-cycles labeled with k and $k + 1$ as well as 1 and N, and $N - 2$ cycles of length two labeled with i and $i + 1$, where $1 \leq i \leq N - 1$, $i \neq k$. For the other singular set, $\widetilde{S_2}$, there is one edge labeled with 1 between two pairs of vertices: $\tilde{v}_1'$ and $\tilde{v}_4'$, as well as $\tilde{v}_2'$ and $\tilde{v}_3'$. Additionally, there are $N-1$ edges between two other pairs of vertices, $\tilde{v}_1'$ and $\tilde{v}_2'$ as well as $\tilde{v}_3'$ and $\tilde{v}_4'$; these edges are labeled with all integers between 2 and N. In total, there are again two four-cycles with edges labeled with 1 and 2 as well as 1 and N, and $N-2$ cycles of length two labeled with i and $i+1$ where $2 \leq i \leq N-1$.

By the conditions imposed on $\mathcal{D}_\Gamma$, the four-cycles labeled by k and $k+1$ in $\widetilde{S_1}$ and 1 and 2 in $\widetilde{S_2}$ have the same set of jester hats glued to them. Additionally, the two-cycles labeled with 1 and 2 in $\widetilde{S_1}$ and k and $k+1$ in $\widetilde{S_2}$ have the same set of jester hats glued to them. For other integers I where $1 \leq I \leq N$ and $I \neq 1, k$, for every cycle in $\widetilde{S_1}$ labeled with I and $I+1$, there is a corresponding cycle of the same length labeled with I and $I+1$ in $\widetilde{S_2}$. As a result, for each cycle in $\widetilde{S_1}$ with a set of jester hats glued to it, there is a corresponding cycle in $\widetilde{S_2}$ with the same set of jester hats glued to it. By taking their regular neighborhoods, we can see that both $\widetilde{S_1}$ and $\widetilde{S_2}$ are homotopic to a $(2N-4)$-holed sphere. We can then conclude that $\widetilde{X_1}$ and $\widetilde{X_2}$ are homotopic to the same $(2N-4)$-holed sphere with the same same sets of jester hats glued to their boundary components. However, they are not homeomorphic since the complements of their cut pairs have different numbers of connected components.

If $N \neq 3$, there are two special cases where the above construction does not work. First, note that if $k = N$, the construction does not work because the first cover will be disconnected. Thus, we construct modified non-homeomorphic but homotopic degree four covers. As before, we will call these covers $\widetilde{X_1}$ and $\widetilde{X_2}$ with singular sets $\widetilde{S_1}$ and $\widetilde{S_2}$ that have essential vertices $\tilde{v}_i$ and $\tilde{v_i'}$ respectively, where $1 \leq i \leq 4$. For $\widetilde{S_1}$, construct two subdivided edges labeled with 1 and 2 between $\tilde{v}_1$ and $\tilde{v}_4$ as well as $\tilde{v}_2$ and $\tilde{v}_3$. Then between vertices $\tilde{v}_3$ and $\tilde{v}_4$ as well as $\tilde{v}_1$ and $\tilde{v}_2$, construct $N-2$ edges with labels between 3 and N. For $\widetilde{S_2}$, between $\tilde{v_1'}$ and $\tilde{v_4'}$ as well as $\tilde{v}_2$' and $\tilde{v_3'}$, construct $N-1$ subdivided edges labeled with all integers between 3 and N and one subdivided edge labeled with 1. Then construct a subdivided vertex labelled with 2 between $\tilde{v_3'}$ and $\tilde{v_4'}$ as well as $\tilde{v_1'}$ and $\tilde{v_2'}$.

Second, note that if $k = N-1$, the general construction does not work since $\widetilde{S_1}$ and $\widetilde{S_2}$ are homeomorphic; therefore, $\widetilde{X_1}$ and $\widetilde{X_2}$ are also homeomorphic. In order to fix this, construct new $\widetilde{S_1}$ and $\widetilde{S_2}$ as follows: for $\widetilde{S_1}$, construct 3 edges between $\tilde{v_1}$ and $\tilde{v_2}$ as well as $\tilde{v_3}$ and $\tilde{v_4}$, which are labeled with $1, 2$, and N. Then construct $N-3$ edges between $\widetilde{v_1'}$ and $\widetilde{v_4'}$ as well as $\widetilde{v_2'}$ and $\widetilde{v_3'}$ labeled from 3 to $N-1$. For $\widetilde{S_2}$, follow the construction for the case where $k = N$. Again, in both edge cases ($k = N$, $k = N-1$, $N > 3$), both $\widetilde{S_1}$ and $\widetilde{S_2}$ are homotopic to a $(2N-4)$-holed sphere and have the same sets of jester hats glued to them, but complements of their cut pairs have different numbers of connected components. Thus, $\widetilde{S_2}$ and $\widetilde{S_1}$ are not homeomorphic but $\widetilde{X_1}$ and $\widetilde{X_2}$ are homotopic.

For the case $N = 3$, any of the previous constructions will yield homeomorphic $\widetilde{X_1}$ and $\widetilde{X_2}$, so a different pair of covers is necessary. For this special case, we will construct 16-sheeted covers. Assume without loss of generality that Θ_1 and Θ_3 have the same Euler characteristic vectors. For both $\widetilde{S_1}$ and $\widetilde{S_2}$, label edges $[v_i, v_{i+1}]$ and $[v_i', v_{i+1}']$ respectively with 2 if i is odd and 3 if i is even. Then construct the edges $[\widetilde{v_2}, \widetilde{v_{13}}]$, $[\widetilde{v_3}, \widetilde{v_{12}}]$, $[\widetilde{v_4}, \widetilde{v_{11}}]$, $[\widetilde{v_1}, \widetilde{v_{16}}]$, $[\widetilde{v_{14}}, \widetilde{v_{15}}]$, $[\widetilde{v_9}, \widetilde{v_{10}}]$, $\widetilde{v_7}, \widetilde{v_8}]$, and $[\widetilde{v_5}, \widetilde{v_6}]$ labelled with 1. For $\widetilde{S_2}$, construct edges $[\widetilde{v_1'}, \widetilde{v_{16}'}]$, $[\widetilde{v_2'}, \widetilde{v_{15}'}]$, $[\widetilde{v_{13}'}, \widetilde{v_{14}'}]$, $[\widetilde{v_3'}, \widetilde{v_{12}'}]$, $[\widetilde{v_4'}, \widetilde{v_5'}]$, $[\widetilde{v_6'}, \widetilde{v_7'}]$, $[\widetilde{v_8'}, \widetilde{v_{11}'}]$, and $[\widetilde{v_9'}, \widetilde{v_{10}'}]$ and label them with 1. As a result, $\widetilde{S_1}$ has one 6-cycle, one 4-cycle, and three 2-cycles labelled with 1 and 2 and one 8-cycle, one 4-cycle, and two 2-cycles labelled with 1 and 3. On the other hand, $\widetilde{S_2}$ has the same set of cycle counts with different labels: one 6-cycle, one 4-cycle, and three 2-cycles labelled with 1 and 3 and one 8-cycle, one 4-cycle, and two 2-cycles labelled with 1 and 2. Thus, both $\widetilde{S_1}$ and $\widetilde{S_2}$ are homotopic to 10-holed spheres, and by construction have the same sets of jester hats glued to them; thus, $\widetilde{X_1}$ and $\widetilde{X_2}$ are homotopic. Observe, however, that $\widetilde{S_1}$ and $\widetilde{S_2}$ are not homeomorphic, as desired. This completes the proof for $K = 1$.

Case 2 ($\mathbf{K > 1}$): Suppose $K > 1$ and $K = p_1^{u_1} p_2^{u_2} ... p_t^{u_t} = \prod_{d=1}^{t} p_d^{u_d}$ is the prime factorization of K. We will now construct two non-homeomorphic $2\prod_{d=1}^{t} p_d^{u_d+1}$-sheeted covers $\widetilde{X_1}$ and $\widetilde{X_2}$. First, we construct their singular sets $\widetilde{S_1}$ and $\widetilde{S_2}$, which will both have $2\prod_{d=1}^{n} p_d^{u_d+1}$

essential vertices, which we will label $\tilde{v}_i$ and $\widetilde{v'}_i$ for $1 \le i \le 2\prod_{d=1}^{t} p_d^{u_d+1}$. Note that all the vertex indices in the construction will be taken modulo $2\prod_{d=1}^{t} p_d^{u_d}$.

Construction of $\widetilde{S_1}$: First, suppose that $k \neq N$; if $k = N$, the following construction will be disconnected. For the first singular set $\widetilde{S_1} \subset \widetilde{X_1}$, for even i, construct k edges between $\tilde{v}_i$ and $\tilde{v}_{i+1} \mod 2\prod_{d=1}^{t} p_d^{u_d+1}$. The edges will be labeled with all integers I where $2 \le I \le k+1$. For odd i, construct $N-k$ edges between $\tilde{v}_i$ and $\tilde{v}_{i+1} \mod 2\prod_{d=1}^{t} p_d^{u_d+1}$. One of these edges will be labeled with 1 and the rest with all integers I where $k+2 \le I \le N$. As a result, for $i \neq 1, k+1$, between $\widetilde{v_i}$ and $\widetilde{v_{i+1}}$, there is a copy of a two-cycle that is a double cover of the nonreflection edges in $\mathcal{D}_\Gamma$ labeled with i and $(i+1) \mod N$. We will then have a graph with $\prod_{d=1}^{t} p_d^{u_d+1}$ copies of two-cycles labeled with i and $(i+1)$ for $2 \le i \le N, i \neq k+1$, as well as two $2\prod_{d=1}^{t} p_d^{u_d+1}$-cycles, one labeled with 1 and 2, and one labeled with $k+1$ and $k+2$. For an example, refer to the graph on the right in Figure 3.5, which is an 18-sheeted cover of the singular set from the orbicomplex in Figure 3.2.

We now consider the special case where $k = N$. For even i, construct $N-2$ edges between $\tilde{v}_i$ and $\tilde{v}_{i+1}$ labeled with integers I where $2 \le I \le N-1$. Then for odd i, construct 2 edges between vertices $\tilde{v}_i$ and $\tilde{v}_{i+1}$ labeled with 1 and N. We will again have a graph with $\prod_{d=1}^{t} p_d^{u_d+1}$ copies of two-cycles labeled with i and $(i+1)$ for $i = N$ and $2 \le i \le N-2$. We will also still have two $2\prod_{d=1}^{t} p_d^{u_d+1}$-cycles, one labeled with $N-1$ and N and the other with 1 and 2 as before.

Construction of $\widetilde{S_2}$: For the second singular set $\widetilde{S_2} \subset \widetilde{X_2}$, first partition $\{\tilde{v'}_i\}$ into $p = \prod_{d=1}^{t} p_d$ sets of equal size, so the first set will contain vertices $\tilde{v'}_i$ where $1 \le i \le 2\prod_{d=1}^{t} p_d^{u_d}$, the second set will contain vertices where $2\prod_{d=1}^{t} p_d^{u_d} + 1 \le i \le 4\prod_{d=1}^{t} p_d^{u_d}$, so in general, the nth set will contain $\tilde{v'}_i$ labeled $2n\left(\prod_{d=1}^{t} p_d^{u_d}\right)+1 \le i \le 2(n+1)\prod_{d=1}^{t} p_d^{u_d}$, where $0 \le n \le p-1$.

We first consider the case where $k \neq N$. Construct $k-1$ edges between the pairs of vertices labeled $\tilde{v'}_{2n(\prod_{d=1}^{t} p_d^{u_d})+1}$ and $\tilde{v'}_{2(n+1)\prod_{d=1}^{t} p_d^{u_d}}$, for $0 \leq n \leq p-1$. Label these edges with integers I, where $2 \leq I \leq k$. For all $1 \leq n \leq p$, construct an edge each labeled with $k+1$ between $\tilde{v'}_{2n\prod_{d=1}^{t} p_d^{u_d}}$ and $\tilde{v'}_{2n(\prod_{d=1}^{t} p_d^{u_d})+1}$. Finally, add edges to the remaining vertices with no edges between them. For even i, if there are no edges between $\tilde{v'}_i$ and $\tilde{v'}_{i+1}$, add k edges and label them with integers I where $2 \leq I \leq k+1$. For odd i, add k edges between all $\tilde{v'}_i$ and $\tilde{v'}_{i+1}$ with no edges between them. One of these edges will be labeled with a 1, while the rest are labeled with integers I such that $k+2 \leq I \leq N$. As a result, we will have p cycles of length $2\prod_{d=1}^{t} p_d^{u_d}$ labeled with 1 and 2, $\prod_{d=1}^{t} p_d^{u_d+1}$ two-cycles labeled with i and $i+1$ for $2 \leq i \leq k$ or $k+2 \leq i \leq N$, $\prod_{d=1}^{t} p_d^{u_d}$ two-cycles labeled with k and $k+1$, and a $2p$-cycle labeled with k and $k+1$. For an example, see the graph on the left in Figure 3.5.

Now consider the edge case of $k = N$. This time, we will construct one edge labeled with 1 between the pairs of vertices labeled $\tilde{v'}_{2n(\prod_{d=1}^{t} p_d^{u_d})+1}$ and $\tilde{v'}_{2(n+1)\prod_{d=1}^{t} p_d^{u_d}}$, for $0 \leq n \leq p-1$. For all $1 \leq n \leq p$, construct an edge labeled with N between $\tilde{v'}_{2n\prod_{d=1}^{t} p_d^{u_d}}$ and $\tilde{v'}_{2n(\prod_{d=1}^{t} p_d^{u_d})+1}$. For the other pairs of vertices, for even i, if there are no edges between $\tilde{v'}_i$ and $\tilde{v'}_{i+1}$, add 2 edges and label them with 1 and N. For odd i, add $N-2$ edges between all $\tilde{v'}_i$ and $\tilde{v'}_{i+1}$ labeled with integers I such that $2 \leq I \leq N-1$. As a result, we will have p cycles of length $2\prod_{d=1}^{t} p_d^{u_d}$ labeled with 1 and 2, $\prod_{d=1}^{t} p_d^{u_d+1}$ each of two-cycles labeled with i and $i+1$ for $2 \leq i \leq N-2$ or $i = N$, and a $2p$-cycle labeled with 1 and N.

Observe that $\widetilde{S_1}$ and $\widetilde{S_2}$ are not homeomorphic since $\widetilde{S_1}$ has many more cut pairs. In particular, any two essential vertices will form a cut pair in $\widetilde{S_1}$, but $\widetilde{S_2}$ only has $\binom{2p}{2}$ pairs of cut pairs. As a result, $\widetilde{X}_1$ and $\widetilde{X}_2$ are not homeomorphic. However, note that $\widetilde{S_1}$ and $\widetilde{S_2}$ are homotopic; take a regular neighborhood of both graphs to obtain a $2\prod_{d=1}^{t} p_d^{u_d+1}(N-2)+1$-holed sphere.

We then examine what orbifolds are glued to $\widetilde{S_1}$. Recall that $r_{s,b}$ denotes the number of reflection edges of the orbifold corresponding to the bth branch of $\Theta_s \subset \Gamma$. By Lemma 2.1, we can see that there will be a disk with $\prod_{d=1}^{t} p_d^{u_d+1}(r_{s,b}-3)+2$ cone points glued

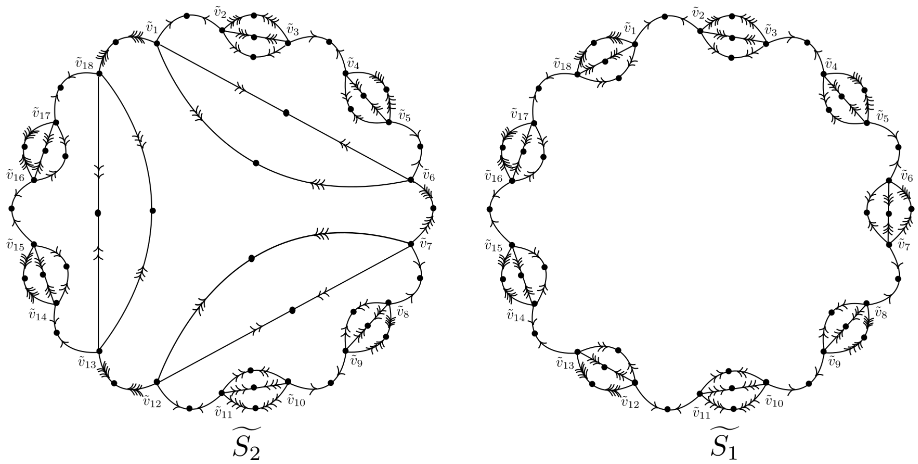

Figure 3.5: Two homotopic but non-homeomorphic covers of the singular set of W_Γ from Figure 3.2. Instead of labeling the edges with numbers, we use different numbers of arrows.

along its boundary circle to the $2\prod_{d=1}^{t} p_d^{u_d}$-cycles labeled by 1 and 2 and $k+1$ and $k+2$ for $s = 1, k+1$ for the case $k \neq N$; if $k = N$, the second $2\prod_{d=1}^{t} p_d^{u_d}$-cycle is labeled with $N-1$ and N and $s = N-1$. Additionally, there will be a disk with $r_{i,b} - 1$ cone points glued to each of the $\prod_{d=1}^{t} p_d^{u_d+1}$ two-cycles labeled by i and $i+1 \mod N+1$ for $2 \leq i \leq k$ or $k+2 \leq i \leq N$ if $k \neq N$ ($2 \leq i \leq N-2$ or $i = N$ if $k = N$).

Next, we list the orbifolds glued to $\widetilde{S_2}$. First, there will be p copies of disks with $\prod_{d=1}^{t} p_d^{u_d}(r_{1,b} - 3) + 2$ cone points glued to the $2\prod_{d=1}^{t} p_d^{u_d}$-cycles with edges labeled with 1 and 2. Second, there will be one copy of a disk with $\prod_{d=1}^{t} p_d^{u_d+1}(r_{k+1,b} - 3) + 2$ cone points glued to the $2\prod_{d=1}^{t} p_d^{u_d+1}$ cycle labeled by $k+1$ and $k+2$ if $k \neq N$ and $N-1$ and N if $k = N$. There will also be $\prod_{d=1}^{t} p_d^{u_d+1}$ copies of a disk with $r_{i,b} - 1$ cone points glued to each 2-cycle labeled by i and $i+1 \mod N$ for $2 \leq i \leq k-1$ or $k+1 \leq i \leq N$ if $k \neq N$ and $2 \leq i \leq N-2$ if $k = N$. Finally, there are $\prod_{d=1}^{t} p_d^{u_d}$ copies of a disk with $r_{k,b} - 1$ cone points glued to each two-cycle labeled by k and $k+1$ as well as one copy of a disk with $p(r_{k,b} - 3) + 2$ cone points glued to the $2p$-cycle labeled by k and $k+1$. Then for both orbicomplex covers, we have the same collection of orbifolds glued to $\widetilde{S_1}$ and $\widetilde{S_2}$:

- $\prod_{d=1}^{t} p_d^{u_d+1}$ copies of each disk with $r_{i,b} - 1$ cone points, where $2 \leq i \leq k-1$ or

$k+1 \leq i \leq N$ if $k \neq N$ and $2 \leq i \leq N-2$ if $k = N$;

- One copy of a disk with $\prod_{d=1}^{t} p_d^{u_d+1}(r_{k+1,b}-3)+2$ cone points;
- $\prod_{d=1}^{t} p_d^{u_d} + p = \prod_{d=1}^{t} p_d^{u_d+1}$ copies of a disk with $r_{k,b}-1 = \prod_{d=1}^{t} p_d^{u_d}(r_{1,b}-3)+2$ cone points since there are $\prod_{d=1}^{t} p_d^{u_d}$ copies of two-cycles labeled with k and $k+1$ as well as p copies of $2\prod p_d^{u_d}$-cycles labeled with 1 and 2 in $\widetilde{S_2}$ and $\prod_{d=1}^{t} p_d^{u_d+1}$ copies of two-cycles labeled with k and $k+1$ in $\widetilde{S_1}$;
- One copy of the disk with $p(r_{k,b}-3)+2 = \prod_{d=1}^{t} p_d^{u_d+1}(r_{1,b}-3)+2$ cone points since there is one $2p$-cycle labeled with k and $k+1$ in $\widetilde{S_2}$ and one $2\prod_{d=1}^{t} p_d^{u_d+1}$-cycle labeled with 1 and 2 in $\widetilde{S_1}$.

As a result, since $\widetilde{X}_1$ and $\widetilde{X}_2$ have the same sets of jester hats glued to homotopic graphs, we have found finite covers that are homotopic but not homeomorphic. This proves Theorem 3.1.6.

□

Next, we prove a Lemma which immediately implies Theorem 3.1.9.

Lemma 3.3.4. *Suppose* Γ_1 *and* Γ_2 *form a non-isomorphic permuted pair (see Definition 3.1.7). Then there exist finite-sheeted covers of* D_{Γ_1} *and* D_{Γ_2} *that are homotopic but not homeomorphic.*

Proof. Note that $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$ have the same number of generalized Θ graphs glued together, so the two-sheeted covers of their singular sets will be two isomorphic generalized Θ graphs S_1 and S_2. We will now construct two non-homeomorphic double covers of S_1 and S_2, which we will call $\widetilde{S_1}$ and $\widetilde{S_2}$.

Since Γ_1 and Γ_2 are not isomorphic, $N > 3$, where N as before denotes the number of generalized Θ graphs glued together in Γ_1 and Γ_2. Indeed, suppose that Γ_1 and Γ_2 are cycles of three generalized Θ graphs, Θ_1, Θ_2, and Θ_3. Without loss of generality, suppose that in Γ_1, Θ_i has essential vertices v_i and v_{i+1} and in Γ_2, Θ_1 has essential vertices v_1' and

v'_2, Θ_2 has essential vertices v'_1 and v'_3, and Θ_3 has essential vertices v'_2 and v'_3. Then there is a graph isomorphism $f : \Gamma_1 \to \Gamma_2$, defined by $f(v_1) = v'_2$, $f(v_2) = v'_1$, and $f(v_3) = v'_3$ (with f defined on the valence 2 vertices in the natural way). We point this out since the construction detailed below will not work for $N = 3$.

To construct $\widetilde{S_1}$, fix a cyclic ordering on a set of four vertices $\{\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4\}$. Construct edges between $\tilde{v}_1$ and $\tilde{v}_2$, and between $\tilde{v}_3$ and $\tilde{v}_4$, labeled with all integers I such that $1 \leq I \leq N-1$. Then construct one edge between $\tilde{v}_2$ and $\tilde{v}_3$, as well as $\tilde{v}_4$ and $\tilde{v}_1$ and label those edges with N.

According to the assumptions, for every generalized Θ graph Θ_i between vertices v_i and v_{i+1} in Γ_1, there is an isomorphic generalized Θ graph Θ'_j in Γ_2 between v'_j and $v'_{j+1} \in \Gamma_2$. We can therefore assume there exists some $\Theta'_j \subset \Gamma_2$ that is isomorphic to $\Theta_1 \subset \Gamma_1$, and some $\Theta'_k \subset \Gamma_2$ that is isomorphic to $\Theta_N \subset \Gamma_1$. Without loss of generality, assume that $j < k$. Construct $\widetilde{S_2}$ with vertices $\{\tilde{v'}_1, \tilde{v'}_2, \tilde{v'}_3, \tilde{v'}_4\}$ and construct edges between $\tilde{v}_1$ and $\tilde{v'}_2$ as well as $\tilde{v'}_3$ and $\tilde{v'}_4$ labeled with all integers I such that $j+1 \leq I \leq k$. Then construct edges between $\tilde{v'}_2$ and $\tilde{v'}_3$ as well as $\tilde{v'}_4$ and $\tilde{v'}_1$ labeled with all integers I such that $k+1 \leq I \leq N$ or $1 \leq I \leq j$.

In the Davis orbicomplexes, the isomorphic Θ graphs $\Theta_1 \subset \Gamma_1$ and $\Theta'_j \subset \Gamma_2$ give rise to identical sets of orbifolds glued to edges labeled 1 and 2 in D_{Γ_1} and j and $j+1$ in D_{Γ_2}. As a result, if there is a cycle in X_1 labeled with 1 and 2 and a cycle of the same length in X_2 labeled with j and $j+1$, the sets of jester hats glued to them will be identical. The same holds for $\Theta_N \subset \Gamma_1$ and $\Theta'_k \subset \Gamma_2$. Note that in $\widetilde{S_1}$, there are two four-cycles with jester hats glued to them, one labeled with 1 and N and the other with N and $N-1$. The other cycles with jester hats glued to them are all two-cycles labeled with I and $I+1$ where $1 \leq I \leq N-2$. On the other hand, $\widetilde{S_2}$ is a generalized Θ graph with two cycles of length four that have jester hats glued to them- one labeled with k and $k+1$ and the other with j and $j+1$. The other cycles with jester hats glued to them are labeled with I and $I+1$ where $I \neq j, k$ are two-cycles. Note that the regular neighborhood of both $\widetilde{S_1}$ and $\widetilde{S_2}$ is the $(2N-2)$-holed sphere $S_{0,2N-2}$, and there is a bijective correspondence between

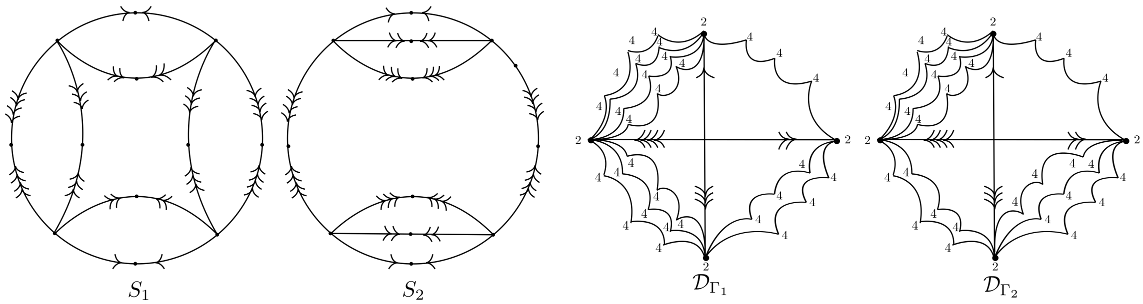

Figure 3.6: Here, the defining graphs Γ_1 and Γ_2 are the permuted pair from Figure 3.1, so they satisfy the conditions of Lemma 3.3.4. The singular sets S_i of X_i, which are four-sheeted covers of $\mathcal{D}_{\Gamma_i}$ for $i = 1, 2$ are shown on the left. Note that X_1 and X_2 are homotopic but not homeomorphic.

sets of jester hats glued to boundary components of $S_{0,2N-2} \subset \widetilde{S_1}$ and sets of jester hats glued to boundary components of $S_{0,2N-2} \subset \widetilde{S_2}$. Thus, X_1 and X_2 are homotopic but not homeomorphic because $\widetilde{S_1}$ and $\widetilde{S_2}$ are not homeomorphic. For an example, refer to Figure 3.6.

