HOMOMORPHISMS TO 3–MANIFOLD GROUPS

DANIEL GROVES, MICHAEL HULL, AND HAO LIANG

ABSTRACT. We prove foundational results about the set of homomorphisms from a finitely generated group to the collection of all fundamental groups of compact 3–manifolds and answer questions of Agol–Liu [2] and Reid–Wang–Zhou [34].

CONTENTS

1. INTRODUCTION

There is a well-developed structure theory for compact 3–manifolds based on the Geometrization Theorem of Perelman (conjectured by Thurston). In this paper we develop a structure theory for the collection of all maps between 3–manifolds, from the point of view of their fundamental groups. Let $\mathscr{M}^{\pi}$ be the set of all fundamental groups of compact 3–manifolds (see [3] for background). Our main result gives a positive answer to a question of Agol–Liu [2, Question 10.3].

t:main **Theorem A.** *For any finitely generated group G there is a finitely presented group G_0 and a surjection $\alpha\colon G_0 \twoheadrightarrow G$ so that for every $\Gamma \in \mathscr{M}^{\pi}$ the map*

$$\alpha^*\colon \mathrm{Hom}(G,\Gamma) \to \mathrm{Hom}(G_0,\Gamma)$$

induced by precomposition with α is a bijection.

Of course, if G is finitely presented, then one can take $G_0 = G$, and there is nothing to prove. Since finitely generated 3–manifold groups are finitely presented [36], it might appear that Theorem A is not relevant to the study of homomorphisms between 3–manifold

Date: October 21, 2021.

The first author is supported in part by NSF grant DMS-1904913. The third author is supported by NSFC (No.11701581 and No.11521101).

groups. On the contrary, Theorem A is an important tool, and in particular implies the following theorem, answering a question of Reid–Wang–Zhou [34, Question 3.1.(C2)].

t:DCC **Theorem B.** *Let M_i be compact 3–manifolds (possibly with boundary). Every infinite sequence of epimorphisms $\pi_1(M_1) \twoheadrightarrow \cdots \twoheadrightarrow \pi_1(M_n) \twoheadrightarrow \cdots$ contains an isomorphism.*

It is worth remarking that Reid, Wang and Zhou ask their question about closed, orientable, aspherical 3–manifolds, whereas the only assumption we need is that they are compact. Theorem B proves the descending chain condition for the following partial order on 3–manifold groups: $G_1 \geq G_2$ if there is an epimorphism $\phi\colon G_1 \to G_2$. Per Calegari [11], understanding this order is an important question in 3–manifold topology.

Theorem B is proved in Section 2 as a consequence of Theorem A. The proof of Theorem B is by contradiction, and the group G from Theorem A used in this proof arises as a direct limit of the sequence of maps $\pi_1(M_i) \twoheadrightarrow \pi_1(M_{i+1})$.

Partial results about Reid, Wang and Zhou's question were already known. Reid–Wang–Zhou [34, Theorem 3.4] gave a positive answer to this question in case all M_i are (closed, orientable, aspherical) Seifert 3-manifolds. A key step of their proof is showing that epimorphisms between fundamental groups of aspherical Seifert 3-manifolds with the same π_1 rank are realized by non-zero degree maps, which does not hold in general for epimorphisms between closed aspherical 3-manifolds groups of the same rank (See [16]).

Soma gave a positive answer to Reid, Wang and Zhou's question in case all M_i are hyperbolic [42, Theorem 1]. Since every hyperbolic 3-manifold group embeds into $\mathrm{PSL}(2,\mathbb{C})$, Soma's result can be derived using the classical Hilbert Basis Theorem. To make this approach work in the general case one would need to not only show that all 3-manifold groups are linear (which is still open for certain graph manifolds), but also to find a *single* linear group into which *every* 3-manifold group embeds. Whether this can be done is an open question which seems to be out of reach of current methods.

