Aspects of Hyperbolicity in Groups and Spaces

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree
Doctor of Philosophy in the Graduate School of The Ohio State
University

By

Hyeran Cho, M.S.

Graduate Program in Mathematics

The Ohio State University

2025

Dissertation Committee:

Jean-François Lafont, Advisor

James Fowler

Jingyin Huang

Abstract

This paper presents two related research projects that explore different aspects of hyperbolicity in both group-theoretic and geometric contexts.

In part 1, we construct a random model for an n-fold branched cover of a finite 2-complex X. With mild hypotheses on X, we show that as n goes to infinity, a random branched cover is asymptotically almost surely homotopy equivalent to a 2-complex satisfying geometric small cancellation. As a consequence, the fundamental group is asymptotically almost surely Gromov hyperbolic and has small cohomological dimension.

In part 2, we consider geometric and topological properties of cusped spaces and branched coverings of finite-volume manifolds modeled on $\mathbb{H}^m \times \mathbb{H}^n$. Let M be a finite-volume manifold whose universal cover $\tilde{M}$ is isometric to $\mathbb{H}^m \times \mathbb{H}^n$, and let S be a compact, totally geodesic, codimension-two submanifold of M whose lift to $\tilde{M}$ is isometric to $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. We consider manifolds $N = M \setminus S$ and X_d, an n-fold branched covering over S. For a product lattice $\Gamma = \Gamma_m \times \Gamma_n$ that gives rise to M, we prove that N admits a complete, finite-volume, A-regular metric with nonpositive sectional curvature and X_d also admits a nonpositively curved Riemannian metric. More generally, for any lattice $\Gamma \leq \text{Isom}(\mathbb{H}^m \times \mathbb{H}^n)$ and $\epsilon > 0$, we show that N and X_d admit complete Riemannian metrics of almost nonpositive sectional curvature, with volume bounds independent of ϵ.

To Min Ho and Haun

Acknowledgments

First and foremost, I would like to express my heartfelt gratitude to my advisor, Jean-François Lafont, for his invaluable academic guidance and unwavering support throughout my doctoral journey. His mentorship went far beyond mathematics—he has been a steady and fatherly presence during many difficult moments of my time abroad.

I thank my collaborators-Rachel Skipper, for many stimulating discussions and her generous support that allowed me to participate in conferences and present my work. I am equally grateful to Barry Minemyer. It has been a privilege to collaborate with him, and I deeply admire his thoughtful and dedicated approach to research.

I would also like to thank Jim Fowler and Jingyin Huang for kindly serving on both my candidacy and defense committees, and for their time and thoughtful engagement with my work.

Special thanks go to Sang Youl Lee and Young Ho Im at Pusan National University, who first introduced me to the excitement of research and encouraged me to pursue my studies abroad.

To my husband, Min Ho, thank you for being my partner, best friend, and unwavering supporter through every challenge. And to our beloved son Haun, the most precious part of my life—your very existence has given my efforts new meaning.

I am forever grateful to my parents, my sister Yoonjung and her family for their endless love and belief in me. I also cherish the comfort brought by Lang, Romi, and the cheerful encouragement of Relay 3rd. I would like to thank my in-laws, especially my father-in-law, Joong Jae Cho, for his warm support not only as family, but also as a fellow scholar.

To my dear friends Kriti and Katelynn, thank you for always bringing me joy and for being sources of motivation and joy. I am especially thankful to Jonghoo, Nick, Shreeya, and Chaeeun for making this experience more meaningful. Finally, to Bitna, thank you for your lasting friendship and encouragement from afar.

Vita

January 31, 1989 .... Born - Republic of Korea

2015 .... B.S. in Mathematics Education,
Pusan National University

2017 .... M.S. in Mathematics,
Pusan National University

2018-present .... Graduate Teaching Associate,
The Ohio State University.

Publications

Research Publications

H. Cho, SY. Lee, H. Song "Derivation of Schubert normal forms of 2-bridge knots from $(1, 1)$-diagrams". *Journal of Knot Theory and Its Ramifications*, 28(2019), no.8,1950048

Fields of Study

Major Field: Mathematics, specializing in Geometric Group Theory and Topology

Table of Contents

List of Figures

Chapter 1: Hyperbolicity of Random Branched Coverings

1.1 Introduction

The probabilistic method was originally pioneered by Erdös, and was used as a non-constructive approach to showing the existence of interesting examples in combinatorics and graph theory, see e.g. the classic text by Alon—Spencer [1]. The method has been particularly effective in graph theory, with random graphs having evolved into its own field of study. Random models have since been developed to study higher dimensional simplicial complexes (Kahle [28]), random closed surfaces (Brooks—Makover [8]), and random 3-manifolds (Dunfield—Thurston [11]). In the late 1980s, Gromov launched the study of random groups. The two main models for random groups are the density model and the few relator model, see Gromov [18] and the survey article by Ollivier [37].

A common theme in these approaches is that a space is built by attaching spaces together via a random process. In the case of random graphs or simplicial complexes, edges or simplices are added to a vertex set at random. In the setting of random surfaces or 3-manifolds, triangles or handlebodies are glued together via a suitable random process. In the setting of random groups, one can think of relations as 2-cells being randomly attached to a bouquet of circles, with the random group being the fundamental group of the resulting 2-complex.

A different topological construction that is commonly used in low-dimensional topology is that of branched covers. All closed oriented surfaces can be realized as branched covers over the sphere, a fact that is also true for 3-manifolds (Hilden [27] and Montesinos [35]). Branched covers have also been a source of many interesting examples in the geometry of negatively curved manifolds, see e.g. Gromov—Thurston [21], Fine—Premoselli [14], Stover—Toledo [41], Minemyer [33], Guenancia-Hamenstädt [24]. From this viewpoint, it is natural to look for a random model for branched covers. In the present paper, we construct a random model for branched covers of finite polygonal 2-complexes and prove:

Main Theorem. *Let X be an acceptable finite polygonal 2-complex, and fix $\lambda > 0$ an arbitrarily constant. Let $X(\sigma)$ be an n-fold random branched cover of X. Then $X(\sigma)$ is asymptotically almost surely homotopy equivalent to a 2-complex satisfying geometric $C'(\lambda)$-small cancellation.*

By a *polygonal* 2-complex, we mean a 2-dimensional CW-complex where the attaching maps are particularly simple, as described here. The 1-skeleton is metrized by having each edge of length one and having a prescribed orientation. The 2-cells are identified with disks scaled so that their perimeter is an integer. We then subdivide the boundary of the 2-cell into consecutive intervals of length one, and the attaching map is required to map each of these intervals isometrically onto a single edge of the 1-skeleton. For these CW-complexes, the attaching maps for each disk can be described just be enumerating the sequence of edges traversed in the 1-skeleton. Note that we are allowing the possibility of a disk attaching to a single edge loop, or to a pair of edges (so paths of combinatorial length one or two).

By an *acceptable* polygonal 2-complex, we mean one which satisfies the following additional conditions:

1. the 1-skeleton $X^{(1)}$ has fundamental group of rank at least two;

2. the 2-cells have attaching maps that are not proper powers in $\pi_1(X^{(1)})$ (so in particular, are non-trivial), and that are pairwise distinct.

It is easy to see that the Main Theorem **fails** for non-acceptable 2-complexes, so our hypotheses are actually necessary.

It is well-known that the $C'(1/6)$-small cancellation property has strong geometric consequences [23]. As a result, the $\lambda = 1/6$ case of our main theorem immediately implies the following:

Corollary 1.1.1. *Let X be an acceptable finite 2-complex, and let $X(\sigma)$ be a n-fold random branched cover of X. Then we have asymptotically almost surely the following properties hold:*

- $\pi_1(X(\sigma))$ *is Gromov hyperbolic and cubulable;*

- $X(\sigma)$ *is aspherical, hence an Eilenberg–MacLane space* $K(\pi_1(X(\sigma)), 1)$*;*

- *The cohomological dimension of* $\pi_1(X(\sigma))$ *is at most* 2.

1.2 Preliminaries

1.2.1 Branched Coverings

Let us recall the notion of branched covering of smooth manifolds (see e.g. [16]).

Definition 1.2.1. (n-fold branched covering of manifolds) Given a pair of smooth k-manifolds X^k, Y^k, an n-fold branched (or ramified) covering is a smooth, proper map $f : X^k \to Y^k$ exhibiting some particularly simple local form. The critical set $B^{k-2} \subset Y$ is called the branch locus, and we require that it is a smoothly embedded codimension two submanifold. Moreover, $f|_{X \setminus f^{-1}(B)} : X \setminus f^{-1}(B) \to Y \setminus B$ is a covering map of degree n, and for each $p \in f^{-1}(B)$ there are local coordinate charts $U, V \to \mathbb{C} \times \mathbb{R}^{k-2}$ about $p, f(p)$ on which f is given by $(z, x) \mapsto (z^m, x)$ for some positive integer m called the branching index of f at p.

Notice that, when restricted to the branching locus B, the map $f|_{f^{-1}(B)}$ is just an ordinary covering map. The pre-image of B is not assumed to be connected, and indeed, could have multiple connected components. Transverse to the branching locus B, f behaves like the map $z \to z^m$ near the origin – though again, at different pre-image points the value of m might be different.

This definition can be readily extended to the setting of CW-complexes. As we will only need the 2-dimensional case, we will focus on that case.

Definition 1.2.2. (n-fold branched covering of 2-complexes) Given a pair of finite 2-dimensional CW-complexes X, Y, an n-fold branched (or ramified) covering is a continuous map $f : X \to Y$ satisfying the following property. There is a finite subset of points $B \subset Y$, called the branching locus, which satisfies $B \cap Y^{(1)} = \emptyset$ (so B lies in the interior of the 2-cells). Moreover, $f|_{X \setminus f^{-1}(B)} : X \setminus f^{-1}(B) \to Y \setminus B$ is a covering map of degree n, and for each $p \in f^{-1}(B)$ there are local coordinate charts $U, V \to \mathbb{C}$ about $p, f(p)$ on which f is given by $z \mapsto z^m$ for some positive integer m, called the branching index of f at p.

Note that, in the case where a 2-cell in Y contains more than one branch point, connected components of its preimage might no longer be homeomorphic to a disk. However, if there is a single branch point inside a 2-cell, then each component in its preimage will be homeomorphic to a disk.

1.2.2 Small Cancellation Conditions

Small cancellation has been a useful tool in combinatorial group theory since the 1970s, see e.g. the references [30], [22], [23], and [25]. There are various notions of small cancellation. Here, we start by recalling the metric small cancellation (or classical small cancellation) with respect to a group presentation. Roughly speaking, this condition says that any common subword between two relators in a presentation is short compared to the length of the relators.

Let X be a symmetric generating set for a group Γ, i.e. X contains all elements of a generating set S and their inverses. We call an element of S a *letter*. A *word* w is finite string of letters $w = s_1 \dots s_m$. We consider w as an element of the free group F with the generating set S. Then each element of F other than the identity 1 has a unique representation as a *reduced word* $w = s_1 \dots s_n$ in which no two successive letters $s_i s_j$ form an inverse pair $s_i s_i^{-1}$. The integer n is the *length* of w, which we denote by $|w|$. A reduced word is called *cyclically reduced* if s_n is not the inverse of s_1. If there is no cancellation in the product $z = u_1 \dots u_n$, we write $z \equiv u_1 \dots u_n$.

A subset R of F is called *symmetrized* if all elements of R are cyclically reduced and for each r in R, all cyclically reduced cyclic permutations of both r and r^{-1} also belong to R. Suppose that $r_1 \equiv bc_1$ and $r_2 \equiv bc_2$ are distinct elements of R. If b is the maximal such subword then it is called a *piece relative to the set* R or simply a *piece*.

Definition 1.2.3. We say that R satisfies the small cancellation condition $C'(\lambda)$ if for $r \in R$ with $r \equiv bc$ where b is a piece, then $|b| < \lambda|r|$. In this case, we also say that the presentation satisfies $C'(\lambda)$. Also, for a group Γ, if there is a presentation that satisfies $C'(\lambda)$, we say that Γ is $C'(\lambda)$ group.

The following is well known:

Proposition 1.2.4. *[22] If a finitely presented group* Γ *satisfies* $C'(\frac{1}{6})$*, then* Γ *is word hyperbolic.*

Geometric consequences of the small cancellation hypothesis are further studied in [30], typically via the group's presentation 2-complex, as well as van Kampen diagrams and their mapping to 2-complexes. This allows the small cancellation condition to be reformulated geometrically, and the results to be generalized to the setting of polygonal 2-complexes.

Recall that the attaching maps for the 2-cells in a polygonal 2-complexes are given by a (cyclic) sequence of directed edges from the 1-skeleton. For each 2-cell D in a polygonal 2-complex, the boundary ∂D is a cycle graph, and one can label the edges of ∂D according to the directed edges they map to in the 1-skeleton.

We will consider combinatorial subpaths b in ∂D which are injective on their interior (so at most, agree at the two endpoints). A *subpiece* is a subpath in the boundary of a pair of disks ∂D, $\partial D'$, whose labels, including orientation, are identical. Note that D' can possibly be the same disk D, but with distinct initial vertices for the two subpaths of the boundary. In that case a subpiece would be contained in the self-intersection of the attaching map on ∂D. A *piece* is a subpiece in the boundary of a pair of disks $\partial D, \partial D'$ which is maximal under containment.

Remark 1.2.5. Maximality gives us some insight on the local behavior of the labels on D, D' near the endpoints. If the endpoints of the piece are distinct vertices in ∂D, $\partial D'$, then maximality tells us that at the initial (and terminal) vertex of the path b, the previous (resp. following) edges of ∂D and $\partial D'$ have to be distinct.

The other possibility is that the endpoints of the piece b coincide in D (for example). Note that this means the label on D' contains an entire copy of the label on D.

We say that D satisfies $C'(\lambda)$ if for any piece b of D, we have

$$\frac{\ell(b)}{\ell(\partial D)} < \lambda$$

where ℓ is the combinatorial path length. We say that a 2-complex satisfies $C'(\lambda)$ if each of its 2-cells satisfy $C'(\lambda)$. This provides us with a geometric notion of small cancellation, and results on groups satisfying small cancellation (established via analysis of the presentation 2-complex) readily generalize to 2-complexes satisfying geometric small cancellation.

In our later constructions, we will consider certain finite covers of the 1-skeleton of X, with certain lifts of attaching maps. Given a 2-cell D with attaching map $\alpha : \partial D \to X^{(1)}$, we have an associated map $\tilde{\alpha} : \mathbb{R} \to X^{(1)}$ obtained by composing the universal covering map $\pi : \widetilde{\partial D} \to \partial D$ with the attaching map. We can identify $\widetilde{\partial D}$ with $\mathbb{R}$ equipped with its standard simplicial structure. The map $\tilde{\alpha}$ is then described by the bi-infinite, periodic word obtained by lifting the edge labels from ∂D to $\widetilde{\partial D} \cong \mathbb{R}$. Since we will be interested in studying small cancellation properties associated to some of these covers, we now formulate a notion that is slightly more general than a piece.

A *sub-overlap* between two disks D, D' is a pair of finite combinatorial subpaths $\mathbf{p} \subset \widetilde{\partial \mathbf{D}}$ and $\mathbf{p}' \subset \widetilde{\partial \mathbf{D}'}$ on which the lifted attaching maps $\tilde{\alpha} : \widetilde{\partial D} \to X^{(1)}$ and $\tilde{\beta} : \widetilde{\partial D'} \to X^{(1)}$ coincide. Note that the paths $\mathbf{p}$ and $\mathbf{p}'$ are not required to respect the orientation on the real line. Sub-overlaps are considered equivalent if they differ by translation by the $\pi_1(\partial D)$ and $\pi_1(\partial D')$ actions and reversing the orientation on both $\mathbf{p}$ and $\mathbf{p}'$. We also allow the case where $D = D'$, in which case we also require either the orientation on $\mathbf{p} \subset \widetilde{\partial \mathbf{D}}$ and $\mathbf{p}' \subset \widetilde{\partial \mathbf{D}'}$ to be different; or the starting points of the overlaps to be in distinct orbits of the $\pi_1(\partial D)$-action (i.e. correspond to distinct initial points in ∂D).

Definition 1.2.6. An *overlap* is a sub-overlap which is maximal under containment The *overlap ratio* of D with D' is defined to be

$$o(D, D') = \sup_{\mathbf{p}} \frac{\ell(\mathbf{p})}{\ell(\partial D)}$$

where the supremum is over all overlaps $(\mathbf{p}, \mathbf{p}')$ between D and D'. The overlap ratio of a 2-cell D is then defined to be $o(D) := \sup_{D'} o(D, D')$, and the overlap ratio of the polygonal 2-complex X is defined by $o(X) = \sup_D o(D) = \sup_{D,D'} o(D, D')$.

Observe that a piece whose length is strictly smaller than the length of both ∂D and $\partial D'$ is automatically an overlap. In particular, for $\epsilon < 1$ the $C'(\epsilon)$ small cancellation condition is implied by the statement that $o(X) < \epsilon$. On the other hand, when the overlap ratio of a pair satisfies $o(D, D') \geq 1$, this just means that the label for ∂D is entirely contained in the label for $\partial D'$.

Lemma 1.2.7. *Let X be an acceptable polygonal 2-complex, $(\mathbf{p}, \mathbf{p}')$ an overlap between disks D, D', and set $M = \max\{\ell(\partial D), \ell(\partial D')\}$. Then the length of the overlap is bounded above by $\ell(\mathbf{p}) < M^2 + M$. In particular, for any pair D, D' of 2-cells, the overlap ratio $o(D, D')$ is finite. It follows that for any finite acceptable polygonal 2-complex, the overlap ratio $o(X)$ is finite.*

Proof. Let us consider the case where $D \neq D'$, and assume that the overlap $(\mathbf{p}, \mathbf{p}')$ has length

$$\ell(\mathbf{p}) \geq M^2 + M.$$

Since the lifted attaching maps $\bar{\alpha}, \bar{\beta}$ coincide on $\mathbf{p}, \mathbf{p}'$, these paths have identical edge labelings. Moreover, these labeled paths can be viewed as subpaths of the labeled bi-infinite paths $\widetilde{\partial D}, \widetilde{\partial D'}$. These bi-infinite labeled paths are periodic with respect to the $\pi_1(\partial D)$-action and $\pi_1(\partial D')$-action, which are translations by $\ell(D), \ell(D')$ respectively. Since the common subpath contains fundamental domain for both translations, one can apply the Euclidean algorithm to find a subpath of length $r = GCD(\ell(D), \ell(D'))$ that tiles both fundamental domains. To see this, consider the initial subpath $\mathbf{q} \subset \mathbf{p}$ of length r. Since $r = GCD(\ell(D), \ell(D'))$, the Euclidean algorithm provides us with an integral solution to Bézout's identity

$$r = A\ell(D) + B\ell(D')$$

and the solution satisfies $|A| \leq \ell(D'), |B| \leq \ell(D)$. Note that exactly one of the integers A, B is positive, the other is negative. Assume without loss of generality that A is positive. Viewing $\mathbf{q} \subset \mathbf{p}'$ as the initial segment of $\mathbf{p}'$, and using $\pi_1(D')$-periodicity of $\widetilde{\partial D'}$, we can translate A times along $\widetilde{\partial D'}$. Since the length of $\mathbf{p}'$ is $> M^2 + M$, the translate of q is still contained within $\mathbf{p}'$. We now switch to viewing that translate as contained in $\mathbf{p}$, and use $\pi_1(D)$-periodicity of $\widetilde{\partial D}$ to translate B times along $\widetilde{\partial D}$. This has the effect of translating $\mathbf{q}$ by exactly r, and hence the initial portion of $\mathbf{p}$ of length $2r$ consists of two copies of $\mathbf{q}$. We can iterate this process M/r-times, noting that the hypothesis that $\ell(\mathbf{p}) > M^2 + M$ guarantees that the forward and backward translates from Bézout's identity land within the common subword $\mathbf{p}, \mathbf{p}'$. If $r < M$, this tells us that the larger of the two words is a proper power, contradicting the definition of acceptable 2-complex. On the other hand, if $r = M$ we get that $\ell(D) = \ell(D')$, and the two disks D, D' are attached along the same map, which again contradicts the definition of acceptable 2-complex.

Next we consider the case where $D = D'$, i.e. self-overlaps. Then one has that the bi-infinite label on $\widetilde{\partial D}$ is periodic with respect a translation by $\ell(D)$. Since the subwords $\mathbf{p}, \mathbf{p}'$ differ by a translation by some $0 < k < \ell(D)$, the subword $\mathbf{p}$ is also k-periodic. Then as before we can apply the Euclidean algorithm to obtain a solution to Bézout's identity, and use combinations of $\ell(D)$-translations and k-translations to see that the $\mathbf{p}$ is actually periodic with period $r = GCD(k, \ell(D)) < \ell(D)$. This implies the attaching map for D is a proper power, which again contradicts X acceptable.

Finally, since X has only finitely many pairs of 2-cells, the supremum of the overlap ratios will still be finite. □

Corollary 1.2.8. *Any acceptable polygonal* 2*-complex* X *only contains finitely many equivalence classes of overlaps* $(\mathbf{p}, \mathbf{p}')$.

Proof. For a given pair of 2-cells D, D', we can count the overlap pairs $(\mathbf{p}, \mathbf{p}')$. Up to the action of $\pi_1(\partial D)$, the initial point of the path $\mathbf{p}$ can be chosen in a fixed fundamental domain $F \subset \widetilde{\partial D}$, where F is a combinatorial interval of length $\ell(\partial D)$. Thus there are at most $\ell(\partial D)$ possible initial vertices for $\mathbf{p}$. From the Lemma, there is also a uniform bound of $\max\{\ell(\partial D), \ell(\partial D')\}$ on the length of $\mathbf{p}$. Thus there are at most finitely many possibilities for the path $\mathbf{p}$. A symmetric argument shows that there are at most finitely many choices for $\mathbf{p}'$, hence finitely many possibilities for the pairs $(\mathbf{p}, \mathbf{p}')$. Since the acceptable 2-complex X has a finite number of 2-cells, the corollary follows. □

1.3 Random Branched Coverings

In this section, we describe our random model for branched coverings of finite 2-complexes, with the motivating example being the case of the presentation 2-complex of a finitely presented group. We then establish a few basic properties concerning the behavior of 2-cells in our random model.