□

As a segue into our next section, we make the following remark.

Remark 3.3.5. Finding necessary and sufficient conditions for topological rigidity is a very nuanced task. In our setting, by Theorem 3.1.9, we know a topologically rigid set cannot contain the finite-sheeted covers of $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$, where Γ_1, Γ_2 are strongly repetitive or form a permuted pair (see Definition 3.1.7). Unfortunately, simply excluding finite-sheeted covers of all $\mathcal{D}_\Gamma$ where Γ is repetitive and part of a permuted pair is not sufficient for constructing a topologically rigid set. For example, for $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$ from Figure 3.7, Γ_1 and Γ_2 are not composed of the same set of generalized Θ graphs glued together. Furthermore, for both defining graphs, there do not exist pairs of commensurable Euler characteristic vectors of generalized Θ graphs in Γ_1 and Γ_2. Nevertheless, there exist eight- and four-sheeted covers of $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$ respectively that are homotopic but not homeomorphic.

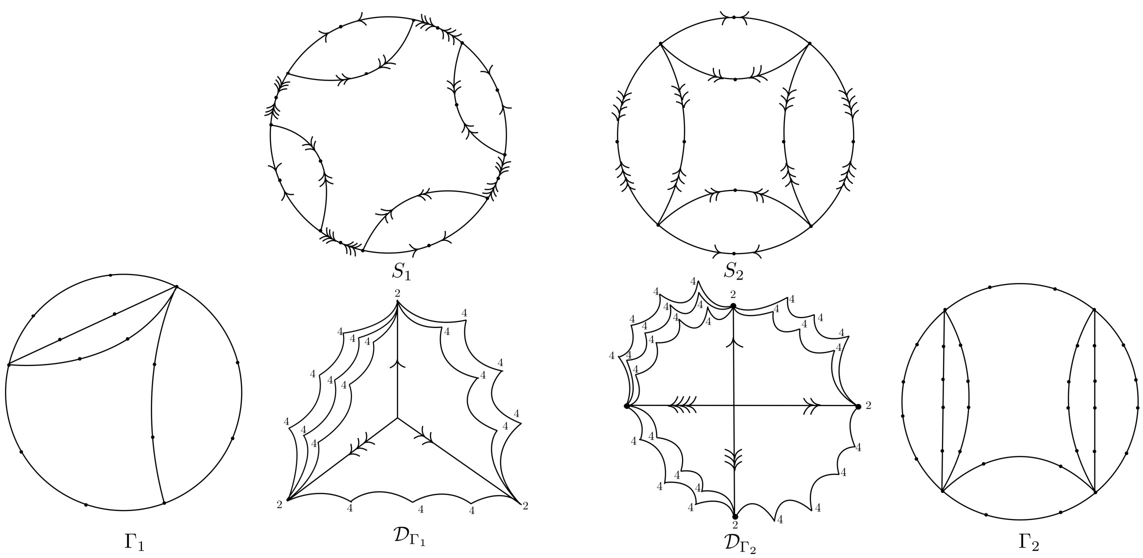

Figure 3.7: S_1 and S_2 are non-homeomorphic singular sets of homotopic eight- and four-sheeted covers of $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$. Thus, neither Γ_1 nor Γ_2 is repetitive and Γ_1 and Γ_2 do not form a permuted pair, yet any set $\mathcal{X}''$ that contains all finite-sheeted covers of $\mathcal{D}_{\Gamma_1}$ and $\mathcal{D}_{\Gamma_2}$ is not topologically rigid.

3.4 A topologically rigid set

In this section, we introduce a class of finite-sheeted covers of Davis orbicomplexes that is topologically rigid. Since topological rigidity is difficult to achieve, many assumptions are necessary, so our class does not contain the complete set of finite-sheeted covers of Davis orbicomplexes. In particular, to find rigid classes, we not only need restrictions on the defining graphs but also on the singular sets of the covers.

A key tool in the proof of Theorem 3.4.6 is Lafont's topological rigidity result for simple, thick hyperbolic surface amalgams (see Theorem 2.0.2). In order to use Theorem 2.0.2, we impose a restriction on the finite covers of Davis orbicomplexes we are examining, which we state below.

Assumption 3.4.1. X is homotopic to an orbicomplex Y that consists of jester hats glued along their boundaries to the boundaries of an h-holed genus g surface $S_{g,h}$. Furthermore, in Y, each boundary of $S_{g,h}$ has at least one jester hat glued to it.

Recall that a graph Γ is said to be *3-convex* if every edge between its essential vertices has at least 3 subdivisions. Note that if a cycle of generalized Θ graphs Γ is 3-convex, then W_Γ is hyperbolic since Γ is square-free. Although the converse is not true, we impose the 3-convexity condition in our proof to ensure our construction of hyperbolic surface amalgam lifts of Davis orbicomplexes will work.

We now use Assumption 3.4.1 to prove a lemma that will be important in the proof of our main result of the section.

Lemma 3.4.2. *Let X_1 and X_2 be finite covers of Davis orbicomplexes D_{Γ_1} and D_{Γ_2}, where Γ_1 and Γ_2 are 3-convex and satisfy Assumption 3.4.1, and suppose $\pi_1(X_1) \cong \pi_1(X_2)$. Then the isomorphism $f : \pi_1(X_1) \to \pi_1(X_2)$ induces a bijection f_* between jester hats of X_1 and X_2 and for a jester hat $\mathcal{O}_1 \subset X_1$, if $f_*(\mathcal{O}_1) = \mathcal{O}_2 \subset X_2$, then $\mathcal{O}_1$ and $\mathcal{O}_2$ are homeomorphic. Furthermore, if S_1 and S_2 are singular subsets of X_1 and X_2 respectively, if $\gamma_1 \subset S_1$ is the boundary component of $\mathcal{O}_1$, then $f_*(\gamma_1) = \gamma_2$ where γ_2 is the boundary component of $\mathcal{O}_2$.*

Proof. Suppose X_1 and X_2 are two orbicomplexes with 3-convex defining graphs where $\pi_1(X_1) \cong \pi_1(X_2)$. We first use the construction from Proposition 3.2 of [Sta18]: let X_i be a finite cover of $\mathcal{D}_{\Gamma_i}$ for $i = 1, 2$. Each jester hat in X_i with p cone points lifts to a orbifold with $2(p-2)$ cone points and two boundary components, which in turn has a two-sheeted cover $S_{g,4}$, where $g = \frac{2(p-2)}{2} - 1$. Then, we glue each boundary component of $S_{g,4}$ to a copy of S_i, the singular set of X_i, to obtain a torsion-free four-sheeted cover. We will call these torsion-free covers $\widehat{X_1}$ and $\widehat{X_2}$. See Figure 3.8 for an illustration of the construction.

Since abstract commensurability is an equivalence relation, $\pi_1(\widehat{X_1})$ and $\pi_1(\widehat{X_2})$ are abstractly commensurable, so there exist finite-sheeted covers $\mathcal{Y}_1$ and $\mathcal{Y}_2$ such that

$$\pi_1(\widehat{X_1}) \geq \pi_1(\mathcal{Y}_1) \cong \pi_1(\mathcal{Y}_2) \leq \pi_1(\widehat{X_2}).$$

Note that $\widehat{X_1}$ and $\widehat{X_2}$ are homotopic to hyperbolic surface amalgams; take the regular neighborhoods of their singular sets. As a result, finite-sheeted covers of $\widehat{X_1}$ and $\widehat{X_2}$ are also homotopic to hyperbolic surface amalgams by Nielsen-Schreier, so it follows that $\mathcal{Y}_1$

and $\mathcal{Y}_2$ are homotopic to hyperbolic surface amalgams Y_1 and Y_2. By Corollary 3.5 of [Laf07b], a homotopy ϕ between hyperbolic surface amalgams induces a bijection between homeomorphic chambers (hyperbolic manifolds with boundary) of Y_1 and Y_2. Additionally, for a chamber $C_1 \subset Y_1$ with boundary component γ_1, if $\phi(C_1) = C_2 \subset Y_2$, then $\phi(\gamma_1) = \gamma_2$ where γ_i is a boundary component of C_i for $i = 1, 2$. Using these results, we can conclude that surfaces with boundary in $\mathcal{Y}_1$ are mapped bijectively, and homeomorphically, to surfaces in $\mathcal{Y}_2$, as the maps are preserved under homotopy. The homotopy lifting property then gives us the statement of the lemma.

Alternatively, observe that the Davis complex Σ_Γ, the universal cover of $\mathcal{D}_\Gamma$, is CAT(0) and thus contractible. Since W_Γ acts freely on Σ_Γ, it follows that $\mathcal{D}_\Gamma$ is a classifying space of W_Γ, or equivalently a $K(W_\Gamma, 1)$ space. Furthermore, finite covers of $\mathcal{D}_\Gamma$ are quotients of Σ_Γ by a free action as well, so X_1 and X_2 are classifying spaces for the same finite-index subgroup of W_Γ. As a result, since $\pi_1(X_1) \cong \pi_1(X_2)$, X_1 and X_2 are homotopy equivalent by Whitehead's Theorem. This allows us to construct a shorter tower of covers since a homotopy between X_1 and X_2 induces a homotopy between $\widehat{X_1}$ and $\widehat{X_2}$, which are homotopic to hyperbolic surface amalgams. We can then apply Theorem 2.0.2 to obtain our result. □ □

We stress that Assumption 3.4.1 is key for the proof of Lemma 3.4.2. For example, let $S_{g,n}$ be a genus g surface with n boundary components. Recall the *graph genus* of a graph G is the minimal genus of an orientable surface into which G can be embedded. In general, a cover of a Davis orbicomplex X is homotopic to an orbicomplex consisting of jester hats identified along their boundaries to a set of simple closed curves C on $S_{g,n}$ since every graph has a genus (see [Whi01]). Analyzing X can be difficult since there is no guarantee of the four-sheeted hyperbolic surface amalgam cover constructed in Lemma 3.4.2, as the lifts of C may not be a disjoint union of circles. For example, consider Figure 3.9, which depicts a six-sheeted cover of a Davis orbicomplex $\mathcal{D}_\Gamma$. The singular set of X is the complete bipartite graph $K_{3,3}$, which is homotopic to $S_{1,3}$. Since the jester hats are not identified along disjoint circles, Lafont's rigidity result is not available for use and the

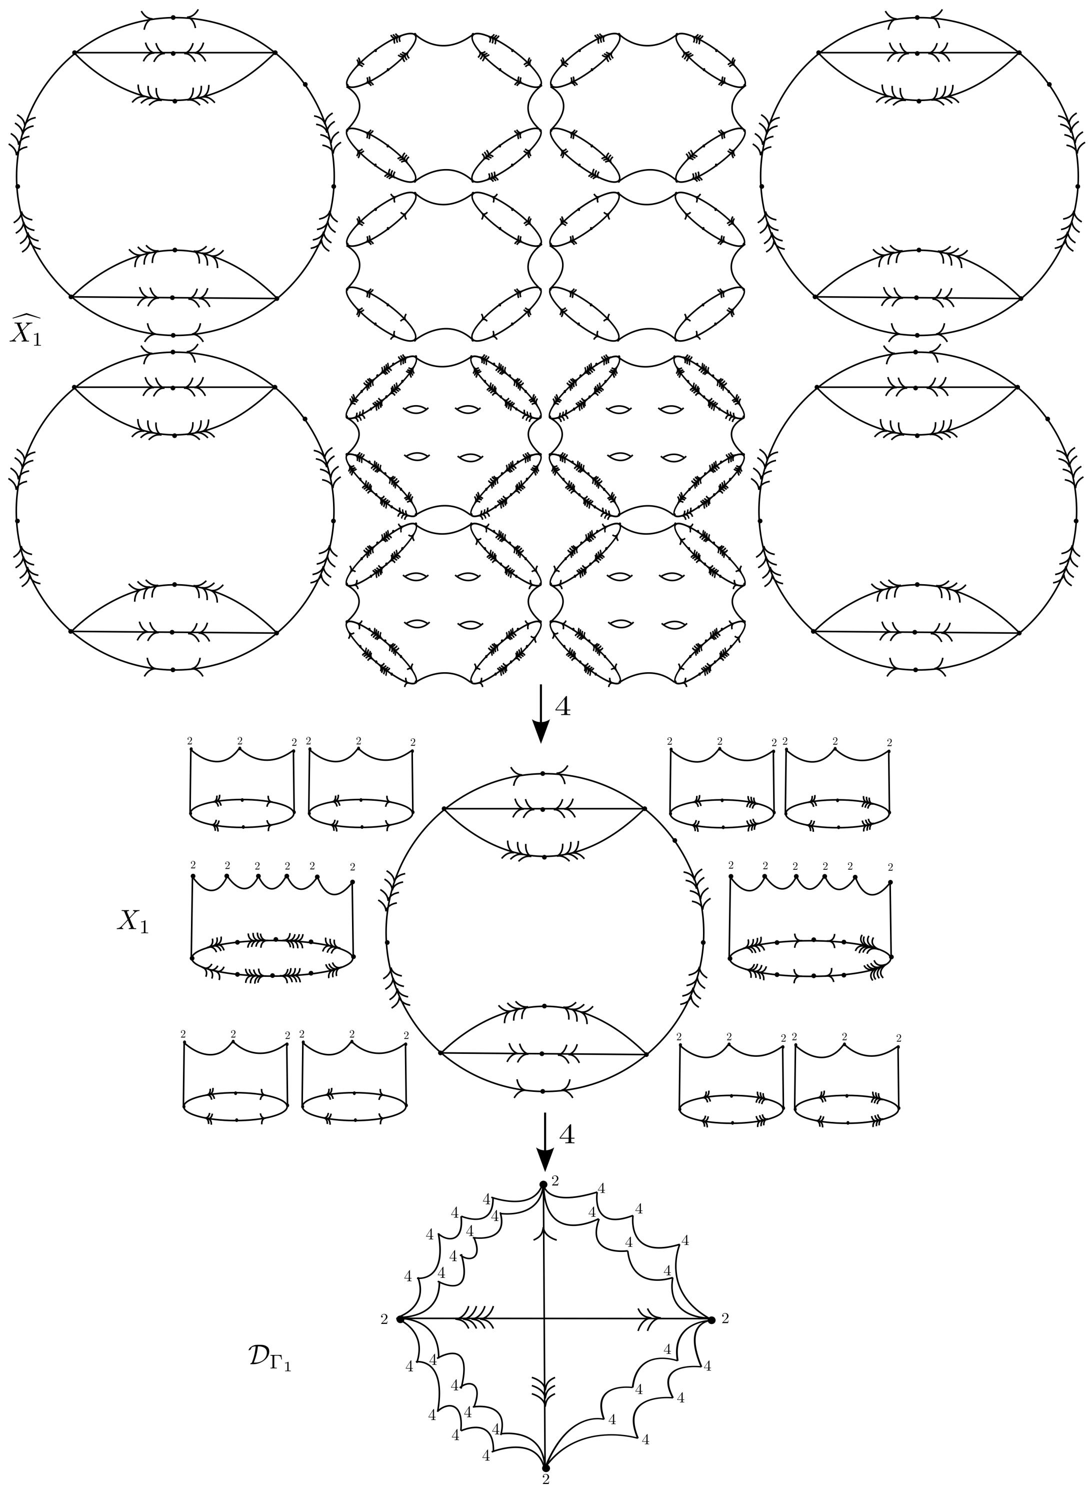

Figure 3.8: A tower of covers constructed in the proof of Lemma 3.4.2. Note that $\widehat{X_1}$ is homotopic to a hyperbolic surface amalgam.

proof for Lemma 3.4.2 does not work.

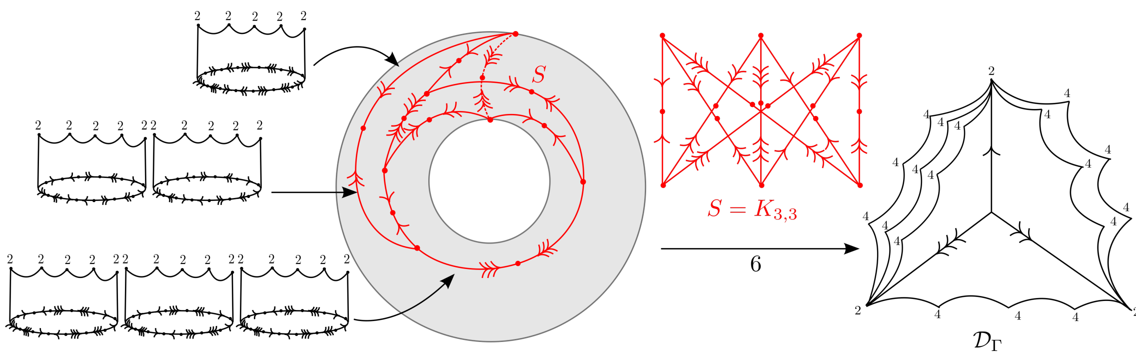

Figure 3.9: An example of a finite-sheeted cover of a Davis orbicomplex $\mathcal{D}_\Gamma$ that does not satisfy Assumption 3.4.1. The singular set S is not necessarily planar; in this example, S (depicted in red) embeds on a torus.

In order to determine whether two finite-sheeted covers of Davis orbicomplexes are homeomorphic, which we need to do to determine topological rigidity, we need to check that their singular sets are homeomorphic. Unfortunately, since finite covers of the singular sets are graphs, determining topological rigidity therefore requires solving a graph isomorphism problem, which has a high computational complexity. Recall that a *complete graph invariant* is combinatorial tool for determining whether a pair of graphs in a family of graphs is isomorphic. To simplify our problem, we will define a family of singular sets $\mathcal{S}$ of finite-sheeted covers of a Davis complex $\mathcal{D}_\Gamma$ with an easily computable complete graph invariant, which we now introduce.

Definition 3.4.3 (Cycle count vectors)**.** Consider X, a $2d$-sheeted cover of a Davis orbicomplex, where $d > 0$ is any arbitrary integer. Recall that for a cycle of generalized Θ graphs Γ, its associated Davis orbicomplex $\mathcal{D}_\Gamma$ has a singular set that is a star graph with N edges, which we will label with integers $i = 1, 2, ..., N$. If an edge $e = [v, w]$ is labeled by i, we will write $e = [v, w]_i$. Fix a cyclic labeling of the edges, which will lift to a labeling in any cover of the singular set. For $1 \leq i \leq N$, let $x_i = (x_{i,1}, x_{i,2}, ..., x_{i,d})$ be a vector where $x_{i,j}$ is the number of cycles in S of length $2j$ that are labeled with i and $i+1$. Recall that as usual, we are counting the edges between two essential vertices as one edge. Note that x_i is a vector of length d since the possible cycle lengths of a $2d$-sheeted cover will range from 2 to $2d$. For an example, refer to Figure 3.6. Let the labeling of an edge of S_i be the number of arrows seen on the edge, so $N = 3$. For S_1, the set of cycle count vectors is $x_1 = (2, 0)$, $x_2 = (0, 1)$, $x_3 = (2, 0)$, and $x_4 = (0, 1)$ since there are two 2-cycles labeled with 1 and 2, two 2-cycles labeled with 3 and 4, one 4-cycle labeled with 2 and 3, and one 4-cycle labeled with 4 and 1. Similarly, for S_2, the set of cycle count vectors is $x_1 = (2, 0)$, $x_2 = (2, 0)$, $x_3 = (0, 1)$ and $x_4 = (0, 1)$. As we will see soon, since the cycle count vectors are different, S_1 and S_2 are not isomorphic.

We now define a family of singular sets of finite-sheeted covers of a Davis complex $\mathcal{D}_\Gamma$. Recall that the double cover of the singular set of $\mathcal{D}_\Gamma$ is itself a generalized Θ graph Θ_N with N branches, where N is the number of generalized Θ graphs in the defining graph Γ (see Construction 4.4.3). A double cover of Θ_N is a cycle of four (possibly trivial) generalized Θ graphs S' with valence N vertices. Note there exists some $a, b \in \mathbb{Z}_{\geq 0}$ and $a + b = N$ such that adjacent essential vertices of S' either have a or b branches between them. For example, in Figure 3.6, all the essential vertices in S_1 and S_2 have valence $N = 4$ since the original defining graph consisted of four generalized Θ graphs glued together. In S_1, $a = 2$ and $b = 2$, and in S_2, $a = 1$ and $b = 3$.

Construction 3.4.4 (A special class of singular sets $\mathcal{S}$)**.** To begin our construction, take S', a four-sheeted cover of the singular set of $\mathcal{D}_\Gamma$, which is a cycle of four generalized Θ graphs each with either a or b branches. Then arbitrarily choose two adjacent vertices of valence greater than two, v_i and v_{i+1} with $n_i = a$ or b edges between them. Delete some fixed number of (subdivided) edges j between them, where $1 \leq j \leq n_i$, add two essential vertices u_i and u_{i+1}, and add j edges between v_i and u_i as well as v_{i+1} and u_{i+1}. Finally, add $N - j$ edges between u_i and u_{i+1}. Then arbitrarily choose two other adjacent essential vertices and repeat the process any finite number of times. See Figure 3.10 for an example of an element of $\mathcal{S}$; at each step, two edges are added between the new essential vertices.

Let $S \in \mathcal{S}$ be a graph that can be constructed from the process described above. By construction, S covers the original singular set of the Davis orbicomplex. Notice $\mathcal{S}$

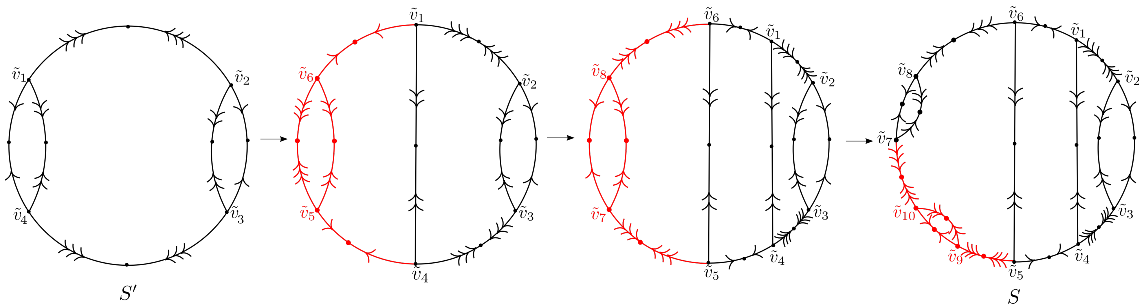

Figure 3.10: Construction of a $S \in \mathcal{S}$ from a cycle of generalized Θ graphs. The edges and vertices added at each step are depicted in red.

describes singular sets, not defining graphs, so the graphs in $\mathcal{S}$ are not restricted to cycles of generalized Θ graphs. We now introduce a second assumption:

Assumption 3.4.5. The singular set S of X is an element of the set $\mathcal{S}$ defined in Construction 3.4.4.

We are now ready to introduce a topologically rigid class finite covers of Davis orbicomplexes.

Theorem 3.4.6. *Suppose Γ is 3-convex and not repetitive (see Definition A.0.3). Let $\mathcal{X}''$ contain all finite-sheeted covers of $\mathcal{D}_\Gamma$ that satisfy Assumptions 3.4.1 and 3.4.5. Then $\mathcal{X}''$ is topologically rigid.*

We now give an outline of the proof of Theorem 3.4.6 (for the actual proof, see the end of the section). We first show that the statement of the theorem reduces to a graph isomorphism problem on singular sets of two finite covers of Davis orbicomplexes (see Lemma 3.4.9). We then show that if two finite covers of Davis orbicomplexes satisfying the conditions listed in Theorem 3.4.6 are homotopic, then they have the same cycle count vectors. Finally, in Lemma 3.4.11, we show that cycle count vectors are a complete graph invariant for singular sets in $\mathcal{S}$ detailed in Construction 3.4.4.

Remark 3.4.7. Recall that if X_1 and X_2 are homeomorphic, then W_{Γ_1} and W_{Γ_2} are commensurable. Given that Assumption 3.4.1 is true, the converse is very much false. Figure 1.2 of [DST18] gives some examples of pairs of defining graphs $\{\Gamma_i, \Gamma'_i\}$ $(i = 1, 2, 3)$

of commensurable RACGs, which we will now reference. We can check that none of the finite-sheeted covers of $\mathcal{D}_{\Gamma_i}$ and $\mathcal{D}_{\Gamma_i'}$ can be homotopic by Lemma 3.4.2, even though the RACGs W_{Γ_i} and $W_{\Gamma_i'}$ are commensurable. To see the full commensurability classification for cycles of generalized Θ graphs, refer to Theorem 1.12 of [DST18].