In order to solve Reid, Wang and Zhou's question in full generality, a new approach was needed. We avoid questions of linearity and assumptions of positive degree by using Theorem A as a replacement for the Hilbert Basis Theorem. From the point of view of equations over groups, a homomorphism $h\colon G \to H$ between groups corresponds to a solution in H to a system of equations (where the generators of G are variables, and the relators of G give equations). In this language, Theorem A is an analogue of the Hilbert Basis Theorem, uniformly for all groups in $\mathscr{M}^\pi$. In the language of [17] Theorem A says that $\mathscr{M}^\pi$ is an *equationally noetherian* family of groups.

Many other people have studied collections of maps between 3–manifolds or between 3–manifold groups. Examples of previous work include [35, 9, 22, 33, 40, 41, 46, 13, 14, 7, 8, 27, 2]. For an overview of the state of the art in 2002 about positive degree maps, see Wang's ICM address [45]. Highlights of work since then have been Agol and Liu's solution to Simon's Conjecture (about epimorphisms between knot groups) [2], and Liu's proof that for fixed $k \geq 1$, a given closed 3-manifold admits a degree k map onto only finitely many other closed 3–manifolds [27]. However, as far as we are aware, all previous work in this area has made assumptions about the 3–manifolds or the types of maps considered, whereas Theorem A makes no assumptions beyond compactness.

Since the Borel Conjecture is true in dimension 3 (see for instance [24, Section 5]), we also obtain:

Corollary C. *Let $(M_i)_{i\in\mathbb{N}}$ be a sequence of closed, aspherical 3–manifolds and let $f_i\colon M_i \to M_{i+1}$ be π_1–surjective maps. There exists N so f_i is homotopic to a homeomorphism for all $i \geq N$.*

We now turn to our approach to proving Theorem A. We fix a finitely generated group G and consider the collection of all homomorphisms from G to Γ, for all $\Gamma \in \mathscr{M}^{\pi}$. Our basic approach, inspired by Sela [38, 39] and his work on limit groups, etc., is to consider sequence of homomorphisms and try to extract limits. In this theory, the goal is usually to find a limiting $\mathbb{R}$-tree from a *divergent* sequence of homomorphisms. At this point, Sela's shortening argument is used to reduce to the case of non-divergent homomorphisms which are usually constant modulo conjugation.

In order to find limiting $\mathbb{R}$–trees, one needs actions on δ–hyperbolic spaces. For 3–manifold groups, there are the Bass–Serre trees arising from the Kneser–Milnor decomposition or the geometric decomposition. This is the context of previous work by the first two authors [17], who used limiting $\mathbb{R}$–trees and a version of Sela's shortening argument to reduce many questions to the case where a sequence of homomorphisms does not diverge on these trees. This is the starting point of this paper. However, this 'non-divergent' case is still highly nontrivial. We do obtain a limiting simplicial action on a tree in this setting, which we call a *generalized geometric decomposition*. For each vertex group V of this splitting, there is a corresponding sequence of pieces in the geometric decompositions. For each such piece, we construct a new space called the *collapsed space* which is variant of a construction of Farb from [15] (Section 5). This sequence of spaces induces a new limiting action of V on an $\mathbb{R}$-tree (see Section 7), and the collapsed spaces are carefully tuned so that the limiting actions behaves sufficiently well and so we can repeat the shortening argument for this new action (see Section 9).

Even having applied the shortening argument twice, and having a new more restrictive notion of non-divergent, the non-divergent sequences are still far from being constant. Thus, we have yet another layer to our analysis. In this case, we construct a new splitting of our limit group L dual to a tree that we construct as a quotient of the Cayley graph of L. This splitting has an algebraic property which we call edge-twisted (see Definition 6.5). We prove some technical results about edge-twisted splittings Appendix A, and applying these to our splitting of L allows us to conclude that each vertex group in the generalized geometric decomposition of L is itself a limit group coming from a sequence of homomorphisms to geometric 3-manifolds. Finally, we can complete the proof by using the results of third author [25] to understand these limit groups over geometric 3-manifolds.

This paper is in some sense a culmination of the prior work [17, 25] of the authors. On the other hand, we expect the tools in this paper to be useful for many other questions about maps between 3–manifold groups. In particular, although Theorems A and B are immediate from the Hilbert Basis Theorem in the case of Kleinian groups, we believe that the tools built in this paper will be useful for other questions about maps between and to Kleinian groups, and we intend to pursue this direction in future work.