1.3.1 Branched Coverings of presentation 2-complexes

Let us first focus on the setting of a presentation 2-complex. Let $\Gamma = \langle u_1, \ldots, u_t \mid r_1, \ldots, r_s \rangle$ be a finite presentation of a group Γ, and let X be the presentation 2-complex for the fixed group presentation above. The complex consists of a single vertex v along with oriented loops $x_1, \ldots, x_t$ corresponding to each generator, and 2-disks $D_1, \ldots, D_s$ that are attached by the attaching maps $r_1, \ldots, r_s$ corresponding to the relators. For such complexes, being an acceptable 2-complex just means that the finitely presented group has at least two generators, that no relation is a proper power.

Remark 1.3.1. Conversely, any finite polygonal 2-complex that has a single vertex can be viewed as a presentation 2-complex for its fundamental group. These spaces have the advantage of having a canonical basepoint for the fundamental group. To deal with the general case of a finite connected polygonal 2-complex X, we can contract a spanning tree T for the 1-skeleton to obtain a 1-vertex 2-complex X', and use the branched cover model for covers of X'. The details can be found in Section...

Next let us describe an n-fold branched covering of X, with branching locus B consisting of the set of centers of the 2-disks $D_1, \ldots, D_s$. Consider a not necessarily connected n-fold covering of the 1-skeleton $X^{(1)}$. Since the presentation 2-complex X has a single vertex, an n-fold covering of $X^{(1)}$ has n vertices, which we will label v_1 to v_n. Under the covering map, the pre-image of each directed loop x_i will consist of n directed edges which we denote $x_{i1}, \ldots, x_{in}$. Our labeling convention is to label the lifted edge x_{ij} to originate at the lifted vertex v_j. Then each vertex v_j has some lifted edge x_{ik} coming in, and the lifted edge x_{ij} going out. Note the situation where $j = k$ corresponds to the case where the lifted edge starting at v_j is a loop at the vertex.

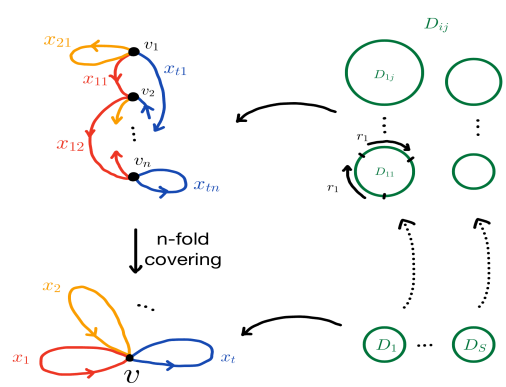

Figure 1.1: Branched cover of a presentation 2-complex

Given such a covering of $X^{(1)}$, we then attach a disk to each of the lifts of the attaching map $r_1, \ldots, r_s$. Note that each disk D_i in X has its boundary labeled

r_i, and for each vertex in the cover we have a path starting at that vertex and following the letters of r_i. If this path in the cover ends at a vertex different from the starting vertex, we follow the letters of r_i again, repeating the process until the word ends at the starting vertex. In this process, we get a closed path and we attach a disk boundary along this path. We call the disks D_{ij}, for suitable indexing set for j, and refer to D_{ij} as a *lift* of the disk D_i. Note that some lifts have the same attaching map up to cyclic permutation. We consider these as the same disk and do not include repetitions. The number of lifts that are considered as a single disk is denoted by $\mathrm{ind}(D_{ij})$ and called the *index* of the disk D_{ij}. Attaching all lifts of disks $D_1, \ldots, D_s \subset X$, we obtain a *branched covering* of X where the branching locus is the set of the centers $c_i \subset D_i$ of all of the disks in the original 2-complex. See the figure above for an illustration of this process.

Remark 1.3.2. The branching index of each center $c_{ij} \subset D_{ij}$ is the same as the index of the disk D_{ij}. The combinatorial lengths of the attaching maps are related via the simple formula

$$\ell(\partial D_{ij}) = ind(D_{ij})\ell(\partial D_i) = ind(D_{ij})|r_i|$$

where ℓ is the combinatorial length.

1.3.2 Random model for branched coverings

We now proceed to define our random model for branched coverings of a presentation 2-complex, which allows us to randomly pick a degree n branched cover of the 2-complex X. We call this the *random labeled branched cover model.* Note that, as detailed in the previous section, each degree n branched cover of the 2-complex X determines, and is determined by, an ordinary degree n cover of the 1-skeleton $X^{(1)}$.

Since in our special case $X^{(1)}$ is a bouquet of t circles, it is easy to describe the degree n covers of this graph.

Labeling the n pre-images of the single vertex v by labels $V = \{v_1, \ldots, v_n\}$, covering space theory tells us that associated to each loop x_i in X, we have a permutation σ_i of the vertex set V. The collection of permutations is determined by the finite cover, and conversely, determines the cover up to label preserving isomorphism. Thus there is a bijection between the set of degree n branched covers, and elements in the product of t copies of the symmetric group $\mathrm{Sym}(n)$.

A random n-fold covering of $X^{(1)}$ can now be generated by choosing t random permutations $\sigma = (\sigma_1, \ldots, \sigma_t)$ with uniform distribution, where each $\sigma_i \in \mathrm{Sym}(n)$ corresponds to each generator $u_i, i = 1, \ldots, t$. For a generator u_i and its corresponding loop x_i in $X^{(1)}$, a permutation σ_i represents the preimage of x_i in the n-fold covering space. More precisely, if the permutation σ_i maps the integer a to the integer b, then there exists an oriented pre-image of the edge x_i that joins the vertex v_a to the vertex v_b. For a finite presentation of a group Γ, the random choice of $\sigma = (\sigma_1, \ldots, \sigma_t)$ where $\sigma_i \in \mathrm{Sym}(n)$ completely determines a n-fold covering of $X^{(1)}$, and thus determines the n-fold branched covering after lifting the attaching maps. We denote the branched covering by $X(\sigma)$.

Example 1.3.3. Let $\Gamma = \langle a, b \mid a^{-1}b^2ab^{-1} \rangle$ and X be its presentation 2-complex. Let $n = 3$ and $\sigma_a = (123), \sigma_b = (12)(3) \in \mathrm{Sym}(3)$. From the choice of $\sigma = (\sigma_a, \sigma_b)$, we have a 3-fold covering of $X^{(1)}$. A disk $D \subset X$ is attached to $X^{(1)}$ along the relation $a^{-1}b^2ab^{-1}$, so our label on $\partial(D)$ is $a^{-1}b^2ab^{-1}$. The attaching map of D has two connected lifts D_1 and D_2, which are attached via the maps

$$\partial D_1 = a_3^{-1}b_3^2a_3b_2^{-1}a_1^{-1}b_1b_2a_1b_1^{-1},$$

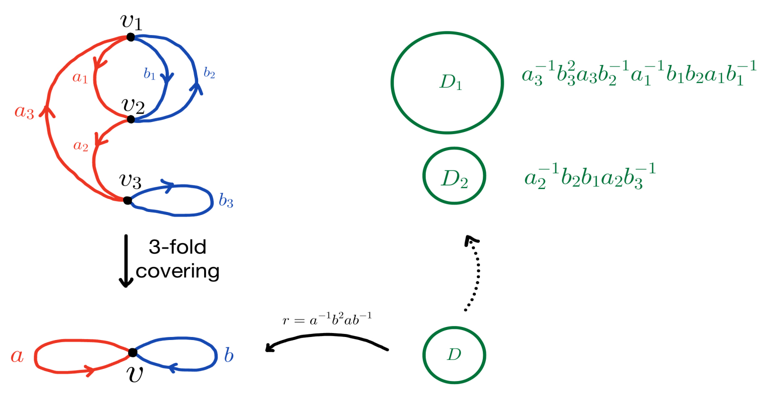

Figure 1.2: A branched covering of degree 3 (Example 1.3.3)

$$\partial D_2 = a_2^{-1}b_2b_1a_2b_3^{-1}.$$

Then the obtained 2-complex $X(\sigma)$ would be a 3-fold branched covering of X with branching locus $\{c_1\} \cup \{c_2\}$ where c_1 (resp. c_2) is the center of the disk D_1 (resp. D_2). The branching index at c_1 is 2 and at c_2 is 1. An illustration of the cover $X(\sigma)$ is given in Figure 1.2.

Consider the fundamental group $\pi_1(X(\sigma))$ in this case. Choose a_1, a_2 as a spanning tree of the 1-skeleton. Then the generators of $\pi_1(X(\sigma))$ will be a_3, b_1, b_2 and b_3. We have two relators $a_3^{-1}b_3^2a_3b_2^{-1}b_1b_2b^{-1}$ and $b_2b_1b_3^{-1}$, obtained by collapsing the spanning tree $a_1 \cup a_2$. We can remove the generator b_3 and the relator $b_2b_1b_3^{-1}$ by a Tietze transformation as $b_3 = b_2b_1$. Therefore, the fundamental group of the branched cover $X(\sigma)$ is the one relator group with presentation

$$\pi_1(X(\sigma)) \cong \langle a_3, b_1, b_2 \mid a_3^{-1}b_2b_1b_2b_1a_3b_2^{-1}b_1b_2b_1^{-1}\rangle. \tag{1.3.1}$$

Example 1.3.4. Let $\Gamma = \langle a, b \mid aba^{-1}b^{-1}\rangle$ and X be its presentation 2-complex. Let $n = 3$ and as before, let $\sigma_a = (123), \sigma_b = (12) \in \mathrm{Sym}(3)$. From the choice of

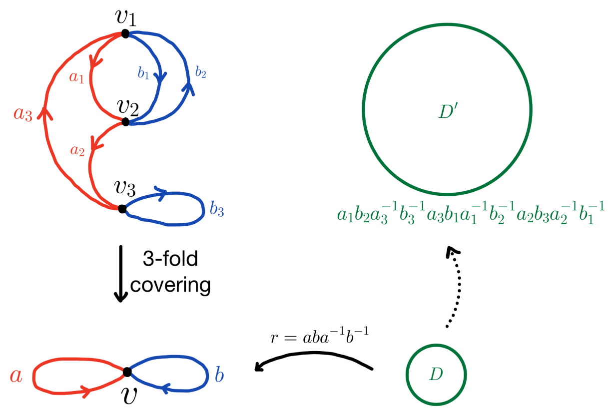

Figure 1.3: A branched covering of degree 3 (Example 1.3.4)

$\sigma = (\sigma_a, \sigma_b)$, we have a 3-fold covering of $X^{(1)}$. A disk $D \subset X$ is attached to $X^{(1)}$ along the relation $aba^{-1}b^{-1}$ so we the label on $\partial(D)$ is $aba^{-1}b^{-1}$. In this case the attaching map of D has a unique connected lift D', with label

$$\partial D' = a_1b_2a_3^{-1}b_3^{-1}a_3b_1a_1^{-1}b_2^{-1}a_2b_3a_2^{-1}b_1^{-1}.$$

See Figure 1.3. Thus the branched cover 2-complex $X(\sigma)$ will be a 3-fold branched covering of X with a unique branch point $\{c'\}$ where c' is the center of the disk D'. The branching index at c' is 3.

Note that the original group $\Gamma = \langle a, b \mid aba^{-1}b^{-1}\rangle$ is a surface group of genus 1, and X is homeomorphic to a torus. Recall that a branched cover of an oriented surface is again an oriented surface, and we can determine which surface by considering the Euler characteristic of the branched covering. Looking again at Figure 1.3, we see that $X(\sigma)$ has a CW-structure with three vertices, six edges, and a single 2-cell, thus giving us $\chi(X(\sigma)) = -2$. Since $\chi(X(\sigma)) = 2 - 2g$ where g is the genus, we conclude that $X(\sigma)$ will be a surface of genus 2.

This can also be seen directly from the fundamental group $\pi_1(X(\sigma))$. Choose a_3, b_1 as a spanning tree of the 1-skeleton, then the generators of $\pi_1(X(\sigma))$ will be a_1, a_2, b_2 and b_3. After collapsing the spanning tree $a_3 \cup b_1$, the attaching map gives rise to the single relation $a_1 b_2 b_3^{-1} a_1^{-1} b_2^{-1} a_2 b_3 a_2^{-1}$. This gives us the presentation

$$\pi_1(X(\sigma)) \cong \langle a_1, a_2, b_2, b_3 \mid a_1 b_2 b_3^{-1} a_1^{-1} b_2^{-1} a_2 b_3 a_2^{-1} \rangle \tag{1.3.2}$$

Let $\alpha_1 = a_1$, $\beta_1 = b_2$, $\alpha_2 = a_2^{-1} b_2 a_1$, and $\beta_2 = b_3$. By applying Tietze transformations, we obtain the presentation

$$\pi_1(X(\sigma)) \cong \langle a_1, a_2, b_2, b_3 \mid a_1 b_2 b_3^{-1} a_1^{-1} b_2^{-1} a_2 b_3 a_2^{-1} \rangle \tag{1.3.3}$$

$$\cong \langle \alpha_1, \beta_1, \alpha_2, \beta_2 \mid [\alpha_1, \beta_1][\alpha_2, \beta_2] \rangle \tag{1.3.4}$$

which is the standard presentation of the surface group of genus 2.

We now have, for each natural number n, a model that randomly produces a degree n branched cover $X(\sigma)$ of the finite 2-complex X. We will be interested in understanding topological and geometric properties of $X(\sigma)$, as n gets large.

Definition 1.3.5. Given an event $E = E_n$ depending on a parameter n, E holds *asymptotically almost surely* if it holds with probability $1 - o(1)$. Thus the probability of success goes to 1 in the limit as $n \to \infty$.

1.3.3 Connectedness of branched covers

Our random model associates to t-tuples of elements in the symmetric group, chosen independently with uniform distribution, a corresponding branched cover. Our goal is now to translate interesting properties of the branched cover $X(\sigma)$ into properties of the t-tuple in the symmetric group. We can then hope to leverage our understanding of random elements in symmetric groups to analyze whether or not the property holds for random branched covers in our model. As an easy example, let us consider the connectedness of the branched cover.

Lemma 1.3.6. *The branched cover $X(\sigma)$ is connected if and only if the subgroup generated by the permutations $\sigma_1, \ldots, \sigma_t$ acts transitively on the vertex set.*

Proof. Connectedness of $X(\sigma)$ is completely determined by connectedness of its 1-skeleton. Given an edge path joining a pair of vertices, one can read off the corresponding product of permutations (and their inverses) taking the initial vertex to the terminal vertex. So if $X(\sigma)$ is connected, then $\langle \sigma_1, \ldots, \sigma_t \rangle$ acts transitively on the vertex set. Conversely, if the subgroup acts transitively on the vertex set, then given any two vertices in $X(\sigma)$ we can find a product of permutations (and their inverses) taking one of these vertices to the other. This then gives us a sequence of edges connecting the two vertices, showing the 1-skeleton of $X(\sigma)$ is connected. □

So understanding connectedness of our random branched covers is completely equivalent to understanding when a randomly selected t-tuple of elements in Sym(n) generate a transitive subgroup. This is a classically studied problem, and we have the following result of Dixon [10]:

Proposition 1.3.7. *The proportion of ordered pairs (σ_a, σ_b), where $\sigma_a, \sigma_b \in Sym(n)$ generate a transitive subgroup of $Sym(n)$ is $1 - \frac{1}{n} + O(\frac{1}{n^2})$ as $n \to \infty$.*

Translating this result back to our model gives us:

Corollary 1.3.8. *Let $\Gamma = \langle u_1, \ldots, u_t \mid r_1, \ldots, r_s \rangle$, with corresponding presentation 2-complex X. If $t \geq 2$, then $X(\sigma)$ is asymptotically almost surely connected.*

Proof. The $t = 2$ case follows immediately from Proposition 1.3.7. For $t \geq 3$, it follows easily from the two generator case, because adding an additional generator just adds more edges to an already connected graph. Equivalently, if the first two elements σ_1, σ_2 already generate a transitive group, then adding additional generators $\sigma_3, \ldots, \sigma_t$ does not change transitivity of the action. □

Remark 1.3.9. For the connectedness in Corollary 1.3.8, we had to assume that there is more than one generator. For the single generator case, random coverings are **not** asymptotically almost surely connected. In particular, the cover will be connected if and only if the chosen permutation is an n cycle. Thus, the probability of connectedness is $\frac{(n-1)!}{n!}$, which goes to 0 when $n \to \infty$.

This is the reason for our requirement that the rank of the 1-skeleton is ≥ 2 in our **Main Theorem** (see definition of acceptable 2-complex). Nevertheless, it is obvious that for a single generator case, any random branched covering has a Gromov hyperbolic fundamental group. Indeed, the presentation 2-complex has 1-skeleton consisting of a single loop. Thus any covering of the 1-skeleton is just a disjoint union of cycles. It follows that each connected component will have fundamental group that is generated by a single element, hence cyclic. It will be isomorphic to $\mathbb{Z}$ if there are no relations, and isomorphic to a (potentially larger) finite cyclic group if there is at least one relation. In either case, the fundamental group is an (elementary) Gromov hyperbolic group.

1.3.4 Disks in random branched covering spaces

Let X be the presentation 2-complex for the presentation

$$\Gamma = \langle u_1, \dots, u_t \mid r_1, \dots, r_s \rangle$$

where $t \geq 2$ and let $f : X(\sigma) \to X$ be the n-fold branched covering obtained from the t-tuple of permutations $\sigma = (\sigma_1, \dots, \sigma_t)$, where $\sigma_i \in \mathrm{Sym}(n), i = 1 \dots t$. Recall that each disk $D_i \subset X$ corresponding to the relation r_i has finitely many lifts $D_{ij} \subset X(\sigma)$ in the branched cover. For each lift $D_{ij} \subset X(\sigma)$, the index of D_{ij}, given by $\mathrm{ind}(D_{ij})$, is the branching index of f at the center of D_{ij}. Clearly, the sum $\sum_{j=1} \mathrm{ind}(D_{ij}) = n$ for each $i = 1, \dots, s$.

We are interested in studying the small cancellation properties for a random branched cover. We can view a piece b in X (respectively $X(\sigma)$) as a combinatorial path in the 1-skeleton $X^{(1)}$ (resp. $X(\sigma)^{(1)}$). It is tempting to use the covering map $\rho : X(\sigma)^{(1)} \to X^{(1)}$ to compare pieces in $X(\sigma)$ with pieces in X. Unfortunately, the image of piece in $X(\sigma)$ might not be a piece in X, as it might map to a path that has length greater than the boundary of the disks. Similarly, the connected lift of a piece in X might not be a piece in $X(\sigma)$, as it might be properly contained in an overlap in X. For this reason, it is more convenient to work with overlaps. As overlaps are defined in terms of the universal cover of the attaching loops, these behave better with respect to the branched covering map ρ restricted to the one skeleton. Indeed, we have the following:

Lemma 1.3.10. *Let X be an acceptable polygonal 2-complex, $X(\sigma)$ a branched cover of X, and $\rho : X(\sigma) \to X$ the branched covering map. Then:*

- *if* $(\mathbf{p}, \mathbf{p}')$ *is an overlap for the pair of* 2*-cells* D, D' *in* X*, then every lift of* $(\mathbf{p}, \mathbf{p}')$ *is an overlap in* $X(\sigma)$*;*

- *if* $(\mathbf{p}, \mathbf{p}')$ *is an overlap in* $X(\sigma)$ *for the pair of disks* $\hat{D}, \hat{D}'$*, then the pair defines an overlap for the image pair of disks* $D = \rho(\hat{D}), D' = \rho(\hat{D}')$ *in* X*.*

Proof. Recall that the overlap $(\mathbf{p}, \mathbf{p}')$ is actually a pair of (equivalence classes of) subpaths $\mathbf{p} \subset \widetilde{\partial \mathbf{D}}$ and $\mathbf{p}' \subset \widetilde{\partial \mathbf{D}'}$, with the property that the lifted attaching maps $\tilde{\alpha} : \widetilde{\partial D} \to X^{(1)}$, $\tilde{\beta} : \widetilde{\partial D'} \to X^{(1)}$ coincide on the subpaths $\mathbf{p}, \mathbf{p}'$. Given any pre-image v_i of the unique vertex $v \in X^{(1)}$, covering space theory tells us we can lift the maps $\tilde{\alpha}, \tilde{\beta}$ to maps $\bar{\alpha} : \widetilde{\partial D} \to X(\sigma)^{(1)}, \bar{\beta} : \widetilde{\partial D'} \to X(\sigma)^{(1)}$ based at the vertex v_i. Since the original maps $\tilde{\alpha}, \tilde{\beta}$ coincide on the subpaths $\mathbf{p}, \mathbf{p}'$, the lifted maps will have the same property. Moreover, since the original maps $\tilde{\alpha}, \tilde{\beta}$ *differ* on the two edges immediately preceding (respectively following) the subpaths $\mathbf{p}, \mathbf{p}'$, the same property holds for the lifted maps. Finally, we note that the lifted maps $\bar{\alpha}, \bar{\beta}$ are periodic, as they will cover one of the connected lifts $\hat{\alpha}, \hat{\beta}$ of the attaching maps α, β (the lifts based at the vertex v_i). These lifted attaching maps $\hat{\alpha} : \widehat{\partial D} \to X(\sigma)^{(1)}$, $\hat{\beta} : \widehat{\partial D'} \to X(\sigma)^{(1)}$ are defined on finite covers $\widehat{\partial D} \to \partial D$ and $\widehat{\partial D'} \to \partial D'$. These define a pair of 2-cells $\hat{D}, \hat{D}'$ in $X(\sigma)$. We conclude that the pair $(\mathbf{p}, \mathbf{p}')$ is an overlap for the 2-cells $\hat{D}, \hat{D}'$.