Definition 3.4.8. We say that a graph homeomorphism $\bar{f} : G_1 \to G_2$ is *label-preserving* if for all $[v, w]_i \in G_1$, $\bar{f}([v, w]_i) \subset G_2$ is also labeled by i.

We now prove the following useful result involving singular sets of finite covers of Davis orbicomplexes.

Lemma 3.4.9. *Let X_1 and X_2 satisfy the assumptions from Lemma 3.4.2. Then every label-preserving graph homeomorphism $\bar{f} : S_1 \to S_2$ between the singular subsets of X_1 and X_2 will induce a homeomorphism $f : X_1 \to X_2$.*

Proof. Suppose the singular sets $S_1 \subset X_1$ and $S_2 \subset X_2$ are homeomorphic. Then the vertices $v_j \in V(S_1)$ with valence greater than two map bijectively to the vertices $v_j' \in V(S_2)$ with the same valence, and if $\bar{f}(v_1) = v_1'$ and $\bar{f}(v_2) = v_2'$, then for every edge $[v_1, v_2]_i$ labeled with i, $\bar{f}([v_1, v_2]_i) = [v_1', v_2']_i$. As a result, every cycle $\gamma_1 \subset S$ labeled with i and $i + 1 \pmod N$ is bijectively mapped to a cycle of the same length in $\gamma_2 \subset S_2$ labeled with i and $i + 1 \pmod N$. By Lemma 3.4.2, for every jester hat $\mathcal{O}_1$ with c cone points glued to γ_1, there must also be a jester hat $\mathcal{O}_2$ with c cone points glued to $\gamma_2 = \bar{f}(\gamma_1)$. As a result, $\bar{f}$ induces a homeomorphism $f : X_1 \to X_2$ where $f(S_1) = S_2$ and $f(\mathcal{O}_1) = \mathcal{O}_2$. □ □

Recall that a graph G with genus g can be embedded into a genus g surface S_g. The edges of the graph will divide S_g into regions called *faces*. Let $|V|$ denote the number of vertices of G, $|E|$ the number of edges, and $|F|$ the number of faces. For planar graphs, we can calculate the number of faces using Euler's formula, $2 = |V| - |E| + |F|$. In general, using the definition of Euler characteristic for simplicial complexes, for a graph with genus g, $2 - 2g = |V| - |E| + |F|$.

Lemma 3.4.10. *Let $\mathcal{X}''$ be the set of finite-sheeted covers of a single Davis orbicomplex $\mathcal{D}_\Gamma$, where Γ is 3-convex and not repetitive. Suppose $X_1, X_2 \in \mathcal{X}''$ are homotopic, and their singular sets S_1 and S_2 satisfy the conditions listed in Theorem 3.4.6. Then $x_i = x_i'$ for all $1 \leq i \leq N$, where x_i and x_i' are the cycle count vectors of X_1 and X_2 respectively.*

Proof. First, we show that if Γ consists of N essential vertices, and X_i are d_i sheeted covers of $\mathcal{D}_\Gamma$ for $i = 1, 2$, then $d_1 = d_2$ necessarily. As usual, we will denote S_i to be the singular set of X_i. Since X_1 and X_2 are homotopic, by Lemma 4.3, they consist of the same sets of jester hats identified along their boundary components to some F-holed genus g surface $S_{g,F}$. Note that F is also the number of faces of both $S_1 \subset X_1$ and $S_2 \subset X_2$. Note the number of vertices in S_i is d_i and the number of edges is $\frac{d_i N}{2}$ for $i = 1, 2$, so using the definition of Euler characteristic, we have:

$$d_1 - \frac{d_1 N}{2} + F = 2 - 2g = d_2 - \frac{d_2 N}{2} + F \implies d_1(2-N) = d_2(2-N).$$

Thus, $d_1 = d_2$ necessarily. Then X_1 and X_2 are finite sheeted covers of the same degree of the same Davis orbicomplex $\mathcal{D}_\Gamma$.

Let $\Theta = \Theta(n_1, n_2, ..., n_k)$ and $\Theta' = \Theta(n_1', n_2', ..., n_{k'}')$ be two arbitrary Θ graphs in Γ. Recall that by Lemma 3.3.1, the number of cone points c_i of a jester hat corresponding to b_i, the ith branch of Θ, is $\frac{d}{2}(r_i - 3) + 2$ where d is the index of the cover the jester hat corresponds to. Similarly, the number of cone points c_j' of a jester hat corresponding to b_j', the jth branch of Θ', is $\frac{d}{2}(r_j' - 3) + 2$. Note that Γ is not repetitive, so there do not exist K, L such that $K(\frac{1-n_1}{4}, \frac{1-n_2}{4}, ..., \frac{1-n_k}{4}) = L(\frac{1-n_1'}{4}, \frac{1-n_2'}{4}, ..., \frac{1-n_{k'}'}{4})$. So in particular, $(n_1 - 1, n_2 - 1, ..., n_k - 1) \neq (n_1' - 1, n_2' - 1, ..., n_{k'}' - 1)$ and since n_i and n_j' are equal to $r_i - 2$ and $r_j' - 2$ respectively, it follows that $\{\frac{d}{2}(r_i - 3) + 2\}_{1 \leq i \leq k} \neq \{\frac{d}{2}(r_j' - 3) + 2\}_{1 \leq j \leq k'}$. Thus, in any finite-sheeted cover of D_Γ, there are different sets of jester hats glued to cycles with different labels since the sets of cone point counts are different. In particular, jester hats in X_1 that are lifts of orbifolds corresponding to a Θ graph in Γ must map to a collection of jester hats in X_2 that are lifts of orbifolds corresponding to the same Θ graph. In order for the sets of jester hats to be the same, cycles of length $2j$ labelled with

i and $i+1$ must map to cycles of the same length and also labeled with i and $i+1$. As a result, $x_i = x'_i$. □ □

Lemma 3.4.11. *Let $\mathcal{X}''$ be the class of finite-sheeted covers of Davis orbicomplexes from Theorem 3.4.6. Suppose $S_1, S_2 \in \mathcal{S}$ are the singular sets of two $2d$- sheeted covers $X_1, X_2 \in \mathcal{X}''$. If $x_i = x'_i$ for all $1 \leq i \leq N$, then there exists a label-preserving homeomorphism $\bar{f} : S_1 \to S_2$.*

Proof. We use induction on the degree of the covers, $2d$. Note that since the cycle count vectors are the same, the covers must be of the same degree. For the base case, suppose $d = 2$, so $2d = 4$. By Assumption 3.4.1, each cycle in both S_1 and S_2 must be attached to at least one jester hat. Thus, cycles of odd length are not allowed since a jester hat cannot be attached to such a cycle. Thus, the only possible cycle lengths of S_1 and S_2 are 2 and 4. By construction, every edge in S_1 must be attached to part of the boundary of at least one jester hat, so every edge in S_1 must be included into at least one cycle. In order for S_1 to be a connected graph with four vertices satisfying the property that every edge is part of a at least one cycle, there must be at least one four-cycle in S_1. The same holds for S_2. Since three-cycles are not allowed and each edge in S_1 will be included in a cycle, the only possible edges in S_1 are of the form $[v_i, v_{i+1}]$, where as before, i is taken modulo 4. The same holds for S_2. Therefore, the only possible S_1 and S_2 are graphs with k edges between v_1 and v_2 as well as v_3 and v_4 (where $1 \leq k < N$), and $N - k$ edges between v_2 and v_3 as well as v_4 and v_1 (namely, a cycle of 4 subdivided Θ graphs- see S_1 and S_2 in Figure 3.7 for examples of such graphs). Note that such graphs will have two four-cycles and $(2N - 4)$ two-cycles. As a result, the only 4-sheeted covers that satisfy Assumption 3.4.1 of Theorem 3.4.6 consist of two sets of jester hats glued to four-cycles, and the rest of the sets of jester hats glued to two-cycles.

If $x_i = x'_i$ for all $1 \leq i \leq N$, we know that for both S_1 and S_2, there is one four-cycle labeled with j and $j + 1$ for some $1 \leq j \leq N$, and another four-cycle labeled with k and $k + 1$ for $k \neq j$. Without loss of generality, suppose $j < k$. Note that in S_1, if an edge labeled $j + 1$ is between v_i and v_{i+1}, then an edge labeled k must necessarily also be between v_i and v_{i+1}. Otherwise, the edge labeled with $k+1$ must be between v_i and v_{i+1}, so an edge labeled with k is between v_i and v_{i-1}. In this case, the labels on the edges between v_i and v_{i+1} will range from $j+1$ to $k+1$, so one of the edges must be labeled with k. However, there is already a k edge between v_i and v_{i-1}, which is impossible. The same argument can be used for S_2 if we replace v_i with v_i'. We can thus see that in S_1, for all $1 \leq i \leq N$, if edges labeled with all integers between $j+1$ and k connect v_i and v_{i+1} (v_i' and v_{i+1}' in S_2), then the edges connecting v_i and v_{i-1} (v_i' and v_{i-1}' in S_2) are labeled with integers $1 \leq l \leq N$ such that $l \leq j$ or $l \geq k+1$. We can then construct a homeomorphism $\bar{f} : S_1 \to S_2$ where $f(v_i) = v_i'$ for all $1 \leq i \leq N$. We can easily check that edge labels and vertex adjacencies are preserved under $\bar{f}$, completing the base case.

Suppose the lemma holds for d-sheeted covers. By Assumption 3.4.5, there exist $\{v_i, v_{i+1}\} \in V(S_1)$ and $\{v_i', v_{i+1}'\} \in V(S_2)$ with the same number of edges and the same set of labels between them. Additionally, the only other edges attached to v_i and $v_{i+1} \in V(S_1)$ are also attached to v_{i-1} and v_{i+2} respectively; the same holds for $v_i', v_{i+1}' \in V(S_2)$. Note that for arbitrary $d > 0$, S_1 and S_2 must both have $(2d+2)$-cycles for the same reason the four-sheeted covers in the base case necessarily have 4-cycles: each edge of S_1 and S_2 is necessarily attached to a jester hat, and thus by Assumption 3.4.1 necessarily belongs to a cycle. Suppose there are no cycles in S_1 and S_2 of length $2d+2$. Then S_1 and S_2 would be disconnected since they are graphs with $2d+2$ vertices. Thus, there must be at least one $(2d+2)$-cycle in S_1 and S_2, which we will label with m and $m+1$. Note that if v_i and v_{i+1} have j edges between them, then the other two pairs of vertices $\{v_i, v_{i-1}\}$ and $\{v_{i+1}, v_{i+2}\}$ must also necessarily have $N-j$ edges between them, and the edges between the two pairs of vertices have the same set of labels. Delete v_i and v_{i+1} and the edges they are adjacent to, and construct the j deleted edges between v_{i+2} and v_{i-1} to create the singular set T_1 of a $2d$-sheeted cover of $\mathcal{D}_{\Gamma_1}$. Construct a singular set T_2 of $2d$-sheeted cover of $\mathcal{D}_{\Gamma_2}$ in the same way. Let y_i and y_i' be the new set of cycle count vectors. Note that in total, for both $2d$-sheeted covers, we have deleted and added the same set of cycles with the same set of edge labels, so $y_i = y_i'$ for all $1 \leq i \leq N$. Then by the inductive hypothesis, there exists a homeomorphism $\bar{g} : T_1 \to T_2$.

We then can extend $\bar{g}$ to $\bar{f} : S_1 \to S_2$. Suppose $\bar{g}(v_{i+2}) = v'_k$ for some $v'_k \in T_2$ (note that k is not necessarily equal to i). Then $\bar{g}(v_{i-1})$ maps to an adjacent vertex v_l such that there are j edges between v'_k and v'_l with the same labelings as the edges between v_i and v_{i+1} in the original S_1. Construct two vertices u_k and u_l in T_2 and v_{i-1} and v_{i+2} in T_1, and delete the $N - j$ edges between $v_k, v_l \in V(T_2)$ and $v_{i-1}, v_{i+2} \in V(T_1)$ that have the same labels as the edges added to construct T_i from S_i for $i = 1, 2$. Then reconstruct the j deleted edges between $u_k, u_l \in V(T_2)$ and $v_{i-1}, v_{i+2} \in V(T_1)$ as well as $N - j$ edges $[v'_k, u_k]$ and $[v'_l, u_l]$ in T_2 and $[v_{i-1}, v_i]$ and $[v_{i+1}, v_{i+2}]$ in T_1. Call the new graphs U_1 and U_2, but note that U_1 is identical to S_1. Additionally, U_2 is homeomorphic to S_2 since they are the same graph up to a relabeling of vertices. Let $\bar{g}(v) = \bar{f}(v)$ for all $v \in V(S_1)$, $\bar{f}(v_i) = u_k$ and $\bar{f}(v_{i-1}) = u_l$, which also determines the maps between the newly added edges, giving us a label-preserving homeomorphism $\bar{f} : S_1 \to S_2$. □ □

We now have all the tools to prove Theorem 3.4.6.

Theorem 3.4.6. It suffices to show that for $X_1, X_2 \in \mathcal{X}$, $\pi_1(X_1) \cong \pi_1(X_2)$ implies $X_1 \overset{\text{homeo}}{\cong} X_2$. As a result of Lemma 3.4.9, in order to show $X_1 \overset{\text{homeo}}{\cong} X_2$, it suffices to show there exists a label-preserving graph homeomorphism between the singular sets of X_1 and X_2 respectively. Since subdivisions of a graph belong to the same homeomorphism class, we can delete the valence two vertices from the singular sets to obtain S_1 and S_2, and compare the graphs. By Lemma 3.4.10, $x_i = x'_i$ for $1 \leq i \leq N$, where x_i and x'_i are the cycle count vectors of S_1 and S_2 defined in Definition 3.4.3. Then by Lemma 3.4.11, we can conclude $S_1 \overset{\text{homeo}}{\cong} S_2$, as desired. □ □

Chapter 4

Marked Length Spectrum Rigidity

4.1 Introduction

Recall that 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 4.1.1 (Marked length spectrum)**.** The *marked length spectrum of a metric space* (M, g) is the class function

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

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 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 $\varphi : (M, g_0) \to (M, g_1)$ that is isotopic to the identity.

In this section, we equip a simple, thick negatively-curved surface amalgam (see Definition 2.0.1) X with a metric in a class we denote as $\mathcal{M}_{\leq}$, following the notation from [CL19b]. For the precise definition of $\mathcal{M}_{\leq}$, we refer the reader to section 4.2.1, but roughly speaking, metrics in $\mathcal{M}_{\leq}$ are CAT(-1) piecewise Riemannian metrics with an additional condition that limits pathological behavior around the gluing curves of X. We now precisely state the main result of the paper:

Theorem 4.1.2. *Suppose (X, g_1) and (X, g_2) are simple, thick negatively-curved surface amalgams 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 $\phi : (X, g_1) \to (X, g_2)$ 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 [FK65]). Recall that Burns and Katok then conjectured that closed negatively curved manifolds of all dimensions are marked length spectrum rigid in [BK85]. A major breakthrough occurred when Croke and Otal independently proved the conjecture in dimension two (see [Ota90], [Cro90]). 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 [CFF92]. In another direction, Hersonksy and Paulin extended Otal's result to the case of negatively curved metrics with finitely many cone singularities in [HP97]. 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 [DLR10], a result extended by Bankovic and Leininger in [BL18], 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 nonpositively curved metrics on surfaces with finitely many cone singularities in [Con18]. 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 [GM18]).

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 [Ham99]). 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 [CM87]) and Alperin and Bass (see [AB87]) leads to a marked length spectrum rigidity result for Culler-Vogtmann Outer Space, built to study the group $\mathrm{Out}(\mathbb{F}_n)$ in analogy to the relationship between the Teichmüller space of S, $\mathscr{T}(S)$, and the mapping class group, $\mathrm{Mod}(S)$. Generalizing the work from [CM87] and [AB87], Constantine and Lafont prove a version of marked length spectrum rigidity of compact, non-contractible one-dimensional spaces in [CL19a]. In [CL19b], they show that certain cocompact quotients of Fuchsian buildings, including those with piecewise hyperbolic metrics, are marked length spectrum rigid.

Outline of the Chapter. In Section 4.2, we define precisely the class of metrics $\mathcal{M}_{\leq}$ and review some fundamental facts about geodesic currents and CAT(-1) spaces. Furthermore, since the boundary of the universal cover of a surface amalgam is much more complicated than that of a surface, we also define an intersection number which works well in our setting. In Section 4.3, we patch together isometries constructed in [Ota90] to prove the theorem for surface amalgams 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 4.4, we prove the general case by constructing finite-sheeted covers and doubles of the surface amalgams to reduce to the base case examined on Section 4.3. We devote a large part of this section to showing the doubles of surface amalgams with the same marked length spectra will also have the same marked length spectra. In order to do so, we construct a Liouville current which completely defines lengths of closed geodesics by taking the product of the usual Liouville current on closed surfaces with a Markov probability measure to account for the branching behavior around the gluing geodesics. We also prove that intersection numbers completely determine a current, a fact whose proof requires careful treatment due to the fractal-like branching behavior of the universal cover.

4.2 Preliminaries

4.2.1 The class of metrics $\mathcal{M}_{\leq}$

We first precisely define the class of metrics $\mathcal{M}_{\leq}$ under consideration in Theorem 4.1.2 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 [HST20]);

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 $\epsilon > 0$. Furthermore, we require that the natural extension of g to the double of X (see Section 4.4.2) is also a negatively curved smooth Riemannian metric.

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 4.1.2 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 [Ota90] in the proof of Lemma 4.3.8, one would instead use Theorem C from [HP97] or Corollary 2 from [Con18]. 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.

We now discuss some properties of $\mathcal{M}_{\leq}$ that will be useful in the proof of Theorem 4.1.2.

Properties of $\mathcal{M}_{\leq}$

First, we remark that metrics in $\mathcal{M}_{\leq}$ are CAT(-1):

Remark 4.2.1. If (X, g) is a negatively-curved surface amalgam where $g \in \mathcal{M}_{\leq}$, then (X, g) is 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 CAT(κ) (see [BH99] Theorem 1A.6). As a result, the restriction of g to C is CAT(-1) since C is a 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 CAT(κ) spaces glued isometrically along a convex, complete metric subspace $A \subset X_1 \cap X_2$, then $X_1 \sqcup_A X_2$ is CAT(κ) (see Theorem 2.11.1 in [BH99]). As a result, since we have endowed each chamber in a surface amalgam X with a CAT(-1) metric, (X, g) will also be CAT(-1), which shows the remark.

The 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 4.3. Note that 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 4.3.

Remark 4.2.2. Another way to show that (X, g) where $g \in \mathcal{M}_{\leq}$ is Gromov hyperbolic is to observe that any surface amalgam (a metric space satisfying only Conditions 1 and 2 from Definition 2.0.1) 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 [BF92]).

4.2.2 The Universal Cover of a Simple, Thick Surface Amalgam

Next, we describe the universal cover of a simple, thick surface amalgam, which we will be working with extensively in Sections 4.3 and 4.4. Roughly speaking, $\widetilde{X}$ 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 $\widetilde{X}$ *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 $\widetilde{X}$ that project to gluing curves under the covering map *branching geodesics*.

While in general, $\widetilde{X}$ 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 $\widetilde{X}$ 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 $\widetilde{X}$, 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 4.1 for an illustration of $\widetilde{X}$.

Remark 4.2.3. Some readers may notice similarities between $\widetilde{X}$ and a two-dimensional

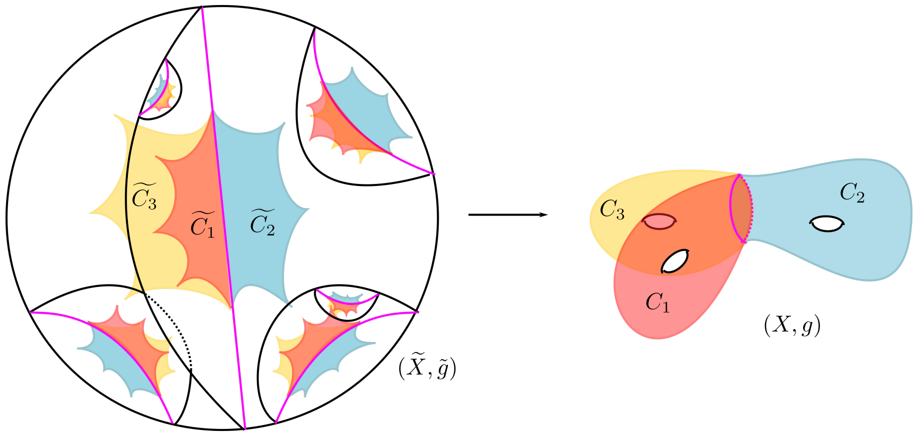

Figure 4.1: The universal cover of a simple, thick surface amalgam (X, g).

hyperbolic building. Even the terminology of surface amalgams 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 (see Definition 4.2.13), from [CL19b]. However, for buildings, boundaries of chambers are not totally geodesic in the space, so their walls are not totally geodesic. The proof methods from [CL19b] 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 $\widetilde{X}$. See Section 1.4.3 of [Laf02] for a more complete list of differences between buildings and surface amalgams.

4.2.3 Visual Boundaries

In this section, we will assume that $(\widetilde{X}, \widetilde{g})$ 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 [KB02]. We say that two geodesic rays in $\widetilde{X}$, $\alpha_1 : \mathbb{R}_{\geq 0} \to \widetilde{X}$, and $\alpha_2 : \mathbb{R}_{\geq 0} \to \widetilde{X}$ 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 $\partial^\infty(\widetilde{X}, \widetilde{g})$, fix any basepoint $x_0 \in (\widetilde{X}, \widetilde{g})$ 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 4.2.4 (Visual Boundary)**.** We define equivalence classes of such geodesic rays based at x_0 as the *visual boundary of* $\widetilde{X}$, $\partial^\infty(\widetilde{X}, \widetilde{g})$.

When $\widetilde{g}$ is CAT(-1), there is a unique geodesic in $(\widetilde{X}, \widetilde{g})$ between any two points in $\partial^\infty(\widetilde{X}, \widetilde{g})$ (see Proposition 1.4 of [BH99]), so the space of geodesics in $\widetilde{X}$ is identified with $\partial^\infty(\widetilde{X}, \widetilde{g}) \times \partial^\infty(\widetilde{X}, \widetilde{g}) \setminus \Delta$.

We define the cross ratio as follows:

Definition 4.2.5 (Cross ratio and Möbius maps)**.** If $(\widetilde{X}, \widetilde{g})$ has boundary $\partial^\infty(\widetilde{X}, \widetilde{g})$, one can define the *cross ratio* of a four-tuple of distinct boundary points $(\xi, \xi', \eta, \eta') \in (\partial^\infty(\widetilde{X}))^4$:

$$[\xi\xi'\eta\eta'] = \lim_{(a,a',b,b') \to (\xi,\xi',\eta,\eta')} \frac{1}{2}(\widetilde{g}(a,b) + \widetilde{g}(a',b') - \widetilde{g}(a,b') - \widetilde{g}(a',b)) \qquad (4.1)$$

See Figure 4.2 for an example in $\mathbb{H}^2$).

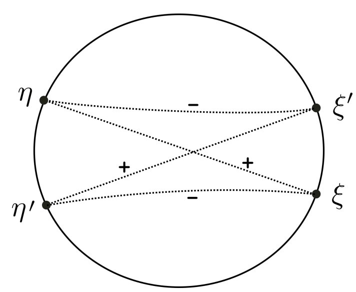

Figure 4.2: An illustration of the calculation of $[\xi\xi'\eta\eta']$ in $\mathbb{H}^2$.

A map $\partial^\infty f : \partial^\infty(\widetilde{X}, \widetilde{g_1}) \to \partial^\infty(\widetilde{X}, \widetilde{g_2})$ 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:

$$[abcd] = [cdab] \text{ and } -[abdc] = [abcd] = -[bacd] \qquad (4.2)$$

We are now ready to define the Gromov Product:

Definition 4.2.6 (Gromov Product)**.** Let $x \in (\widetilde{X}, \widetilde{g})$ and $a, b \in \partial^\infty(\widetilde{X})$. Then the *Gromov Product* of a and b with respect to a basepoint $x \in \widetilde{X}$ is:

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

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

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

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 $\partial^\infty(\widetilde{X}, \widetilde{g})$. Indeed, we can define a countable neighborhood base for any $p \in \partial^\infty(\widetilde{X})$ and $d \in \mathbb{N}$ as follows:

$$B(p, d) = \{q \in \partial^\infty(\widetilde{X}) | \langle p, q \rangle_x > d\} \tag{4.3}$$

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

Definition 4.2.7 (Visual metric)**.** Let $(\widetilde{X}, \widetilde{g})$ be a proper CAT(-1) space. Suppose $a, b \in \partial^\infty(\widetilde{X}, \widetilde{g})$. Then for some fixed basepoint $x \in (\widetilde{X}, \widetilde{g})$, we can assign $\partial^\infty(\widetilde{X}, \widetilde{g})$ a *visual metric*:

$$\widetilde{g}_{\infty,x}(a, b) = \begin{cases} e^{-\langle a,b \rangle_x} \text{ if } a \neq b \\ 0 \text{ otherwise} \end{cases} \tag{4.4}$$

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

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

The Visual Boundary of a Surface Amalgam

We now present some useful properties of the visual boundary of a simple, thick surface amalgam $(\widetilde{X}, \widetilde{g})$. As before, we will assume all the chambers in X are negatively curved.