We now outline the contents of this paper. In Section 2 we outline the proof of Theorem A and use this to derive Theorem B. In Section 2 Theorem A is reduced to four results, proved later in the paper. In Section 3 we prove Theorem 2.4, which allows us to focus on a restricted class of 3–manifold groups. In Section 4 we consider various splittings of limit groups. Later sections are dedicated to the proof of the main technical result, Theorem 2.7.

s:outline of proof

2. Outline of proofs

In this section we explain the proof of Theorem A. Specifically, we explain the outline of the proof, and also state four results: Theorem 2.4, Theorem 2.5, Theorem 2.6, and Theorem 2.7 (proved in later sections). Assuming these results, we give a complete proof of Theorem A. We deduce Theorem B from Theorem A.

Our first reduction uses the basic structure of compact 3–manifolds and 3–manifold groups to restrict the class of 3–manifolds we need to consider. We refer to [3] for background about 3–manifolds, their fundamental groups, etc.

d:gen type

Definition 2.1. *Let $\mathscr{M}_{\mathrm{Gen}}$ be the collection of* 3*–manifolds M so that:*

(1) M is closed, orientable, and irreducible;
(2) All geometric pieces are hyperbolic or large Seifert-fibered; and
(3) The base orbifold of each Seifert-fibered piece is orientable.

Let $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ denote the collection of fundamental groups of elements of $\mathscr{M}_{\mathrm{Gen}}$.

Throughout this paper, we fix a non-principal ultrafilter ω and use the concepts of ω–limits, etc. See [44] for the background and basic results about ultrafilters, ultralimits, etc. Fix a non-principal ultrafilter ω which is a finitely additive probability measure $\omega\colon 2^{\mathbb{N}} \to \{0,1\}$ so $\omega(F) = 0$ for any finite set F. A statement $P(i)$ depending on an index i holds *ω–almost surely* if $\omega(\{i \mid P(i) \text{ holds}\}) = 1$.

def:limit group

Definition 2.2. *Suppose $\mathscr{G}$ is a family of groups, G is a finitely generated group, and (ϕ_i) is a sequence of homomorphisms from* $\mathrm{Hom}(G,\mathscr{G})$*. Associated to (ϕ_i) is the* stable kernel $\mathrm{Ker}^{\omega}(\phi_i) = \{g \in G \mid \ \omega\text{–almost surely } \phi_i(g) = 1\}$ *and a $\mathscr{G}$–*limit group *defined by $L := G/\mathrm{Ker}^{\omega}(\phi_i)$. Let ϕ_∞ denote the quotient map $G \twoheadrightarrow L$. The sequence (ϕ_i)* ω–factors through the limit *if ϕ_i ω–almost surely factors through ϕ_∞.*

The notion of an *equationally noetherian* family of groups was introduced by the first and second author in [17, Definition A], and a characterization was given in [17, Theorem 3.5]. For the purposes of this paper, the formulation from [17, Theorem 3.5] is most useful to us, and we use this as a definition. See [17] for more details about these notions.

def:eqchar

Definition 2.3. *The family $\mathscr{G}$ is equationally noetherian if for every finitely generated group G, every sequence from* $\mathrm{Hom}(G,\mathscr{G})$ *ω–factors through the limit.*

Our first reduction towards the proof of Theorem A is the following, proved in Section 3.

t@@reductiontheorem@data

thmt@@reductiontheorem

Theorem 2.4. *If $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is equationally noetherian then so is $\mathscr{M}^{\pi}$.*

Our proof of Theorem A proceeds by verifying the criterion from Definition 2.3. One feature a sequence $(\phi_i\colon\ G \to \Gamma_i)$ from $\mathrm{Hom}(G,\mathscr{M}^{\pi}_{\mathrm{Gen}})$ might have is being *$\mathscr{T}$–divergent* (see Definition 4.4 – roughly speaking the actions of ϕ_i on the trees dual to the geometric decompositions diverge). The main result of [17] implies the following theorem, proved in Section 4.