Conversely, if we have an overlap $(\mathbf{p}, \mathbf{p}')$ in $X(\sigma)$ for a pair of 2-cells $\hat{D}, \hat{D}'$, we can project the overlap via the covering map ρ. More precisely, from the construction of $X(\sigma)$, the attaching maps $\hat{\alpha} : \partial \hat{D} \to X(\sigma)^{(1)}$, $\hat{\beta} : \partial \hat{D}' \to X(\sigma)^{(1)}$ are lifts of the attaching maps $\alpha : \partial D \to X^{(1)}$, $\beta : \partial D' \to X^{(1)}$ for the pair of 2-cells D, D' in X. This means there are finite covering maps $\pi : \partial \hat{D} \to \partial D$, $\pi' : \partial \hat{D}' \to \partial D'$, and commutative diagrams $\alpha \circ \pi = \rho \circ \hat{\alpha}$, $\beta \circ \pi' = \rho \circ \hat{\beta}$.

The covering map π allow us to identify the universal covers of $\partial\hat{D}$ and ∂D, via the lift of the covering map $\widetilde{\pi} : \widetilde{\partial\hat{D}} \to \widetilde{\partial D}$ (and similarly for π', $\partial\hat{D}'$, and $\partial D'$). With this identification, the map $\bar{\alpha} : \widetilde{\partial\hat{D}} \to X(\sigma)^{(1)}$ descends to a map $\tilde{\alpha} := \rho \circ \bar{\alpha} \circ \tilde{\pi}^{-1} : \widetilde{\partial D} \to X^{(1)}$. Similarly, we have a map $\tilde{\beta} := \rho \circ \bar{\beta} \circ (\tilde{\pi}')^{-1} : \widetilde{\partial D'} \to X^{(1)}$. Since the maps $\bar{\alpha}, \bar{\beta}$ agree on the subpaths $\mathbf{p}, \mathbf{p}'$ but differ on the immediately preceding (and immediately following) edges, the same property is true for the composite maps $\rho \circ \bar{\alpha}, \rho \circ \bar{\beta}$. Using the identification $\tilde{\pi}, \tilde{\pi}'$ of universal covers, we can view $\mathbf{p}, \mathbf{p}'$ as subpaths in $\widetilde{\partial D}, \widetilde{\partial D'}$. This shows that $(\mathbf{p}, \mathbf{p}')$ defines an overlap for the 2-cells D, D', completing the proof of the Lemma.

□

An immediate consequence of the lemma is the following

Corollary 1.3.11. *Let X be an acceptable polygonal 2-complex, and $X(\sigma)$ a branched cover of X. If the 2-cell $\bar{D}$ in $X(\sigma)$ is an index k branched cover of the 2-cell D in X, then the overlap ratios are related by $o(\bar{D}) = o(D)/k$. In particular, $o(X(\sigma)) \leq o(X)$.*

Proof. From the lemma, we see that the lengths of overlaps for D coincide with the lengths of overlaps for $\bar{D}$. Since $\ell(\partial\bar{D}) = k \cdot \ell(\partial D)$, the result follows. □

As we remarked earlier, the small cancellation condition $C'(\lambda)$ (where $0 < \lambda < 1$) for a 2-cell D is implied by $o(D) < \lambda$. So we immediately obtain:

Corollary 1.3.12. *If a 2-cell in X satisfies the $o(D) < \lambda$, then all of its lifts in $X(\sigma)$ satisfy $C'(\lambda)$.*

We denote by R_L (respectively R_S) the length of the longest (resp. shortest) relation in the presentation Γ. Let us introduce some constants associated to the presentation 2-complex X.

In view of Lemma 1.2.7, there is a uniform bound on the length of overlaps in X. We introduce the parameter $\mathcal{O} := R_L^2 + R_L$, which serves as a global upper bound on the length of overlaps in X. In view of Lemma 1.3.10, $\mathcal{O}$ also serves as an upper bound on the length of overlaps in **any** of the branched covers $X(\sigma)$. Since all the relators in Γ have length $\geq R_S$, we also obtain the upper bound on the overlap ratio $o(X) < \frac{\mathcal{O}}{R_S}$. We are looking for branched covers with overlap ratio less than λ. To this end, let us introduce the critical index $I := \frac{\mathcal{O}}{\lambda R_S}$. From Corollary 1.3.11, we know that any disk D in a branched cover $X(\sigma)$ whose index is $\geq I$ automatically has $o(D) < \lambda$, thus satisfies $C'(\lambda)$.

Definition 1.3.13. Given an acceptable polygonal 2-complex X, and $\lambda \in (0, 1)$, a disk D in a branched cover $X(\sigma)$ is called a λ-*good disk* if its index is greater than or equal to I. A disk that is not a good-disk is called a λ-*worrisome disk*. We will typically be working with a fixed value of λ, and refer to λ-good disks as *good disk* and λ-worrisome disk as a *worrisome disk*.

All the λ-good disks satisfy $C'(\lambda)$ small cancellation in $X(\sigma)$. However, this is not true for λ-worrisome disks, which may or may not satisfy the $C'(\lambda)$ small cancellation condition.

For the relators $r_1, \ldots, r_s$ and a random choice of permutations $\sigma = (\sigma_1, \ldots, \sigma_t)$, $\sigma_i \in \mathrm{Sym}(n)$, we define another type of permutation $r_i(\sigma) \in \mathrm{Sym}(n)$ that represents the structure of lifts of the disk D_i. In $X(\sigma)$, we first attach a disk for a lift of D_i of the relation r_i starting from the vertex v_1. Let $v_1 = v_{i(1)}$ and let $v_{i^j(1)}$ be the vertex that is arrived at after following the letters of r_i a total of j times. Let $k \geq 1$ be the smallest such that $v_{i^k(1)} = v_1$. Then $(v_{i(1)} \cdots v_{i^k(1)})$ forms a cycle of length k with entries corresponding to the indices of the vertices appearing in the procedure. We repeat the same process at a vertex that is not already obtained in a previous cycle until there are no vertices remaining. Finally we get a permutation of n elements and denote it $r_i(\sigma)$. The cycle lengths of $r_i(\sigma)$ have a one-to-one correspondence with the indices of lifts of D_i. The permutation $r_i(\sigma)$ is the result of applying the *word map* $r_i : \mathrm{Sym}(n) \times \cdots \times \mathrm{Sym}(n) \to \mathrm{Sym}(n)$ (see [26]) to the t-tuple σ of permutations. We will use the following result from [26, Corollary 1.5]:

Proposition 1.3.14. *Let $k \geq 2$ be a fixed integer. If the permutation r_i is not a proper power, then the expected number of cycles of length k in $r_i(\sigma)$ is $\frac{1}{k} + \mathbf{O}(n^{-\pi(r_i)})$, where $\pi(r_i) \geq 2$ is the primitive rank defined in [26].*

Let $L_n(k)$ be the total number of cycles of length at most k in all of the permutations $r_1(\sigma), \dots, r_s(\sigma) \in Sym(n)$. Note that, by hypothesis, none of the r_i are proper powers, so applying Proposition 1.3.14, we have

$$\mathbb{E}\big(L_n(k)\big) = s(1 + \frac{1}{2} + \cdots + \frac{1}{k}) + \mathbf{O}(n^{-\pi(r_1)} + \cdots + n^{-\pi(r_s)}).$$

As a result, when $n \to \infty$ the expected value $\mathbb{E}(L_n(k)) \to s(1 + \frac{1}{2} + \cdots + \frac{1}{k})$. Knowing the asymptotics of the expected value allows us to deduce information about the tails of the probability distributions.

Lemma 1.3.15. *For any $\epsilon > 0$, there exists $N, m \in \mathbb{N}$ such that if $n \geq N$, then*

$$\mathbb{P}\big(L_n(k) \leq m\big) > 1 - \frac{\epsilon}{2}.$$

Proof. Suppose not. Then there exists $\epsilon > 0$ with the property that for each m, we can find a sequence $n_j \to \infty$ such that $\mathbb{P}\big(L_{n_j}(k) \geq m+1\big) > \frac{\epsilon}{2}$. We can now estimate the expectation of $L_{n_j}(k)$ from below:

$$\begin{aligned}\mathbb{E}\big(L_{n_j}(k)\big) &= \sum_{i=0}^{\infty} i\mathbb{P}\big(L_{n_j}(k) = i\big) \\ &\geq \sum_{i=m+1}^{\infty} (m+1)\cdot\mathbb{P}\big(L_{n_j}(k) = i\big) \\ &= (m+1)\cdot\mathbb{P}\big(L_{n_j}(k) \geq m+1\big) > (m+1)\cdot\frac{\epsilon}{2}\end{aligned}$$

For a fixed choice of m, this estimate holds for all the n_j in the sequence. Applying this to the specific case where $m = \frac{2s}{\epsilon}\cdot\big(1+\frac{1}{2}+\cdots+\frac{1}{k}\big)$, we obtain an infinite sequence of integers $n_j \to \infty$ where

$$\begin{aligned}\mathbb{E}\big(L_{n_j}(k)\big) &> (m+1)\cdot\frac{\epsilon}{2} \\ &> s\big(1+\frac{1}{2}+\cdots+\frac{1}{k}\big)+\frac{\epsilon}{2}\end{aligned}$$

On the other hand, we saw earler that the expected value $\mathbb{E}(L_n(k))$ converges to $s(1+\frac{1}{2}+\cdots+\frac{1}{k})$ as $n\to\infty$, giving us a contradiction. □

Note that in the proof above, we are not assuming the existence of limiting distribution of $L_n(k)$, but using the expected value of $L_n(k)$ for finite $n \in \mathbb{N}$ and its limit as $n\to\infty$. We can interpret Proposition 1.3.14 and Lemma 1.3.15 in terms of random branched coverings.

Corollary 1.3.16. *Let X be an acceptable 2-complex. Then given any integer k, and $\epsilon > 0$, we can find an integer $M = M(k,\epsilon)$ with the following property: for any n sufficiently large, with probability $\geq 1-\frac{\epsilon}{2}$ a random degree n branched cover $X(\sigma)$ contains $\leq M$ disks of index $\leq k$.*

Next we turn our attention to topological properties of disks

Lemma 1.3.17. *Let X be an acceptable 2-complex, and m a given integer. Then asymptotically almost surely the random branched covers of X have all disks of index m that are injectively embedded.*

Proof. We will count all possible random branched coverings and see how many of them contain non-injective lifts of index m. Since the number of generators is t and $|\mathrm{Sym}(n)| = n!$, the number of choices for permutations $\sigma = (\sigma_1, \ldots, \sigma_t)$ is $(n!)^t$. In other words, there are $(n!)^t$ random branched coverings of X.

Let r be a relator. By the Proposition 1.3.14, the expected number of lifts of r of index m in a random branched covering is $\frac{1}{m} + \mathbf{O}(n^{-\pi(r)})$. Since there are $(n!)^t$ random branched coverings, the total number of lifts of r of index m in all possible random branched coverings is $(n!)^t \left(\frac{1}{m} + \mathbf{O}(n^{-\pi(r)})\right)$.

Now, in all possible branched coverings, we count the total number of injective lifts of the relator r of index m. Let $r = w_1 \ldots w_{|r|}$ where each

$$w_i \in \{u_1, \ldots, u_t, u_1^{-1}, \ldots, u_t^{-1}\},$$

the symmetric generating set. For such a lift $D \subset X(\sigma)$ of r, since it has index m, the length of ∂D is $m|r|$. We label the $m|r|$ edges of ∂D as in Example 1.3.3. Recall that in a branched covering, there are tn total edges denoted u_{ij} where $i = 1, \ldots, t$ and $j = 1, \ldots, n$. To be injective, ∂D has to contain $m|r|$ distinct vertices. Assume that n is sufficiently large so that $n \geq m|r|$. Choosing the $m|r|$ distinct vertices in order amounts to $n(n-1)\ldots(n-m|r|+1)$ possibilities. Once we choose the labels on the vertices of ∂D, the labels on the oriented edges along ∂D will be determined by $r = w_1 \ldots w_{|r|}$. Thus there are $n(n-1)\ldots(n-m|r|+1)$ different ways of labeling the boundary of an injective lift D. Again, we remark that these lifts may occur in different branched covers. Note that since cyclic permutations of the labeling on disk's boundary represent the same lift, there are $\frac{1}{m}n(n-1)\dots(n-m|r|+1)$ different ways of labeling the boundary up to cyclic relabeling.

Now we count all possible branched coverings $X(\sigma)$ that contain the choice of ∂D with a given injective labeling. Let ℓ_i be the number of occurrences of a lift of the generator u_i along ∂D. Obviously, $\sum_{i=1}^{t} \ell_i = \ell(\partial D) = m|r|$. For the remaining $n - l_i$ lifts of u_i that are not contained in the labeled ∂D, there are $n - \ell_i$ possible initial vertices for the lifts, and $n - \ell_i$ possible ending vertices for the lifts. Hence the number of ways to place the remaining lifts of u_i boils down to choosing a pairing between the possible initial vertices and terminal vertices. There are $(n - \ell_i)!$ such pairings. Ranging over all the edges, we obtain $(n-\ell_1)!\dots(n-\ell_t)!$ labeled branched coverings that contain the lift D with the prescribed (injective) labeling on ∂D.

Thus in all possible branched coverings, the total number of injective lifts of r of index m will be

$$\frac{\big[n(n-1)\dots(n-m|r|+1)\big]\cdot\big[(n-\ell_1)!\dots(n-\ell_t)!\big]}{m}$$

and therefore the total number of non-injective lifts of r in all possible branched coverings is

$$\frac{(n!)^t}{m} - \frac{\big[n(n-1)\dots(n-m|r|+1)\big]\cdot\big[(n-\ell_1)!\dots(n-\ell_t)!\big]}{m} + (n!)^t\mathbf{O}(n^{-\pi(r)}).$$

The total number of branched coverings that contain a non-injective lifts will be bounded above by the total number of non-injective lifts in all branched coverings. Let us denote by $P_n(r)$ by the probability that a branched covering contains a non-injective lift of r of index m. Then using the above estimate we obtain the upper bound

$$P_n(r) \leq \frac{(n!)^t - n(n-1)\dots(n-m|r|+1)(n-\ell_1)!\dots(n-\ell_t)!}{m(n!)^t} + \mathbf{O}(n^{-\pi(r)}).$$

With this upper bound in hand, we can easily compute the limit of $P_n(r)$ as $n \to \infty$:

$$\begin{aligned}\lim_{n\to\infty} P_n(r) &\leq \lim_{n\to\infty} \frac{(n!)^t - n(n-1)\dots(n-m|r|+1)(n-\ell_1)!\dots(n-\ell_t)!}{m(n!)^t}\\ &= \frac{1}{m} - \frac{1}{m}\cdot \lim_{n\to\infty} \frac{n(n-1)\dots(n-m|r|+1)}{\left(\prod_{i=0}^{\ell_1-1}(n-i)\right)\dots\left(\prod_{i=0}^{\ell_t-1}(n-i)\right)}.\end{aligned}$$

Note that in the last term, the numerator and the denominator are both $m|r|$-degree monic polynomials in n. As $n \to \infty$, the ratio tends to one, and as a result $\lim_{n\to\infty} P_n(r) = 0$. This tells us that in a random branched covering covering, all lifts of the single relation r of index m are embedded asymptotic almost surely.

The probability that a branched covering contains a non-injective index m lift of one of the finitely many relators $r_1, \dots, r_s$ is less than $\sum_{i=1}^{s} P(r_i)$. Since the sum goes to zero as $n \to \infty$, we conclude that every lift of an r_i index m are embedded in a random branched covering asymptotic almost surely. □

A similar argument can be used to control intersections of disks.

Lemma 1.3.18. *Let X be an acceptable 2-complex, and M a given integer. Then asymptotically almost surely in the random branched covers of X all disks of index at most M are pairwise disjoint.*

Proof. Consider a pair r, r' of relations. We will start by considering disks D, D' that are lifts of the relations r, r' of index $\leq M$.

If r and r' don't have a common piece, any lifts D and D' have no common edges on their boundaries. So it is sufficient to consider the case that r and r' have a piece in common, i.e. $r = pw$ and $r' = pw'$. We first consider the case that p is the only maximal piece contained in both r and r'. Then if D and D' are lifts of r and r' respectively, then by Remark 1.3.11, the subpath $\partial D \cap \partial D'$ has the same length $|p|$ and we call the subpath by p for convenience.

We want to compute the probability that two lifts of index at most M intersect in some branched covering and show that it goes to 0 as $n \to \infty$. Among $(n!)^t$ number of all possible branched coverings, we count the total number of pairs (D, D') such that $\partial D \cap \partial D' \neq \emptyset$. The idea is very similar to the proof of Lemma 1.3.17. Note that by Lemma 1.3.17, any lift of index $\leq M$ is injective asymptotically almost surely. Let m (resp. m') be the index of D (resp.D').

Since for a labeling of D (resp. D'), there are m (resp. m') cyclic permutations that actually represent the same disk.

When we first label ∂D with the edges in $X(\sigma)$, there are $n(n-1)\dots(n-m|r|+1)$ number of ways of injective labeling. But for a labeling on ∂D, there are m different cyclic permutations that actually represent the same disk. Thus we have $\frac{1}{m}n(n-1)\dots(n-m|r|+1)$ ways of distinct injective labelings up to cyclic permutation. Note that there are m possible ways to choose $p \subset \partial D$ that intersects with $\partial D'$. Once we choose a path $p \subset \partial D$, we label $\partial D'$ so that $\partial D \cap \partial D' = p$.

When we label $\partial D'$, the key idea is that there is no choice for the $|p|+1$ number of vertices that are contained in p. So we choose $m'|r'|-|p|-1$ number of vertices in order so that $\partial D'$ contains $m'|r|$ different vertices and intersect with ∂D only on $|p|$. There are $(n-m|r|)(n-m|r|-1)\dots(n-m|r|-m'|r'|+|p|+2)$ number of ways to label $\partial D'$ that are all cyclically distinct. Therefore we have

$$n(n-1)\dots(n-m|r|-m'|r'|+|p|+2) \tag{1.3.5}$$

number of injective labeling on D and D' that intersect on the piece p. Here, the numerator has $m|r|+m'|r'|-|p|-1$ number of factors.

Now, let a_i be number of u_i generator on the label of $\partial D \cup \partial D'$. Obviously, we have

$$\sum_{i=1}^{t} a_i = m|r|+m'|r'|-|p|. \tag{1.3.6}$$

For a pair of labeled D and D', we count the number of possible branched coverings that contain D, D'. By using the same labeling process as 1.3.17 on remaining edges to construct a branched coverings, we have $(n-a_1)!(n-a_2)!\dots(n-a_t)!$ number of branched coverings containing D and D'. Thus for given indices m and m', we have

$$n(n-1)\dots(n-m|r|-m'|r'|+|p|+2)(n-a_1)!\dots(n-a_t)! \tag{1.3.7}$$

number of branched coverings that contain D and D' of the given indices.

The total number of pairs of intersecting lifts (D, D') of index $\leq M$ in all possible branched coverings is

$$\sum_{m,m'\leq M} n(n-1)\dots(n-m|r|-m'|r'|+|p|+2)(n-a_1)!\dots(n-a_t)! \tag{1.3.8}$$

Here, in the summation, the numerator has $nt-1$ number of factors.

In the general case, for two relations r and r' that have more than one maximal piece on the intersection, then the number of pairs of intersecting lifts of r and r' of index$\leq M$ in all possible branched coverings would be less than or equal to the number in the Equation (1.3.8).

The number of branched coverings that contains index $\leq M$ lifts of r and r' that intersect is less than or equal to the number of pairs of intersecting lifts (D, D') of index $\leq M$ in all possible branched coverings. We denote the probability that a branched covering contains intersecting lifts of r and r' of index $\leq M$ by $P(r, r')$ and it is less than or equal to

$$\sum_{m,m' \leq m} \frac{n(n-1)\dots(n - m|r| - m'|r'| + |p| + 2)(n - a_1)!\dots(n - a_t)!}{(n!)^t}. \tag{1.3.9}$$

Since (1.3.9) goes to 0 as $n \to \infty$, we see that in a branched covering, all lifts of r and r' of index $\leq M$ are pairwise disjoint asymptotic almost surely.

Now, the probability that a random branched covering contains pair of intersecting lifts of any two relators is less than or equal to $\sum_{r,r' \in \{r_1,\dots,r_k\}} P(r, r')$. Since the summation goes to 0 as $n \to \infty$, we conclude that every lifts of index $\leq M$ pairwise disjoint asymptotically almost surely. □

Corollary 1.3.19. *All worrisome-disks are pairwise disjoint in $X(\sigma)$ asymptotically almost surely.*

1.4 Proof of the Main Theorem

Let $\Gamma = \langle u_1, \dots, u_t \mid r_1, \dots, r_s \rangle$ be a finitely presented group and X the presentation 2-complex where $t \geq 2$. Let $f : X(\sigma) \to X$ be a n-fold random branched covering obtained from random permutations $\sigma = (\sigma_1, \dots, \sigma_t)$ where $\sigma_i \in \mathrm{Sym}(n), i = 1 \dots t$.

We are now ready to prove the **Main Theorem** in the special case of the presentation 2-complex. For the convenience of the reader, we restate the special case.