Bass-Serre theory provides a useful characterization of points on $\partial^\infty(\widetilde{X})$. 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 [DST18] 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 [Mal10], which uses Bowditch's construction to characterize boundary points for *geometric amalgams of free groups* (e.g. $\pi_1(X)$). Given a point in $x \in \partial(\widetilde{X})$, Bowditch characterizes x by $\mathrm{val}(x)$, the number of topological ends of $\partial^\infty(\widetilde{X})\setminus\{x\}$. In particular, for a simple, thick surface amalgam with negatively-curved chambers, $\mathrm{val}(x) = n \geq 3$ if x is a branching geodesic attached to n distinct boundary components and $\mathrm{val}(x) = 2$ otherwise. Bowditch defines separate equivalence classes for $M(2) = \{x \in \partial^\infty(\widetilde{X}) : \mathrm{val}(x) = 2\}$ and $M(3+) = \{x \in \partial^\infty(\widetilde{X}) : \mathrm{val}(x) \geq 3\}$:

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

2. *For points in* $M(2)$: We say $x \sim y$ if $x = y$ or the number of connected components of $\partial^\infty(\widetilde{X}) \setminus \{x, y\}$ is equal to 2.

In both cases, equivalent pairs of points form a cut pair of $\partial^\infty(\widetilde{X})$. 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 $\sim$-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 $\partial^\infty(\widetilde{X}) \setminus \{a, b\}$, 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 $\partial^\infty(\widetilde{X})$:

Proposition 4.2.8. *[Proposition 1.3 of [Bow98] and Lemma 10 of [KK00]] If $\pi_1(X)$ acts on Bass-Serre tree T, then each point $x \in \partial^\infty(\widetilde{X})$ corresponds to exactly one of the following equivalence classes of points in $\partial^\infty(\widetilde{X})$:*

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 $\partial^\infty(T)$.*

In other words, $\approx$-classes in $M(3+)$ correspond to vertices in category (1) in Proposition 4.2.8 while vertices in category (2) correspond to $\sim$-classes of $M(2)$.

Using Proposition 4.2.8, we can topologically characterize points in $\partial^\infty(\widetilde{X})$. As mentioned earlier, points in category (1) correspond to endpoints of branching geodesics, lifts of gluing curves in $\partial^\infty(\widetilde{X})$. 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, $\partial^\infty(\widetilde{X}) \setminus \{x, y\}$ 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, $p : \widetilde{X} \to X$ is the universal covering map, and $\mathcal{BG}$ is the full set of branching geodesics in $\widetilde{X}$. 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 $\widetilde{X}$. Thus, the unique points in category (3) correspond to geodesic rays that intersect infinitely many branching geodesics in $\widetilde{X}$.

We now point out that points on $\partial^\infty(\widetilde{X})$ lie on boundaries of embedded disks, a fact stated in the proof of Lemma 3.1 in [Laf07b]. For the convenience of the reader, we provide a proof as well.

Lemma 4.2.9. *Given a geodesic $\gamma \in \widetilde{X}$ with endpoints $p, q \in \partial^\infty(\widetilde{X})$, there exists an isometrically embedded apartment containing γ. That is, p and q lie in an embedded $S^1 \subset \partial^\infty(\widetilde{X})$. As a consequence, the boundary $\partial^\infty(\widetilde{X})$ 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 $\widetilde{X}$, 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 $\widetilde{X}$, 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 $\widetilde{X}$. 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 $\widetilde{X}$. 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 $\partial^\infty(\widetilde{X})$. 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 4.3).

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 $\partial^\infty(\widetilde{X})$ is path-connected, we refer the reader to [Laf07b].

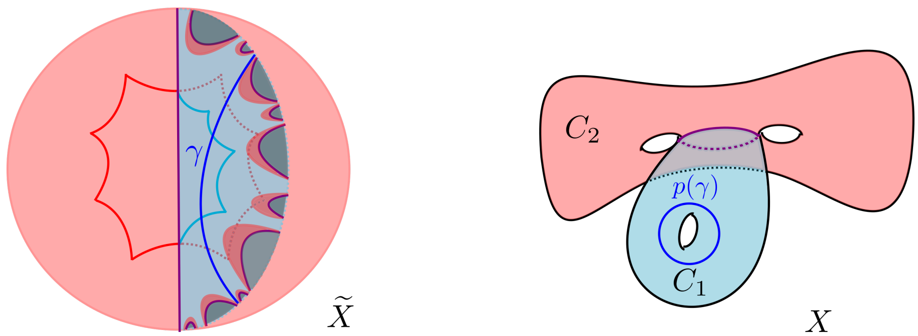

Figure 4.3: An example of how to choose an apartment containing γ. In this example, $p(\gamma)$ is completely contained in the torus with one boundary component C_1, which lifts to a totally geodesic subspace of a half-disk. We can fill in the missing portions of the half-disk (shaded in gray) with portions 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.

□

Lemma 4.2.9 allows us to define intervals in $\partial^\infty(\widetilde{X})$:

Definition 4.2.10 (Interval on $\partial^\infty(\widetilde{X})$)**.** Suppose X is a simple, thick negatively-curved surface amalgam. An *interval* on $\partial^\infty(\widetilde{X})$ with endpoints $a, b \in \partial^\infty(\widetilde{X})$ is an interval on the boundary of an apartment containing the unique geodesic in $\widetilde{X}$ with endpoints a and b.

4.2.4 Geodesic Currents

We now briefly introduce geodesic currents, originally defined by Bonahon (see [Bon86] and [Bon88]), and some important facts about them necessary for the rest of the paper. For more details on the theory of geodesic currents, we refer the reader to [AL17] and [ES22]).

Suppose now that $X = \widetilde{X}/\Gamma$. In our setting, $\Gamma = \pi_1(X)$ (where again, X is proper and CAT(-1) or Gromov hyperbolic). Let $\mathscr{G}(\widetilde{X})$ be the space of unpointed, unoriented geodesics of $\widetilde{X}$ endowed with the topology of the Hausdorff distance on compact sets of $\widetilde{X}$. Note that $\mathscr{G}(\widetilde{X})$ is identified with $(\partial^\infty(\widetilde{X}) \setminus \Delta)^2/[\xi,\eta] \sim [\eta,\xi]$, and that Γ acts naturally on $\mathscr{G}(\widetilde{X})$.

Definition 4.2.11 (Geodesic current)**.** A *geodesic current* on X is a locally finite, positive regular Borel measure on $\mathscr{G}(\widetilde{X})$ that is invariant under the action of Γ.

We will use $\mathscr{C}(X)$ to denote the space of currents on X equipped with the weak convergence topology. One key example of a geodesic current is a *counting current* associated with each primitive $\alpha \in \pi_1(X)$. Throughout the paper, following convention, we will abuse notation by using α to also denote the counting current associated with α. Given a measurable $E \subset \mathscr{G}(X)$, we measure E with α by simply counting the number of points of intersection of E with $\{\widetilde{\alpha}\}$, the set of all lifts of α in $\widetilde{X}$. In other words:

$$\alpha(E) = |E \cap \{\widetilde{\alpha}\}|$$

In the case of surfaces, counting currents are especially important because they give important data for determining a geodesic current via intersection numbers, which we will define next.

Transversality and Intersection Numbers

Recall that for a surface S, there is a bilinear, symmetric, continuous function, the *intersection number* $i : \mathscr{C}(S) \times \mathscr{C}(S) \to \mathbb{R}_{\geq 0}$ defined by:

$$i(\mu,\nu) = \int_{D(\mathscr{G}(\widetilde{S}))/\Gamma} d\mu d\nu$$

where $D(\mathscr{G}(\widetilde{S}))/\Gamma$ denotes pairs of unoriented, unparametrized bi-infinite geodesics that intersect transversely in $\widetilde{S}$ (see [Bon86]). Transversality of geodesics is well-defined in $\widetilde{S}$

since $\partial^\infty \widetilde{S}$ is a circle; namely, two geodesics $\xi = (x_1, y_1)$ and $\eta = (x_2, y_2)$ are *transverse* or *linked* if x_2 and y_2 lie on different connected components of $\partial^\infty(\widetilde{S}) \setminus \{x_1, y_1\}$. For the universal cover of a surface amalgam X, we use Definition 6.4 of [CL19b] to define transversality of geodesics in $\widetilde{X}$.

Definition 4.2.12 (Transversality of geodesics for Surface Amalgams)**.** We say that ξ and η, two geodesics in $\widetilde{X}$, are *transverse* if there exists some apartment $A \subset \widetilde{X}$ such that ξ and η are transverse (linked) in A.

Note that transversality is well-defined for an apartment A since apartments are isometrically embedded disks in $\widetilde{X}$.

As before, let $D(\mathscr{G}(\widetilde{X})) \subset \mathscr{G}(\widetilde{X}) \times \mathscr{G}(\widetilde{X})$ denote the pairs of bi-infinite, unoriented, unparametrized geodesics that intersect transversely in $\widetilde{X}$. In the case of closed surfaces, intersection numbers recover information about the lengths of closed geodesics; we want a similar result for surface amalgams. As a result, following [CL19b], we modify the intersection number for surface amalgams by defining the function:

$$\varpi : D(\mathscr{G}(\widetilde{X})) \to \mathbb{R}_{\geq 0}, \varpi(\xi, \eta) = \begin{cases} \prod_{w \in \mathcal{W}(\xi,\eta)} n(w) - 1 \text{ if } \mathcal{W}(\xi, \eta) \neq \varnothing \\ 1 \text{ otherwise} \end{cases}$$

where $\mathcal{W}(\xi, \eta)$ denotes the set of branching geodesics intersecting ξ and η transversely at two distinct points (in particular, if $\xi \cap \eta$ lies on a branching geodesic γ, then $\gamma \notin \mathcal{W}(\xi, \eta)$), and $n(w)$ denotes the number of chambers adjacent to $w \in \mathcal{W}(\xi, \eta)$. See Figure 4.4 for an example of branching geodesics in $\mathcal{W}(\xi, \eta)$.

Constantine and Lafont define $\mathcal{W}(\xi, \eta)$ to be the set of *walls*, combinatorial paths created from sides of polygons that tessellate disks (also called "apartments") in a piecewise negatively curved Riemannian Fuchsian building, that transversely intersect $(\xi, \eta) \in D(\mathscr{G}(\widetilde{X})$. They prove that given any transverse (ξ, η) in a Fuchsian building X equipped with a negatively curved metric, $\mathcal{W}(\xi, \eta)$ is finite (see Lemma 7.4 of [CL19b]) and the function ϖ is a π_1-invariant measurable function on $D(\mathscr{G}(\widetilde{X}))$ (see Proposition 7.7 and

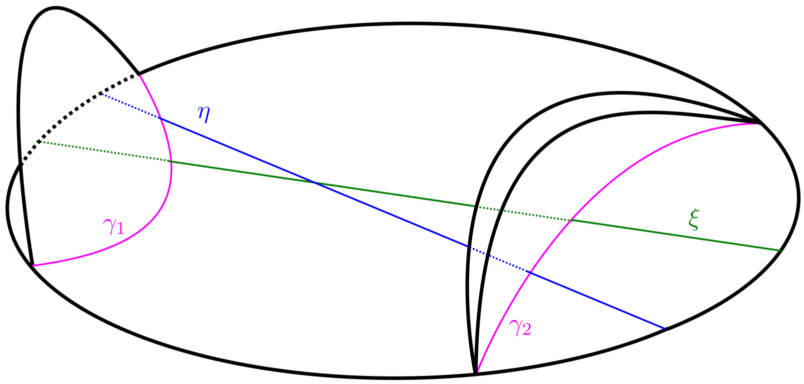

Figure 4.4: An example of transverse ξ and η. Note that $\gamma_1, \gamma_2 \in \mathcal{W}(\xi, \eta)$ since they intersect both ξ and η transversely. If there are no other geodesics in $\mathcal{W}(\xi, \eta)$, then $\varpi(\xi, \eta) = (3-1)(4-1) = 6$.

Lemma 7.8 of [CL19b]). Although the vocabulary is different, as Fuchsian buildings have walls instead of disjoint branching geodesics, the proof ideas also hold in the setting of surface amalgams. We can thus define intersection numbers for surface amalgams similarly to the way they were defined for Fuchsian buildings:

Definition 4.2.13 (Intersection numbers for Surface Amalgams)**.** The intersection number $i(\mu, \nu)$ associated with two geodesic currents μ and ν on a surface amalgam is a bilinear, symmetric function $i : \mathscr{C}(X) \times \mathscr{C}(X) \to \mathbb{R}$ defined by:

$$i(\mu, \nu) = \int_{\left((\xi,\eta)\in D\left(\mathscr{G}(\widetilde{X})\right)\right)/\Gamma} \varpi(\xi, \eta) d\mu d\nu$$

Remark 4.2.14. The existence of a continuous intersection number suitable for our setting is not needed for the proofs in this paper, but is nonetheless an interesting open question. Constantine and Lafont show in Proposition 10.1 of [CL19b] that in the special case that ν is the Liouville current (see Definition 4.4.5), $i(\cdot, L_g)$ is continuous with respect to the weak-$*$ topology on $\mathscr{C}(X)$.

4.3 The Base Case

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

Proposition 4.3.1. *Suppose X is a simple surface amalgam 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 $S \subset X$. Then if (X, g_1) and (X, g_2) have the same marked length spectrum, then there exists an isometry $\phi : (X, g_1) \to (X, g_2)$ that is induced by the boundary homeomorphism $\partial^\infty f : \partial^\infty(\widetilde{X}, \widetilde{g_1}) \to \partial^\infty(\widetilde{X}, \widetilde{g_2})$ discussed above.*

We first provide an outline of the proof. First, in the proof of Proposition 4.3.2 of Section 4.3.1, we show that the boundary homeomorphism $\partial^\infty f : \partial^\infty(\widetilde{X}, \widetilde{g_1}) \to \partial^\infty(\widetilde{X}, \widetilde{g_2})$ induced by the identity map is Möbius (see Definition 4.2.5). The proof of Proposition 4.3.2 requires the ergodicity of the geodesic flow map, which can be deduced from general theory developed by Kaimanovich in [Kai94] for Gromov hyperbolic spaces (see the Appendix). Using ergodicity of the geodesic flow map, the proof of Lemma 4.3.5 shows the cross ratio of any distinct $(a, b, c, d) \in (\partial^\infty(\widetilde{X}))^4$ 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 4.3.2, we patch together isometries $\phi_S : (S, g_1|_S) \to (S, g_2|_S)$ constructed by Otal in [Ota90] on every closed subsurface $S \subset X$ to construct a global isometry $\phi : (X, g_1) \to (X, g_2)$. 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.

4.3.1 The boundary map is Möbius

In this section, we prove the following proposition by adapting Otal's proof of Theorem 1 in [Ota90]. Note that in this section, the metrics are only assumed to be CAT(-1).

Proposition 4.3.2. *Suppose X is a simple, thick CAT(-1) surface amalgam endowed with a pair of metrics g_1 and g_2. 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 surface amalgams, 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 $\gamma : I \to X$ has *speed* $s \geq 0$ if for every $t \in I$, there exists a neighborhood $U \subset I$ 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 4.3.3 (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 $\mathscr{G}(X) \times \mathbb{R}$, where as before, $\mathscr{G}(X)$ 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 4.3.4 (Geodesic flow)**.** The *geodesic flow* on SX is the map $\phi_t : SX \to SX, \phi_t(\xi)(s) = \xi(t + s)$, 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 surface amalgam, 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 surface amalgams, the usual proof of the ergodicity of the geodesic flow map on surfaces does not generalize to the setting of surface amalgams 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 surface amalgams equipped with CAT(-1) metrics are proper Gromov hyperbolic spaces.

Next, we introduce a lemma whose proof is adapted from Otal's proof of Theorem 1 in [Ota92]:

Lemma 4.3.5. *Suppose X is simple, thick surface amalgam endowed with a CAT(-1) metric. Then given a 4-tuple of distinct points $(a, b, c, d) \in (\partial^\infty(\widetilde{X}))^4$, $[abcd]$ can be approximated arbitrarily well by lengths of closed geodesics.*

Proof. Fix a 4-tuple of distinct points $(a, b, c, d) \in (\partial^\infty(\widetilde{X}, \widetilde{g}))^4$. By the Birkhoff Ergodic Theorem (see Section A.0.3), we can conclude that there exists $v \in SX$ with dense orbit under the geodesic flow map ϕ_t on SX. Then there exist $n_i \in \mathbb{N}$, sufficiently large numbers so the basepoints of $\phi_{n_1}(\widetilde{v_1}), \phi_{-n_2}(\widetilde{v_1}), \phi_{n_3}(\widetilde{v_2})$, and $\phi_{-n_4}(\widetilde{v_2})$ are arbitrarily close to a, b, c, and d respectively (here, $\widetilde{v_1}$ and $\widetilde{v_2}$ are lifts of v in $S\widetilde{X}$). In other words, $[\phi_{n_1}(\widetilde{v_1})(0)\phi_{-n_2}(\widetilde{v_1})(0)\phi_{n_3}(\widetilde{v_2})(0)\phi_{-n_4}(\widetilde{v_2})(0)]$ approximates $[abcd]$ arbitrarily well.

Fix any $u \in S\widetilde{X}$. Then since the orbit of v is dense in $SX = S\widetilde{X}/\Gamma$, for every sufficiently large $m_1 \in \mathbb{N}$ there exists some $\gamma_1 \in \Gamma$ such that $\gamma_1 \cdot \phi_{m_1}(\widetilde{v_1})$ lies arbitrarily close to $\gamma_1 \cdot u$ in $S\widetilde{X}/\Gamma$. Similarly, there exists some $\gamma_3 \in \Gamma$ and sufficiently large $m_3 \in \mathbb{N}$ such that $\gamma_3 \cdot \phi_{m_3}(\widetilde{v_2})$ and $\gamma_3 \cdot u$ are arbitrarily close. In the end, we can choose $l_i >> m_i, n_i$ $(i = 1, 3)$ and corresponding $\gamma_1, \gamma_3 \in \Gamma$ with long translation axes such that $\gamma_1 \cdot \phi_{l_1}(\widetilde{v_1})$ lies arbitrarily close to both $\gamma_1 \cdot u$ and a and $\gamma_3 \cdot \phi_{l_3}(\widetilde{v_2})$ lies arbitrarily close to both $\gamma_3 \cdot u$ and c (see Figure 4.5).

Let $\ell_{\gamma_3^{-1}\gamma_1}$ be the length of the fundamental domain of the action of $\gamma_3^{-1} \circ \gamma_1$ on its translation axis. Note that $\ell_{\gamma_3^{-1}\gamma_1}$, the length of a closed geodesic corresponding to $\gamma_3^{-1} \circ \gamma_1$, is approximated arbitrarily well by $\widetilde{g}(\gamma_3 \cdot u, \gamma_1 \cdot u)$ since by construction, $\gamma_3 \cdot u$ and $\gamma_1 \cdot u$ lie arbitrarily close to endpoints of a geodesic $(a, c) \in (\widetilde{X}, \widetilde{g})$ and $\gamma_1 \cdot u = (\gamma_1 \circ \gamma_3^{-1})(\gamma_3 \cdot u)$. Additionally, $\widetilde{g}(\gamma_3 \cdot u, \gamma_1 \cdot u)$ approximates $\widetilde{g}(a, c)$ arbitrarily well, so $\ell_{\gamma_3^{-1}\gamma_1}$ approximates $\widetilde{g}(a, c)$ arbitrarily well. We apply a similar argument to approximate the other lengths by

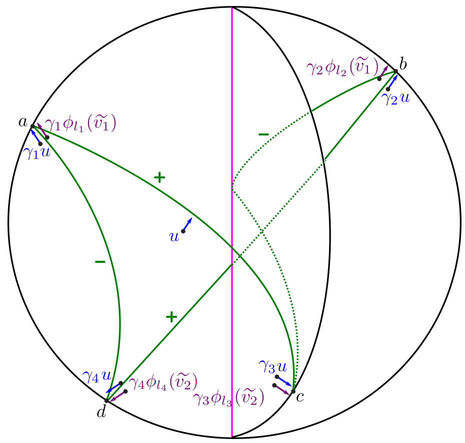

Figure 4.5: An illustration of the setup of Lemma 4.3.5.

lengths of translation axes (and thus closed geodesics in X) to obtain the desired result.

□

We can now prove Proposition 4.3.2. By Lemma 4.3.5, 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:

$$|[abcd] - (\widetilde{g_1}(x_i, y_i) + \widetilde{g_1}(x_i', y_i') - \widetilde{g_1}(x_i, y_i') - \widetilde{g_1}(x_i', y_i))| < \frac{1}{i}$$

By construction, (x_i, y_i, x_i', y_i') converges to (a, b, c, d) in $(\widetilde{X}, \widetilde{g_1})$ 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 $(\widetilde{X}, \widetilde{g_2})$ since f is continuous. Thus, by definition of cross ratio, we have that:

$$[abcd] = \lim_{i\to\infty} \widetilde{g_1}(x_i, y_i) + \widetilde{g_1}(x_i', y_i') - \widetilde{g_1}(x_i, y_i') - \widetilde{g_1}(x_i', y_i)$$

and

$$\lim_{i\to\infty} \widetilde{g_2}(f(x_i, y_i)) + \widetilde{g_2}(f(x_i', y_i')) - \widetilde{g_2}(f(x_i, y_i')) - \widetilde{g_2}(f(x_i', y_i)) = [\partial^\infty f(a)\partial^\infty f(b)\partial^\infty f(c)\partial^\infty f(d)]$$

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, $\widetilde{g_1}(x_i, y_i) = \ell_{g_1}(\alpha_i) = \ell_{g_2}(\alpha_i) = \widetilde{g_2}(f(x_i, y_i))$, and a similar statement holds for the other distances as well. It then follows that $[abcd] = [\partial^\infty f(a)\partial^\infty f(b)\partial^\infty f(c)\partial^\infty f(d)]$. □

4.3.2 Proof of Proposition 4.3.1

In order to prove Proposition 4.3.1, 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 4.3.6. *Suppose U_i and V_i are complete, convex, locally compact metric spaces, where $i = 1, 2$. Suppose $\phi_i : U_i \to V_i$ are (invertible) isometries, and $\phi_1|_{U_1 \cap U_2} = \phi_2|_{U_1 \cap U_2}$. Then there exists an isometry $\phi : U_1 \cup U_2 \to V_1 \cup V_2$ such that $\phi|_{U_i} = \phi_i|_{U_i}$ for $i = 1, 2$.*

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 4.6).

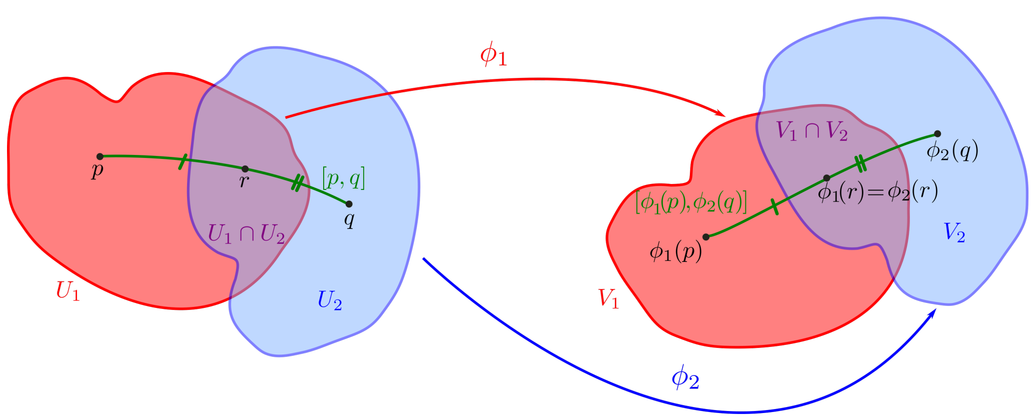

Figure 4.6: An illustration of the setup of Lemma 4.3.6.

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 $i = 1, 2$,

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

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

The following is a technical lemma about the cross-ratio defined in Equation 4.1.

Lemma 4.3.7. *Suppose $(\widetilde{S}, \widetilde{g}|_{\widetilde{S}})$ and $(\widetilde{S}', \widetilde{g}|_{\widetilde{S}'})$ are two lifts of $(S, g|_S)$ and $(S', g|_{S'})$ respectively that meet at a branching geodesic $\widetilde{\gamma} \subset (\widetilde{X}, \widetilde{g})$. Then for every $p \in \widetilde{\gamma}$, 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 $p \in \partial^\infty(\widetilde{X})$. 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 $(a, c) \subset (\widetilde{S}, \widetilde{g}|_{\widetilde{S}})$ that transversely intersects $\widetilde{\gamma}$ at p in $\widetilde{S}$. Suppose $\widetilde{\gamma}$ divides $\widetilde{S}'$ into two half planes, which we will call H_1' and H_2'. Consider a map $h : \partial^\infty(H_1') \to \widetilde{\gamma}$ that maps a point $x \in \partial^\infty(H_1')$ to the point $h(x) = \widetilde{\gamma} \cap (x, c)$ where (x, c) denotes the geodesic arc with endpoints x and c. Note that h extends continuously to a map $\widetilde{h} : \partial^\infty(H_1') \to \widetilde{\gamma} \cup \{q, r\}$, where $q, r \in \partial^\infty(\widetilde{S}) \cap \partial^\infty(\widetilde{S}')$ are endpoints of $\widetilde{\gamma}$ in the following way: $\widetilde{h}(q) = q$ and $\widetilde{h}(r) = r$. As we vary x along $\partial^\infty(H_1')$, $h(x)$ varies continuously along $\widetilde{\gamma}$, so by an Intermediate Value Theorem argument, since

$\widetilde{\gamma}$ and $\partial^\infty(H_1')$ are both connected, it follows that $\widetilde{h}(x) = p$ for some $x \in \partial^\infty(H_1')$ (see Figure 4.7). We set b equal to this x.