thmt@@groveshull@data

thmt@@groveshullGH

Theorem 2.5. *Suppose that for every finitely generated group G, every sequence from* $\mathrm{Hom}(G,\mathscr{M}^{\pi}_{\mathrm{Gen}})$ *which is not $\mathscr{T}$–divergent ω–factors through the limit. Then for every finitely generated group G* every *sequence from* $\mathrm{Hom}(G,\mathscr{M}^{\pi}_{\mathrm{Gen}})$ *ω–factors through the limit.*

Thus, we fix a non-$\mathscr{T}$–divergent sequence $(\phi_i\colon\ G \to \Gamma_i)$ defining an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group L. The G–actions on the trees associated to the Γ_i converge to an L–action on a (simplicial) tree, inducing the *geometric decomposition* of L (Definition 4.15). A key set of subgroups are *stably parabolic* subgroups (Definition 4.5). These are abelian (Lemma 4.6) and edge groups of the geometric decomposition are stably parabolic. The following is proved in Section 4.

thmt@@fingenedges@data

thmt@@fingenedges

Theorem 2.6. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a non–$\mathscr{T}$–divergent sequence (ϕ_i) and suppose the edge groups of the geometric decomposition of L (wrt (ϕ_i)) are finitely generated. Then (ϕ_i) ω–factors through the limit.*

At this point, the following result (together with Theorems 2.4, 2.5 and 2.6) completes the proof of Theorem A.

collapsingdivergent@data

thmt@@collapsingDivergent

Theorem 2.7. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a non-$\mathscr{T}$–divergent sequence (ϕ_i). All stably parabolic subgroups of L are finitely generated.*

Our approach to proving Theorem 2.7 is to analyze the geometric decomposition of L using the *collapsed spaces* defined in Section 5. This leads to the notion of a *$\mathscr{C}$–divergent* sequence (see Definition 6.1). In Section 6 we prove Theorem 6.2 (modulo the technical Theorem 6.7 proved in Appendix A), which deals with the case where the sequence is not $\mathscr{C}$–divergent. After build-up in Sections 5, 7 and 8, we prove Theorem 2.7 in Section 9.

We finish by explicitly deducing Theorem A from the results in this section. By Theorem 2.4 it suffices to prove $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is equationally noetherian. Since edge groups of the geometric decomposition are stably parabolic, Theorems 2.7 and 2.6 prove that non–$\mathscr{T}$–divergent sequences ω–factor through the limit. Theorem 2.5completes the proof of Theorem A.

2.1. **Proof of Theorem B.** We now deduce Theorem B from Theorem A. Suppose that $(\Gamma_i)_{i=1}^{\infty}$ is a sequence from $\mathscr{M}^{\pi}$, and that for each $i \geq 1$ there is a surjection $\tau_i\colon \Gamma_i \to \Gamma_{i+1}$. Define (ρ_i) from $\mathrm{Hom}(\Gamma_1, \mathscr{M}^{\pi})$ by $\rho_i = \tau_i \circ \cdots \circ \tau_1\colon \Gamma_1 \to \Gamma_{i+1}$. By Theorem A $\mathscr{M}^{\pi}$ is equationally noetherian, so by Definition 2.3 (ρ_i) ω–factors through the limit map ρ_∞.

Now, $\mathrm{Ker}(\rho_{i+1}) \subseteq \mathrm{Ker}(\rho_i)$, so

$$\mathrm{Ker}(\rho_\infty) = \bigcup_{i=1}^{\infty} \mathrm{Ker}(\rho_i).$$

By Theorem A for ω–almost all j the homomorphism ρ_j factors through ρ_∞. Fix such a j. Then $\mathrm{Ker}(\rho_\infty) \subseteq \mathrm{Ker}(\rho_j)$, and since $\mathrm{Ker}(\rho_j) \subset \mathrm{Ker}(\rho_\infty)$, the two are equal. Thus, the ascending sequence of kernels stabilizes after ρ_j, and for all $k \geq j$ the map τ_k is an isomorphism, completing the proof.

2.2. **Notation and conventions.** If X is a metric space and $x, y \in X$ then we denote a geodesic between x and y by $[x, y]$.