Main Theorem. *Let X be the presentation 2-complex associated to the finite presentation $\Gamma = \langle u_1, \dots, u_t \mid r_1, \dots, r_s \rangle$. We assume that the relations are cyclically reduced, and that none of the relations are proper powers. Let $\lambda \in (0,1)$ be a given constant, and let $X(\sigma)$ be an n-fold random branched cover of X. Then $X(\sigma)$ is asymptotically almost surely homotopy equivalent to a 2-complex satisfying geometric $C'(\lambda)$-small cancellation.*

Proof. Observe that the presentation 2-complex X is an acceptable 2-complex. The parameters $\mathcal{O}$ and I are defined as in Section 1.3.4. Observe that all the parameters we have introduced so far only depend on the initial complex X, and the small cancellation parameter λ that we want to achieve.

Now given an $\epsilon > 0$, we want to show that for all sufficiently large n, a random branched cover $X(\sigma)$ is homotopy equivalent to a 2-complex satisfying geometric small cancellation with probability $> 1 - \epsilon$. We now introduce some parameters that also depend on the given ϵ. From Corollary 1.3.16, Lemma 1.3.17, and Lemma 1.3.18, we can choose an N sufficiently large, so that for all $n > N$, with probability $> 1 - \epsilon$, a random branched cover $X(\sigma)$ has the following properties:

1. the number of disks in $X(\sigma)$ of index $\leq I$ is $\leq M$ (where M is fixed parameter provided by Corollary 1.3.16);
2. all disks of index $\leq K$ are injective;
3. all disks of index $\leq K$ are pairwise disjoint.

Where here the constant K in statements (2) and (3) is given by the value:

$$K := R_S^{-1}(1 + \lambda^{-1})\mathcal{O}\big(M^2 I (R_L \mathcal{O})^2\big) + R_S^{-1}\lambda^{-1}\mathcal{O}$$

Thus the choice of K depends (via M) on the choice of ϵ. To complete the proof, it suffices to show that these branched covers are homotopy equivalent to a complex satisfying $C'(\lambda)$-small cancellation.

Observe that, if $X(\sigma)$ has **no** worrisome disks, then all disks in $X(\sigma)$ have overlap ratio $< \lambda$ and $X(\sigma)$ itself satisfies $C'(\lambda)$-small cancellation. But in general, most of the $X(\sigma)$ will contain some disks of index $\leq I$. To analyze these branched covers, we partition the disks in $X(\sigma)$ according to their index:

- *small disks* are those with index $\leq I$,
- *medium disks* are those with index $> I$ but $\leq K$
- *large disks* have index $> K$.

By the definition of worrisome disks, the set of small disks are exactly the same as the set of worrisome disks. Since these are the disks which might have overlap ratio $> \lambda$, we construct a new space $Y(\sigma)$ by collapsing each of the small disks in $X(\sigma)$ to a point. There is a natural map quotient map $q : X(\sigma) \to Y(\sigma)$.

Fact 1: The map q is a homotopy equivalence.

To check that the quotient map is a homotopy equivalence, recall that quotienting out a contractable subcomplex from a CW-complex yields a homotopy equivalence. From property (2), small disks are embedded, hence have image in $X(\sigma)$ that are homeomorphic to $\mathbb{D}^2$. From property (3), small disks are pairwise disjoint. Collapsing them one by one yields a finite sequence of homotopy equivalences from $X(\sigma)$ to $Y(\sigma)$.

Next we need to establish that $Y(\sigma)$ satisfies $C'(\lambda)$-small cancellation. It suffices to check that all the overlap ratios of disks in $Y(\sigma)$ are $< \lambda$. Since disks in $Y(\sigma)$ are images of disks in $X(\sigma)$, we will use the same terminology of "medium" and "large" disks in $Y(\sigma)$. There will not be any "small" disks in $Y(\sigma)$ as those disks are collapsed to points.

Fact 2: If $\hat{D}$ is a medium disk in $Y(\sigma)$, then it has overlap ratio $o(\hat{D}) < \lambda$.

The disk $\hat{D}$ is the image of a medium disk D in $X(\sigma)$. Since the index of D is greater than the critical index I, we have $o(D) < \lambda$. From property (3), the disk D is disjoint from all the small disks. So the quotient map q leaves D and all edges incident to the curve ∂D unchanged. It follows that $o(\hat{D}) = o(D) < \lambda$, as desired.

This leaves us with checking the overlap ratio of large disks in $Y(\sigma)$. In order to do this, we need to give a lower bound on the length of the large disk, and an upper bound on the length of the overlaps in the large disk. As before, we let $\hat{D}$ be a large disk in $Y(\sigma)$, which is the image of a large disk D in $X(\sigma)$. The boundary $\partial\hat{D}$ is obtained from ∂D by collapsing the subpaths that are images of overlaps with small disks.

Note that, since $K > \mathcal{O}$, the length of ∂D exceeds the length of any of the overlaps in $X(\sigma)$. So if $(\mathbf{p}, \mathbf{p}')$ is an overlap with $\mathbf{p} \subset \widetilde{\partial D}$, we can instead view $\mathbf{p}$ as an embedded path in ∂D. We know from Corollary 1.2.8 that there are only finitely many overlaps in $X(\sigma)$, so we can list out all the overlaps $(\mathbf{p}, \mathbf{p}')$ between $\hat{D}$ (so $\mathbf{p} \subset \partial D$) and small disks. This gives us a finite list of overlaps $\{(\mathbf{p_1}, \mathbf{p_1'}), \ldots (\mathbf{p_k}, \mathbf{p_k'})\}$, cyclically ordered according to the initial vertex of the paths $\mathbf{p_i} \subset \partial D$. Each $\mathbf{p_i'}$ lies in $\widetilde{\partial D_i}$ where D_i is a small disk in $X(\sigma)$. The boundary $\partial\hat{D}$ is obtained from ∂D by collapsing each of the intervals $\mathbf{p} \subset \partial D$ to points.

Fact 3: For any large disk D, there are at most $\leq M^2 I(R_L\mathcal{O})^2$ many overlaps $(\mathbf{p}, \mathbf{p}')$, where $\mathbf{p} \subset \partial D$ and $\mathbf{p}' \subset \widetilde{\partial D_i}$ is in any of the short disks D_i.

Applying Lemma 1.3.10, any such overlap covers an overlap $(\mathbf{q}, \mathbf{q}')$ in X. There are $\leq M$ short disks in $X(\sigma)$, hence $\leq M$ images of short disks in X. Since each image disk E in X has at most $\leq R_L\mathcal{O}$ paths that could serve as $\mathbf{q}'$, the total number of possible image overlaps $(\mathbf{q}, \mathbf{q}')$ in X is bounded above by $M(R_L\mathcal{O})^2$.

Lastly, given a candidate image overlap $(\mathbf{q}, \mathbf{q}')$ in X, we need to check how many pre-images in $X(\sigma)$ correspond to overlaps between D and one of the small disks. The image of the common path $\mathbf{q}'$ in $X^{(1)}$ lifts to n paths inside $X(\sigma)^{(1)}$, each of them lying on some lift of the disk E. However, we know that there are at most $\leq M$ lifts that are small disks, and as each of them have index $\leq I$, there are $\leq MI$ lifts of $\mathbf{q}'$ along small disks. Since there is a bijection between the lifts of $\mathbf{q}'$ and those of $\mathbf{q}$ (they define the same path in $X^{(1)}$), we see that the pair $(\mathbf{q}, \mathbf{q}')$ has at most $\leq MI$ lifts that are overlaps between the given disk D and one of the short disk lifts of E. Combining this with the estimate on the number of possible projected overlaps in X from the previous paragraph, **Fact 3** follows.

From the upper bound on the number of overlaps, we can deduce a lower bound on the length of the boundary $\partial\hat{D}$ for the quotient disk. The other ingredient we will need is to compute an upper bound on the size of the overlaps for the quotient disk $\hat{D}$. We have:

Fact 4: Any overlap for $\hat{D}$ has length $\leq \big(M^2 I(R_L\mathcal{O})^2 + 1\big)\mathcal{O}$.

To see this, let us consider an overlap between $\hat{D}$ and some other disk $\hat{E}$ in the quotient space $Y(\sigma)$. These are images of disks D, E inside $X(\sigma)$, and we would like to relate the overlaps in $X(\sigma)$ between D, E with those in $Y(\sigma)$ between $\hat{D}, \hat{E}$. Observe that if we have an overlap in $X(\sigma)$ with the property that the two edges preceding and following survive in the quotient space, then the image will be an overlap in $Y(\sigma)$. But if the preceding and/or following edges are in the subsets being collapsed, then we can potentially lose the "witness" to the start/end of the overlap. In that case, in the quotient space the overlap could continue, as the subsets where they differed can be collapsed down to points. This would result in a potentially longer overlap in $Y(\sigma)$, obtained by concatenating two overlaps in $X(\sigma)$.

Now the only way such a concatentation can occur is if the overlap in X started and ended on part of a short disk. More precisely, along the disk D we have a collection of paths $\mathbf{p_i} \subset \partial D$ that come from overlaps with small disks. From **Fact 3** there are at most $\leq M^2 I(R_L \mathcal{O})^2$ such overlaps. So at most $\leq M^2 I(R_L \mathcal{O})^2 + 1$ concatenations can occur. Since overlaps in $X(\sigma)$ have length at most $\mathcal{O}$, **Fact 4** follows.

Finally, with **Fact 3** and **Fact 4** in hand, it is straightforward to estimate the overlap ratio of $\hat{D}$. Indeed, $\hat{D}$ is the image of the long disk D in $X(\sigma)$ under the quotienting map. Since D is a long disk in $X(\sigma)$, it is a branched cover of a disk in X, with index $\geq K$. So the length of D is bounded below by $\geq KR_S$. When computing the length of $\hat{D}$, we see from **Fact 3** that at most $\leq M^2 I(R_L \mathcal{O})^2$ overlaps with short disks get collapsed to points. Since each of these overlaps has length $\leq \mathcal{O}$, we get the lower bound

$$
\begin{aligned}
\ell(\partial \hat{D}) &\geq KR_S - M^2 I(R_L \mathcal{O})^2 \mathcal{O} \\
&\geq \lambda^{-1} \mathcal{O} M^2 I(R_L \mathcal{O})^2 + \lambda^{-1} \mathcal{O}
\end{aligned}
$$

where the second inequality follows from the chosen value of K. Now using the estimate for the overlap length from **Fact 4** we get:

$$o(\hat{D}) \leq \frac{\big(M^2 I(R_L \mathcal{O})^2 + 1\big)\mathcal{O}}{\lambda^{-1}\mathcal{O}M^2 I(R_L\mathcal{O})^2 + \lambda^{-1}\mathcal{O}} = \lambda$$

as desired. Since this estimate holds for any long disk in our $X(\sigma)$ satisfying conditions (1), (2), (3), it concludes the proof of the theorem.

□

1.5 Concluding Remarks

We end our paper with some general remarks on topics related to our random models and our main theorem.

1.5.1 Multiple vertex case

The attentive reader will notice that our **Main Theorem** is stated for acceptable 2-complexes, but that in our proofs we work exclusively with the special case of a presentation 2-complex. In fact, the two statements are equivalent, as we now explain.

Given an arbitrary finite acceptable 2-complex X, we can take a spanning tree T in the 1-skeleton of X, and create a new 2-complex Z by collapsing T to a point. By construction, the 1-skeleton $Z^{(1)}$ is a bouquet of circles, so Z has a single vertex. The quotient map $\phi : X \to Z$ is a homotopy equivalence, since it is obtained by collapsing the contractable set T. Each polygonal 2-cell in X gives rise to a polygonal 2-cell in Z. Moreover, the restriction of ϕ to the 1-skeleton is a homotopy equivalence between $X^{(1)}$ and $Y^{(1)}$, so provides an isomorphism $\phi_\# : \pi_1(X^{(1)}) \to \pi_1(Y^{(1)})$. It follows that an attaching map α for a disk in X is a proper power in $\pi_1(X^{(1)})$ if and only if the corresponding attaching map $\phi \circ \alpha$ for a disk in Z is a proper power. Similarly a pair of disks have identical attaching map in $\pi_1(X^{(1)})$ if and only if the corresponding attaching maps in $\pi_1(Z^{(1)})$ are identical. This shows that Z is also an acceptable 2-complex, but with a single vertex, so can be viewed as a presentation 2-complex.

Finally, our model for random branched covers of Z are obtained by taking ordinary degree n covers of the 1-skeleton $Z^{(1)}$, and inducing a branched cover by attaching disks along all the connected lifts of an attaching map (see Section 1.3.1 and Section 1.3.2). From covering space theory, all information on lifting is encoded in the fundamental group of the 1-skeletons. Hence the group isomorphism $\phi_\#$ allows you to obtain a corresponding finite cover of $X^{(1)}$, and a homotopy equivalence between this finite cover and the 1-skeleton of the branched cover $Z(\sigma)$. Under this homotopy equivalence, we can transfer the lifts of the attaching maps to the finite cover of $X^{(1)}$ and form a corresponding branched cover $X(\sigma)$ of X. By construction, there is then an induced homotopy equivalence $X(\sigma) \simeq Z(\sigma)$. It follows that topological results about the random model can be transferred from the presentation 2-complex case to the general case of acceptable polygonal 2-complexes.

1.5.2 Relations that are proper powers

One of the conditions in our definition of acceptable 2-complexes is that none of the attaching maps for the 2-cells represent proper powers in $\pi_1(X^{(1)})$. This property was used in the proof of Lemma 1.2.7, showing that acceptable 2-complexes have finite overlap ratio. There is however a more fundamental reason for requiring this property. We introduced in Definition 1.2.1 the notion of a branched cover for a 2-complex X, and in Section 3.2 explained how to associate, to each element $\sigma \in \mathrm{Sym}(n)^k$ a space $X(\sigma)$. Let us consider this construction when some of the relations are proper powers.

Example 1.5.1. Consider the finite group $\mathbb{Z}_2 \cong \langle x \mid x^2 \rangle$. The presentation two complex X has a single loop labelled x, and a single attached by the degree two map on the circle. So X is homeomorphic to the projective plane $\mathbb{R}P^2$, a closed non-orientable surface. Any branched cover of X would also have to be a closed surface.

Now consider the space $X(\sigma)$, where $\sigma = (12 \ldots n)$. The 1-skeleton of $X(\sigma)$ is a cycle of length n. There is only one lift of the attaching map of degree 2 to this 1-skeleton. This is either a degree one map if n is even, or a degree two map if n is odd. So the space $X(\sigma)$ is either homeomorphic to $\mathbb{R}P^2$ (if n odd) or to $\mathbb{D}^2$ (if n even). Of these $X(\sigma)$, only the n odd case produces a branched cover.

The example above might cause some concern about our model for random branched covers. However, for acceptable 2-complexes, we can check that we do indeed obtain branched covers.

Lemma 1.5.2. *Let X be an arbitrary presentation 2-complex. If none of the relations are proper powers, then $X(\sigma)$ is a branched cover.*

Proof. We first note that there is a natural map $X(\sigma) \to X$ induced by the covering map on the 1-skeleton, extended to a branched cover from each 2-cell in $X(\sigma)$ to its image 2-cell in X. It is straightforward to check that every point in X that is **not** the center of a 2-cell has exactly n pre-images, where n is degree of the covering on the 1-skeleton. By construction, in the interior of the 2-cells the map is a branched covering. So we are left with checking the local topology along pre-images of edges.

In a polygonal 2-complex, the local topology at a point inside an edge is easy to describe. Let k denote the number of occurrences of that edge (and its inverse) along the boundary labels of all the attached 2-disks. Then locally, a closed neighborhood of the point is homeomorphic to the product $I \times C_k$, where C_k is the cone over a discrete set of k points. Under this identification the edge corresponds to the product of I times the cone point.

So to decide whether the canonical map $X(\sigma) \to X$ is a covering map, it suffices to compare, for a loop $e \subset X$ and a pre-image edge $e_i \subset X(\sigma)$, the number of occurrences of those edges along labels of disks in X and $X(\sigma)$ respectively. Take one of the local branches of the neighborhood $I \times C_k$ of the edge e in X. This corresponds to a unique edges in a disk D with label e. We can lift the boundary of this disk starting at that labeled edge e, with initial lift e_i. This will yield a lifted disk D' inside $X(\sigma)$. If we can construct such a lift for each edge labeled e in the boundary labels of disks in X, then we would get the exact same number of local branches around e_i in $X(\sigma)$. The problem that arises, as illustrated in the previous example, is that one could potentially start lifting from *distinct* edges labeled e in the boundary of a disk D, whose lifts give you the identical path in $X(\sigma)^{(1)}$. In that case, instead of having *two* local branches from the two distinct lifts, we will only obtain *one* local branch, as our construction of $X(\sigma)$ associates one disk to this closed loop.

Next we notice that this situation can only happen if the label on ∂D has a partial rotational symmetry. Namely, if s is the distance along ∂D where the two e labels occur, we note that rotating the labeling by s gives back the identical labeling. Since $1 \leq s < \ell(\partial D)$, the bi-infinite word on $\widetilde{\partial D}$ is now periodic under both a $\ell(\partial D)$-translation, and an s-translation. Arguing as in the proof of Lemma 1.2.7, one sees that the bi-infinite word is in fact periodic with period $GCD(s, \ell(\partial D)) < \ell(\partial D)$.

This implies that the relation corresponding to the disk D is a proper power, a contradiction. $\square$

Chapter 2: Codimension Two Complements in Hyperbolic Products

2.1 Introduction

For any $m, n \geq 2$, suppose that M is a finite-volume manifold with universal cover isometric to the product of hyperbolic spaces, $\mathbb{H}^m \times \mathbb{H}^n$, and containing a (possibly disconnected) embedded, compact, codimension two totally geodesic submanifold S whose lifts to $\mathbb{H}^m \times \mathbb{H}^n$ are isometrically embedded copies of $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. Let $N = M \setminus S$ and, for some integer $d > 2$, suppose that the ramified branched cover X_d of M about S is a smooth manifold. In this paper we aim to study geometric and topological properties of N and X_d.

The reason for considering submanifolds of codimension two in our definition of N is that their removal significantly alters the fundamental group. If the codimension were higher, then we would have that $\pi_1(N) = \pi_1(M)$, whereas if S had codimension one then there is a well-known simple relationship between $\pi_1(N)$, $\pi_1(M)$. In the codimension two setting, this connection is more complicated and rich. In the present setting, it is straightforward that $\pi_1(N)$ is an *overlattice* in the sense of Gromov [19, Pg. 126], that is, $\pi_1(N)$ surjects onto the lattice $\pi_1(M)$, with infinitely generated kernel. We include the study of the branched cover X_d since many of our geometric constructions will apply to these manifolds as well.

Geometry of (M, S)

From a geometric point of view we look to answer the following two questions:

1. Does N admit a complete, finite volume, nonpositively curved metric which turns S into a cusp of N?

2. Does X_d admit a nonpositively curved Riemannian metric?

There exists a lattice $\Gamma < \text{Isom}(\mathbb{H}^m \times \mathbb{H}^n)$ such that $M = (\mathbb{H}^m \times \mathbb{H}^n)/\Gamma$. It turns out, maybe somewhat surprisingly, that the difficulty in answering questions (1) and (2) depends on properties of Γ. The simplest case is when Γ splits as a product. That is, when there exist lattices $\Gamma_m < \text{Isom}(\mathbb{H}^m)$ and $\Gamma_n < \text{Isom}(\mathbb{H}^n)$ such that $\Gamma = \Gamma_m \times \Gamma_n$. In this case we can prove the following, where statement (2) in Theorem 2.1.1 is due to Fornari and Schroeder [15].

Theorem 2.1.1. *Suppose that* $\Gamma = \Gamma_m \times \Gamma_n$ *splits as a product, and let* M, S, N, *and* X_d *be as defined above. Then*

1. *N admits a complete, finite volume, A-regular*[1]*, nonpositively curved Riemannian metric, and*

2. *X_d admits a nonpositively curved Riemannian metric.*

Our proof of Theorem 2.1.1 (1) uses the confluence of an idea of Fornari–Schroeder [15] and a straightforward adaptation of the A-regular metric constructed by Belegradek in [4]. The latter is similar to the "funnel" construction of Avramidi and

[1] See Definition 2.3.9

Pham in [2]. But, unfortunately, the arguments for Theorem 2.1.1 are insufficient for when Γ is not a product (including the argument for (2) in [15]). The reason for this is as follows. Our proof constructs a nonpositively curved metric $\tilde{g}$ on the universal cover $\mathbb{H}^m \times \mathbb{H}^n$. But there exist points arbitrarily far away from $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$ where $\tilde{g}$ does not agree with the product hyperbolic metric $\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n$. We will show that, when Γ splits as a product, $\tilde{g}$ still descends to a well-defined metric g on M (Proposition 2.3.4). But there is no reason to believe that g will be well-defined when Γ is not a product. See Remarks 2.3.5 and 2.4.5 for more information.

Corollary 2.1.2. *Suppose that $\Gamma = \Gamma_m \times \Gamma_n$ splits as a product, and let M, S, and N be as defined above. Then N is aspherical.*

Remark 2.1.3. In order to avoid confusion, throughout this paper a metric $\tilde{g}$ with a "tilde" will denote a metric on the universal cover, whereas we drop the "tilde" and use g for the corresponding metric on our manifold. We will sometimes also drop the "tilde" in subscripts and in calculations when the metric should be clear from context.

The other two cases are when Γ is reducible and when Γ is irreducible (see Definition 2.2.1). If Γ is reducible then it contains a product lattice as a finite index subgroup (Remark 2.2.2). This gives us the following Corollary.

Corollary 2.1.4. *Suppose that $\Gamma < Isom(\mathbb{H}^m \times \mathbb{H}^n)$ is reducible, and let M and S be as defined above. Then M admits a finite cover $\overline{M}$ with lift $\overline{S}$ of S such that*

1. *$\overline{N} = \overline{M} \setminus \overline{S}$ admits a complete, finite volume, A-regular, nonpositively curved Riemannian metric, and*

2. $\overline{X}_d$*, the d-fold ramified cover of* $\overline{M}$ *about* $\overline{S}$*, admits a nonpositively curved Riemannian metric.*

The latter statement, of course, assumes that the branched cover $\overline{X}_d$ is a smooth manifold.