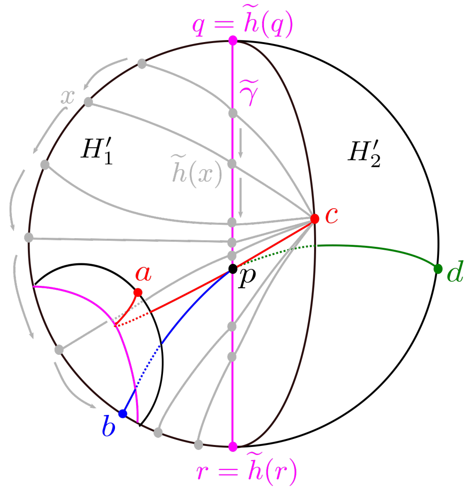

Figure 4.7: By an intermediate value theorem argument, there exists some $x \in \partial^\infty H_1'$ such that $\widetilde{h}(x) = p$ for a given point $p \in \widetilde{\gamma}$.

Note that (b, c) is a geodesic that passes through p. Continue the geodesic ray (b, p) into $H_2' \subset \widetilde{S}'$ to obtain a bi-infinite geodesic $(b, d) \subset \widetilde{S}'$. 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) such that:

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

□

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 4.3.8. *Suppose $S, S' \in \mathcal{S}$ are two closed surfaces in X that are identified along some set of gluing curves $\{\gamma_1, ..., \gamma_n\}$. Then there exist isometries $\phi_S : (S, g_1|_S) \to (S, g_2|_S)$*

and $\phi_{S'} : (S', g_1|_{S'}) \to (S', g_2|_{S'})$ *where* $\phi_S|_{\gamma_i} = \phi_{S'}|_{\gamma_i}$ *for all* $1 \le i \le n$.

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 $\partial^\infty(\widetilde{X}, \widetilde{g_1})$ and $\partial^\infty(\widetilde{X}, \widetilde{g_2})$.

Suppose $\widetilde{S}$ and $\widetilde{S'}$ are apartments that are arbitrary lifts of S and S' in $\widetilde{X}$ such that $\widetilde{S} \cap \widetilde{S'} \neq \varnothing$. We can then use the restrictions of the boundary homeomorphism $\partial^\infty f|_{\partial^\infty(\widetilde{S})}$ and $\partial^\infty f|_{\partial^\infty(\widetilde{S'})}$ to construct isometries $\phi_S : (S, g_1|_S) \to (S, g_2|_S)$ and $\phi_{S'} : (S', g_1|_{S'}) \to (S', g_2|_{S'})$ 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 [Ota90], 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 $x = \lim_{i \to \infty} x_i$. 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 $a_i, b_i, c_i, d_i \to x$, 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 4.8). Then since $\partial^\infty f$ is continuous, $\lim_{i \to \infty} \partial^\infty f(a_i) = \partial^\infty f(x)$, and the same is true for b_i, c_i, and d_i. Then:

$$\lim_{i \to \infty} \widetilde{\phi_S}(x_i) = \lim_{i \to \infty} \widetilde{\phi_S}((a_i, b_i) \cap (c_i, d_i)) = \lim_{i \to \infty} (\partial^\infty f(a_i), \partial^\infty f(b_i)) \cap (\partial^\infty f(c_i), \partial^\infty f(d_i)) = \partial^\infty f(x)$$

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:

$$\partial^\infty \widetilde{\phi_S}(p) = \partial^\infty f|_{\widetilde{S} \cap \widetilde{S'}}(p) = \partial^\infty \widetilde{\phi_{S'}}(p) \tag{4.5}$$

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 4.5, $\widetilde{\phi_S}(\widetilde{\gamma})$ 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 points in the CAT(-1)

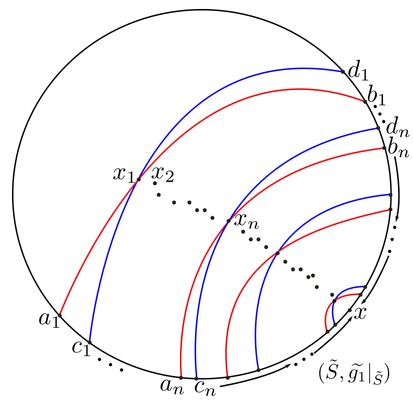

Figure 4.8: Given a sequence of points $\{x_i\}_{i=1}^n$ converging radially to x, one can find a sequence of geodesics ξ_i and η_i with endpoints a_i, b_i, c_i, d_i converging to x. The images of these endpoints under the boundary homeomorphism $\partial^\infty f$ will also converge to $\partial^\infty f(x) = \lim_{i\to\infty} \widetilde{\phi_S}(x_i)$.

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 4.3.2, if an isomorphism of fundamental groups preserves the marked length spectrum of two CAT(-1) spaces, then the induced map at infinity is Mobius. As a consequence, $\partial^\infty f$ is Mobius. By Lemma 4.3.7, 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 4.9). 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 $a_i, b_i \in \xi$, $a'_i, b'_i \in \xi'$, $(a_i, b_i) \to (a, b)$ and $(a'_i, b'_i) \to (a', b')$. The Parallelogram Law for CAT(κ) spaces (see Exercise 1.16 of [BH99]) gives us the following inequality:

$$\widetilde{g_2}(\widetilde{\phi_S}(a'_i), \widetilde{\phi_S}(b_i)) + \widetilde{g_2}(\widetilde{\phi_S}(a_i), \widetilde{\phi_S}(b'_i)) \neq \widetilde{g_2}(\widetilde{\phi_S}(a_i), \widetilde{\phi_S}(b_i)) + \widetilde{g_2}(\widetilde{\phi_{S'}}(a'_i), \widetilde{\phi_{S'}}(b'_i)) \quad (4.6)$$

We therefore conclude that:

$$
\begin{aligned}
[\partial^\infty f(a)\partial^\infty f(a')\partial^\infty f(b)\partial^\infty f(b')] &= \lim_{i\to\infty}[\widetilde{\phi_S}(a_i)\widetilde{\phi_{S'}}(a_i')\widetilde{\phi_S}(b_i)\widetilde{\phi_{S'}}(b_i')] \\
&= \lim_{i\to\infty}\widetilde{g_2}(\widetilde{\phi_S}(a_i),\widetilde{\phi_S}(b_i)) + \widetilde{g_2}(\widetilde{\phi_{S'}}(a_i'),\widetilde{\phi_{S'}}(b_i')) \\
&- \widetilde{g_2}(\widetilde{\phi_S}(a_i),\widetilde{\phi_{S'}}(b_i')) - \widetilde{g_2}(\widetilde{\phi_S}(b_i),\widetilde{\phi_{S'}}(a_i')) \\
&\overset{(4.6)}{\neq} 0 = [aa'bb']
\end{aligned}
$$

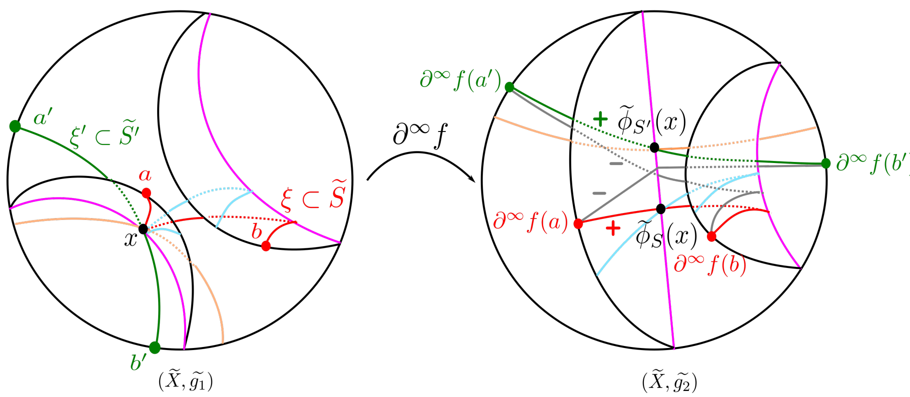

Figure 4.9: If $\widetilde{\phi}_S(x) \neq \widetilde{\phi}_{S'}(x)$, then $[\partial^\infty f(a)\partial^\infty f(a)\partial^\infty f(b_i)\partial^\infty(b_i')] \neq 0$, but a, a', b, and b' were chosen so that $[aa'bb'] = 0$, a contradiction since $\partial^\infty f$ is Mobius. As before, the pink geodesics are branching.

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. □

Lemma 4.3.9. *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, $\phi_S : (S, g_1|_S) \to (S, g_2|_S)$ and $\phi_{S'} : (S', g_1|_{S'}) \to (S', g_2|_{S'})$ are the isometries constructed by Otal.*

Proof. By Lemma 4.3.8, 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 $\widetilde{x}$ 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 $\widetilde{\gamma}_1$, a lift of a gluing curves $\gamma \in S \cap S'$.

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

In order to construct ξ, consider a geodesic arc $\alpha = [\widetilde{x}, \widetilde{\gamma_1}(a)]$ joining $\widetilde{x}$ to some point $\widetilde{\gamma_1}(a)$ on $\widetilde{\gamma_1} \subset \widetilde{X}$. Furthermore, we require that $\widetilde{\gamma_1}(a)$ is chosen so that α lies entirely inside $\widetilde{C_i}$. To find ξ, extend α in $\widetilde{S}$ via the exponential map to a geodesic lying entirely inside a single lift of S, $\widetilde{S}$, and construct $\xi' \subset \widetilde{S}'$ similarly. Along the same vein, construct η by extending $\beta = [\widetilde{x}, \widetilde{\gamma_1}(b)]$ on $\widetilde{S}$, where $\widetilde{\gamma_1}(b)$ is another point on the image of $\widetilde{\gamma_1}$ and β also lies entirely inside $\widetilde{C_i}$. Similarly, we can extend β in $\widetilde{S}'$ to obtain η'. See Figure 4.10 for an example of a choice of α and β given some $\widetilde{x}$.

We thus have two pairs of geodesics $\xi, \eta \subset \widetilde{S}$, and $\xi', \eta' \subset \widetilde{S}'$. 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 $[\widetilde{\gamma_2}(c), \widetilde{\gamma_1}(a)]$ (where $\widetilde{\gamma_2}$ is a lift of a gluing curve in $S \cap S'$) or a geodesic ray with endpoint $\widetilde{\gamma_1}(a)$ and limit point $y \in \partial^\infty(p^{-1}(\{C_i\}))$. Similarly, η and η' agree either on a geodesic arc $[\widetilde{\gamma_3}(d), \widetilde{\gamma_1}(b)]$ or a geodesic ray with endpoint $\widetilde{\gamma_1}(b)$ and limit point $z \in \partial^\infty(p^{-1}(\{C_i\}))$. For example, in Figure 4.10, $\xi \subset \widetilde{S}$ and $\xi' \subset \widetilde{S}'$ agree on a geodesic arc $[\widetilde{\gamma_1}(a), \widetilde{\gamma_2}(c)]$ while $\eta \subset \widetilde{S}$ and $\eta' \subset \widetilde{S}'$ agree on a geodesic ray $[\widetilde{\gamma_1}(b), y_2)$.

Suppose that ξ and ξ' agree on a geodesic arc $[\widetilde{\gamma_2}(c), \widetilde{\gamma_1}(a)] \subset p^{-1}(\{C_i\})$. By Proposition 4.3.8, we know that we can define points $x_1 := \widetilde{\phi_S}(\widetilde{\gamma_1}(a)) = \widetilde{\phi_{S'}}(\widetilde{\gamma_1}(a))$ and $x_2 := \widetilde{\phi_S}(\widetilde{\gamma_2}(c)) = \widetilde{\phi_{S'}}(\widetilde{\gamma_2}(c))$. 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,

$$\widetilde{\phi_S}(\xi \cap \xi') = \widetilde{\phi_S}([\widetilde{\gamma_1}(a), \widetilde{\gamma_2}(c)]) = [x_1, x_2] = \widetilde{\phi_{S'}}([\widetilde{\gamma_1}(a), \widetilde{\gamma_2}(c)]) = \widetilde{\phi_{S'}}(\xi \cap \xi')$$

Suppose on the other hand that ξ and ξ' agree on a geodesic ray $[\widetilde{\gamma_1}(a), y)$. Again, we can set $x_1 := \widetilde{\phi_S}(\widetilde{\gamma_1}(a)) = \widetilde{\phi_{S'}}(\widetilde{\gamma_1}(a))$. 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 4.3.8, we determined that $\partial^\infty \widetilde{\phi_S}|_{\widetilde{S \cap S'}} = \partial^\infty f|_{\widetilde{S \cap S'}} =$

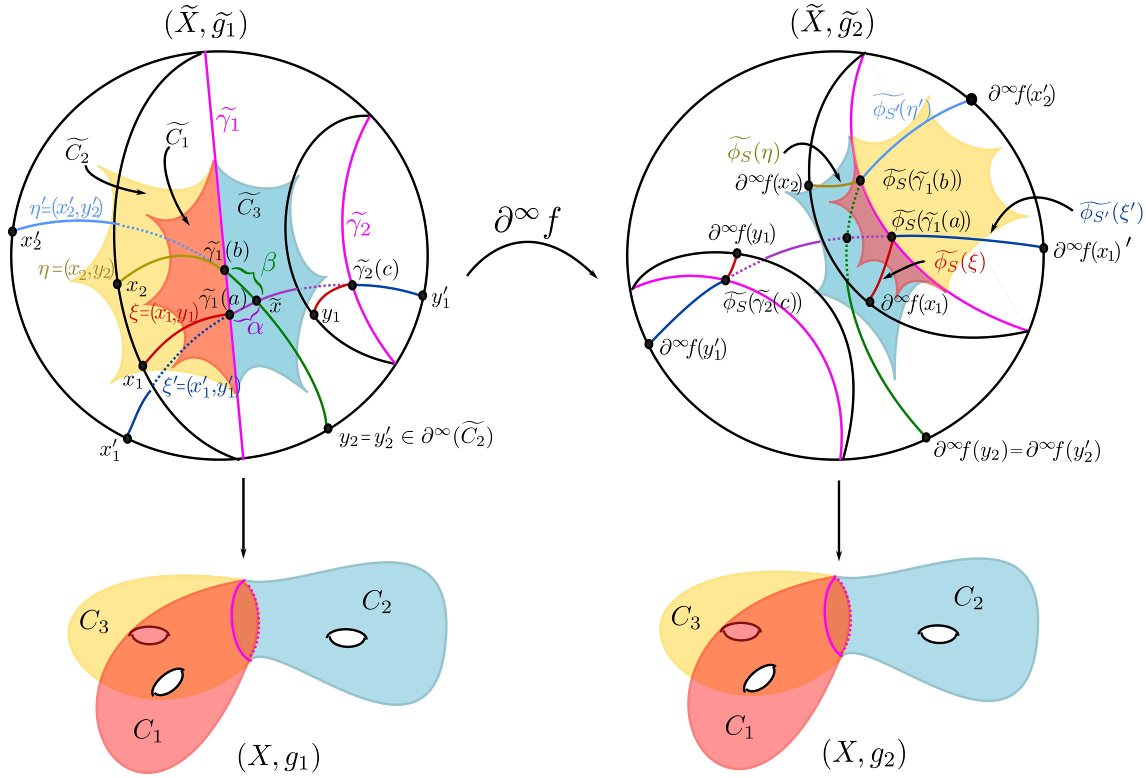

Figure 4.10: We extend $\alpha = (\widetilde{x}, \widetilde{\gamma_1}(a))$ in the picture on the left to $\xi \subset \widetilde{S}$ and $\xi' \subset \widetilde{S}'$. Similarly, we extend $\beta = (\widetilde{x}, \widetilde{\gamma_1}(b))$ where $b \neq a$ to $\eta \subset \widetilde{S}$ and $\eta' \subset \widetilde{S}'$. While ξ and ξ' agree on a geodesic arc $[\widetilde{\gamma_2}(c), \widetilde{\gamma_1}(a)]$, η and η' agree on a geodesic ray $[\widetilde{\gamma_1}(b), y)$.

$\partial^\infty \widetilde{\phi_{S'}}|_{\widetilde{S \cap S'}}$. 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,

$$\widetilde{\phi_S}(\xi \cap \xi') = \widetilde{\phi_S}([\widetilde{\gamma_1}(a), y)) = [x_1, y_1) = \widetilde{\phi_{S'}}([\widetilde{\gamma_1}(a), y)) = \widetilde{\phi_{S'}}(\xi \cap \xi')$$

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 $\widetilde{x} \in \widetilde{\phi_S}(\widetilde{S} \cap \widetilde{S}') = \widetilde{\phi_{S'}}(\widetilde{S} \cap \widetilde{S}')$ by doing some case work:

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 $p_1 = \widetilde{\gamma_1}(a)$ and $p_2 = \widetilde{\gamma_1}(b)$ lie on a branching geodesic $\widetilde{\gamma_1}$ and $r, s \in \partial^\infty(p^{-1}(\{C_i\}))$. Additionally, suppose $\widetilde{\gamma_1}$ has endpoints p and q, $p = \lim_{t \to -\infty} \widetilde{\gamma_1}(t)$,

and $q = \lim_{t\to\infty} \widetilde{\gamma}_1(t)$. Consider an apartment A in $\widetilde{X}$ that contains $\widetilde{\gamma}_1$, r, and s (e.g. $\widetilde{S}$ or $\widetilde{S}'$). 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 $\widetilde{\gamma}_1$ will be preserved under $\partial^\infty f|_A$, it follows that the geodesic rays $[\widetilde{\phi}(p_2), \partial^\infty f(s)) = [\widetilde{\phi}(p_2), \partial^\infty \widetilde{\phi}(s))$ and $[\widetilde{\phi}(p_1), \partial^\infty f(r)) = [\widetilde{\phi}(p_1), \partial^\infty \widetilde{\phi}(r))$ will necessarily intersect (see Figure 4.11).

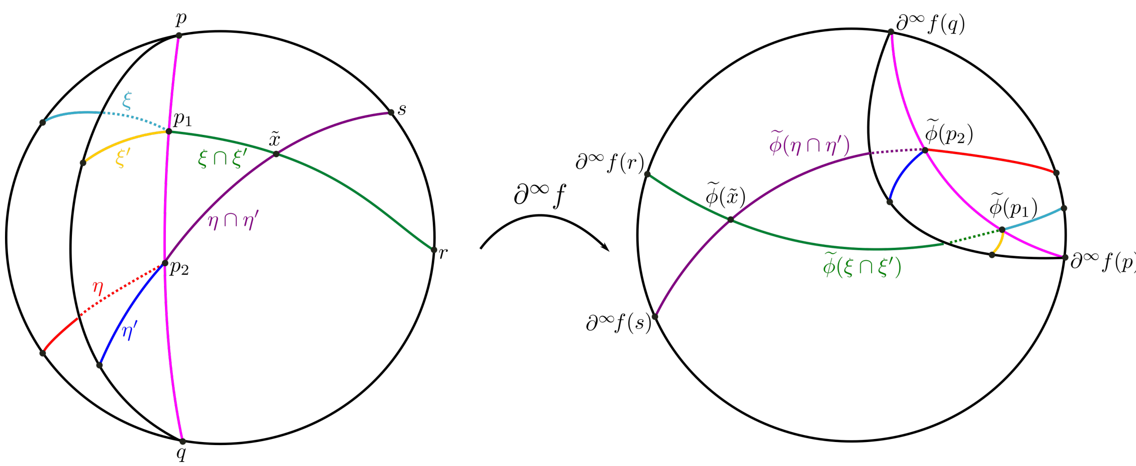

Figure 4.11: An example where $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic rays.

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 $\widetilde{\gamma}_1(a) = p_1 \in \xi \cap \xi'$ and $\widetilde{\gamma}_1(b) = p_2 \in \eta \cap \eta'$ where $a < b$. Additionally, suppose that $\widetilde{\gamma}_2(c) = p_3$ and $\widetilde{\gamma}_2(d) = p_4$ where $d < c$, and $\lim_{t\to-\infty} \widetilde{\gamma}_2(t) = s$ and $\lim_{t\to\infty} \widetilde{\gamma}_2(t) = r$. Again, one can find some apartment A (e.g. $\widetilde{S}$ or $\widetilde{S}'$) 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 $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ is, as before, preserved, it follows that $[\widetilde{\phi}(p_1), \widetilde{\phi}(p_3)]$ and $[\widetilde{\phi}(p_2), \widetilde{\phi}(p_4)]$ intersect (see Figure 4.12).

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 $p_1 = \widetilde{\gamma}_1(a)$ and $p_2 = \widetilde{\gamma}_1(b)$

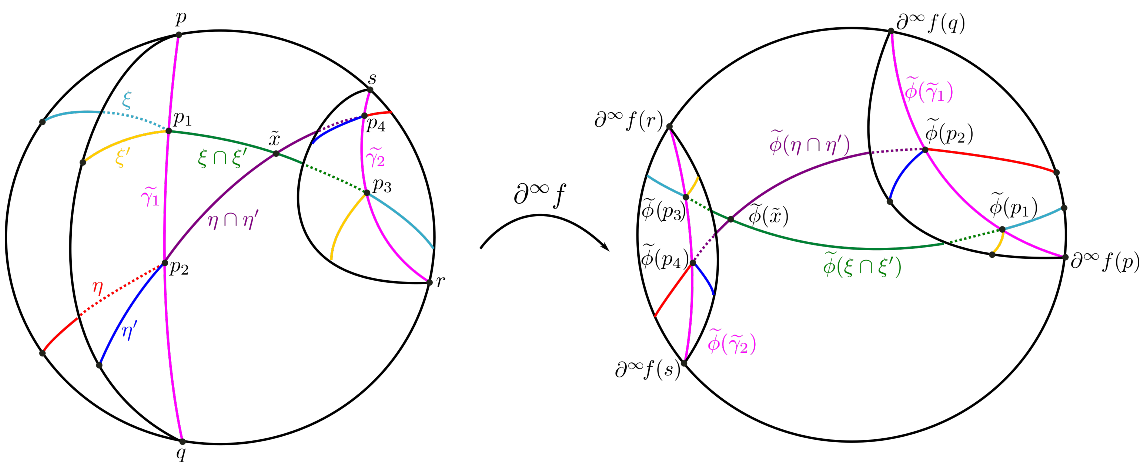

Figure 4.12: An example where $\xi \cap \xi'$ and $\eta \cap \eta'$ are both geodesic arcs.

for some $a < b$ where $\widetilde{\gamma_1}$ has endpoints $p = \lim_{t\to-\infty} \widetilde{\gamma_1}(t)$ and $q = \lim_{t\to\infty} \widetilde{\gamma_1}(t)$. Suppose $p_3 = \widetilde{\gamma_2}(c)$ where $c \in \mathbb{R}$ and the endpoints of $\widetilde{\gamma_2}$ are s and t. Again, choose some apartment A (e.g. $\widetilde{S}$ or $\widetilde{S}'$) that contains all five points ($p, q, r, s,$ 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 $\widetilde{\gamma_1}$ is fixed, it follows that $[\widetilde{\phi}(p_1), \partial^\infty f(r)) = [\widetilde{\phi}(p_1), \partial^\infty \widetilde{\phi}(r))$ and $[\widetilde{\phi}(p_2), \widetilde{\phi}(p_3)]$ will intersect. (See Figure 4.13.)

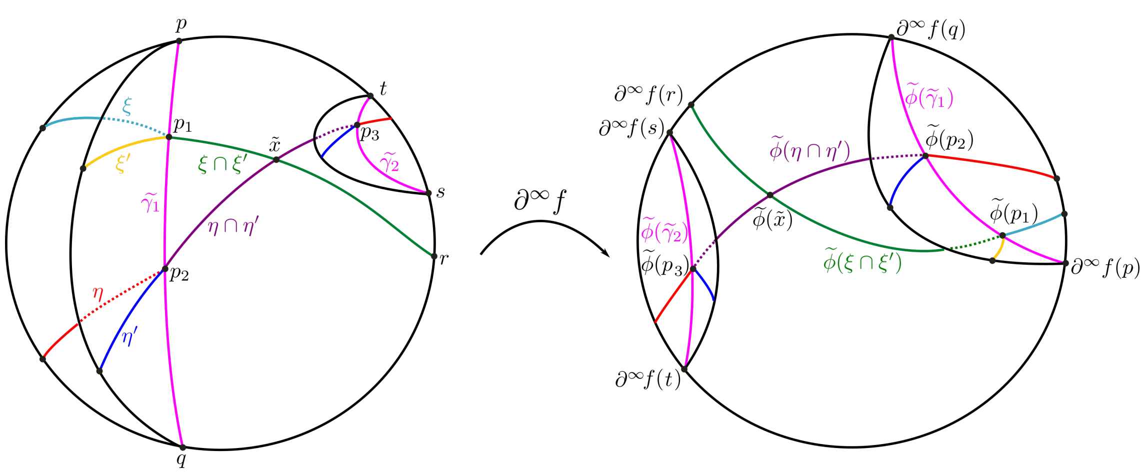

Figure 4.13: An example where $\xi \cap \eta$ is a geodesic ray and $\xi' \cap \eta'$ is a geodesic arc.