A graph of groups is *minimal* if there are no proper invariant subtrees of the Bass–Serre tree. Let G be a group, and $\mathscr{A}$ and $\mathscr{H}$ families of subgroups closed under conjugation. An *$(\mathscr{A}, \mathscr{H})$–splitting* of G is a splitting of G in which all edge groups are in $\mathscr{A}$ and all groups in $\mathscr{H}$ are elliptic.

Throughout this paper a *δ–hyperbolic space* is a geodesic metric space in which all geodesic triangles are *δ–thin* in the sense of [10, Definition III.H.1.16]. We remark that if a geodesic metric space has ν–slim triangles (in the sense of [10, Definition III.H.1.1]) then it has 4ν–thin triangles.

s:reduction

3. Reduction to $\mathscr{M}^{\pi}_{\mathrm{Gen}}$

The purpose of this section is to record some basic facts about (limit groups over the family of) 3–manifold groups, and in particular to prove Theorem 2.4.

3.1. **Equationally Noetherian families.** In this subsection, we record some facts about equationally noetherian families not particular to 3–manifold groups. The first follows from [4, Theorem B1] and [17, Lemma 3.9].

lem:linear **Lemma 3.1.** *For any field K and $n \geq 1$, if $\mathscr{G}$ is a family of groups so each $\Gamma \in \mathscr{G}$ is a subgroup of $GL(n,K)$, then $\mathscr{G}$ is equationally noetherian.*

lem:relfplim **Lemma 3.2.** [17, Lemma 3.20] *Suppose (ϕ_i) is a sequence from* $\mathrm{Hom}(G,\mathscr{G})$ *and the associated limit group is finitely presented relative to subgroups $\{P_1,\dots,P_n\}$. Suppose for each P_j, there is $\widetilde{P}_j \leq G$ so that $\phi_\infty(\widetilde{P}_j) = P_j$ and that $\phi_i|_{\widetilde{P}_j}$ ω–almost surely factors through $\phi_\infty|_{\widetilde{P}_j}$. Then (ϕ_i) ω–factors through the limit.*

Theorem 1 of [5] states that if H is a finite index subgroup of G and H is equationally noetherian, then G is equationally noetherian. We generalize this to a family of groups.

Definition 3.3. *Let $\mathscr{G}$ be a family of groups and let $n \geq 1$. The class $\mathscr{G}^{\leq n}$ is the set of groups G containing a subgroup from $\mathscr{G}$ of index at most n.*

prop:fi **Proposition 3.4.** *Suppose that $\mathscr{G}$ is an equationally noetherian family of groups. For any $n \geq 1$ the family $\mathscr{G}^{\leq n}$ is equationally noetherian.*

Proof. Let G be finitely generated and let $(\phi_i \colon G \to \Gamma_i)$ be a sequence with $\Gamma_i \in \mathscr{G}^{\leq n}$. By Definition 2.3 it suffices to show that ω–almost surely $\mathrm{Ker}(\phi_\infty) \subseteq \mathrm{Ker}(\phi_i)$. So, let $\Gamma_i' \in \mathscr{G}$ be so that $|\Gamma_i : \Gamma_i'| \leq n$. Let H be the intersection of all subgroups of G of index at most n. Then $\phi_i(H) \subseteq \Gamma_i'$, so $\phi_i|_H \in \mathrm{Hom}(H,\mathscr{G})$. Since $\mathscr{G}$ is equationally noetherian by Definition 2.3 ω–almost surely $\mathrm{Ker}(\phi_\infty|_H) \subseteq \mathrm{Ker}(\phi_i|_H)$.

Let $g \in G$. If $\phi_\infty(g) \notin \phi_\infty(H)$, then for any $h \in H$ $gh \notin \mathrm{Ker}(\phi_\infty)$. Suppose now that $\phi_\infty(g) \in \phi_\infty(H)$. Fix $k \in H$ such that $\phi_\infty(g) = \phi_\infty(k)$. Now, ω–almost surely $gk^{-1} \in \mathrm{Ker}(\phi_i)$ and $\mathrm{Ker}(\phi_\infty|_H) \subseteq \mathrm{Ker}(\phi_i|_H)$. For any such i and for any $h \in H$, if $gh \in \mathrm{Ker}(\phi_\infty)$ then $kh \in \mathrm{Ker}(\phi_\infty) \subseteq \mathrm{Ker}(\phi_i)$. Hence $gh = gk^{-1}kh \in \mathrm{Ker}(\phi_i)$. Putting together these two cases, we see that for any $g \in G$ ω–almost surely $\mathrm{Ker}(\phi_\infty) \cap gH \subseteq \mathrm{Ker}(\phi_i)$. Repeating this argument over a finite of coset representatives of H in G proves that ω–almost surely $\mathrm{Ker}(\phi_\infty) \subseteq \mathrm{Ker}(\phi_i)$. □