Remark 2.1.5. Suppose Γ is reducible but not a product. Then N has a finite cover $\overline{N} = (\mathbb{H}^m \times \mathbb{H}^n)/\overline{\Gamma}$ with $\overline{\Gamma}$ a product, and so $\overline{N}$ admits a nonpositively curved metric. It is tempting to think that one can construct an action of the finite group $G = \Gamma/\overline{\Gamma}$ on N to define this nonpositively curved metric on N. But such an action would induce a G-action on the space of nonpositively curved metrics on N via the pullback metric, and one can find a G-equivariant nonpositively curved metric on N exactly when this induced action has a fixed point. It is unknown to the authors if this is possible, but results like those in [13] make this seem unlikely in general.

When Γ is irreducible, the best that we can prove is that N and X_d admit an "almost nonpositively curved metric". We say that a Riemannian manifold Y with finite volume is *almost nonpositively curved* if, given any $\varepsilon > 0$, Y admits a Riemannian metric with all sectional curvatures bounded above by ε and with volume bounded above by some constant C_{vol} independent of ε. Our precise statement is as follows.

Theorem 2.1.6. *Suppose that* $\Gamma < Isom(\mathbb{H}^m \times \mathbb{H}^n)$*, and let* $\varepsilon > 0$*. Let* M, S, N*, and* X_d *be as defined above.*

1. *The manifold* N *admits a complete Riemannian metric of almost nonpositive curvature. Moreover, one can choose* $C_{vol} = vol(M, \mathfrak{h}_m \times \mathfrak{h}_n) + \xi$ *for any* $\xi > 0$*.*

2. *The branched cover* X_d *is almost nonpositively curved. In addition, one can choose* $C_{vol} = d \cdot vol(M, \mathfrak{h}_m \times \mathfrak{h}_n) + \xi$ *for any* $\xi > 0$*.*

Remark 2.1.7. Our definition of *almost nonpositive curvature* for a general manifold is not equivalent to the definition that is frequently used for compact manifolds (see, for example, [17] or [3]). Using Chern-Weil theory and an argument similar to that in [9], one can show that a 4-manifold with almost nonpositive curvature (in our sense) has nonnegative Euler characteristic. In particular, Theorem 2.1.6 is nontrivial in that it cannot apply to any 4-manifold with negative Euler characteristic.

The above Theorem is true for any lattice Γ, but we are mainly interested in Γ irreducible. In this case, we must have $m = n$. Also, for statement (1) above note that the product hyperbolic metric $\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n$ descends to a nonpositively curved metric on N, but this metric fails to be complete at S. Lastly, note that we do not make any assumptions about the normal injectivity radius of S in M (besides that it is positive since S is compact). Theorem 2.1.6 asserts that N and X_d admit metrics with all curvatures less than ε independent of the size of the normal injectivity radius of S.

We prove Theorem 2.1.6 in the standard way: we write the product metric $\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n$ on $\mathbb{H}^m \times \mathbb{H}^n$ in polar coordinates about $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$ (equation (2.3.7)), we consider the corresponding warped-product metric $\tilde{\lambda}$ (2.3.8), we calculate formulas for the sectional curvature tensor of $\tilde{\lambda}$ (Theorem 2.3.6), and then we attempt to find coefficient functions so that $\tilde{\lambda}$ has bounded volume, nonpositive curvature, and satisfies our desired properties. The difficulty, of course, is that at each point of M there are many 2-planes with curvature 0 with respect to the metric $\mathfrak{h}_m \times \mathfrak{h}_n$. So, in one direction, there is no room for error in this warping procedure. In light of [40], it is certainly possible that one or both of N and X_d do not admit a nonpositively curved metric with the desired properties, but we do not know how to prove this.

We now describe how our results above fit into the existing literature. Given M as above, any embedded, compact, connected, totally geodesic, codimension two submanifold S of M, has universal cover $\tilde{S}$ given by exactly one of the following three types:

1. $\widetilde{S}$ is isometric to either $\mathbb{H}^{m-2} \times \mathbb{H}^n$ or $\mathbb{H}^m \times \mathbb{H}^{n-2}$,
2. $\widetilde{S}$ is isometric to $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$,
3. $m = n = 2$ and $\widetilde{S}$ is isometric to $\mathbb{H}^2$ which (up to isometry) has been diagonally embedded in $\mathbb{H}^2 \times \mathbb{H}^2$.

We first consider complete, finite volume metrics on N that turn S into a cusp. In Case (1), N admits a complete, finite-volume, non-positively curved, A-regular metric by a straightforward extension of a result of Belegradek [5] to the product case. Our stated results pertain to case (2). In Case (3), it is presently unknown if N admits such a metric. A similar study for complex hyperbolic manifolds is contained in [4] and [32]. However, in that setting, the consideration of products of complex hyperbolic manifolds is not interesting, as the only type of codimension two totally geodesic real submanifolds not included in case (1) are induced by diagonal embeddings of $\mathbb{CH}^1$ into $\mathbb{CH}^1 \times \mathbb{CH}^1$, which is precisely Case (3) above. Lastly, note that quaternionic (dimension > 1) and Cayley hyperbolic manifolds do not have totally geodesic submanifolds with real codimension two.

We now turn our attention to metrics on branched cover manifolds. The fact that branched covers X of manifolds M from Case (1) above admit a non-positively curved metric is a straightforward extension of the famous Gromov-Thurston construction [20]. Case (2) is again considered above, where we include the work of Fornari and

Schroeder [15]. Conversely, Stadler [40] has shown that such a result does not hold in the setting of Case (3). More precisely, if M is modeled on $\mathbb{H}^2 \times \mathbb{H}^2$ and S lifts to the diagonal in $\mathbb{H}^2 \times \mathbb{H}^2$ (up to isometry), then any d-fold branched cover X over M does not admit a non-positively curved Riemannian metric. This suggests that, in the setting of Case (3), $M \setminus S$ might not admit a complete, finite-volume, non-positively curved Riemannian metric, and also lends evidence that Theorem 2.1.6 may be optimal. In the complex hyperbolic setting, similar branched covers were recently constructed in important work of Stover–Toledo [42], and interesting negatively curved metrics were constructed on collections of these manifolds in [43] and [34].

2.2 Lattices in $\mathrm{Isom}(\mathbb{H}^m \times \mathbb{H}^n)$

In this section, we recall some basic facts about locally symmetric orbifolds M whose universal cover is isometric to $\mathbb{H}^m \times \mathbb{H}^n$, for $m, n \geq 2$.

Throughout we let q_m denote the *standard quadratic form* on $\mathbb{R}^{m+1}$, which is defined by the formula

$$q_m(x_1, \dots, x_{m+1}) = x_1^2 + \cdots + x_m^2 - x_{m+1}^2. \tag{2.2.1}$$

The hyperboloid model of hyperbolic space is then given by

$$\mathbb{H}^m = \left\{(x_1, \dots, x_{m+1}) \in \mathbb{R}^{m+1} \mid q_m(x_1, \dots, x_{m+1}) = -1,\ x_{m+1} > 0\right\}, \tag{2.2.2}$$

and the orientation preserving group of isometries, $\mathrm{Isom}^+(\mathbb{H}^m)$, of this model is given by

$$\mathrm{SO}_0(m,1) = \left\{A \in \mathrm{SL}_{m+1}(\mathbb{R}) \mid q_m(Ax) = q_m(x),\ \forall x \in \mathbb{R}^{m+1}\right\}^\circ,$$

where the notation on the righthand side of this equation denotes the identity component. In particular,

$$\mathrm{SO}_0(m,1)\times\mathrm{SO}_0(n,1)\cong\mathrm{Isom}^+(\mathbb{H}^m)\times\mathrm{Isom}^+(\mathbb{H}^n)\leq_{f.i.}\mathrm{Isom}^+(\mathbb{H}^m\times\mathbb{H}^n),\quad(2.2.3)$$

where the notation means that this is a subgroup of finite index. For ease of notation, we let $G_m = \mathrm{SO}_0(m,1)$ throughout the remainder of this section.

Recall that two discrete subgroups $\Gamma, \Gamma' < \mathrm{Isom}^+(\mathbb{H}^m \times \mathbb{H}^n)$ are *commensurable* (in the wide sense) if there exists $g \in \mathrm{Isom}^+(\mathbb{H}^m \times \mathbb{H}^n)$ such that $\Gamma \cap g\Gamma' g^{-1}$ is a finite index subgroup of both Γ and $g\Gamma' g^{-1}$. Geometrically this is equivalent to $(\mathbb{H}^m\times\mathbb{H}^n)/\Gamma$ and $(\mathbb{H}^m \times \mathbb{H}^n)/\Gamma'$ having a common finite cover. The notion of commensurability forms an equivalence relation and we will primarily be interested in commensurability classes of subgroups (equivalently, orbifolds). As such, Equation (2.2.3) implies that we may focus our attention on subgroups $\Gamma < G_m \times G_n$.

We call $\Gamma < G_m \times G_n$ a *lattice* if Γ is discrete and the quotient $(G_m \times G_n)/\Gamma$ has finite Haar measure, equivalently, if the orbifold $(\mathbb{H}^m \times \mathbb{H}^n)/\Gamma$ has finite-volume. If a lattice Γ is additionally torsion-free, then the quotient $(\mathbb{H}^m\times\mathbb{H}^n)/\Gamma$ is a finite-volume manifold. Such lattices are always finitely generated [38, Cor 13.15] and therefore Selberg's lemma [39, Lem 8] allows one to find a finite index subgroup $\Gamma' \leq \Gamma$ which is torsion-free and consequently a finite-volume manifold $(\mathbb{H}^m\times\mathbb{H}^n)/\Gamma'$ which finitely covers $(\mathbb{H}^m \times \mathbb{H}^n)/\Gamma$. Owing to this, for the time being, we postpone the issue of requiring lattices to be torsion-free.

Definition 2.2.1. We say that a lattice $\Gamma < G_m \times G_n$ is *irreducible* if its projection to each of G_m, G_n is dense and *reducible* otherwise.

Remark 2.2.2. It is straightforward to construct reducible lattices by letting $\Gamma = \Gamma_1 \times \Gamma_2$, where $\Gamma_1 < G_m$ and $\Gamma_2 < G_n$ are lattices, and indeed all reducible lattices have a finite index subgroup of this form [38, Ch. 5].

It is also well-known that irreducible lattices exist in $G_m \times G_n$ if and only if G_m and G_n are isotypic, that is, if their complexifications are isogenous (see for instance [36, Ch. 5F]). By a straightforward dimension argument, one sees that this is possible if and only if $m = n$. Therefore when $m \neq n$, necessarily any lattice $\Gamma < G_m \times G_n$ is reducible. When $m = n$ irreducible lattices are necessarily arithmetic by seminal work of Margulis [31, Ch. IX], as $G_m \times G_n$ is a higher rank Lie group.

Then we consider families of irreducible arithmetic lattices in $G_m \times G_n$ whenever $m = n \geq 2$. In fact, as it requires minimal extra work, we will describe how to construct infinitely many commensurability classes of arithmetic lattices in $(G_m)^k$ for any $k \in \mathbb{N}$ and therefore finite-volume orbifolds of the form $(\mathbb{H}^m \times \cdots \times \mathbb{H}^m)/\Gamma$. Throughout this subsection, fix a choice of $k \in \mathbb{N}$.

Let K be a totally real number field with $d = [K : \mathbb{Q}]$ and denote by $\sigma_i : K \to \mathbb{R}$ the Galois embeddings of K for $1 \leq i \leq d$. We will always require that $d \geq k$ in the sequel. Let q be a K-quadratic form on the K-vector space K^{m+1}, then we say that the pair (q, K) is *admissible* if 1) K is the minimal field of definition of q and 2) the forms $q^{\sigma_1}, \ldots, q^{\sigma_k}$ each have signature $(n, 1)$ over $\mathbb{R}$, and q^{σ_i} is positive definite for any $i > k$.

Given a field E and a quadratic form q on the vector space E^{m+1}, let B_q denote the matrix associated to q in the standard basis. Then the *special orthogonal group associated to* (E, q) is given by

$$\mathrm{SO}_q(E) = \{A \in \mathrm{SL}_{m+1}(E) \mid A^T B_q A = B_q\}.$$

Of particular importance is when (q, K) is an admissible pair, then Sylvester's law of inertia gives that

$$\mathrm{SO}_{q^{\sigma_i}}(\mathbb{R}) \cong \begin{cases} \mathrm{SO}(m,1), & 1 \le i \le k \\ \mathrm{SO}(m+1), & i > k \end{cases}. \tag{2.2.4}$$

In particular, we have the embedding

$$\psi : \mathrm{SO}_q(K) \hookrightarrow \underbrace{\mathrm{SO}(m,1) \times \cdots \times \mathrm{SO}(m,1)}_{k \text{ times}} \times \prod_{i>k} \mathrm{SO}(m+1),$$
$$g \mapsto (\sigma_1(g), \ldots, \sigma_k(g), \sigma_{k+1}(g), \ldots, \sigma_d(g))$$

where we abusively suppress the post-composition with the isomorphisms from Equation (2.2.4). If $\mathcal{O}_K$ denotes the ring of integers of K and

$$\Lambda = \mathrm{SO}_q(\mathcal{O}_K) = \left\{ A \in \mathrm{SL}_{m+1}(\mathcal{O}_K) \mid A^T B_q A = B_q \right\},$$

then $\psi(\Lambda)$ is discrete in $\mathrm{SO}(m,1)^k \times \mathrm{SO}(m+1)^{d-k}$. Indeed, this follows from the fact that $\mathcal{O}_K$ is discrete in $\mathbb{R}^d$ via the similar map induced by the Galois embeddings. An important theorem of Borel–Harish-Chandra [7] moreover shows that $\psi(\Lambda)$ is a lattice in this product.

Letting π denote the composition of ψ with the projection from this product to the $\mathrm{SO}(m,1)^k$ factor, it follows that $\Gamma = \pi(\psi(\Lambda)) \cap (G_m)^k$ is a lattice in $(G_m)^k$. Indeed, the kernel of π is the compact group $\prod_{i>k} \mathrm{SO}(m+1)$. It is also straightforward to check from the setup that the projection to each G_m factor of $(G_m)^k$ is dense and therefore Γ is an irreducible lattice in $(G_m)^k$.

Definition 2.2.3. We call any lattice in $\mathrm{Isom}^+(\mathbb{H}^m \times \cdots \times \mathbb{H}^m)$ that is commensurable with some Γ as constructed above an *irreducible arithmetic lattice of simplest type.* Given such a lattice, we call the pair (q, K) used to construct Γ the *arithmetic data* associated to the lattice.

A couple of facts follow from this definition which we list as remarks.

Remark 2.2.4. Given two arithmetic lattices of simplest type Γ, Γ' with associated arithmetic datum (q, K), (q', K') then Γ and Γ' are commensurable if and only if $K \cong K'$ and q and q' are in the same K-similarity class. The latter means that q is K-isometric to $\lambda q'$ for some $\lambda \in K^*$. In particular, the number field K and the K-similarity classes of q is a complete invariant of the commensurability class.

Remark 2.2.5. By Godement's compactness criterion [36, Prop 5.26] an arithmetic lattice of simplest type Γ with associated arithmetic data (q, K) is cocompact if and only if q is anisotropic, that is, if there is no non-zero $v \in K^{m+1}$ for which $q(v) = 0$. Otherwise, we say that q is isotropic. These notions are commensurability invariants. Moreover, the Hasse–Minkowski theorem shows that q is isotropic if and only if q is isotropic over every completion of K. Therefore, whenever $d > k$ the form q is necessarily anisotropic since there is at least one completion associated to a Galois embedding for which q is positive definite (and hence anisotropic).

Using the previous subsections, for each pair $m, n \geq 2$, we now provide infinitely many commensurability classes of manifolds to which Theorems 2.1.1 and 2.1.6 applies.

Theorem 2.2.6 (N.Miller)**.** *For every $m, n \in \mathbb{N}$ such that $m, n \geq 2$, there exists infinitely many commensurability classes of finite-volume manifolds with universal cover isometric to $\mathbb{H}^m \times \mathbb{H}^n$ and containing an embedded compact totally geodesic submanifold with universal cover isometric to $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. Moreover, when $m = n$, these manifolds may be taken to be either reducible or irreducible.*

2.3 Riemannian Metric for Product Lattices

In this section we first provide details for a construction that was quickly outlined in the second and third paragraphs of Section 3 of [15]. We provide full details here for two reasons. The first is because this construction is critical to our proof for Theorem 2.1.1 (1). The second reason is to illustrate why this argument only works for product lattices. We then prove Theorem 2.1.1 (1), using curvature formulas for a warped product metric (Theorem 2.3.6) that will play a much larger role in Section **??**.

2.3.1 Warped Product Metric on Product Lattices

Consider the manifold $\mathbb{R}^n \cong \mathbb{R}^{n-1} \times (-\infty, \infty)$. If we endow $\mathbb{R}^{n-1}$ with the hyperbolic metric $\tilde{\mathfrak{h}}_{n-1}$ and let x_n denote the variable in the last coordinate, then it is well-known (see for example [5]) that the hyperbolic metric $\tilde{\mathfrak{h}}_n$ on $\mathbb{R}^n$ can be written as

$$\tilde{\mathfrak{h}}_n = \cosh^2(x_n)\tilde{\mathfrak{h}}_{n-1} + dx_n^2. \tag{2.3.1}$$

Now, let $\sigma_n : \mathbb{R} \to (0, \infty)$ be a positive, real-valued function of x_n. Replacing $\cosh(x_n)$ with $\sigma_n(x_n)$ in equation (2.3.1) yields the *warped-product metric*

$$\tilde{g}_n = \sigma_n^2(x_n)\tilde{\mathfrak{h}}_{n-1} + dx_n^2. \tag{2.3.2}$$

The components of the sectional curvature tensor of $\tilde{g}_n$ will, of course, be functions of σ_n. Formulas for the sectional curvature tensor of $\tilde{g}_n$ are again well-known (see for example [5, Eqns. 3.1, 3.2]). We list the relevant formulas after establishing an appropriate frame. All calculations in the following paragraph are with respect to the hyperbolic metric $\tilde{\mathfrak{h}}_n$.

Let $\varphi : \mathbb{R}^n \to \mathbb{R}^{n-1} \times \{0\}$ be the orthogonal projection. Let $q \in \mathbb{R}^n$ and $p = \varphi(q)$. Choose an orthonormal frame $(\check{X}_1, \ldots, \check{X}_{n-1})$ about p for $T_p\mathbb{R}^{n-1}$. Since $\varphi : \mathbb{R}^n \to \mathbb{R}^{n-1}$ is a submersion, $d\varphi : T_q\mathbb{R}^n \to T_p\mathbb{R}^{n-1}$ is an epimorphism. Moreover, letting $\mathcal{H}$ denote the orthogonal complement to $\partial x_n := \partial/\partial x_n$ in $T_q\mathbb{R}^n$, we have that $d\varphi|_{\mathcal{H}} : \mathcal{H} \to T_p\mathbb{R}^{n-1}$ is an isomorphism. Letting $X_i = d\varphi|_{\mathcal{H}}^{-1}(\check{X}_i) \in \mathcal{H}$ defines an orthogonal frame $(X_1, \ldots, X_{n-1})$ for $\mathcal{H}$ about q. This frame is *not* orthonormal since $\tilde{\mathfrak{h}}_n(X_i, X_i) = \cosh^2(x_n)$ for each i. So the collection of vectors

$$Y_i = \frac{1}{\cosh(x_n)} X_i \quad \text{for } 1 \leq i \leq n-1, \qquad Y_n = \frac{\partial}{\partial x_n} \tag{2.3.3}$$

forms an orthonormal frame for $T_q\mathbb{R}^n$ about q with respect to the metric $\tilde{\mathfrak{h}}_n$.

The corresponding orthonormal frame for $\tilde{g}_n$ is

$$Y_i = \frac{1}{\sigma_n} X_i \quad \text{for } 1 \leq i \leq n-1, \qquad Y_n = \frac{\partial}{\partial x_n}. \tag{2.3.4}$$

Let $K_n(A, B)$ denote the sectional curvature of the 2-plane spanned by the vectors A and B with respect to $\tilde{g}_n$. Formulas for the sectional curvatures of $\tilde{g}_n$ with respect to the frame in Equation (2.3.4) are

$$K_n(Y_i, Y_j) = -\frac{1}{\sigma_n^2} - \left(\frac{\sigma_n'}{\sigma_n}\right)^2 \quad \text{for } 1 \leq i, j \leq n-1. \tag{2.3.5}$$

$$K_n(Y_i, Y_n) = -\frac{\sigma_n''}{\sigma_n} \quad \text{for } 1 \leq i \leq n-1. \tag{2.3.6}$$

Additionally, there are no mixed terms. The frame in Equation (2.3.4) diagonalizes the curvature operator with respect to $\tilde{g}_n$.

Before stating Proposition 2.3.2, we state Lemma 2.3.1, a smoothing lemma that we will use at several points throughout this paper including in the proof of Proposition 2.3.2. Lemma 2.3.1 is intuitively clear, but the proof of a more general version of this lemma can be found in [4, App. A].