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 $\widetilde{\phi_S}(\xi \cap \eta) = \widetilde{\phi_S}(\widetilde{x})$, $\widetilde{\phi_{S'}}(\xi' \cap \eta') = \widetilde{\phi_{S'}}(\widetilde{x})$

and $\widetilde{\phi_S}$ and $\widetilde{\phi_{S'}}$ are injective, it follows that:

$$\widetilde{\phi_{S'}}(\widetilde{x}) = \widetilde{\phi_{S'}}(\xi\cap\xi'\cap\eta\cap\eta') = \widetilde{\phi_{S'}}(\xi\cap\xi')\cap\widetilde{\phi_{S'}}(\eta\cap\eta') = \widetilde{\phi_S}(\xi\cap\xi')\cap\widetilde{\phi_S}(\eta\cap\eta') = \widetilde{\phi_S}(\xi\cap\xi'\cap\eta\cap\eta') = \widetilde{\phi}_S(\widetilde{x})$$

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

The proof of Proposition 4.3.1 then follows easily:

Proof of Proposition 4.3.1. As before, consider a minimal collection of closed surfaces $\mathcal{S}$ that covers X. Consider any arbitrary $S, S' \in \mathcal{S}$. By Lemmas 4.3.8 and 4.3.9, ϕ_S and $\phi_{S'}$ pointwise agree on $S \cap S'$. Since X is complete, convex, and compact, by Lemma 4.3.6, one can thus patch the collection of isometries $\{\phi_S\}_{S\in\mathcal{S}}$ together to obtain a global isometry $\phi : (X, g_1) \to (X, g_2)$. □

4.4 General Case

We now tackle the general case of Theorem 4.1.2. Let (X, g) be a simple, thick surface amalgam with $g \in \mathcal{M}_{\leq}$. We first show that the marked length spectrum of any finite-sheeted cover of (X, g) is determined completely by the marked length spectrum of (X, g). Most of the section is devoted to proving Proposition 4.4.3, which says the marked length spectrum of the double of (X, g) is also determined by the marked length spectrum of (X, g). In order to do so, we define a Liouville current that works for surface amalgams. As part of the proof of Proposition 4.4.3, we also argue that the Liouville current will completely determine intersection numbers, which requires some care since the visual boundary of a surface amalgam is much more complex than that of a surface (see Section 4.2.3), so the classical arguments from the surface case (as seen in [Ota90], [HP97], [Con18], etc.) do not generalize. Finally, we construct finite-sheeted covers and doubles of (X, g) which fall under the base case, thus proving the theorem.

4.4.1 Constructing Covers

Given a metric space (X, g), there is a natural metric $\widehat{g}$ defined on a finite-sheeted cover $\widehat{X}$ of X; endow each copy of X in $\widehat{X}$ with g. This section is dedicated to deriving some useful facts about covers of simple, thick Riemannian surface amalgams that we will use to prove Theorem 4.1.2.

First, we observe that the marked length spectrum of simple, thick surface amalgam completely determines the marked length spectrum of any of its finite-sheeted covers, which follows from an argument from covering space theory.

Lemma 4.4.1. *Suppose X is a CAT(-1) metric space, and (X, g_1) and (X, g_2) have the same marked length spectrum. Then for any finite-sheeted cover $\widehat{X}$, it follows that $(\widehat{X}, \widehat{g_1})$ and $(\widehat{X}, \widehat{g_2})$ will have the same marked length spectrum.*

Proof. Let $\widetilde{\gamma} \in \pi_1(\widehat{X})$ be a geodesic in $(\widehat{X}, \widehat{g})$, a degree d cover of (X, g). Since $\widetilde{\gamma}$ is a closed curve, by the Galois correspondence it projects under the covering map $p : \widehat{X} \to X$ to a closed curve $\gamma \in \pi_1(X)$. Recall that for CAT(-1) spaces, every local geodesic is also a geodesic (see Proposition 1.4 of [BH99]). Observe that γ is a local geodesic since p is a local isometry, so γ is also a geodesic. Since $\widetilde{\gamma}$ consists of k isometric copies of γ concatenated together to form a closed geodesic, it follows that follows that $\ell_{\widehat{g}}(\widetilde{\gamma}) = k\ell_g(\gamma)$ for some $1 \leq k \leq d$, $k \in \mathbb{N}$. Thus, for any $\widetilde{\gamma} \in \pi_1(\widehat{X})$, $\ell_{\widehat{g_1}}(\widetilde{\gamma}) = k\ell_{g_1}(\gamma) = k\ell_{g_2}(\gamma) = \ell_{\widehat{g_2}}(\widetilde{\gamma})$. The lemma then follows. □

We then prove the following purely topological lemma:

Lemma 4.4.2. *Let X be a surface amalgam. Then there exists a finite-sheeted cover of X, $\widehat{X}$, such that for every chamber $C \subset \widehat{X}$, no two boundary components of C are glued together in $\widehat{X}$.*

Proof. Suppose that C is a chamber in X, and let $\{\gamma_i\}_{i=1}^n$ denote the set of gluing curves in X. Let p_i^C denote the number of boundary components of C attached to a gluing curve γ_i of X adjacent to C. Furthermore, let $B_i^C = \{b_{i,1}^C, b_{i,2}^C, ..., b_{i,p_i^C}^C\}$ denote the collection of boundary components of C glued to γ_i. We will construct $\widehat{C}$, a $P^C = \prod_{i=1}^{n} p_i^C$ sheeted cover of C such that no two boundary components of $\widehat{C}$ are glued together.

First, if $p_1^C > 1$, we construct a p_1^C-sheeted cover. Take p_1^C copies of C labeled from 1 to p_1^C. We will use C_{j_1} to denote these copies of C $(1 \leq j_1 \leq p_1^C)$. For all $1 \leq k \leq p_1^C$, glue $b_{1,k}^{C_{j_1}}$ in C_{j_1} to $b_{1,k+1}^{C_{j_1+1}}$ in C_{j_1+1} (where k and j_1 are taken modulo p_1^C and $1 \leq k, j_1 \leq p_1^C$). Notice that if we choose any two copies of C, e.g. C_{j_1} and C_{k_1}, every boundary component of C_{j_1} is glued to exactly one boundary component of C_{k_1} and vice versa (see Figure 4.14).

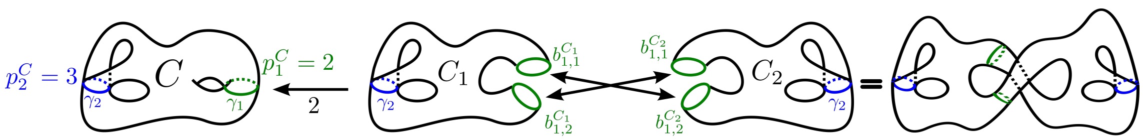

Figure 4.14: Constructing a p_1^C-sheeted cover of a chamber C, which in this case consists of a five-holed sphere glued to two different gluing curves.

Then, we take p_2^C copies of the p_1-sheeted cover and use C_{j_1,j_2} to denote the j_2th copy of C_{j_1}. Next, similar to before, for each j_1, we glue $b_{2,l}^{C_{j_1,j_2}}$ in C_{j_1,j_2} to $b_{2,l+1}^{C_{j_1,j_2+1}}$ in C_{j_1,j_2+1}, where as before, $1 \leq l \leq p_2^C$ indexes the boundary components of each copy of C attached to γ_2. In general, we construct a $P_m^C = p_1^C p_2^C ... p_m^C$-sheeted cover of C by taking p_m^C copies of the $P_{m-1}^C = p_1^C p_2^C ... p_{m-1}^C$-sheeted cover and gluing together boundary components of $C_{j_1 j_2 j_3 \cdots j_{m-1} j_m}$ with the same j_i values where $1 \leq j_i \leq m-1$ but different j_m values. As before, for all the covers, we glue boundary component $b_{m,k}^{C_{j_1 j_2 \cdots j_{m-1} j_m}}$ in $C_{j_1 j_2 \cdots j_{m-1} j_m}$ to $b_{m,k+1}^{C_{j_1 j_2 \cdots j_{m-1} j_m+1}}$ in $C_{j_1 j_2 \cdots j_{m-1} j_m+1}$ where $1 \leq k \leq p_m^C$. The resulting tower of covers will yield $\widehat{C}$. Notice that by construction, no two boundary components of C are glued together, as claimed. See Figure 4.15 for an example of this construction.

Let $\{C_i\}_{i=1}^N$ be the collection of chambers in X. Perform the construction above to obtain a collection of $P_i = P^{C_i}$-sheeted covers $\{\widehat{C_i}\}_{i=1}^N$. Consider the least common multiple of $P_1, P_2, ..., P_N$, which we will call P. Notice that in each $\widehat{C_i}$, there are P_i copies of each gluing adjacent to C_i in X. For each $\widehat{C_i}$, take P/P_i copies of $\widehat{C_i}$ to obtain $(\frac{P}{P_i})(P_i) = P$ copies of each gluing curve adjacent to C_i. Thus, one can glue together the collections of P/P_i copies of $\widehat{C_i}$ to construct a finite-sheeted cover of X. See Figure 4.16

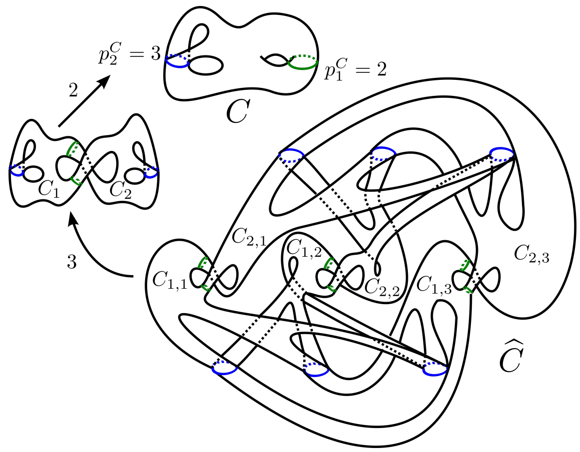

Figure 4.15: Given a chamber C, a $P^C = 6$ sheeted cover of C is constructed. Note that the boundary components of $C_{1,1}$, $C_{1,2}$, and $C_{1,3}$ are all glued together while the boundary components of $C_{2,1}$, $C_{2,2}$, and $C_{2,3}$ are glued together.

for an example of this construction. □

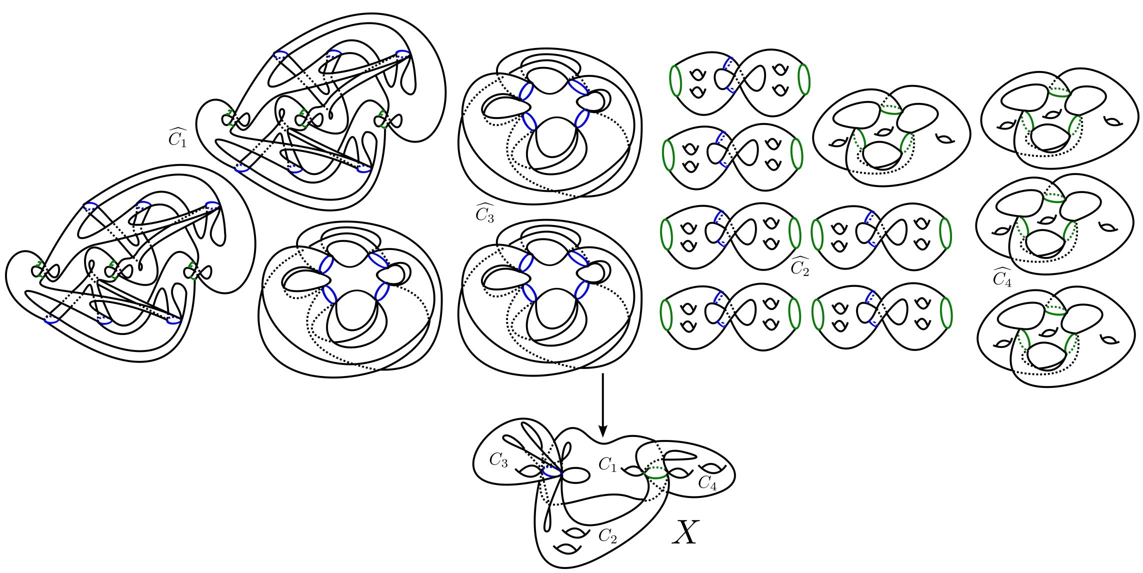

Figure 4.16: The complete collection of chambers required for a 12-sheeted cover of X, the simple, thick surface amalgam from Figure 2.1 following the algorithm detailed in Lemma 4.4.2. Note that there are 12 copies of each gluing curve, so it is possible to glue the copies of chambers together to obtain $\widehat{X}$.

4.4.2 Doubling Argument

Given a surface amalgam, one can consider the double DC of each chamber in $C \subset X$. If we identify the collection of doubles of all the chambers DC along the branching geodesics, we obtain the *double of a surface amalgam along its gluing curves*, which we will call DX (see Figure 4.17 for an illustration).

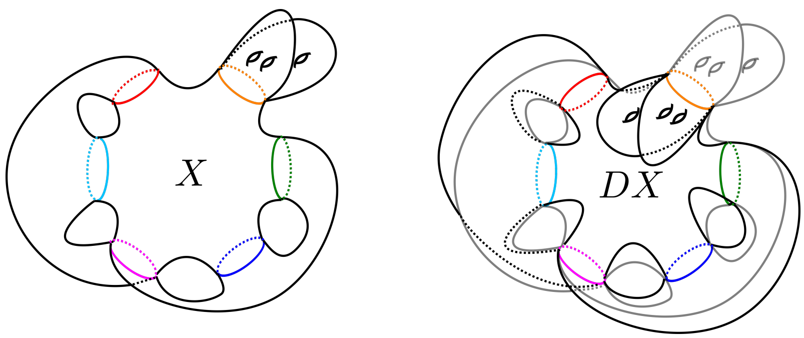

Figure 4.17: The double of X along its gluing curves, DX, is shown on the right. Note that the original copy of X (shown in gray) is isometrically embedded into DX. Although X is not thick, the notion of a double is well-defined for all surface amalgams.

Given (X, g), there is a natural extension of g to a metric on DX, which we will denote by $\bar{g}$. This section will be devoted to proving the following proposition:

Proposition 4.4.3. *Let (X, g_1) and (X, g_2) be two simple, thick surface amalgams endowed with $g_i \in \mathcal{M}_{\leq}$, $i = 1, 2$. Let $(DX, \overline{g_1})$ and $(DX, \overline{g_2})$ denote the doubles of X along branching geodesics with g_1 and g_2 extended to metrics on DX as described above. Then (X, g_1) and (X, g_2) have the same marked length spectrum if and only if $(DX, \overline{g_1})$ and $(DX, \overline{g_2})$ have the same marked length spectrum.*

In order to prove the proposition, we first introduce the Liouville current.

The Liouville Current

In the case of closed surfaces, the Liouville current encodes important information about the length spectra. We define an analogous current, which we will also call the Liouville current, that encodes information about the length spectra of surface amalgams. We follow Ballmann and Brin's construction in [BB95] for any locally finite, finite dimensional and boundaryless complex with a piecewise smooth Riemannian metric. The proof of the following lemma reviews their construction.

Lemma 4.4.4. *Let (X,g) be a negatively curved surface amalgam with $g \in \mathcal{M}_{\leq}$. Then there exists a geodesic current on (X,g) (i.e. π_1-invariant Radon measure on $\mathscr{G}(\widetilde{X})$) that restricts to the usual Liouville current on the interior of chambers of X.*

Proof. We first begin with the original construction by Ballmann and Brin in [BB95], whose argument we repeat here for convenience. Let $x \in \gamma$, where γ is a gluing geodesic of X. Let $S'_x C$ be the open hemisphere of unit tangent vectors at x pointing inside a chamber C adjacent to γ. 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 γ (see Figure 4.18). Ballmann and Brin define a Liouville measure μ on $S' = \bigcup_{x \in \{\gamma_j\}_{j=1}^n} S'_x$, where $\{\gamma_j\}_{j=1}^n$ denotes the set of gluing geodesics of X.

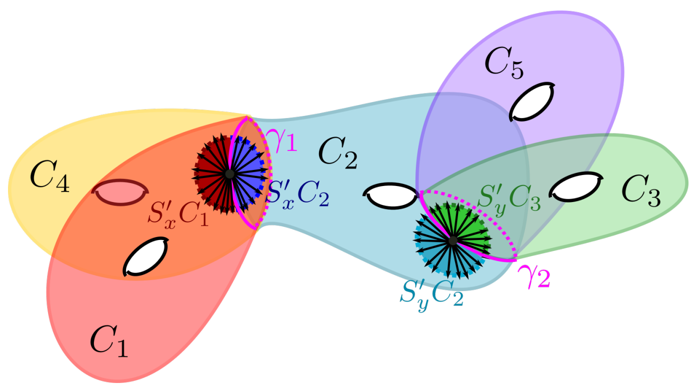

Figure 4.18: Given two points $x \in \gamma_1$ and $y \in \gamma_2$, we illustrate examples of hemispheres of unit tangent vectors based at x and y. Here, $S'_x = S'_x C_1 \cup S'_x C_2 \cup S'_x C_4$ and $S'_y = S'_y C_2 \cup S'_y C_3 \cup S'_y C_5$.

Suppose v is based at $x \in \gamma$ and points into a chamber C. Let $\theta(v)$ be the angle between v and the normal to γ based at x pointing into C (see Figure 4.19). Recall that S'_x is a collection of open hemispheres of unit tangent vectors and thus can be identified with an open semicircle by mapping each vector to its endpoint. As a result, we can put a Lebesgue measure λ_x, the usual measure on a collection of open semicircles, on S'_x. Finally, let dx be the volume element on $\{\gamma_j\}_{j=1}^n$. Then define the Liouville measure on

S' as:

$$d\mu(v) = \cos(\theta(v))d\lambda_x dx \tag{4.7}$$

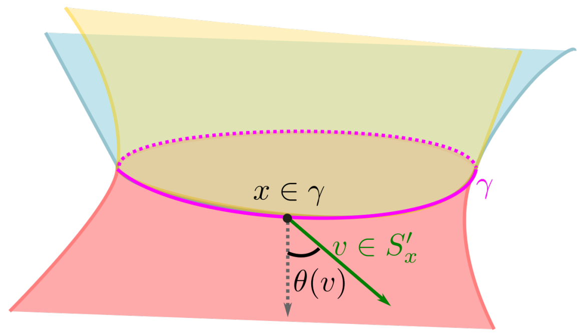

Figure 4.19: Given a vector $v \in S'_x$, $\theta(v)$ is the angle between v and the unit normal of the gluing curve x emanates from.

Ballmann and Brin also observe that $d\overline{L_g} = d\mu \times dt$, a measure on the set of unit-speed parameterized geodesics that transversely intersect $\{\gamma_j\}_{j=1}^n$, is invariant under geodesic flow and flip using arguments from billiards dynamics. Furthermore, they observe $\overline{L_g}$ is the usual Liouville measure on the interiors of the chambers of X. Recall from the discussion in Section A.0.1 that that there is a bijective correspondence between $\mathcal{M}_{\text{flip-flow}}(X)$, the set of geodesic flow and flip invariant Radon measures on the generalized unit tangent bundle SX, and $\mathscr{C}(X)$, π_1-invariant Radon measures on $\mathscr{G}(\widetilde{X})$ for any Gromov hyperbolic space (including negatively curved surface amalgams), and $\overline{L_g}$ corresponds to a geodesic current $L_g = \mu$. □

The Liouville Current in Local Coordinates. Inspired by Constantine and Lafont (see [CL19b]), we now give a local expression for the the current $L_g \in \mathscr{C}(X)$. Let (X, g) be a simple, thick surface amalgam equipped with a metric $g \in \mathcal{M}_{\leq}$. Let $\sigma \subset (X, g)$ be a geodesic segment.

Suppose $v \in S'C$, where $S'C$ denotes the set of unit tangent vectors pointing into a particular chamber C. We can use v to extend a geodesic segment $\sigma \subset C$ until σ reaches a (possibly the same) gluing curve $\gamma' \in \{\gamma_j\}_{j=1}^n$ at some point $x \in \gamma'$. Following the notation in [BB95], we let $F(v) \in S'$ denote the set of vectors emanating from x which can geodesically extend σ (see Figure 4.20); note that $|F(v)|$ is exactly one less than the number of chambers adjacent to γ', as σ can be extended into any chamber adjacent to γ' barring the one it came from. Let W denote the set of all $(w_n)_{n\in\mathbb{Z}}$, bi-infinite sequences of vectors such that $w_{n+1} \in F(w_n)$.

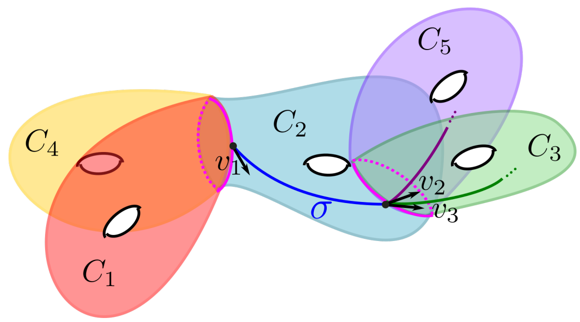

Figure 4.20: Given $v_1 \in S'_x C_2$, there are two ways to continue σ, the geodesic segment determined by v_1: along $v_1 \in S'_x C_3$ or along $v_2 \in S'_x C_5$, so $|F(v)| = 2$.

Let $V \subset S'$ be the set of all vectors that can be included into some $(w_n)_{n\in\mathbb{Z}} \in W$, which Ballmann and Brin observe to be a full measure subset of S'. We can view V as the state space of a Markov Chain with transition probabilities on $v, w \in V$ defined by:

$$p(v,w) = \begin{cases} \frac{1}{|F(v)|} & \text{if } w \in F(v) \\ 0 & \text{otherwise} \end{cases} \tag{4.8}$$

We can thus parameterize the mass of unoriented, unparameterized geodesics $\mathscr{G}(\sigma) \subset \mathscr{G}(\widetilde{X})$ that intersect a geodesic segment $\sigma \subset \widetilde{X}$ at a positive angle with local coordinates $(x, \theta(v), (\widetilde{w_n})_{n\in\mathbb{N}})$. More precisely, given $\xi \in \mathscr{G}(\sigma)$, we have:

- x is the point of intersection of the element $\xi \in \mathscr{G}(\widetilde{X})$ with σ;
- $\theta(v) \in (0, \pi)$ is the angle between the normal to σ at x and a unit tangent vector of ξ at x (with fixed orientations for ξ and σ);
- $(\widetilde{w_n})_{n\in\mathbb{N}}$ denotes a bi-infinite sequence of unit vectors based at branching geodesics along which ξ continues. Note that these are lifts of the unit vectors in $(w_n)_{n\in\mathbb{N}} \in W$

described above. We will let $\widetilde{W}$ denote the set of all possible $(\widetilde{w_n})_{n\in\mathbb{N}}$.

It is then natural to locally define L_g on the triple $(x, \theta(v), (\widetilde{w_n})_{n\in\mathbb{N}})$:

Definition 4.4.5 (Liouville current for Riemannian Surface Amalgams)**.** Let (X, g) be a simple, thick surface amalgam equipped with $g \in \mathcal{M}_{\leq}$. Then given a geodesic segment $\sigma \subset (\widetilde{X}, \widetilde{g})$, one can use the *Liouville current* L_g to measure the mass of unparameterized, unoriented geodesics $(x, \theta(v), (\widetilde{w_n})_{n\in\mathbb{N}})$ intersecting σ at a nonzero angle with the local coordinates:

$$dL_g = \cos(\theta) d\theta dx d\nu$$

Where dx is the volume element on $\{\gamma\}_{j=1}^{n}$ and ν is the Markov probability measure given by the transition probabilities from Equation 4.8.

Note that $d\theta d\nu$ is equivalent to the Lebesgue measure $d\lambda_x$ on S'_x, a collection of unit tangent vectors based at $x \in X$. Indeed, the Lebesgue measure on a unit semicircle is $d\theta$, as the arclength of a sector of a unit semicircle is determined completely by the angle measure of the sector in radians. Thus, the local coordinates from Definition 4.4.5 agree with the coordinates given by Ballmann and Brin in equation 4.7.

Intersection numbers determine lengths of geodesics

We now associate intersection numbers with the length spectrum of a simple, thick surface amalgam (X, g) where $g \in \mathcal{M}_{\leq}$, taking inspiration from Constantine and Lafont's methods from [CL19b].

If I is a geodesic segment in $\widetilde{X}$, let $\mathscr{G}(I)$ denote the compact subset of geodesics of $\mathscr{G}(\widetilde{X})$ that intersect I but are not tangent to I. Let $\widetilde{\alpha}$ be some lift of α in $\widetilde{X}$ and let I_α be the fundamental domain on $\widetilde{\alpha}$ of the action by α; note that I_α has length $\ell_g(\alpha)$. First, we prove a useful lemma that relates $L_g(\mathscr{G}(I_\alpha))$ and $i(\alpha, L_g)$ (see Definition 4.2.13).