Lemma 2.3.1 (Smoothing Lemma)**.** *Let* $f_1 : [a_1, c] \to \mathbb{R}$ *and* $f_2 : [c, a_2] \to \mathbb{R}$ *be smooth functions which satisfy* $f_i'' > 0$ *for* $i = 1, 2$*,* $f_1(c) = f_2(c)$*, and* $f_1'(c) \leq f_2'(c)$*. If* $f : [a_1, a_2] \to \mathbb{R}$ *denotes the concatenation of* f_1 *and* f_2*, then there exists* $\delta > 0$ *and a smooth function* $f_\delta : [a_1, a_2] \to \mathbb{R}$ *such that* $f_\delta'' > 0$ *and both* $f_\delta = f$ *and* $f_\delta' = f'$ *at the endpoints* a_1 *and* a_2*.*

Proposition 2.3.2. *Let* $\eta > 0$*. Then there exist* ε_1 *and* ε_2 *satisfying* $0 < \varepsilon_1, \varepsilon_2 < (1/2)\eta$ *and a non-positively curved metric* g *on* $\mathbb{R}^{n-1} \times [-\eta, \eta]$ *which satisfies*

$$\tilde{g} = \tilde{h}_{n-1} + dx_n^2 \quad \text{over } \mathbb{R}^{n-1} \times (-\varepsilon_1, \varepsilon_1),$$

$$\tilde{g} = \tilde{\mathfrak{h}}_n \quad \text{over } \mathbb{R}^{n-1} \times ([-\eta, -\eta + \varepsilon_2) \cup (\eta - \varepsilon_2, \eta]) .$$

Remark 2.3.3. Proposition 2.3.2 is easily seen to be true for any $\varepsilon_1, \varepsilon_2$ which satisfy the given inequality, however, in what follows we only need the weaker statement above.

Proof. Let $\varepsilon_1 = \varepsilon_2 = (1/8)\eta$. Define the continuous, piecewise-smooth function $\hat{\sigma}_n$ over $[0, \eta]$ by

$$\hat{\sigma}_n = \begin{cases} 1, & \text{if } 0 \leq x \leq \frac{1}{4}\eta \\ 1 + \frac{\cosh\left(\frac{3}{4}\eta\right) - 1}{\frac{1}{2}\eta}\left(x - \frac{1}{4}\eta\right), & \text{if } \frac{1}{4}\eta \leq x \leq \frac{3}{4}\eta \\ \cosh(x), & \text{if } \frac{3}{4}\eta \leq x \leq \eta \end{cases}$$

The middle piece of this function is linear and interpolates between the values 1 and $\cosh(\frac{3}{4}\eta)$ over the closed interval $[(1/4)\eta, (3/4)\eta]$. Note that, away from the breakpoints $(1/4)\eta$ and $(3/4)\eta$, this function is smooth and satisfies $\hat{\sigma}_n'' \geq 0$.

One sees immediately that $\lim_{x\to(1/4)\eta^-}\hat{\sigma}'_n(x) = 0 < \lim_{x\to(1/4)\eta^+}\hat{\sigma}'_n(x)$. For the other breakpoint, note that

$$\lim_{x\to(3/4)\eta^-}\hat{\sigma}'_n(x) = \frac{\cosh\left(\frac{3}{4}\eta\right)-1}{\frac{1}{2}\eta},$$
$$\lim_{x\to(3/4)\eta^+}\hat{\sigma}'_n(x) = \sinh\left(\frac{3}{4}\eta\right).$$

Using Taylor series it is an easy exercise to see that, again, $\lim_{x\to(3/4)\eta^-}\hat{\sigma}'_n(x) < \lim_{x\to(3/4)\eta^+}\hat{\sigma}'_n(x)$. The smooth function σ_n is obtained by applying Lemma 2.3.1 to $\hat{\sigma}_n$ about both break points. By our above calculations we know that $\sigma''_n(x) \geq 0$ for all $x \in [0,\eta]$. Finally, since 1 and $\cosh(x)$ are even functions, we can define σ_n in a symmetric manner over $[-\eta, 0]$. □

Suppose $\Gamma < \mathrm{Isom}(\mathbb{H}^m \times \mathbb{H}^n)$ is a product lattice. So there exist $\Gamma_m < \mathrm{Isom}(\mathbb{H}^m)$ and $\Gamma_n < \mathrm{Isom}(\mathbb{H}^n)$ such that $\Gamma = \Gamma_m \times \Gamma_n$. Let $X_m = \mathbb{H}^m/\Gamma_m$ and $X_n = \mathbb{H}^n/\Gamma_n$. Then

$$M = (\mathbb{H}^m \times \mathbb{H}^n)/\Gamma = (\mathbb{H}^m/\Gamma_m) \times (\mathbb{H}^n/\Gamma_n) = X_m \times X_n.$$

The universal cover of (each component of) our codimension two submanifold S is isometric to $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. Thus, S naturally splits as a product $Y_m \times Y_n$, where $Y_m \subset X_m$ and $Y_n \subset X_n$ are both codimension one totally geodesic submanifolds. The hyperbolic metrics $\tilde{\mathfrak{h}}_m$ and $\tilde{\mathfrak{h}}_n$ descend to metrics on X_m and X_n, of course. But by Proposition 2.3.2 the manifolds X_i $(i = m, n)$ admit nonpositively curved metrics g_i which, locally about Y_i, are isometric to $\mathfrak{h}_{i-1} \times e_1$ (where $\tilde{e}_j$ denotes the Euclidean metric on $\mathbb{R}^j$). The metric $g = g_m \times g_n$ is thus a nonpositively curved metric on M which, for some $r > 0$, is isometric to $\mathfrak{h}_{m-1} \times \mathfrak{h}_{n-1} \times e_2$ on the interior of the r-tube about S.

Since $(M, \mathfrak{h}_m \times \mathfrak{h}_n)$ is assumed to have finite volume, each $(X_i, \mathfrak{h}_i)$ has finite volume. But g_i and $\mathfrak{h}_i$ agree on X_i outside of a compact set. Therefore (X_i, g_i) has finite volume, and thus $(M, g_m \times g_n)$ likewise has finite volume. We summarize all of the above in the following Proposition.

Proposition 2.3.4. *Let $\Gamma < Isom(\mathbb{H}^m \times \mathbb{H}^n)$ be a product lattice, let $M = (\mathbb{H}^m \times \mathbb{H}^n)/\Gamma$, and let $S \subset M$ be a totally geodesic, codimension two submanifold whose universal cover is isometric to $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. Then there exists a metric g on M which is nonpositively curved, has finite volume, and locally about S is isometric to $\mathfrak{h}_{m-1} \times \mathfrak{h}_{n-1} \times e_2$.*

Remark 2.3.5. The construction above only works for product lattices. At the level of the universal cover, the metric $\tilde{g}_i$ differs from $\tilde{\mathfrak{h}}_i$ on some strip about $\mathbb{H}^{i-1}$ $(i = m, n)$. The product of these strips will contain points arbitrarily far from $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$. So, if Γ does not split as a product of lattices, then there is no guarantee that the metric $\tilde{g} = \tilde{g}_m \times \tilde{g}_n$ will descend to a well-defined metric on M.

There are several easy arguments to prove Theorem 2.1.1 (1) using Proposition 2.3.4. But to prove Theorem 2.1.6 we will need the curvature formulas in Theorem 2.3.6 below. So we develop these formulas in this subsection, and then use them to prove Theorem 2.1.1 in the next subsection.

Consider the product $\mathbb{R}^m \times \mathbb{R}^n \cong \mathbb{R}^{m-1} \times (-\infty, \infty) \times \mathbb{R}^{n-1} \times (-\infty, \infty)$. By Equation (2.3.1), we have that the product hyperbolic metric can be written as

$$\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n = \cosh^2(x_m)\tilde{\mathfrak{h}}_{m-1} + dx_m^2 + \cosh^2(y_n)\tilde{\mathfrak{h}}_{n-1} + dy_n^2.$$

If we define

$$x_m = r\cos(\theta), \quad y_n = r\sin(\theta),$$

for $r \in [0, \infty)$ and $\theta \in \mathbb{S}^1$, then the product hyperbolic metric becomes

$$\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n = \cosh^2(r\cos(\theta))\tilde{\mathfrak{h}}_{m-1} + \cosh^2(r\sin(\theta))\tilde{\mathfrak{h}}_{n-1} + r^2 d\theta^2 + dr^2. \tag{2.3.7}$$

The associated warped-product metric is

$$\tilde{\lambda} = \alpha^2(r,\theta)\tilde{\mathfrak{h}}_{m-1} + \beta^2(r,\theta)\tilde{\mathfrak{h}}_{n-1} + f^2(r)d\theta^2 + dr^2, \tag{2.3.8}$$

where α and β are positive functions of r and θ and f is a positive function of r. In Theorem 2.3.6 we work out formulas for the components of the sectional curvature tensor of λ. These formulas will, of course, be functions of α, β, and f.

As in Section 2.3.2, we must first define an orthonormal frame for λ. For the hyperbolic metrics in λ, we use the half-space model and write

$$\mathfrak{h}_{m-1} = \frac{1}{x_{m-1}^2}dx_1^2 + \cdots + \frac{1}{x_{m-1}^2}dx_{m-1}^2, \tag{2.3.9}$$

$$\mathfrak{h}_{n-1} = \frac{1}{y_{n-1}^2}dy_1^2 + \cdots + \frac{1}{y_{n-1}^2}dy_{n-1}^2. \tag{2.3.10}$$

Let $X_i' = \frac{\partial}{\partial x_i}$, for $1 \leq i \leq m-1$ and $Y_j' = \frac{\partial}{\partial y_j}$, for $1 \leq j \leq n-1$. Then $\{X_1', \ldots, X_{m-1}'\}$ and $\{Y_1', \ldots, Y_{n-1}'\}$ will be an orthogonal frame for the orthogonal complement to $\{\partial\theta, \partial r\}$ defined in an identical manner as in Subsection 2.3.2.1. We define the orthonormal frame $\{X_1, \ldots, X_{m-1}, Y_1, \ldots, Y_{n-1}, Z_\theta, Z_r\}$ by

$$X_i = \frac{x_{m-1}}{\alpha(r,\theta)}X_i' \quad \text{for } 1 \leq i \leq m-1, \tag{2.3.11}$$

$$Y_i = \frac{y_{n-1}}{\beta(r,\theta)}Y_i' \quad \text{for } 1 \leq i \leq n-1, \tag{2.3.12}$$

$$Z_\theta = \frac{1}{f(r)}\frac{\partial}{\partial\theta}, \tag{2.3.13}$$

$$Z_r = \frac{\partial}{\partial r}. \tag{2.3.14}$$

2.3.2 Sectional Curvature Formula

In what follows, for the functions $\alpha(r,\theta)$ and $\beta(r,\theta)$, we will use the standard notation

$$\alpha_r = \frac{\partial \alpha}{\partial r}, \quad \alpha_\theta = \frac{\partial \alpha}{\partial \theta},$$

and similarly for β and for second order partial derivatives. It is important to note that subscripts only stand for partial derivatives when they are used with the functions α and β.

Theorem 2.3.6. *Let $\tilde{\lambda} = \alpha^2(r,\theta)\tilde{\mathfrak{h}}_{m-1} + \beta^2(r,\theta)\tilde{\mathfrak{h}}_{n-1} + f^2(r)d\theta^2 + dr^2$ be the metric on $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times \mathbb{S}^1 \times [0,\infty)$ defined in Equation (2.3.8). Up to the symmetries of the curvature tensor, the non-zero components of the curvature tensor are:*

$$K_\lambda(X_i, X_j) = -\frac{1}{\alpha^2}\left(1 + \alpha_r^2 + \frac{\alpha_\theta^2}{f^2}\right) \quad \textit{where } i \neq j, \tag{2.3.15}$$

$$K_\lambda(Y_i, Y_j) = -\frac{1}{\beta^2}\left(1 + \beta_r^2 + \frac{\beta_\theta^2}{f^2}\right) \quad \textit{where } i \neq j, \tag{2.3.16}$$

$$K_\lambda(X_i, Z_\theta) = -\frac{1}{\alpha f}\left(\alpha_r f' + \frac{\alpha_{\theta\theta}}{f}\right), \tag{2.3.17}$$

$$K_\lambda(Y_i, Z_\theta) = -\frac{1}{\beta f}\left(\beta_r f' + \frac{\beta_{\theta\theta}}{f}\right), \tag{2.3.18}$$

$$K_\lambda(X_i, Z_r) = -\frac{\alpha_{rr}}{\alpha}, \tag{2.3.19}$$

$$K_\lambda(Y_i, Z_r) = -\frac{\beta_{rr}}{\beta}, \tag{2.3.20}$$

$$K_\lambda(X_i, Y_j) = -\frac{1}{\alpha\beta}\left(\alpha_r\beta_r + \frac{\alpha_\theta\beta_\theta}{f^2}\right), \tag{2.3.21}$$

$$K_\lambda(Z_r, Z_\theta) = -\frac{f''}{f}. \tag{2.3.22}$$

For non-zero mixed terms, we only have

$$< R_\lambda(X_i, Z_r)X_i, Z_\theta >_\lambda = \frac{1}{\alpha f}\left(\frac{\alpha_\theta f'}{f} - \alpha_{r\theta}\right), \tag{2.3.23}$$
$$< R_\lambda(Y_i, Z_r)Y_i, Z_\theta >_\lambda = \frac{1}{\beta f}\left(\frac{\beta_\theta f'}{f} - \beta_{r\theta}\right).$$

Proof. The authors calculated the above formulas using two different methods and obtained the same results each way. We will describe both methods, and illustrate their use by giving details for a few of the calculations. The first method will be used to calculate one of the formulas in Equation (2.3.15), and the second method to compute one of the formulas in Equation (2.3.23).

The first method is to calculate everything directly using local coordinates and Christoffel symbols. Recall that, for a metric g and its coefficient matrix $[g_{ij}]$, the formulas for the curvature tensor with orthonormal frame $\{X_1, \ldots, X_n\}$ are

$$< R(X_i, X_j)X_k, X_s >= R_{ijks} = \sum_l R^l_{ijk} g_{ls}, \tag{2.3.24}$$

where

$$R^s_{ijk} = \sum_l \Gamma^l_{ik}\Gamma^s_{jl} - \sum_l \Gamma^l_{jk}\Gamma^s_{il} + \frac{\partial}{\partial x_j}\Gamma^s_{ik} - \frac{\partial}{\partial x_i}\Gamma^s_{jk}, \tag{2.3.25}$$

and

$$\Gamma^m_{ij} = \frac{1}{2}\sum_l \left(\frac{\partial}{\partial x_i} g_{jl} + \frac{\partial}{\partial x_j} g_{li} + \frac{\partial}{\partial x_l} g_{ij}\right) g^{lm}. \tag{2.3.26}$$

Note that, in Equation (2.3.26), we are using the standard notation $[g^{ij}] = [g_{ij}]^{-1}$.

Now, for the metric $\tilde{\lambda} = \alpha^2(r,\theta)\tilde{\mathfrak{h}}_{m-1} + \beta^2(r,\theta)\tilde{\mathfrak{h}}_{n-1} + f^2(r)d\theta^2 + dr^2$, and its orthogonal frame $\{X'_1, \ldots, X'_{m-1}, Y'_1, \ldots, Y'_{n-1}, \frac{\partial}{\partial r}, \frac{\partial}{\partial \theta}\}$, the $(m+n)\times(m+n)$-matrix of the coefficient functions for the metric λ is

$$\left[\tilde{\lambda}_{ij}\right] = \begin{bmatrix} \frac{\alpha^2}{x_{m-1}^2} & 0 & \cdots & & & & & & 0 \\ 0 & \ddots & & & & & & & \vdots \\ \vdots & & \frac{\alpha^2}{x_{m-1}^2} & & & & & & \\ & & & \frac{\beta^2}{y_{n-1}^2} & & & & & \\ & & & & \ddots & & & & \\ & & & & & \frac{\beta^2}{y_{n-1}^2} & & & \\ & & & & & & f^2 & 0 \\ 0 & & & & & & 0 & 1 \end{bmatrix} \tag{2.3.27}$$

In this computation, we label $m+n$ indices not just from 1 to $m+n$, but label $x_1, \ldots, x_{m-1}$ for the first $m-1$ indices, $y_1, \ldots, y_{n-1}$ for the next $n-1$ indices, r for the $(m+n-1)$th index, and θ for the last $(m+n)$th index. From the formula in Equation (2.3.26), we have non-zero entries of the Christoffel symbols

$$\begin{array}{llll} \Gamma^{x_i}_{x_i x_{m-1}} = -\frac{1}{x_{m-1}}, & \Gamma^{x_i}_{x_i r} = \frac{\alpha_r}{\alpha}, & \Gamma^{x_i}_{x_i \theta} = \frac{\alpha_\theta}{\alpha}, & \text{for } 1 \le i \le m-2. \\ \Gamma^{x_{m-1}}_{x_i x_i} = \frac{1}{x_{m-1}}, & \Gamma^{x_{m-1}}_{x_{m-1} r} = \frac{\alpha_r}{\alpha}, & \Gamma^{x_{m-1}}_{x_{m-1}\theta} = \frac{\alpha_\theta}{\alpha}, & \text{for } 1 \le i \le m-1. \\ \Gamma^{y_i}_{y_i y_{n-1}} = -\frac{1}{y_{n-1}}, & \Gamma^{y_i}_{y_i r} = \frac{\beta_r}{\beta}, & \Gamma^{y_i}_{y_i \theta} = \frac{\beta_\theta}{\beta}, & \text{for } 1 \le i \le n-2. \\ \Gamma^{y_{n-1}}_{y_i y_i} = \frac{1}{y_{n-1}}, & \Gamma^{y_{n-1}}_{y_{n-1} r} = \frac{b_r}{\beta}, & \Gamma^{y_{n-1}}_{y_{n-1}\theta} = \frac{\beta_\theta}{\beta}, & \text{for } 1 \le i \le n-1. \\ \Gamma^{\theta}_{x_i x_i} = -\frac{\alpha\alpha_\theta}{x_{m-1}^2 f^2}, & \Gamma^{\theta}_{y_i y_i} = -\frac{\beta\beta_\theta}{f^2 y_{n-1}^2}, & \Gamma^{\theta}_{\theta r} = \frac{f'}{f}, & \text{for all } i. \\ \Gamma^{r}_{x_i x_i} = -\frac{\alpha\alpha_r}{x_{m-1}^2}, & \Gamma^{r}_{y_i y_i} = -\frac{\beta\beta_r}{y_{n-1}^2}, & \Gamma^{r}_{\theta\theta} = -ff', & \text{for all } i. \end{array}$$

and recall that the Christoffel symbols are symmetric i.e., $\Gamma^i_{jk} = \Gamma^i_{kj}$.

By applying the formulas in Equations (2.3.24) and (2.3.25), we obtain the components for the $(4,0)$-curvature tensor. Let us illustrate this process by explicitly working out the computation for $K_\lambda(X_i, X_j)$ where $1 \le i, j \le m-1$ and $i \ne j$. The other tensors can be computed similarly. First note that

$$K_\lambda(X_i, X_j) =< R(X_i, X_j)X_i, X_j >= \frac{x_{m-1}^4}{\alpha^4} R_{x_i x_j x_i x_j} = \frac{x_{m-1}^4}{\alpha^4} \sum_l R^l_{x_i x_j x_i} \lambda_{l x_j}.$$

From the matrix in Equation (2.3.27) we have,

$$\sum_l R^l_{x_i x_j x_i} \lambda_{l x_j} = R^{x_j}_{x_i x_j x_i} \frac{\alpha^2}{x_{m-1}^2}.$$

We apply the formulas in Equations (2.3.25) and (2.3.26), and simplify as the matrix Γ^{x_i} has zero entries except $\Gamma^{x_i}_{x_i x_{m-1}}$,$\Gamma^{x_i}_{x_i r}$, and $\Gamma^{x_i}_{x_i \theta}$ up to symmetry.

$$\begin{aligned}
R^{x_j}_{x_i x_j x_i} &= \sum_l \Gamma^l_{x_i x_i} \Gamma^{x_j}_{x_j l} - \sum_l \Gamma^l_{x_j x_i} \Gamma^{x_j}_{x_i l} + \frac{\partial}{\partial x_j} \Gamma^{x_j}_{x_i x_i} - \frac{\partial}{\partial x_i} \Gamma^{x_j}_{x_j x_i} \\
&= \Gamma^{x_{m-1}}_{x_i x_i} \Gamma^{x_j}_{x_j x_{m-1}} + \Gamma^r_{x_i x_i} \Gamma^{x_j}_{x_j r} + \Gamma^\theta_{x_i x_i} \Gamma^{x_j}_{x_j \theta} \\
&= \frac{1}{x_{m-1}} \frac{-1}{x_{m-1}} - \frac{\alpha \alpha_r}{x_{m-1}^2} \frac{\alpha_r}{\alpha} - \frac{\alpha \alpha_\theta}{f^2 x_{m-1}^2} \frac{\alpha_\theta}{\alpha} \\
&= -\frac{1}{x_{m-1}^2} \left(1 + \alpha_r^2 + \frac{\alpha_\theta^2}{f^2} \right).
\end{aligned}$$

Therefore we have

$$\begin{aligned}
K_\lambda(X_i, X_j) &= \frac{x_{m-1}^4}{\alpha^4} R^{x_j}_{x_i x_j x_i} \frac{\alpha^2}{x_{m-1}^2} \\
&= -\frac{1}{\alpha^2} \left(1 + \alpha_r^2 + \frac{\alpha_\theta^2}{f^2} \right).
\end{aligned}$$

The second (equivalent) method to calculate the formulas in Theorem 2.3.6 is to calculate the Lie brackets for the orthonormal basis in Equations (2.3.11)–(2.3.14), use these to calculate the Levi-Civita connection, and then use the connection to calculate the components of the sectional curvature tensor.

The relevant Lie brackets needed for Equation (2.3.23) are

$$[X_i, Z_\theta] = \frac{\alpha_\theta}{\alpha f} X_i, \quad [X_i, Z_r] = \frac{\alpha_r}{\alpha} X_i, \quad [Z_\theta, Z_r] = -\frac{f'}{f} Z_\theta.$$

Using the formula for an orthonormal basis

$$\langle \nabla_A B, C \rangle = -\frac{1}{2} \left(\langle [B, C], A \rangle + \langle [A, C], B \rangle + \langle [B, A], C \rangle \right),$$

one calculates

$$\nabla_{X_i} X_i = -\frac{\alpha_r}{\alpha} Z_r - \frac{\alpha_\theta}{\alpha f} Z_\theta, \quad \nabla_{Z_r} X_i = 0, \quad \nabla_{Z_r} Z_r = 0, \quad \nabla_{Z_r} Z_\theta = 0.$$

Let us quickly note that no Lie brackets other than what are listed above contribute to the above formulas for the connection.

Finally, we compute

$$\begin{aligned}\langle R_\lambda(X_i, Z_r) X_i, Z_\theta\rangle_\lambda &= \langle \nabla_{Z_r} \nabla_{X_i} X_i - \nabla_{X_i} \nabla_{Z_r} X_i + \nabla_{[X_i, Z_r]} X_i, Z_\theta\rangle_\lambda \\ &= \langle \nabla_{\frac{\partial}{\partial r}} \left(-\frac{\alpha_r}{\alpha} Z_r - \frac{\alpha_\theta}{\alpha f} Z_\theta \right) - \nabla_{X_i}(0) + \frac{\alpha_r}{\alpha} \nabla_{X_i} X_i, Z_\theta\rangle_\lambda\end{aligned}$$

$$\begin{aligned}&= \langle \frac{-\alpha_{rr}\alpha + (\alpha_r)^2}{\alpha^2} Z_r - \frac{\alpha_{r\theta} f \alpha - \alpha_\theta (f'\alpha + f\alpha_r)}{\alpha^2 f^2} Z_\theta + \frac{\alpha_r}{\alpha}\left(-\frac{\alpha_r}{\alpha} Z_r - \frac{\alpha_\theta}{\alpha f} Z_\theta \right), Z_\theta\rangle_\lambda \\ &= \frac{f'}{\alpha f^2}\alpha_\theta - \frac{1}{\alpha f}\alpha_{r\theta} = \frac{1}{\alpha f}\left(\frac{\alpha_\theta f'}{f} - \alpha_{r\theta} \right).\end{aligned}$$

Lastly, note that the above calculation also gives the formula for $K_\lambda(X_i, Z_r)$.

□

Corollary 2.3.7. *Using the notation of Theorem 2.3.6, when $\alpha = \beta = 1$ the only nonzero components of the curvature tensor for $\tilde{\lambda}$ are*

$$K_\lambda(X_i, X_j) = -1, \quad K_\lambda(Y_i, Y_j) = -1, \quad K_\lambda(Z_r, Z_\theta) = -\frac{f''}{f}.$$

2.3.3 Proof of Theorem 2.1.1

To prove Theorem 2.1.1 (1), we need to warp the product hyperbolic metric near S in order to turn S into a cusp of N. We do this by defining functions α and β on $\mathbb{S}^1 \times (-\infty, \varepsilon)$ and f on $(-\infty, \varepsilon)$ in Equation (2.3.8) for some sufficiently small $\varepsilon > 0$. As the radial parameter r approaches $-\infty$, this has the geometric effect of turning S into a cusp of N.

Remark 2.3.8. By Fubini's Theorem, and since $\mathbb{S}^1$ is compact, a sufficient condition for the metric $\tilde{\lambda}$ in Equation (2.3.8) to descend to a metric on $N = M \setminus S$ with finite-volume is that all three functions α, β, and f are bounded for $r \in (-\infty, \varepsilon)$, and at least one of the functions approaches 0 exponentially as $r \to -\infty$.

Definition 2.3.9. A Riemannian metric g with curvature tensor R is *A-regular* if there exists a sequence of positive numbers $\{A_k\}$ such that, for each $k \geq 0$, the k^{th} covariant derivative of R satisfies $|\nabla^k R|_{C^0} < A_k$.

Note that when $k = 0$, this implies that the sectional curvature of g is bounded. Obviously, any metric on a compact manifold is A-regular, and the hyperbolic metric is known to be A-regular. Therefore, in our setting, we just need to show that our constructed metric is A-regular on the region about S where we turn S into a cusp.

Proof of Theorem 2.1.1 (1). We want to construct a smooth metric on $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times \mathbb{S}^1 \times (-\infty, \infty)$ which will descend to a well-defined metric on N. We use Corollary 2.3.7 and Proposition 2.3.4, together with a change in coordinates.

Let $\eta > 0$ be the normal injectivity radius of S in M and $\eta' = (1/2)\eta$. By Proposition 2.3.4 there exists a finite volume, nonpositively curved metric g_2 on N which is isometric to $\mathfrak{h}_{m-1} \times \mathfrak{h}_{n-1} \times e_2$ on the η' neighborhood of S. This metric lifts to a metric $\tilde{g}_2$ with the same properties on $\mathbb{R}^{m+n-2} \times \mathbb{S}^1 \times (0, \infty)$.

We now define a metric $\tilde{g}_1$ on $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times \mathbb{S}^1 \times (-\infty, \epsilon_1)$ where $\epsilon_1 = \frac{1}{2}\eta'$. From Equation (2.3.8), by putting $\alpha = 1$ and $\beta = 1$, we have a metric

$$\tilde{\lambda} = \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + f^2(r)d\theta^2 + dr^2. \tag{2.3.28}$$

We want to construct a function $f(r)$ such that $f, f'' > 0$, and $f \to 0$ exponentially as $r \to -\infty$. Let $\epsilon_2 = \frac{1}{2}\epsilon_1$. Then there exists $c < \min\{1, \epsilon_2\}$ and an $\epsilon > 0$ so that $\epsilon e^c = c$. Let $\frac{1}{2}c = \epsilon_3$ and define two functions $f_1(r)$ and $f_2(r)$ by

$$f_1(r) = \epsilon e^r, \quad r \in [\epsilon_3, c], \tag{2.3.29}$$

$$f_2(r) = r, \quad r \in [c, \epsilon_2]. \tag{2.3.30}$$

Notice that $f_1(c) = \epsilon e^c = c = f_2(c)$ and $f_1'(c) = \epsilon e^c = c \leq 1 = f_2'(c)$. Therefore, by Lemma 2.3.1, there exists a smooth function f on $[\epsilon_3, \epsilon_2]$ such that $f'' > 0$, $f_1(\epsilon_3) = f(\epsilon_3)$, and $f_2(\epsilon_2) = f(\epsilon_2)$.

We thus have a smooth metric $\tilde{g}_1$ defined on $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times \mathbb{S}^1 \times (-\infty, \epsilon_1)$ given by

$$\tilde{g}_1 = \begin{cases} \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + (\epsilon e^r)^2 d\theta^2 + dr^2, & r \in (-\infty, \epsilon_3] \\ \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + f^2(r) d\theta^2 + dr^2, & r \in [\varepsilon_3, \varepsilon_2] \\ \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + r^2 d\theta^2 + dr^2, & r \in [\epsilon_2, \epsilon_1] \end{cases}. \tag{2.3.31}$$

By Corollary 2.3.7, the sectional curvature of $\tilde{g}_1$ is nonpositive. Note that, on $\mathbb{R}^{m+n-2} \times \mathbb{S}^1 \times [\varepsilon_2, \varepsilon_1]$, $\tilde{g}_1 = \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + e_2$. Thus, $\tilde{g}_1 = \tilde{g}_2$ on this region. We can therefore glue these metrics together to obtain a nonpositively curved metric $\tilde{g}$ defined on the whole of $\mathbb{R}^{m+n-2} \times \mathbb{S}^1 \times (-\infty, \infty)$. Since this metric agrees with $\tilde{g}_1$ outside of the normal injectivity radius of $\tilde{S}$, it descends to a well-defined metric g on N.

Since $\alpha = \beta = 1$ and $f = \varepsilon e^r$ on $(-\infty, 0)$, by Remark 2.3.8 the metric g has finite-volume. For A-regularity, consider the metric g near the cusp. By Corollary 2.3.7 there are only three nonzero terms of the curvature tensor, all identically equal to -1. Therefore all of its derivatives are zero and hence bounded. Thus the metric g is A-regular. $\square$

2.4 Riemannian Metrics for The General Lattices

2.4.1 Riemannian Metrics on Cusped Spaces

Let $E = \mathbb{H}^{m-1} \times \mathbb{H}^{n-1} \times \mathbb{S}^1$, and consider the metric $\tilde{\lambda}$ from equation (2.3.8) defined on $E \times (0, \infty)$. Let $q \in E \times (0, \infty)$, let $\sigma \subset T_q(E \times (0, \infty))$ be a 2-plane, and consider the orthonormal frame about q given by equations (2.3.11) through (2.3.14). Choose an orthonormal basis (A, B) for σ which satisfies that A is orthogonal to Z_r, which is always possible since the orthogonal complement to Z_r has codimension one.

Choose a unit vector X_1 in the $\mathbb{H}^{m-1}$ component of $T_q(E \times (0, \infty))$ and a unit vector Y_1 in the $\mathbb{H}^{n-1}$ component of $T_q(E \times (0, \infty))$ which satisfy that X_1 is parallel to $\text{proj}_{\mathbb{H}^{m-1}}(A)$ and Y_1 is parallel to $\text{proj}_{\mathbb{H}^{n-1}}(A)$. Finally, choose corresponding vectors X_2 and Y_2 such that $\text{proj}_{\mathbb{H}^{m-1}}(\sigma) \subseteq \text{span}(X_1, X_2)$ and $\text{proj}_{\mathbb{H}^{n-1}}(\sigma) \subseteq \text{span}(Y_1, Y_2)$. Then there exist constants $a_1, a_3, a_5, b_1, b_2, \ldots, b_6$ such that

$$A = a_1 X_1 + a_3 Y_1 + a_5 Z_\theta \quad \text{and} \quad B = b_1 X_1 + b_2 X_2 + b_3 Y_1 + b_4 Y_2 + b_5 Z_\theta + b_6 Z_r.$$

Note that, since A and B are orthonormal, we have

$$a_1^2 + a_3^2 + a_5^2 = 1 \qquad b_1^2 + \ldots + b_6^2 = 1 \qquad a_1 b_1 + a_3 b_3 + a_5 b_5 = 0.$$

In the sequel we use the following notation. Let

$$\Omega_1 = X_1 \quad \Omega_2 = X_2 \quad \Omega_3 = Y_1 \quad \Omega_4 = Y_2 \quad \Omega_5 = Z_\theta \quad \Omega_6 = Z_r$$

and

$$R_{ijk\ell} = \langle R(\Omega_i, \Omega_j)\Omega_k, \Omega_\ell \rangle_\lambda.$$

We then compute

$$K_\lambda(\sigma) = \langle R(A,B)A,B\rangle_\lambda = a_1^2b_2^2R_{1212} + (a_1b_3 - a_3b_1)^2R_{1313} + a_1^2b_4^2R_{1414}$$

$$+ a_3^2b_2^2R_{2323} + a_5^2b_2^2R_{2525} + a_3^2b_4^2R_{3434} + a_5^2b_4^2R_{4545} + a_5^2b_6^2R_{5656}$$

$$+ (a_1b_5 - a_5b_1)^2R_{1515} + 2a_1b_6(a_1b_5 - a_5b_1)R_{1516} + a_1^2b_6^2R_{1616} \tag{2.4.1}$$

$$+ (a_3b_5 - a_5b_3)^2R_{3535} + 2a_3b_6(a_3b_5 - a_5b_3)R_{3536} + a_3^2b_6^2R_{3636} \tag{2.4.2}$$

Lines (2.4.1) and (2.4.2) can be rewritten as

$$\begin{bmatrix} a_1b_5 - a_5b_1 & a_1b_6 \end{bmatrix} \begin{bmatrix} R_{1515} & R_{1516} \\ R_{1516} & R_{1616} \end{bmatrix} \begin{bmatrix} a_1b_5 - a_5b_1 \\ a_1b_6 \end{bmatrix} \tag{2.4.3}$$

and

$$\begin{bmatrix} a_3b_5 - a_5b_3 & a_3b_6 \end{bmatrix} \begin{bmatrix} R_{3535} & R_{3536} \\ R_{3536} & R_{3636} \end{bmatrix} \begin{bmatrix} a_3b_5 - a_5b_3 \\ a_3b_6 \end{bmatrix}. \tag{2.4.4}$$

With these calculations we are ready to prove the following proposition, which gives a relatively simple criterion to determine exactly when $\tilde{\lambda}$ is nonpositively curved.

Proposition 2.4.1. *The metric $\tilde{\lambda}$ in* (2.3.8) *is nonpositively curved if and only if both of the following two conditions hold:*

1. *All of the curvatures in equations* (2.3.15) *through* (2.3.22) *are nonpositive.*

2. *All matrices of the form*

$$\begin{bmatrix} K_\lambda(X_i, Z_\theta) & < R_\lambda(X_i, Z_r)X_i, Z_\theta >_\lambda \\ < R_\lambda(X_i, Z_r)X_i, Z_\theta >_\lambda & K_\lambda(X_i, Z_r) \end{bmatrix} \tag{2.4.5}$$

and

$$\begin{bmatrix} K_\lambda(Y_i, Z_\theta) & < R_\lambda(Y_i, Z_r)Y_i, Z_\theta >_\lambda \\ < R_\lambda(Y_i, Z_r)Y_i, Z_\theta >_\lambda & K_\lambda(Y_i, Z_r) \end{bmatrix} \tag{2.4.6}$$

are negative semi-definite.

Proof. First, notice that the matrices in (2.4.5) and (2.4.6) are of the exact same form as the matrices in equations (2.4.3) and (2.4.4).

To prove the proposition, we first assume that conditions (1) and (2) hold. Then, by (2), lines (2.4.1) and (2.4.2) in the calculation for $\langle R(A,B)A,B\rangle_\lambda$ are nonpositive. By (1), the first two lines are also nonpositive. Therefore, the metric $\tilde{\lambda}$ must have nonpositive sectional curvature.

So let us now assume that all sectional curvatures of $\tilde{\lambda}$ are nonpositive. This clearly implies statement (1). To verify statement (2), we consider specific 2-planes σ which correspond to choosing specific coefficients $a_1, a_3, a_5, b_1, b_2, \ldots b_6$. In particular, we begin by specifying that

$$a_3 = a_5 = b_2 = b_3 = b_4 = 0.$$

The only nonzero line in the calculation of $\langle R(A,B)A,B\rangle_\lambda$ with this choice of coefficients is (2.4.1). Note that $a_3 = a_5 = 0$ forces $a_1 = \pm 1$. The vector

$$\begin{bmatrix} a_1 b_5 \\ a_1 b_6 \end{bmatrix} = \pm \begin{bmatrix} b_5 \\ b_6 \end{bmatrix}$$

can obtain any direction in $\mathbb{RP}^1$ for different choices of b_5 and b_6. So the only way to ensure that (2.4.1) is nonpositive for all possible selections of these variables is if the matrix

$$\begin{bmatrix} R_{1515} & R_{1516} \\ R_{1516} & R_{1616} \end{bmatrix}$$

is negative semi-definite. An identical analysis shows that we also need the matrix

$$\begin{bmatrix} R_{3535} & R_{3536} \\ R_{3536} & R_{3636} \end{bmatrix}$$

to be negative semi-definite, verifying the necessity of condition (2). □

Remark 2.4.2. Condition (1) in Proposition 2.4.1 forces the matrices in (2.4.5) and (2.4.6) to each have at least one nonpositive eigenvalue. If one of $K_\lambda(X_i, Z_\theta)$ or $K_\lambda(X_i, Z_r)$ are negative (as opposed to nonpositive), then the matrices of the form (2.4.5) must have at least one negative eigenvalue (and similarly for $K_\lambda(Y_i, Z_\theta)$, $K_\lambda(Y_i, Z_r)$, and (2.4.6)). In this case, matrices of the form (2.4.5) are negative semi-definite if and only if their determinant is nonnegative. So, if $K_\lambda(X_i, Z_\theta) < 0$ or $K_\lambda(X_i, Z_r) < 0$, and condition (1) in Proposition 2.4.1 holds, then $\tilde{\lambda}$ is nonpositively curved if and only the matrices in (2.4.5) and (2.4.6) have nonnegative determinant.

The proof of Proposition 2.4.1 combined with Remark 2.4.2 gives the following Corollary.

Corollary 2.4.3. *Let $\varepsilon > 0$, and suppose that either $K_\lambda(X_i, Z_\theta)$ or $K_\lambda(X_i, Z_r)$ are bounded above by a negative constant that is independent of ε. Also, assume that condition* (1) *in Proposition 2.4.1 is satisfied. Then there exists $\zeta(\varepsilon) > 0$ such that, if all matrices of the form* (2.4.5) *and* (2.4.6) *have determinant greater than $-\zeta$, then the sectional curvature of $\tilde{\lambda}$ is less than ε.*

We are now ready to prove Theorem 2.1.6 (1).

Proof of Theorem 2.1.6 (1). We again let $E = \mathbb{H}^{m-1} \times \mathbb{H}^{n-1} \times \mathbb{S}^1$, and let $\varepsilon > 0$ be given. Let δ denote the normal injectivity radius of S in M. We first define a piecewise-smooth metric $\tilde{g}'$ by prescribing specific values for the functions α, β, and f in equation (2.3.8) on different portions of $E \times \mathbb{R}$. We will show that, for a set of parameters all chosen sufficiently small, the sectional curvatures of $\tilde{g}'$ at all points where it is smooth are bounded above by $\varepsilon/2$. We will then argue that we can C^2-approximate $\tilde{g}'$ by a smooth metric $\tilde{g}$ whose sectional curvatures are as close to those of $\tilde{g}'$ as we like. The metric $\tilde{g}$ will agree with $\tilde{\mathfrak{h}}_m \times \tilde{\mathfrak{h}}_n$ outside of a sufficiently small tubular neighborhood of $\mathbb{H}^{m-1} \times \mathbb{H}^{n-1}$, and so will thus descend to a well-defined metric g on N. An easy argument will show that we can choose all parameters sufficiently small so that $\mathrm{vol}(N, g) < \mathrm{vol}(M, \mathfrak{h}_m \times \mathfrak{h}_n) + \xi$ for any prescribed $\xi > 0$.

The metric $\tilde{g}'$ will depend on a set of small positive parameters η, η_1, η_2, and η_3. The value of η_1 will be the value of r at which the functions $f_1(r) = \eta e^r$ and $f_2(r) = r$ intersect (which is at approximately $r = \eta$). For now, we will assume that all other parameters are given positive numbers that satisfy $\eta_1 < \eta_2 < \eta_3 < \delta$. At the end of the proof we will discuss the order in which we choose these parameters. Also note that, for $r > 0$ small, we have the approximations

$$\cosh(r\cos\theta) \approx 1 + \frac{1}{2}r^2\cos^2\theta$$
$$\cosh(r\sin\theta) \approx 1 + \frac{1}{2}r^2\sin^2\theta.$$

Since all of our calculations below occur for $r \approx 0$, and the condition $K_\lambda < \varepsilon$ is an open condition, we will use these approximations for the remainder of the proof to simplify calculations.

We define the metric $\tilde{g}'$ via (2.3.8) by

1. $f = \eta e^r$ and $\alpha = \beta = 1$ on $E \times (-\infty, \eta_1)$.

2. $f = r$ and $\alpha = \beta = 1$ on $E \times (\eta_1, \eta_2)$.

3. On $E \times (\eta_2, \eta_3)$ we define $f = r$ and

$$\alpha(r,\theta) = 1 + \frac{\eta_3^2}{2(\eta_3 - \eta_2)^2}(r - \eta_2)^2\cos^2\theta$$
$$\beta(r,\theta) = 1 + \frac{\eta_3^2}{2(\eta_3 - \eta_2)^2}(r - \eta_2)^2\sin^2\theta$$

4. On $E \times (\eta_3, \infty)$ we have $f(r) = r$ and

$$\begin{aligned}\alpha(r,\theta) &= 1 + \frac{1}{2} r^2 \cos^2\theta \\ \beta(r,\theta) &= 1 + \frac{1}{2} r^2 \sin^2\theta.\end{aligned}$$

We quickly note that the metric $\tilde{g}'$ is not only piecewise-smooth, but it is also C^0 and is C^1 at $E \times \{\eta_2\}$.

It is a straightforward calculation using Proposition 2.4.1 to show that $\tilde{g}'$ has nonpositive curvature over $E \times (-\infty, \eta_1)$ and $E \times (\eta_1, \eta_2)$. Of course, $\tilde{g}'$ is the product hyperbolic metric on $E \times (\eta_3, \infty)$ and thus has nonpositive curvature over this region. Our real task is to show that, for η_2 chosen sufficiently small[2], $K_{g'} < \varepsilon$ on $E \times (\eta_2, \eta_3)$.

Of the sectional curvatures in equations (2.3.15) through (2.3.22), it is clear that all of these are nonpositive except for $K_{g'}(X_i, Z_\theta)$ and $K_{g'}(Y_i, Z_\theta)$. For $K_{g'}(X_i, Z_\theta)$, we have

$$\begin{aligned}K_{g'}(X_i, Z_\theta) &= -\frac{1}{f\alpha}\left(f'\alpha_r + \frac{\alpha_{\theta\theta}}{f^2}\right) \\ &= -\frac{1}{r\alpha}\left(\frac{\eta_3^2}{(\eta_3-\eta_2)^2}(r-\eta_2)\cos^2\theta + \frac{\eta_3^2}{r(\eta_3-\eta_2)^2}(r-\eta_2)^2(\sin^2\theta - \cos^2\theta)\right) \\ &= \frac{-\eta_3^2}{r\alpha(\eta_3-\eta_2)^2}(r-\eta_2)\left[\left(1 - \frac{r-\eta_2}{r}\right)\cos^2\theta + \frac{r-\eta_2}{r}\sin^2\theta\right] \\ &= \frac{-\eta_3^2}{r^2\alpha(\eta_3-\eta_2)^2}(r-\eta_2)\left[\eta_2\cos^2\theta + (r-\eta_2)\sin^2\theta\right]\end{aligned}$$

which is clearly nonpositive for all values of θ. The calculation for $K_{g'}(Y_i, Z_\theta)$ is identical.