Lemma 4.4.6. *Suppose $\alpha \in \pi_1(X)$ is represented by a primitive geodesic α, and let μ be a geodesic current with full support in $\mathscr{G}(\widetilde{X})$. Then $i(\alpha, \mu) = \mu(\mathscr{G}(I_\alpha))$.*

Proof. Let $\widehat{A_n}(\alpha) = \{d \in \mathscr{G}(\widetilde{X})/\Gamma : d \pitchfork \widetilde{\alpha} \text{ and } \varpi(\widetilde{\alpha}, d) = n\}$, where $d \pitchfork \widetilde{\alpha}$ means that d is transverse to $\widetilde{\alpha}$ in the sense of Definition 4.2.12. Recall that $\mathcal{W}(\widetilde{\alpha}, d)$ is the set of branching geodesics that transversely intersect both $\widetilde{\alpha}$ and d. Let $\mathcal{P}(\widetilde{\alpha}, d)$ denote the set of points of intersection of d with branching geodesics in $\mathcal{W}(\widetilde{\alpha}, d)$. Let:

$$A_n(\alpha) = \{d' \in \mathscr{G}(\widetilde{X})/\Gamma : d' \in \mathscr{G}(I_\alpha) \text{ and } d' \text{ agrees with some } d \in \widehat{A_n}(\alpha) \text{ inside an } \epsilon\text{-neighborhood of } I_\alpha \text{ before possibly diverging from } d \text{ at some subset of points in } \mathcal{P}(\widetilde{\alpha}, d)\}$$

For $d \in \widehat{A_n}(\alpha)$, since $d \in \mathscr{G}(\widetilde{X})/\Gamma$ and $d \pitchfork \widetilde{\alpha}$, it follows that $d \in \mathscr{G}(I_\alpha)$. As a result, we have that $\widehat{A_n}(\alpha) \subset A_n(\alpha)$. Note that $A_n(\alpha)$ and $\widehat{A_n}(\alpha)$ are both in the support of μ, and since there are n times as many geodesics in $A_n(\alpha)$ than in $\widehat{A_n}(\alpha)$, it follows that $\mu(A_n(\alpha)) = n\mu(\widehat{A_n}(\alpha))$. Indeed, geodesics in $A_n(\alpha)$ are precisely those that agree with agree with geodesics in $\widehat{A_n}(\alpha)$ until possibly diverging from d at some subset of $\mathcal{P}(\widetilde{\alpha}, d)$ (see Figure 4.21).

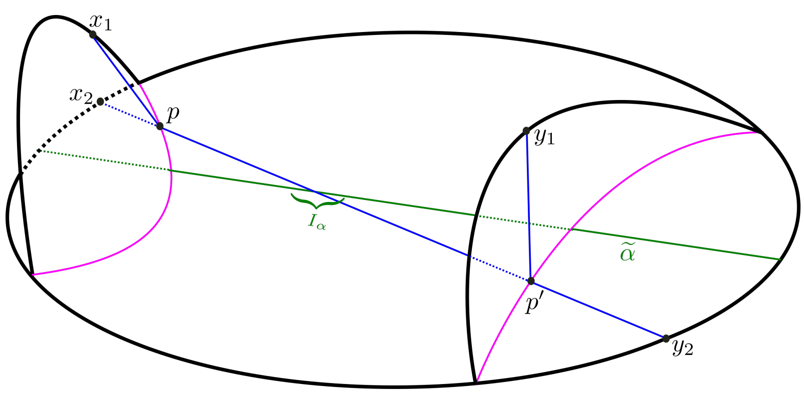

Figure 4.21: While $d = (x_2, y_2) \in \widehat{A_4}(\alpha)$ is the only geodesic in the picture transverse to $\widetilde{\alpha}$, (x_1, y_1), (x_1, y_2), and $(x_2, y_1) \in A_4(\alpha)$, which diverge from (x_2, y_2) at points in $\{p, p'\} = \mathcal{P}(\widetilde{\alpha}, d)$, also intersect I_α at a nonzero angle. All four geodesics are in the support of μ, which suggests $\mu(A_4(\alpha)) = 4\mu(\widehat{A_4}(\alpha))$.

It is clear that the $\widehat{A_n}(\alpha)$ are disjoint for distinct $n \in \mathbb{N}$, and Constantine and Lafont prove that $A_n(\alpha)$ are also disjoint for different $n \in \mathbb{N}$ for Fuchsian buildings; the proof is similar in our setting (see Proposition 8.7 of [CL19b]). Note that $\bigsqcup_{n \in \mathbb{N}} A_n(\alpha) = \mathscr{G}(I_\alpha)$ since every geodesic in $\mathscr{G}(\widetilde{X})/\Gamma$ intersecting I_α at a nonzero angle will agree with some geodesic transverse to $\widetilde{\alpha} \supset I_\alpha$ until possibly diverging from it at branching geodesics.

We are then ready to prove the lemma. Let $p^{-1}(\alpha)$ denote the set of all lifts of α. Note that $\text{supp}(\alpha \times \mu) = \text{supp}(\alpha) \times (\text{supp}(\mu) \cap \bigsqcup_{n\in\mathbb{N}} \widehat{A_n}(\alpha)) = p^{-1}(\alpha) \times \bigsqcup_{n\in\mathbb{N}} \widehat{A_n}(\alpha)$. Thus:

$$
\begin{aligned}
i(\alpha,\mu) &= \int_{D(\mathscr{G}(\widetilde{X}))/\Gamma} \varpi(\xi,\eta)d\alpha d\mu = \sum_{n\in\mathbb{N}} n\left(\int_{\{(\xi,\eta)\in D(\mathscr{G}(\widetilde{X}))/\Gamma \,:\, \varpi(\xi,\eta)=n\}} d\alpha d\mu\right) \\
&= \sum_{n\in\mathbb{N}} n((\alpha\times\mu)(p^{-1}(\alpha)\times\widehat{A_n}(\alpha)) = \sum_{n\in\mathbb{N}} n\mu(\widehat{A_n}(\alpha)) = \sum_{n\in\mathbb{N}} \mu(A_n(\alpha)) \\
&= \mu\left(\bigsqcup_{n\in\mathbb{N}} A_n(\alpha)\right) = \mu(\mathscr{G}(I_\alpha))
\end{aligned}
$$

□

In particular, for the Liouville current L_g from Definition 4.4.5, it follows that for all $\alpha \in \pi_1(X)$, $i(\alpha, L_g) = L_g(\mathscr{G}(I_\alpha))$. We now relate this intersection number with the g-length of α.

Lemma 4.4.7. *Suppose $\alpha \in \pi_1(X)$ is represented by geodesic α and let $\{\gamma_i\}_{i=1}^n$ denote the collection of gluing curves in X. Furthermore, suppose that there are n_i distinct boundary components attached to γ_i. Then:*

$$
i(\alpha, L_g) = \begin{cases} 2\ell_g(\alpha) \text{ if } \alpha \notin \{\gamma_i\}_{i=1}^n \\ n_i\ell_g(\alpha) \text{ otherwise} \end{cases}
$$

Proof. Decompose α into a collection of segments s_i where s_i either lies entirely in the interior of a chamber C or a branching geodesic. We will show that $i(\alpha, L_g) = \sum_i q(s_i)\ell_g(s_i)$, where $\ell_g(s_i)$ denotes the length of s_i, $q(s_i) = 2$ if s_i lies in the interior of a chamber and $q(s_i) = n$ if s_i lies on a gluing curve gluing curve γ_{s_i}, where n is the number of boundary components attached to γ_{s_i}. By Lemma 4.4.6, it suffices to compute $L_g(\mathscr{G}(I_{s_i}))$, where I_{s_i} is some (any) fundamental domain of s_i in $\widetilde{X}$ under the universal covering map.

If s_i is in the interior of a chamber, the proof is the same as the proof in [Ota90] since the Liouville current of (X, g) is identical to the usual Liouville current defined for surfaces. Using local coordinates, we compute:

$$L_g(\mathscr{G}(I_{s_i})) = \int_{x \in I_{s_i}} \int_{\theta=-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\theta) d\theta dx = 2 \int_{x \in I_{s_i}} dx = 2\ell_g(s_i) \tag{4.9}$$

Suppose, on the other hand, that s_i lies along a gluing curve γ_{s_i} and there are n distinct boundary components attached to γ_{s_i}. In the universal cover, a fundamental domain of γ_{s_i} will thus be attached to n distinct chambers in $\widetilde{X}$. First, we compute the measure of the set of all geodesics in $\mathscr{G}(I_{s_i})$ that pass through two distinct, fixed chambers $\widetilde{C}, \widetilde{C'} \subset \widetilde{X}$ adjacent to I_{s_i}. The probability that a geodesic in $\widetilde{X}$ intersecting I_{s_i} will continue into $\widetilde{C'}$ given that it came from $\widetilde{C}$ is $\frac{1}{n-1}$. Recall that since we are measuring the mass of *unoriented* geodesics that pass through I_{s_i}, geodesics originating from $\widetilde{C}$ and entering $\widetilde{C'}$ are identified with those originating from $\widetilde{C'}$ and entering $\widetilde{C}$. The calculation is therefore the same the one seen in Equation 4.9, with an extra factor of $\frac{1}{n-1}$. Since there are $\binom{n}{2}$ choices of $\widetilde{C}$ and $\widetilde{C'}$, the measure of unoriented, unparameterized geodesics passing through I_{s_i} is as follows:

$$L_g(\mathscr{G}(I_{s_i})) = \int_W \int_{x \in I_{s_i}} \int_{\theta=-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\theta) d\theta d\lambda d\nu = \binom{n}{2} \frac{1}{n-1} \int_{x \in I_{s_i}} \int_{\theta=-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\theta) d\theta d\lambda = n\ell_g(s_i)$$

□

In particular, Lemma 4.4.7 implies that $i(\alpha, L_g) = 2\ell_g(\alpha)$ if α transversely intersects the set of gluing geodesics in X. Unlike in [CL19b], the only other case is when α is a gluing geodesic, in which case $i(\alpha, L_g) = n\ell_g(\alpha)$, where n again denotes the number of distinct boundary components attached to α.

Suppose that (X, g_1) and (X, g_2) have the same marked length spectra. From Lemma

4.4.7, we can infer that if $\alpha \in \pi_1(X)$ is not a gluing curve, then:

$$i(\alpha, L_{g_1}) = 2\ell_{g_1}(\alpha) = 2\ell_{g_2}(\alpha) = i(\alpha, L_{g_2})$$

On the other hand, if α is a gluing curve attached to n distinct boundary components, then:

$$i(\alpha, L_{g_1}) = n\ell_{g_1}(\alpha) = n\ell_{g_2}(\alpha) = i(\alpha, L_{g_2})$$

Thus, we can conclude that $i(\alpha, L_{g_1}) = i(\alpha, L_{g_2})$ for all $\alpha \in \pi_1(X)$.

With this information, we now prove a useful fact about geodesic currents $\mu \in \mathscr{C}(X)$ by adapting the proof of Theorem 2 from [Ota90], which says a geodesic current is completely determined by its intersection number data. In order to do so, we need the following Lemma:

Lemma 4.4.8. *Suppose $I \subset \partial^\infty(\widetilde{X})$, $J \subset \partial^\infty(\widetilde{X})$ are intervals (see Definition 4.2.10), and $\overline{I} \cap \overline{J} = \varnothing$. Then there exist finite partitions $\bigcup_{i=1}^{N} I_i$ and $\bigcup_{j=1}^{M} J_j$ of I and J respectively such that for all i and j, I_i and J_j both lie on the boundary of some apartment (disk) in $\widetilde{X}$.*

Proof. We show that one can split I and J each into at most three intervals to satisfy the statement of the lemma. In the following proof, unless otherwise specified, if $p_1, p_2 \in \partial^\infty(\widetilde{X})$ are both in $\overline{I}$ or $\overline{J}$, we use the notation (p_1, p_2) and $[p_1, p_2]$ to denote (the unique) open and closed intervals respectively in $\overline{I} \cup \overline{J}$ with endpoints p_1 and p_2. Furthermore, we will assume that I has endpoints a and b and J has endpoints c and d.

Since $\partial^\infty(\widetilde{X})$ is path-connected (see Section 4.2.3), there is an injective path from a to $\overline{J}$. Let $\gamma_a : [0, 1] \to \partial^\infty(\widetilde{X})$ be an injective path from a to J such that $\gamma_a(0) = a$, $\gamma_a(1) \in \overline{J}$, $\gamma_a(t) \notin \overline{I} \cup \overline{J}$ for all $t \in (0, 1)$, and let $p_a := \gamma_a(1)$. In other words, p_a is the first point in $\overline{J}$ that γ_a reaches. Define γ_b, γ_c, γ_d, and p_b, p_c, and p_d similarly. Then I can be partitioned into $\{I_i\}$ where each with endpoints a_i and b_i, each of which is in the set $\{a, p_c, p_d, b\}$. Similarly, J can be partitioned into $\{J_j\}$ with endpoints c_j and d_j, each of which is in the set $\{c, p_a, p_b, d\}$.

In the special case that p_a, p_b, p_c and p_d are all in $\{a, b, c, d\}$, then I and J themselves are already included in the boundary of an apartment. Indeed, without loss of generality, assume $p_a = c$ and $p_b = d$. Then there is an injective path between a and c as well as one between b and d. Then $[a, c] \cup [c, d] \cup [d, b] \cup [b, a]$, where $[a, c] = \gamma_a([0, 1]) = \gamma_c([1, 0])$ and $[d, b] = \gamma_d([0, 1]) = \gamma_b([1, 0])$, forms the boundary of an apartment.

Otherwise, note that either I and J are both split into two intervals or both split into three intervals. Indeed, if $p_a = c$ (for example), then $p_c = a$, so I and J would be split into at most two intervals with endpoints both in $\{a, p_d, b\}$ or $\{c, p_b, d\}$. If I and J are each split into two or three subintervals, then there are exactly three possible kinds of subintervals that may arise:

- **Case One**: Exactly two of the endpoints of I_i and J_j (i.e. the set $\{a_i, b_i, c_j, d_j\}$) belong in the set $\{a, b, c, d\}$. Note that this can only happen if I_i and J_j each contain exactly one of a, b, c, or d. Without loss of generality, suppose $a_i = a$ and $c_j = c$. There are injective paths γ_a and γ_c from a to p_a and c to p_c respectively; these will trace out part of the boundary of an apartment containing $I_i \cup J_j$. There are two subcases.

  If $p_c \neq a$ (so $p_a \neq c$), then $(a_i, b_i) = (a, b_i) \subseteq (a, p_c)$ and $(c_j, d_j) = (c, d_j) \subseteq (c, p_a)$ by definition of a partition. Then $I_i \subset [a, p_c]$ and $J_j \subset [c, p_a]$, so $[a, p_a] \cup [p_a, c] \cup [c, p_c] \cup [p_c, a]$ is the boundary of an apartment containing $I_i \cup J_j$, where $[a, p_a] = \gamma_a([0, 1])$ and $[c, p_c] = \gamma_c([0, 1])$. On the other hand, if $p_c = a$ so $c = p_a$, then $b_i = p_d$ and $d_j = p_b$ necessarily. Then $[a_i, b_i] \cup [c_j, d_j] = [a, p_d] \cup [c, p_b] \subset [a, b] \cup [c, p_b] \subset [a, c] \cup [c, p_b] \cup [p_b, b] \cup [b, a]$, which forms the boundary of an apartment containing $I_i \cup J_j$. As before, we set $[a, c] = \gamma_a([0, 1])$ and $[p_b, b] = \gamma_b([1, 0])$. See Figure 4.22 for an illustration of these two subcases.

- **Case Two**: Exactly one of the endpoints of I_i and J_j belongs in the set $\{a, b, c, d\}$. Without loss of generality, assume I_i has endpoints a and b_i and J_j has endpoints p_a and p_b. Again, there is an injective path from a to p_a. There is also an injective path from b_i to either c or d; without loss of generality, assume $b_i = p_c$. Note

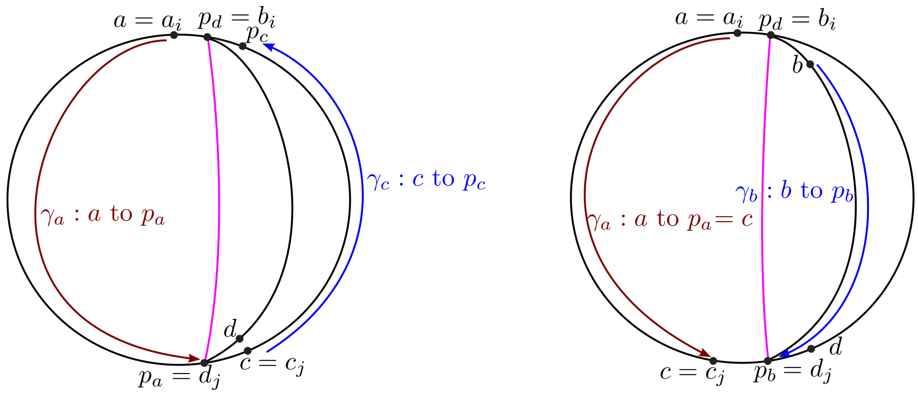

Figure 4.22: Examples of I_i and J_j where exactly one pair of endpoints is in the set $\{a, b, c, d\}$. In the picture on the left, an apartment boundary is formed by $[a, p_a] \cup [p_a, c] \cup [c, p_c] \cup [p_c, a]$ while in the picture on the right, the apartment boundary is formed by $[a, c] \cup [c, p_b] \cup [p_b, b] \cup [b, a]$. In both examples, $I_i \cup J_j$ is included in the boundary of an apartment.

that either (p_b, c) or (p_a, c) contains (p_a, p_b) since (c, d) is partitioned into either $(c, p_a) \cup (p_a, p_b) \cup (p_b, d)$ or $(c, p_b) \cup (p_b, p_a) \cup (p_a, d)$. Without loss of generality, assume $(p_a, p_b) \subset (p_a, c)$. Then $[a, p_a] \cup [p_a, c] \cup [c, p_c] \cup [p_c, a]$ forms the boundary of an apartment containing $I_i \cup J_j$ since $\overline{I_i} = [a, p_c]$ and $\overline{J_j} = [p_a, p_b] \subseteq [p_a, c]$, where $[a, p_a] = \gamma_a([0, 1])$ and $[p_c, c] = \gamma_c([1, 0])$. See the figure on the left of Figure 4.23.

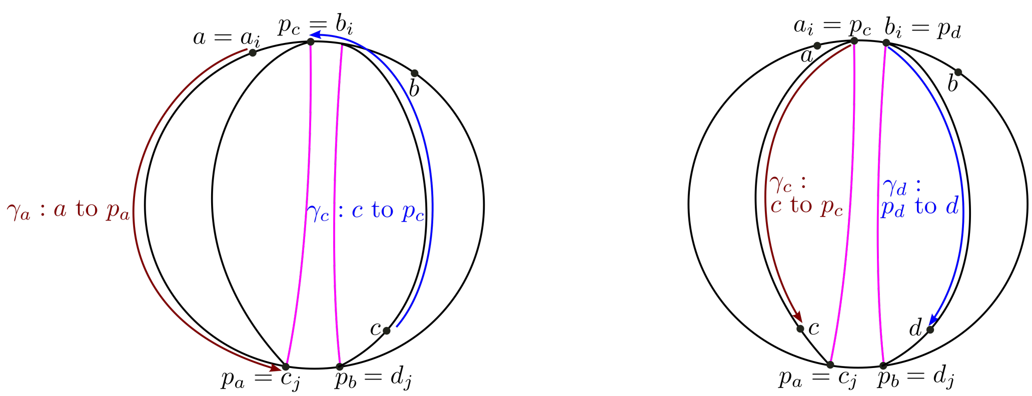

Figure 4.23: Examples of case two (left) and three (right). In both examples, $I_i \cup J_j$ is included in the boundary of an apartment. For the figure on the left, I_i is included in $[a, p_a] \cup [p_a, c] \cup [c, p_c] \cup [p_c, a]$ and for the figure on the right, J_j is included in $[p_c, c] \cup [c, d] \cup [d, p_d] \cup [p_d, p_c]$.

- **Case Three**: None of the four endpoints of I_i and J_j are in $\{a, b, c, d\}$. In other words, I_i has endpoints p_c and p_d and J_j has endpoints p_a and p_b. Note there are injective paths from p_c and p_d to c and d respectively, and $I_i \subset I$. Thus, $\overline{I_i} \cup \overline{J_j} = [p_c, p_d] \cup [p_a, p_b]$ is a subset of the boundary of an apartment $[c, p_c] \cup [p_c, p_d] \cup [p_d, d] \cup [d, c]$ where $[c, p_c] = \gamma_c([0, 1])$ and $[p_d, d] = \gamma_d([1, 0])$. See the figure on the right of Figure 4.5 for an illustration.

□

Lemma 4.4.9. *Let $\mu_1, \mu_2 \in \mathscr{C}(X)$, where X is a Riemannian surface amalgam. Then $\mu_1 = \mu_2$ if $i(\mu_1, \alpha) = i(\mu_2, \alpha)$ for all $\alpha \in \pi_1(X)$.*

Proof. From the discussion in the proof of Lemma 4.3.5, we know there exists $v \in \mathscr{G}(\widetilde{X})/\pi_1(X)$, a unit speed geodesic in X, with dense forward and backward orbit under the geodesic flow map. In particular, the image of v is dense in X, the image of the lifts of v are dense in $\widetilde{X}$, and the endpoints of such lifts are dense in $\mathscr{G}(\widetilde{X})$.

Since the lifts of v and thus their endpoints are countable, it follows that $\mathscr{G}(\widetilde{X})$ is separable since it contains a countable dense subset. As the result, the open sets (and thus the Borel σ-algebra) of $\mathscr{G}(\widetilde{X})$ are generated by products of open balls defined in Equation 4.3 of Section 4.2.3, $B(p_1, d_1) \times B(p_2, d_2)$ with disjoint closure. Recall that $B(p, d)$ describes the set of geodesic rays that 2δ-fellow-travel with the geodesic ray p for distance d.

Observe that each $B(p, d)$ can be covered by a countable collection of intervals. Indeed, each lift $\widetilde{\gamma} \subset \widetilde{X}$ of a gluing geodesic $\gamma \subset X$ is attached to finitely many apartments or totally geodesic subsets of apartments which are universal covers of chambers glued to γ in X. In particular, if one or two chambers glue together at γ to form a closed surface $S \subset X$, then there is an apartment that is a copy of the universal cover of S glued to each $\widetilde{\gamma}$. In all other cases, $\widetilde{\gamma}$ is glued to a copy of the universal cover of a surface with boundary, which is a totally geodesic subset of an apartment. Since there are countably many branching geodesics in $\widetilde{X}$, there are countably many such copies of universal covers of chambers or unions of chambers. Note that the intersection of $B(p, r)$ with the boundary of an apartment can be covered with one interval while the intersection of $B(p, r)$ with the boundary of a copy of the universal cover of a surface with boundary is the complement of a Cantor set, which can be covered with countably many open intervals. Thus, since a countable union of countable sets is still countable, $B(p, r)$ can be covered with countably many intervals.

Since every countable union of sets can be written as a countable disjoint union of subsets of sets in the union, we can write $B(p_1, d_1)$ and $B(p_2, d_2)$ as countable disjoint unions of intervals. In other words, for intervals I_i and J_j, we have the following:

$$\mu(B(p_1, d_1) \times B(p_2, d_2)) = \mu\left(\bigsqcup_{i \in \mathbb{N}} I_i \times \bigsqcup_{j \in \mathbb{N}} J_j \right) = \mu\left(\bigsqcup_{i,j \in \mathbb{N}} (I_i \times J_j) \right) \tag{4.10}$$

By Lemma 4.4.8, for any I_i and J_j, it follows that $I_i = \bigcup_{n=1}^{N(i)} I_{i,n}$ and $J_j = \bigcup_{m=1}^{M(j)} J_{j,m}$ where $I_{i,n}$ and $J_{j,m}$ are disjoint intervals such that for all $1 \leq n \leq N(i)$ and $1 \leq m \leq M(j)$, $I_{i,n} \cup J_{j,m}$ lies on the boundary of a single apartment. As a result, we have that:

$$\mu(I_i \times J_j) = \mu\left(\bigsqcup_{1 \leq n \leq N(i)} I_{i,n} \times \bigsqcup_{1 \leq m \leq M(j)} J_{j,m} \right) = \mu\left(\bigsqcup_{\substack{1 \leq n \leq N(i) \\ 1 \leq m \leq M(j)}} I_{i,n} \times J_{j,m} \right) \tag{4.11}$$

The proof now closely follows the argument by Otal, which we will recall here, with a few modifications. Let $\widetilde{v_1}$, $\widetilde{v_2}$, $\widetilde{v_3}$, and $\widetilde{v_4}$ be four arbitrary lifts of v in $\widetilde{X}$. Suppose $I_{i,n}$ is a geodesic arc in $\partial^\infty(\widetilde{X})$ sharing one endpoint x_n with $\widetilde{v_1}$ and the other, y_n, with $\widetilde{v_2}$ and $J_{j,m}$ a geodesic arc sharing one endpoint w_m with $\widetilde{v_3}$ and the other, z_m, with $\widetilde{v_4}$. Note that since the lifts of v are dense in $\widetilde{X}$, it suffices to find the μ-mass of geodesics with endpoints in all such $I_{i,n}$ and $J_{j,m}$, which we will denote by $\mu(I_{i,n} \times J_{j,m})$.

Let $\mathscr{G}(\gamma)$ denote the set of geodesics that pass through a geodesic $\gamma \subset \widetilde{X}$. Let $\gamma_{x_n,z_m} \subset A$ denote the geodesic in X with endpoints x_n and z_m. Define γ_{y_n,w_n}, γ_{x_n,w_n} and γ_{y_n,z_m} similarly. Note that (see Figure 4.24 for details):

$$\mu(I_{i,n} \times J_{j,m}) = \frac{1}{2}\Big(\mu(\mathscr{G}(\gamma_{x_n,z_m})) + \mu(\mathscr{G}(\gamma_{y_n,w_m})) - \mu(\mathscr{G}(\gamma_{x_n,w_m})) - \mu(\mathscr{G}(\gamma_{y_n,z_m})) \Big) \tag{4.12}$$

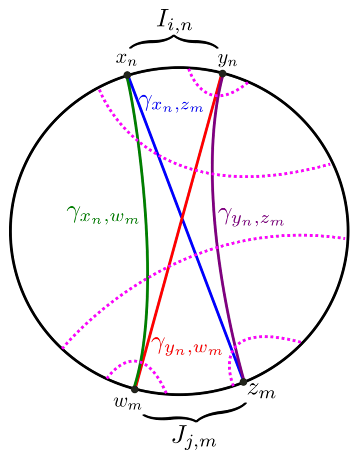

Figure 4.24: An illustration of Equation 4.12. Note that $I_{i,n}$ and $J_{j,m}$ lie on the boundary of some apartment. The pink geodesics represented by dashed lines do not have one endpoint in $I_{i,n}$ and the other in $J_{j,m}$; their mass is thus subtracted from the calculation of the mass of $\mu(I_{i,n} \times J_{j,m})$.