We now move our attention to the matrix in (2.4.5), again with the calculations for (2.4.6) being identical. We first need the intermediate calculations

$$K_{g'}(X_i, Z_r) = -\frac{\alpha_{rr}}{\alpha} = \frac{-\eta_3^2}{(\eta_3-\eta_2)^2\alpha}\cos^2\theta$$

[2]The calculation turns out to not depend on the size of η_3.

and

$$\langle R(X_i, Z_r)X_i, Z_\theta\rangle = \frac{1}{f\alpha}\left(\frac{f'\alpha_\theta}{f} - \alpha_{r\theta}\right)$$
$$= \frac{1}{r\alpha}\left[\frac{-\eta_3^2}{r(\eta_3-\eta_2)^2}(r-\eta_2)^2\sin\theta\cos\theta + \frac{2\eta_3^2}{(\eta_3-\eta_2)^2}(r-\eta_2)\sin\theta\cos\theta\right]$$
$$= \frac{\eta_3^2}{r\alpha(\eta_3-\eta_2)^2}(r-\eta_2)\left(1+\frac{\eta_2}{r}\right)\sin\theta\cos\theta.$$

Letting

$$Q = \begin{bmatrix} K_\lambda(X_i, Z_\theta) & < R_\lambda(X_i, Z_r)X_i, Z_\theta >_\lambda \\ < R_\lambda(X_i, Z_r)X_i, Z_\theta >_\lambda & K_\lambda(X_i, Z_r) \end{bmatrix}$$

we have

$$\mathrm{Det}(Q) = \frac{\eta_3^4(r-\eta_2)}{r^2\alpha^2(\eta_3-\eta_2)^4}\cos^2\theta\left(\eta_2\cos^2\theta + (r-\eta_2)\sin^2\theta\right)$$
$$- \frac{\eta_3^4(r-\eta_2)^2}{r^2\alpha^2(\eta_3-\eta_2)^4}\left(1+\frac{\eta_2}{r}\right)^2\sin^2\theta\cos^2\theta$$
$$= \frac{\eta_3^4(r-\eta_2)}{r^2\alpha^2(\eta_3-\eta_2)^4}\cos^2\theta\left[\eta_2\cos^2\theta + (r-\eta_2)\left(1-\left(1+\frac{\eta_2}{r}\right)^2\right)\sin^2\theta\right]. \qquad (2.4.7)$$

First note that $\mathrm{Det}(Q) = 0$ when $r = \eta_2$. Also, when $\sin\theta = 0$ it is clear that (2.4.7) is nonnegative. The reason why $\tilde{g}'$ does not have nonpositive curvature over $E \times (\eta_2, \eta_3)$ is because the coefficient of $\sin^2\theta$ in (2.4.7) will generally be negative. But notice that, as η_2 approaches 0, $\mathrm{Det}(Q)$ approaches 0 for all values of r and θ. So, for any prescribed value of ζ, we can choose η_2 sufficiently small so that $\mathrm{Det}(Q) > -\zeta$ for all values of $r \in (\eta_2, \eta_3)$ and $\theta \in \mathbb{S}^1$.

Also notice that, as η_2 approaches 0, we have

$$K_{g'}(X_i, Z_\theta) \to -\sin^2\theta \quad \text{and} \quad K_{g'}(X_i, Z_r) \to -\cos^2\theta.$$

Thus, at least one of these values is always nonzero. An identical analysis applies for the matrix in equation (2.4.6). Therefore, by Corollary 2.4.3, we can choose η_2 sufficiently small so that $K_{g'}$ is bounded above by $\varepsilon/2$ over $E \times (\eta_2, \eta_3)$.

We can use Lemma 2.3.1 to smooth $\tilde{g}'$ near $E \times \{\eta_1\}$ in such a way that our smooth metric $\tilde{g}$ will have nonpositive curvature in this region. At $E \times \{\eta_2\}$ the function $\tilde{g}'$ is C^1, and the only obstruction to being C^2 is α_{rr}. One sees immediately from the definition of $\tilde{g}'$ that

$$\lim_{r \to \eta_2^-} \alpha_{rr} = 0$$
$$\lim_{r \to \eta_2^+} \alpha_{rr} = \frac{\eta_3^2}{(\eta_3 - \eta_2)^2} \cos^2 \theta.$$

We can now apply Lemma 2.3.1 to α_{rr} in a small neighborhood of η_2 (and keeping θ fixed), and integrate to obtain an appropriate smoothing of α.

To analyze $E \times \{\eta_3\}$, let

$$\alpha_1(r, \theta) = 1 + \frac{\eta_3^2}{2(\eta_3 - \eta_2)^2}(r - \eta_2)^2 \cos^2 \theta$$

and

$$\alpha_2(r, \theta) = 1 + \frac{1}{2} r^2 \cos^2 \theta.$$

Notice that, as $\eta_2 \to 0$, α_1 smoothly approaches α_2. So, for η_2 chosen sufficiently small, we can find $\mu < (1/4)(\eta_2 - \eta_1)$ and a function $\alpha : \mathbb{S}^1 \times [\eta_2 - \mu, \eta_2 + \mu] \to \mathbb{R}$ such that

- $\alpha = \alpha_1$ on a small neighborhood of $\mathbb{S}^1 \times \{\eta_2 - \mu\}$.
- $\alpha = \alpha_2$ on a small neighborhood of $\mathbb{S}^1 \times \{\eta_2 + \mu\}$.
- $|\alpha_2 - \alpha|_{C^2} < \varepsilon'$ for any prescribed $\varepsilon' > 0$.

Hence, we can smooth $\tilde{g}'$ to obtain a smooth metric $\tilde{g}$ which satisfies $K_g < \varepsilon$.

Let us quickly discuss is how we choose our parameters. We begin by choosing η_3 less than δ. We then choose η_2 sufficiently small so that $K_{g'} < \varepsilon/2$ on $E \times (\eta_2, \eta_3)$ as discussed above, and so that curvatures change by at most $\varepsilon/4$ when smoothing $\tilde{g}'$ to obtain $\tilde{g}$ near $E\times\{\eta_3\}$. Finally, we choose η sufficiently small so that $\eta_1 < (1/2)\eta_2$. We still need to smooth g' in small intervals about η_1 and η_2. We perform all smoothing in intervals which are less than 1/4 of the distance to the next closest parameter. In this way there will be no overlap on the intervals where we smooth $\tilde{g}'$. Finally, we decrease the values for η_2 and η, if necessary, to satisfy the volume conditions below.

The last thing that we need to consider is the volume of (N, g). Suppose the submanifold S has k components, and let S' be one of those components. Since $\alpha = \beta = 1$ for r in $(-\infty, 0)$, by Fubini's Theorem we have that the volume of this portion of the cusp corresponding to S' is

$$2\pi\mathrm{vol}(S')\int_{-\infty}^{0}\eta e^r\,dr = 2\pi\eta\mathrm{vol}(S').$$

We can choose $\eta > 0$ sufficiently small so that this quantity is less than $\xi/(2k)$ for any prescribed $\xi > 0$. Also, the functions α, β, and f are bounded, independent of the parameters, on $(0, \eta_3)$. So by a direct inequality we can bound the volume of N over this region by $\xi/(2k)$ by choosing η_3 sufficiently small. Thus, the metric g has volume at most ξ/k within the η_3-neighborhood of S' (measured with respect to $\mathfrak{h}_m \times \mathfrak{h}_n$), and therefore $\mathrm{vol}(N, g) < \mathrm{vol}(M, \mathfrak{h}_m \times \mathfrak{h}_m) + \xi$.

□

2.4.2 Riemannian Metrics on Branched Coverings

To prove Theorem 2.1.6 (2), we need the following result from [15].

Theorem 2.4.4 ([15] Section 2)**.** *Let (M, g_M) be a nonpositively curved manifold, (S, g_S) a totally geodesic, codimension two submanifold of M, and suppose that, on some r-tube about S, we have that $g_M = g_S + e_2$. Let X_d denote the d-fold cyclic* *branched cover of M about S, and $\phi_d : X_d \to M$ the associated ramified covering map. Then, given any $\delta_1, \delta_2 > 0$ with $0 < \delta_1 < \delta_2 < r$, there exists a smooth, nonpositively curved metric g on X_d such that*

1. *$g = g_M = g_S + e_2$ on the δ_1 neighborhood of S, and*

2. *$g = \phi_d^* g_M$ outside of the δ_2 neighborhood of S. In particular, on the (δ_2, r)-annulus of S, we have $g = g_S + \phi_d^*(e_2)$.*

Note that, in Theorem 2.4.4, we abuse notation and also use S to refer to the ramification locus of X_d.

Remark 2.4.5. Theorem 2.4.4 is essential in Fornari and Schroeder's proof of Theorem 2.1.1 (2) in [15]. But note that, by Remark 2.3.5, Proposition 2.3.2 can only be used to construct a metric on $\mathbb{H}^m \times \mathbb{H}^n$ which will descend to M and satisfies $g_M = g_S + e_2$ whenever Γ splits as a product. This is why the argument in [15] only applies to product lattices.

We can now use Theorem 2.4.4 and a very similar argument as in the proof of Theorem 2.1.6 (1) to prove Theorem 2.1.6 (2).

Proof of Theorem 2.1.6 (2). Let $E = \mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times \mathbb{S}^1$, let $\varepsilon > 0$ be given, and let δ denote the normal injectivity radius of S in M (which is equal to the normal injectivity radius of S in X_d). Let $\tilde{X}_d$ be the universal cover of X_d, which is diffeomorphic to $\mathbb{R}^{m+n} \cong E \times [0, \infty)$, and let $\phi : \tilde{X}_d \to X_d$ denote the associated covering map. Define ϕ_d as in Theorem 2.4.4, and let $\varphi_d = \phi_d \circ \phi$.

Choose δ_1 and δ_2 with $0 < \delta_1 < \delta_2 < (1/2)\delta$, and consider the metric $\tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + \tilde{e}_2$ on $\tilde{X}_d$. This metric satisfies the "local product" assumption in Theorem 2.4.4. So there exists a smooth, nonpositively curved metric $\tilde{g}_1$ on $\tilde{X}_d$ which satisfies

- $\tilde{g}_1 = \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + \tilde{e}_2$ on $E \times [0, \delta_1)$.
- $\tilde{g}_1 = \tilde{\mathfrak{h}}_{m-1} + \tilde{\mathfrak{h}}_{n-1} + d^2\tilde{e}_2$ on $E \times (\delta_2, \infty)$.

Take $\mathbb{S}^1$ and subdivide it into d equidistant subintervals $A_1, \dots, A_d$. The metric $\tilde{g}_1$ restricted to the manifold $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times A_i \times (\delta_2, \infty)$ is equal to $\mathfrak{h}_{m-1} + \mathfrak{h}_{n-1} + e_2$. In what follows we restrict our calculations to $Y_i = \mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times A_i \times (\delta_2, \infty)$, and we perform the same procedures on Y_i for each i so that our metrics glue together to give a well-defined metric on $\tilde{X}_d$.

Choose $\eta_2, \eta_3 > 0$ such that $\delta_2 < \eta_2 < \eta_3 < \delta$. We use equation (2.3.8) to define a metric $\tilde{g}_2'$ on Y_i as follows:

1. $f = r$ and $\alpha = \beta = 1$ on $E \times (\delta_2, \eta_2)$.

2. On $E \times (\eta_2, \eta_3)$ we define $f = r$ and

$$\alpha(r, \theta) = 1 + \frac{\eta_3^2}{2(\eta_3 - \eta_2)^2}(r - \eta_2)^2 \cos^2\theta$$
$$\beta(r, \theta) = 1 + \frac{\eta_3^2}{2(\eta_3 - \eta_2)^2}(r - \eta_2)^2 \sin^2\theta$$

3. On $E \times (\eta_3, \infty)$ we have $f(r) = r$ and

$$\alpha(r, \theta) = 1 + \frac{1}{2}r^2 \cos^2\theta$$
$$\beta(r, \theta) = 1 + \frac{1}{2}r^2 \sin^2\theta.$$

The definition of $\tilde{g}_2'$ above is identical to the last three steps in the definition of $\tilde{g}'$ in the proof of Theorem 2.1.6 (1). Via this same argument, we can smooth $\tilde{g}_2'$ to obtain a metric $\tilde{g}_2$ on Y_i with all curvatures less than ε. This metric $\tilde{g}_2$ descends to a metric g_2 on the complement of the δ_2 neighborhood of S in M.

There exists $\mu > 0$ such that, for each i, the metrics $\tilde{g}_1$ and $\tilde{g}_2$ agree on $\mathbb{R}^{m-1} \times \mathbb{R}^{n-1} \times A_i \times (\delta_2, \delta_2 + \mu)$. We then define the metric $\tilde{g}$ on $\tilde{X}_d$ by

$$\tilde{g} = \begin{cases} \tilde{g}_1 & \text{on } E \times [0, \delta_2] \\ \varphi_d^*(g_2) & \text{on } E \times [\delta_2, \infty) \end{cases}. \tag{2.4.8}$$

Via the discussion above, $\tilde{g}$ is well-defined on the overlap of these regions and thus defines a smooth metric on $\tilde{X}_d$ whose curvatures are bounded above by ε. Since this metric agrees with $\varphi_d^*(\mathfrak{h}_m + \mathfrak{h}_n)$ outside of the δ-neighborhood of S, it descends to a well-defined metric on X_d. An identical argument to the proof of Theorem 2.1.6 (1) shows that, by choosing δ_2 sufficiently small, we can bound the volume of (X_d, g) by $d \cdot \mathrm{vol}(M, \mathfrak{h}_m \times \mathfrak{h}_n) + \xi$ for any $\xi > 0$. □

2.5 Future Works: Topology of (M, S)

The group $\pi_1(N)$ will not be relatively hyperbolic relative to the fundamental groups of its ends. Therefore we cannot directly conclude many standard rigidity results for $\pi_1(N)$, for instance those obtained in [5], [4], and [6]. But when Γ splits as a product, Theorem 2.1.1 allows us to conclude statements about the topological rigidity of $\pi_1(N)$. In particular, by combining Theorem 2.1.1 with results from Farrell–Jones [12] and Lafforgue [29], we conclude the following.

Corollary 2.5.1. *Suppose (M, S) is as in Theorem 2.1.1. If $\bar{N}$ is a compact aspherical manifold with $\pi_1(\bar{N}) \cong \pi_1(N)$, then $\bar{N}$ satisfies Borel's conjecture provided* $\dim(\bar{N}) \geq 5$. *Moreover, if $\pi_1(N)$ satisfies Lafforgue's Rapid Decay property, then it satisfies the Baum-Connes conjecture.*

As an immediate consequence of Corollary 2.1.4, we have that N virtually satisfies the conclusions of Corollary 2.5.1 when Γ is reducible. Using "graph of groups" style arguments, we can prove the following results about $\pi_1(N)$.

Bibliography

[1] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, fourth edition, 2016.

[2] Grigori Avramidi and T. Tâm Nguyễn-Phan. Examples of noncompact nonpositively curved manifolds. *Bull. Lond. Math. Soc.*, 53(4):1061–1071, 2021.

[3] C. Bavard. Courbure presque négative en dimension 3. (Almost negative curvature in dimension 3). *Compos. Math.*, 63:223–236, 1987.

[4] Igor Belegradek. Complex hyperbolic hyperplane complements. *Math. Ann.*, 353(2):545–579, 2012.

[5] Igor Belegradek. Rigidity and relative hyperbolicity of real hyperbolic hyperplane complements. *Pure Appl. Math. Q.*, 8(1):15–51, 2012.

[6] Igor Belegradek and G. Christopher Hruska. Hyperplane arrangements in negatively curved manifolds and relative hyperbolicity. *Groups Geom. Dyn.*, 7(1):13–38, 2013.

[7] Armand Borel and Harish-Chandra. Arithmetic subgroups of algebraic groups. *Ann. of Math. (2)*, 75:485–535, 1962.

[8] Robert Brooks and Eran Makover. Random construction of Riemann surfaces. *J. Differential Geom.*, 68(1):121–157, 2004.

[9] Shiing-Shen Chern. On curvature and characteristic classes of a Riemann manifold. *Abh. Math. Semin. Univ. Hamb.*, 20:117–126, 1955.

[10] John D Dixon. The probability of generating the symmetric group. *Mathematische Zeitschrift*, 110:199–205, 1969.

[11] Nathan M. Dunfield and William P. Thurston. Finite covers of random 3-manifolds. *Invent. Math.*, 166(3):457–521, 2006.

[12] F. T. Farrell and L. E. Jones. Rigidity for aspherical manifolds with $\pi_1 \subset \mathrm{GL}_m(\mathbf{R})$. *Asian J. Math.*, 2(2):215–262, 1998.

[13] F. T. Farrell and P. Ontaneda. The space of nonpositively curved metrics of a negatively curved manifold. *J. Differ. Geom.*, 99(2):285–311, 2015.

[14] Joel Fine and Bruno Premoselli. Examples of compact Einstein four-manifolds with negative curvature. *J. Amer. Math. Soc.*, 33(4):991–1038, 2020.

[15] Susana Fornari and Viktor Schroeder. Ramified coverings with nonpositive curvature. *Math. Z.*, 203(1):123–128, 1990.

[16] Robert E Gompf and András I Stipsicz. *4-manifolds and Kirby calculus*, volume 20. American Mathematical Society, 2023.

[17] M. Gromov. Almost flat manifolds. *J. Differ. Geom.*, 13:231–241, 1978.

[18] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[19] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[20] M. Gromov and W. Thurston. Pinching constants for hyperbolic manifolds. *Invent. Math.*, 89(1):1–12, 1987.

[21] M. Gromov and W. Thurston. Pinching constants for hyperbolic manifolds. In *Collected works of William P. Thurston with commentary. Vol. I. Foliations, surfaces and differential geometry*, pages 655–666. Amer. Math. Soc., Providence, RI, [2022] ©2022. Reprint of [0892185].

[22] Mikhael Gromov. Hyperbolic groups. In *Essays in group theory*, pages 75–263. Springer, 1987.

[23] Mikhail Gromov. Random walk in random groups. *Geometric and Functional Analysis*, 13:73–146, 2003.

[24] Henri Guenancia and Ursala Hamenstädt. Kähler-Einstein metrics of negative curvature. arXiv:2503.02838, 2025.

[25] Vincent Guirardel. Geometric small cancellation. In *Graduate Summer School*, volume 21, pages 55–90. American Mathematical Society, 2012.

[26] Liam Hanany and Doron Puder. Word measures on symmetric groups. *International Mathematics Research Notices*, 2023(11):9221–9297, 2023.

[27] Hugh M. Hilden. Every closed orientable 3-manifold is a 3-fold branched covering space of S^3. *Bull. Amer. Math. Soc.*, 80:1243–1244, 1974.

[28] Matthew Kahle. Topology of random simplicial complexes: a survey. In *Algebraic topology: applications and new directions*, volume 620 of *Contemp. Math.*, pages 201–221. Amer. Math. Soc., Providence, RI, 2014.

[29] V. Lafforgue. K-théorie bivariante pour les algĕbres de banach et conjecture de baum-connes. *Invent. Math.*, 149(1):1–95, 2002.

[30] RC Lyndon. Combinatorial group theory, 1977.

[31] G. A. Margulis. *Discrete subgroups of semisimple Lie groups*, volume 17 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991.

[32] Barry Minemyer. Real hyperbolic hyperplane complements in the complex hyperbolic plane. *Adv. Math.*, 338:1038–1076, 2018.

[33] Barry Minemyer. Kähler manifolds with an almost 1/4-pinched metric. arXiv:2307.15550, 2023.

[34] Barry Minemyer. Kähler manifolds with an almost 1/4-pinched metric. *preprint, arXiv: 2307.15550v2*, 2023.

[35] José M. Montesinos. A representation of closed orientable 3-manifolds as 3-fold branched coverings of S^3. *Bull. Amer. Math. Soc.*, 80:845–846, 1974.

[36] Dave Witte Morris. *Introduction to arithmetic groups*. Deductive Press, [place of publication not identified], 2015.

[37] Yann Ollivier. *A January 2005 invitation to random groups*, volume 10 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro, 2005.

[38] M. S. Raghunathan. *Discrete subgroups of Lie groups*. Ergebnisse der Mathematik und ihrer Grenzgebiete [Results in Mathematics and Related Areas], Band 68. Springer-Verlag, New York-Heidelberg, 1972.

[39] Atle Selberg. On discontinuous groups in higher-dimensional symmetric spaces. In *Contributions to function theory (Internat. Colloq. Function Theory, Bombay, 1960)*, pages 147–164. Tata Inst. Fund. Res., Bombay, 1960.

[40] Stephan Stadler. An obstruction to the smoothability of singular nonpositively curved metrics on 4-manifolds by patterns of incompressible tori. *Geom. Funct. Anal.*, 25(5):1575–1587, 2015.

[41] Matthew Stover and Domingo Toledo. Residual finiteness for central extensions of lattices in $\mathrm{PU}(n, 1)$ and negatively curved projective varieties. *Pure Appl. Math. Q.*, 18(4):1771–1797, 2022.

[42] Matthew Stover and Domingo Toledo. Residual finiteness for central extensions of lattices in $\mathrm{PU}(n, 1)$ and negatively curved projective varieties. *Pure Appl. Math. Q.*, 18(4):1771–1797, 2022.

[43] Fangyang Zheng. Examples of non-positively curved Kähler manifolds. *Comm. Anal. Geom.*, 4(1-2):129–160, 1996.