Using a similar argument to the one in Lemma 4.3.5, one can conclude that there exist $\alpha_i^{nm} \in \pi_1(X)$ where $i = 1, 2, 3, 4$ such that the translation axes of α_i^{nm} approximate the geodesics γ_{x_n,z_m}, γ_{y_n,w_n}, γ_{x_n,w_n} and γ_{y_n,z_m} arbitrarily well. Recall that $i(\mu, \alpha_i^{nm})$ is by definition the μ-mass of geodesics passing through the translation axis of α_i^{nm}. In other words, for any choice of $\epsilon > 0$, there exist α_i^{nm} such that for any $k, M, N \in \mathbb{Z}_{>0}$:

$$\left|\mu(I_{i,n} \times J_{j,m}) - \frac{1}{2}(i(\mu, \alpha_1^{nm}) + i(\mu, \alpha_2^{nm}) - i(\mu, \alpha_3^{nm}) - i(\mu, \alpha_4^{nm}))\right| < \frac{\epsilon}{kMN}$$

So for all $\epsilon > 0$, there exist $\{\alpha_1^{mn}, \alpha_2^{mn}, \alpha_3^{mn}, \alpha_4^{mn}\}_{\substack{1 \leq m \leq M \\ 1 \leq n \leq N}}$ such that:

$$\Big|\mu(I_i \times J_j) - \sum_{\substack{1\leq m\leq M(j)\\ 1\leq n\leq N(i)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big|$$

$$= \Big|\sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(i)}} \mu(I_{i,n} \times J_{j,m}) - \sum_{\substack{1\leq m\leq M(j)\\ 1\leq n\leq N(i)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big|$$

$$\leq \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \Big|\mu(I_{i,n} \times J_{j,m}) - \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big|$$

$$\leq \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \frac{\epsilon}{M(i)N(j)(2^i)(2^j)} = \frac{\epsilon}{(2^i)(2^j)}$$

Finally, combining equations 4.10, 4.11, and the aforementioned inequality, we have that for any $\epsilon > 0$:

$$\Big|\mu(B(p_1,d_1) \times B(p_2,d_2)) - \sum_{i,j\in\mathbb{N}} \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big|$$

$$= \Big|\sum_{i,j\in\mathbb{N}} \mu(I_i \times J_j) - \sum_{i,j\in\mathbb{N}} \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big|$$

$$= \Big|\sum_{i,j\in\mathbb{N}} \Big(\mu(I_i \times J_j) - \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) - i(\mu,\alpha_3^{mn}) - i(\mu,\alpha_4^{mn}))\Big)\Big|$$

$$\leq \sum_{i,j\in\mathbb{N}} \Big|\mu(I_i \times J_j) - \sum_{\substack{1\leq n\leq N(i)\\ 1\leq m\leq M(j)}} \frac{1}{2}(i(\mu,\alpha_1^{mn}) + i(\mu,\alpha_2^{mn}) + i(\mu,\alpha_3^{mn}) + i(\mu,\alpha_4^{mn}))\Big|$$

$$\leq \sum_{i,j\in\mathbb{N}} \frac{\epsilon}{(2^i)(2^j)} = \epsilon$$

This completes the proof since we have shown that we can measure $B(p_1,d_1)\times B(p_2,d_2)$ arbitrarily well using values of $i(\mu,\alpha)$, where $\alpha \in \pi_1(X)$.

□

Proof of Proposition 4.4.3: Suppose (X, g_1) and (X, g_2) have the same marked length spectra. Then by Proposition 4.4.7, $i(L_{g_1}, \alpha) = i(L_{g_2}, \alpha)$ for all $\alpha \in \pi_1(X)$. By Proposition 4.4.9, it then follows that $L_{g_1} = L_{g_2}$. Consider a geodesic $\beta \in \pi_1(DX)$ in $(DX, \overline{g_i})$ for $i = 1$ or 2. Note it is possible to find β_k such that $\ell_{\overline{g_i}}(\beta) = \ell_{\overline{g_i}}\left(\bigcup_{k\in\mathbb{N}} \beta_k\right)$, where each β_k is a geodesic arc lying entirely inside a single chamber of $(DX, \overline{g_i})$. Since each β_k is also a geodesic arc in a copy of a chamber in (X, g_i), β_k can be viewed as a geodesic arc in (X, g_i) as well. Using the fact that $L_{g_1} = L_{g_2}$, which allows us to recover lengths of geodesic arcs in X, it then follows that:

$$\ell_{\overline{g_1}}(\beta) = \ell_{\overline{g_1}}\left(\bigcup_{k\in\mathbb{N}} \beta_k\right) = \sum_{k\in\mathbb{N}} \ell_{\overline{g_1}}(\beta_k) = \sum_{k\in\mathbb{N}} \ell_{g_1}(\beta_k) = \sum_{k\in\mathbb{N}} \ell_{g_2}(\beta_k) = \sum_{k\in\mathbb{N}} \ell_{\overline{g_2}}(\beta_k) = \ell_{\overline{g_2}}\left(\bigcup_{k\in\mathbb{N}} \beta_k\right) = \ell_{\overline{g_2}}(\beta)$$

The key equality is the fourth one, which follows since $L_{g_1} = L_{g_2}$. We then conclude that the marked length spectra of $(DX, \overline{g_1})$ and $(DX, \overline{g_2})$ are the same, as desired.

□

4.4.3 Proof of Theorem 4.1.2

We now have all the ingredients needed for proving Theorem 4.1.2.

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 4.3.1, (X, g_1) and (X, g_2) are isometric via a map isotopic to identity, $\phi : (X, g_1) \to (X, g_2)$.

Otherwise, use the construction from the proof of Lemma 4.4.2 to obtain a finite-sheeted cover $\widehat{X}$ (possibly with index 1) such that no two boundaries components of the same chamber are glued together in $\widehat{X}$. Suppose now that every chamber in $\widehat{X}$ can be included into a closed surface. Then by Proposition 4.3.1, there exists an isometry isotopic to the identity $\widehat{\phi} : (\widehat{X}, \widehat{g_1}) \to (\widehat{X}, \widehat{g_2})$. Recall that $\widehat{\phi}$ is constructed by projecting the $\pi_1(\widehat{X})$-equivariant isometry $\widetilde{\phi} : (\widetilde{X}, \widetilde{g_1}) \to (\widetilde{X}, \widetilde{g_2})$ between copies of the universal cover $\widetilde{X}$ of both X and $\widehat{X}$. By construction, $\widetilde{\phi}$ is also $\pi_1(X)$-equivariant, so there exists some isometry

$\phi : (X, g_1) \to (X, g_2)$ isotopic to the identity as well.

If not every chamber in $\widehat{X}$ can be included into a closed surface, we can consider the double of $\widehat{X}$ along its gluing curves, $D\widehat{X}$ (see Figure 4.25 for an illustration of a case where both taking a finite-sheeted cover and doubling is required to reduce to the base case). Note that by construction, every boundary component of every chamber in $D\widehat{X}$ may now be paired with exactly one boundary component of an isometric chamber, so each chamber in $D\widehat{X}$ can be included into a closed surface. Furthermore, by Proposition 4.4.3, $(D\widehat{X}, \overline{\widehat{g_1}})$ and $(D\widehat{X}, \overline{\widehat{g_2}})$ have the same marked length spectrum. There is thus, again by Proposition 4.3.1, an isometry $\widehat{\phi}_D : (D\widehat{X}, \widehat{g_1}) \to (D\widehat{X}, \widehat{g_2})$ isotopic to the identity. Since restrictions of isometries are also isometries, it follows that the restriction $\widehat{\phi} = \widehat{\phi}_D|_{\widehat{X}} \colon (\widehat{X}, \widehat{g_1}) \to (\widehat{X}, \widehat{g_2})$ is also an isometry isotopic to the identity. Then apply the argument in the previous case to obtain an isometry $\phi : (X, g_1) \to (X, g_2)$ isotopic to the identity, as desired. □

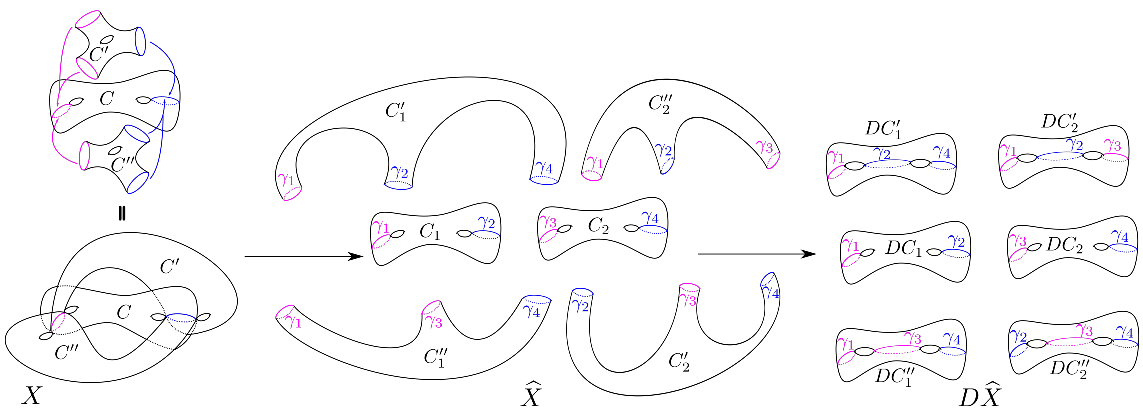

Figure 4.25: Since X does not fall into the base case, as C' and C'' are not included into closed surfaces, we construct a double cover of X based on the algorithm from Section 4.4.1. Note that $\widehat{X}$ also does not fall into the base case since C_1', C_2', C_1'' and C_2'' are each attached to a different triplet of gluing curves and thus are not included into any closed surfaces. After doubling along γ_i for $1 \le i \le 4$, we obtain $D\widehat{X}$, which falls under the base case.

Bibliography

[AB87] R. Alperin and H. Bass, "Length functions of group actions on Λ-trees," in *Combinatorial group theory and topology (Alta, Utah, 1984)*, ser. Ann. of Math. Stud. Vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 265–378.

[AL17] J. Aramayona and C. J. Leininger, "Hyperbolic structures on surfaces and geodesic currents," in *Algorithmic and geometric topics around free groups and automorphisms*, ser. Adv. Courses Math. CRM Barcelona, Birkhäuser/Springer, Cham, 2017, pp. 111–149.

[BB95] W. Ballmann and M. Brin, "Orbihedra of nonpositive curvature," *Inst. Hautes Études Sci. Publ. Math.*, no. 82, 169–209 (1996), 1995.

[BL18] A. Bankovic and C. J. Leininger, "Marked-length-spectral rigidity for flat metrics," *Trans. Amer. Math. Soc.*, vol. 370, no. 3, pp. 1867–1884, 2018.

[BF92] M. Bestvina and M. Feighn, "A combination theorem for negatively curved groups," *J. Differential Geom.*, vol. 35, no. 1, pp. 85–101, 1992.

[Bon86] F. Bonahon, "Bouts des variétés hyperboliques de dimension 3," *Ann. of Math. (2)*, vol. 124, no. 1, pp. 71–158, 1986.

[Bon88] F. Bonahon, "The geometry of Teichmüller space via geodesic currents," *Invent. Math.*, vol. 92, no. 1, pp. 139–162, 1988.

[Bow98] B. H. Bowditch, "Cut points and canonical splittings of hyperbolic groups," *Acta Math.*, vol. 180, no. 2, pp. 145–186, 1998.

[BH99] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature* (Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]). Springer-Verlag, Berlin, 1999, vol. 319, pp. xxii+643.

[BM96] M. Burger and S. Mozes, "CAT(-1)-spaces, divergence groups and their commensurators," *J. Amer. Math. Soc.*, vol. 9, no. 1, pp. 57–93, 1996.

[BK85] K. Burns and A. Katok, "Manifolds with nonpositive curvature," *Ergodic Theory Dynam. Systems*, vol. 5, no. 2, pp. 307–317, 1985, ISSN: 0143-3857.

[Con18] D. Constantine, "Marked length spectrum rigidity in non-positive curvature with singularities," *Indiana Univ. Math. J.*, vol. 67, no. 6, pp. 2337–2361, 2018.

[CL19a] D. Constantine and J.-F. Lafont, "Marked length rigidity for one-dimensional spaces," *J. Topol. Anal.*, vol. 11, no. 3, pp. 585–621, 2019.

[CL19b] D. Constantine and J.-F. Lafont, "Marked length rigidity for Fuchsian buildings," *Ergodic Theory Dynam. Systems*, vol. 39, no. 12, pp. 3262–3291, 2019.

[Coo93] M. Coornaert, "Mesures de Patterson-Sullivan sur le bord d'un espace hyperbolique au sens de Gromov," *Pacific J. Math.*, vol. 159, no. 2, pp. 241–270, 1993, ISSN: 0030-8730.

[CDST18] R. Coulon, R. Dougall, B. Schapira, and S. Tapie, "Twisted Patterson-Sullivan measures and applications to amenability and coverings," *arXiv e-prints*, arXiv:1809.10881, arXiv:1809.10881, Sep. 2018.

[CP08] J. Crisp and L. Paoluzzi, "Commensurability classification of a family of right-angled Coxeter groups," *Proc. Amer. Math. Soc.*, vol. 136, no. 7, pp. 2343–2349, 2008.

[CFF92] C. Croke, A. Fathi, and J. Feldman, "The marked length-spectrum of a surface of nonpositive curvature," *Topology*, vol. 31, no. 4, pp. 847–855, 1992.

[Cro90] C. B. Croke, "Rigidity for surfaces of nonpositive curvature," *Comment. Math. Helv.*, vol. 65, no. 1, pp. 150–169, 1990.

[CM87] M. Culler and J. W. Morgan, "Group actions on **R**-trees," *Proc. London Math. Soc. (3)*, vol. 55, no. 3, pp. 571–604, 1987.

[Dan20] P. Dani, "The large-scale geometry of right-angled Coxeter groups," in *Handbook of group actions. V*, ser. Adv. Lect. Math. (ALM), vol. 48, Int. Press, Somerville, MA, [2020] ©2020, pp. 107–141.

[DST18] P. Dani, E. Stark, and A. Thomas, "Commensurability for certain right-angled Coxeter groups and geometric amalgams of free groups," *Groups Geom. Dyn.*, vol. 12, no. 4, pp. 1273–1341, 2018.

[Dav08] M. W. Davis, *The geometry and topology of Coxeter groups* (London Mathematical Society Monographs Series). Princeton University Press, Princeton, NJ, 2008, vol. 32, pp. xvi+584.

[DLR10] M. Duchin, C. J. Leininger, and K. Rafi, "Length spectra and degeneration of flat metrics," *Invent. Math.*, vol. 182, no. 2, pp. 231–277, 2010.

[ES22] V. Erlandsson and J. Souto, *Mirzakhani's curve counting and geodesic currents* (Progress in Mathematics). Birkhäuser/Springer, Cham, [2022] ©2022, vol. 345, pp. xii+226.

[FM12] B. Farb and D. Margalit, *A primer on mapping class groups* (Princeton Mathematical Series). Princeton University Press, Princeton, NJ, 2012, vol. 49, pp. xiv+472.

[FK65] R. Fricke and F. Klein, *Vorlesungen über die Theorie der automorphen Funktionen. Band 1: Die gruppentheoretischen Grundlagen. Band II: Die funktionentheoretischen Ausführungen und die Andwendungen* (Bibliotheca Mathematica Teubneriana, Bande 3). Johnson Reprint Corp., New York; B. G. Teubner Verlagsgesellschaft, Stuttgart, 1965, vol. 4, Band I: xiv+634 pp., Band II: xiv+668.

[GM18] C. Guillarmou and M. Mazzucchelli, "Marked boundary rigidity for surfaces," *Ergodic Theory Dynam. Systems*, vol. 38, no. 4, pp. 1459–1478, 2018.

[Ham99] U. Hamenstädt, "Cocycles, symplectic structures and intersection," *Geom. Funct. Anal.*, vol. 9, no. 1, pp. 90–140, 1999.

[HP97] S. Hersonsky and F. Paulin, "On the rigidity of discrete isometry groups of negatively curved spaces," *Comment. Math. Helv.*, vol. 72, no. 3, pp. 349–388, 1997.

[HST20] G. C. Hruska, E. Stark, and H. C. Tran, "Surface group amalgams that (don't) act on 3-manifolds," *Amer. J. Math.*, vol. 142, no. 3, pp. 885–921, 2020.

[Kai94] V. A. Kaimanovich, "Ergodicity of harmonic invariant measures for the geodesic flow on hyperbolic spaces.," *Journal für die reine und angewandte Mathematik*, vol. 455, pp. 57–104, 1994.

[KB02] I. Kapovich and N. Benakli, "Boundaries of hyperbolic groups," in *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, ser. Contemp. Math. Vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.

[KK00] M. Kapovich and B. Kleiner, "Hyperbolic groups with low-dimensional boundary," *Ann. Sci. École Norm. Sup. (4)*, vol. 33, no. 5, pp. 647–669, 2000.

[Laf04] J.-F. Lafont, "Rigidity result for certain three-dimensional singular spaces and their fundamental groups," *Geom. Dedicata*, vol. 109, pp. 197–219, 2004.

[Laf06] J.-F. Lafont, "Strong Jordan separation and applications to rigidity," *J. London Math. Soc. (2)*, vol. 73, no. 3, pp. 681–700, 2006.

[Laf07a] J.-F. Lafont, "Rigidity of hyperbolic P-manifolds: A survey," *Geom. Dedicata*, vol. 124, pp. 143–152, 2007.

[Laf07b] J.-F. Lafont, "Diagram rigidity for geometric amalgamations of free groups," *J. Pure Appl. Algebra*, vol. 209, no. 3, pp. 771–780, 2007.

[Laf02] J.-F. R. Lafont, *Rigidity results for singular spaces.* ProQuest LLC, Ann Arbor, MI, 2002, p. 93, Thesis (Ph.D.)–University of Michigan.

[Mal10] W. Malone, *Topics in geometric group theory*. ProQuest LLC, Ann Arbor, MI, 2010, p. 79, Thesis (Ph.D.)–The University of Utah.

[Mos68] G. D. Mostow, “Quasi-conformal mappings in n-space and the rigidity of hyperbolic space forms,” *Inst. Hautes Études Sci. Publ. Math.*, no. 34, pp. 53–104, 1968.

[Ota90] J.-P. Otal, “Le spectre marqué des longueurs des surfaces à courbure négative,” *Ann. of Math. (2)*, vol. 131, no. 1, pp. 151–162, 1990.

[Ota92] J.-P. Otal, “Sur la géometrie symplectique de l’espace des géodésiques d’une variété à courbure négative,” *Rev. Mat. Iberoamericana*, vol. 8, no. 3, pp. 441–456, 1992.

[Ric17] R. Ricks, “Flat strips, Bowen-Margulis measures, and mixing of the geodesic flow for rank one CAT(0) spaces,” *Ergodic Theory Dynam. Systems*, vol. 37, no. 3, pp. 939–970, 2017.

[Sta18] E. Stark, “Topological rigidity fails for quotients of the Davis complex,” *Proc. Amer. Math. Soc.*, vol. 146, no. 12, pp. 5357–5366, 2018.

[Whi01] A. T. White, *Graphs of groups on surfaces* (North-Holland Mathematics Studies). North-Holland Publishing Co., Amsterdam, 2001, vol. 188, pp. xiv+363, Interactions and models.

[Wu23a] Y. Wu, “Marked length spectrum rigidity of surface amalgams,” pp. 1–41, Oct. 2023.

[Wu23b] Y. Wu, “A topologically rigid set of quotients of the Davis complex,” *Geom. Dedicata*, vol. 217, no. 5, Paper No. 82, 20, 2023.

Appendix A

Ergodicity of the geodesic flow map on Gromov hyperbolic spaces

The main purpose of this appendix is to prove the following statement, which is essential to the proof of Lemma 4.3.5:

Proposition A.0.1. *The geodesic flow map on (Gromov) hyperbolic P-manifolds is ergodic.*

The ideas and terminology of this section are mostly taken from [Kai94], 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 [Ric17], but for simplicity, we stick to the ideas in [Kai94].

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), nonelementary, 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 Definition 4.2.3 for details). Then the *(Gromov) hyperbolic compactification* $\overline{X}$ of X may be defined as $\overline{X} = X \cup \partial^\infty(X)$.

Following [Kai94], we impose two additional assumptions:

Assumption A.0.2 (Uniqueness of geodesics)**.** For all $x_1, x_2 \in \overline{X}$, there exists a unique geodesic $[x_1, x_2]$ joining them.

Assumption A.0.3 (Existence of convergent geodesics)**.** If α_1 and α_2 are asymptotic (e.g. they lie within bounded distance of each other), then there exists some $c > 0$ such that:

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

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

A.0.1 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 $\Gamma x = \{\gamma x | \gamma \in \Gamma\}$ 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:

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

where $r : \mathbb{R} \to X$ 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 A.0.4. 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 A.0.4 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 *Poincare series* associated to Γ is the series $P(x,s) = \sum_{\gamma\in\Gamma} e^{sd(x,\gamma.x)}$, 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 [Coo93]). Note that if X is a proper geodesic space and X/Γ is compact, then δ_Γ is finite (see Proposition 1.7 of [BM96]). Furthermore, if the action of Γ on X is non-elementary, then the critical exponent is nonzero. Thus, in the case where X is a simple, thick P-manifold and $\Gamma = \pi_1(X)$, δ_Γ is nonzero 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 the following theorem. 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 A.0.5 ([Coo93], 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 Theoerem 5.4 in [Coo93] 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 [Coo93]) and the Patterson-Sullivan measure is defined as the weak limit of probability measures:

$$\mu_x = \lim_{n\to\infty} \frac{1}{\sum\limits_{\gamma\in B_n} e^{-\delta_\Gamma d(x,\gamma.x)}} \sum_{\gamma\in B_n} e^{-\delta_\Gamma d(x,\gamma.x)} \mathrm{Dirac}_{\gamma.x}$$

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

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

A.0.2 Bowen-Margulis Measures

Given μ_x with density δ, Kaimanovich defines a Γ-invariant measure on $\partial^\infty(X)\times\partial^\infty(X)-\Delta$ (see Section 2.4.1 in [Kai94]):

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

where $a, b \in \partial^\infty(X)$, $x \in X$ and $g_{\infty,x}$ is the visual metric from Definition 4.2.7. Note that Patterson-Sullivan measures with different basepoints are absolutely continuous with respect to each other. We can then denote $\nu \in [\{\nu_x\}_{x\in X}]$ as some measure in this equivalence class of measures. We can thus define a Γ-invariant Radon measure on $\partial^\infty(X)\times\partial^\infty(X)-\Delta$, which we will also denote as ν.

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

$$m = \nu \times dt$$

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

Suppose that $\widetilde{X}$ is a (Gromov) hyperbolic covering space with Γ its deck group. Unsurprisingly, for a Gromov hyperbolic metric space $X = \widetilde{X}/\Gamma$, there is a one-to-one correspondence between Γ-invariant Radon measures on $\partial^\infty(X)\times\partial^\infty(X)-\Delta$ and Radon measures on $SX = S\widetilde{X}/\Gamma$ that are invariant under the geodesic flow (see Theorem 2.2 of [Kai94]). Thus, since ν is Γ-invariant, m_Γ is invariant under geodesic flow.

Definition A.0.6 (Bowen-Margulis measures)**.** The geodesic flow-invariant measure m_Γ defined above is a *Bowen-Margulis measure* on SX/Γ.

Remark A.0.7. 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 discussed in this paper, the Liouville current (see Definition 4.4.5), is an interesting open question.

A.0.3 Proof of Proposition A.0.1

We are now ready to state a key theorem from [Kai94]:

Theorem A.0.8 ([Kai94], Theorem 2.6). *Let $\widetilde{X}$ be a (Gromov) hyperbolic covering space satisfying Assumptions A.0.2 and A.0.3. 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 $SX = S\widetilde{X}/\Gamma$ 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 CAT(-1) metric, recall from Theorem A.0.5 that $\text{supp}(\mu_x) = \Lambda_\Gamma$; therefore, the geodesic flow on SX is ergodic with respect to m_Γ, as desired.

An Application of Proposition A.0.1

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 A.0.9 (Birkhoff Ergodic Theorem). *Let $(Y, \mathscr{B}, \mu)$ be a probability space, and let $T : Y \to Y$ be an ergodic measure-preserving transformation. Let $A \in \mathscr{B}$ be a measurable set of positive measure $\mu(A) > 0$. Then for all $f \in \mathscr{L}^1(Y, \mathscr{B}, \mu)$ and μ-almost everywhere $y \in Y$:*

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

If Γ is a nonelementary 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 [CDST18]) and in particular can be scaled to a probability measure, so we can apply the

Birkhoff Ergodic Theorem. For example, if X is a 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 [CDST18], 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 A.1 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 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 A.1), as every open set of SX has nonzero measure.