Remark 3.5. *The bound on the index in Proposition 3.4 is essential. Indeed, it is straightforward to see that if one considers the family $\mathscr{G}^{\mathrm{fin}}$ of groups which have a finite index subgroup contained in $\mathscr{G}$ then $\mathscr{G}$ may be equationally noetherian while $\mathscr{G}^{\mathrm{fin}}$ is not. This is true, for example, in case $\mathscr{G}$ is the family consisting of the trivial group, since then $\mathscr{G}^{\mathrm{fin}}$ is the family of finite groups, which we have already seen is not equationally noetherian.*

Definition 3.6. *If $\mathscr{G}$ is a family of groups, then let $\mathscr{G}^*$ be the family of all free products of elements of $\mathscr{G}$.*

t:free products **Theorem 3.7.** [17, Cor. C] *If $\mathscr{G}$ is equationally noetherian then so is $\mathscr{G}^*$.*

ss:geometric 3.2. **The geometric decomposition.** If M is a closed, orientable, irreducible 3–manifold then the Geometrization Theorem (see [30, 31, 32, 12, 23, 28]) gives a (possibly trivial) decomposition of M along incompressible tori into geometric pieces. This induces a splitting of $\pi_1(M)$ which we call the *geometric decomposition*. We refer to [3], [37], or [43] for the definition of a geometric 3–manifold, and to [37] for details about each of the 8 Thurston geometries.

Assume M is a closed, orientable, irreducible 3–manifold and also that M is not a geometric 3–manifold (in particular, M is not a torus bundle over a circle). Then the each piece of the geometric decomposition of M is either hyperbolic, large Seifert-fibered, or a twisted I–bundle over a Klein bottle. Replacing M by a double cover as in [49, Lemma 2.4] we may assume there are no twisted I–bundles over a Klein bottle. By a similar argument, there is a double cover so all the base orbifolds of large Seifert-fibered pieces are orientable. Thus:

l:cover to become gen **Lemma 3.8.** *Let M be a closed, orientable, irreducible, non-geometric 3–manifold. Then M has a cover of degree at most 4 which lies in $\mathscr{M}_{\mathrm{Gen}}$.*

An important property of elements of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is the following result, which follows immediately from the proof of [49, Lemma 2.4]. Recall a group action on a tree T is k–*acylindrical* if the stabilizer of any segment in T of length at least $k+1$ is trivial.

l:2-acyl **Lemma 3.9.** *Suppose $\Gamma \in \mathscr{M}^{\pi}_{\mathrm{Gen}}$ and that T_Γ is the Bass-Serre tree dual to the geometric decomposition of Γ. The Γ–action on T_Γ is 2–acylindrical.*

3.3. **Reduction to $\mathscr{M}^{\pi}_{\mathrm{Gen}}$.** We now identify the some pertinent classes of 3–manifolds.

Definition 3.10. *Let $\mathscr{M}^{\pi}_{\mathrm{Ori}}$ be set of the fundamental groups of closed, orientable 3–manifolds. Let $\mathscr{M}^{\pi}_{\mathrm{Irr}}$ be the set of fundamental groups of closed, orientable and irreducible 3–manifolds. Let $\mathscr{M}^{\pi}_{\mathrm{Geo}}$ be the set of fundamental groups of compact geometric 3–manifolds with all boundary components tori.*

c:red to ori **Lemma 3.11.** *If $\mathscr{M}^{\pi}_{\mathrm{Ori}}$ is equationally noetherian then so is $\mathscr{M}^{\pi}$.*

Proof. The fundamental group of any compact 3–manifold with boundary is a subgroup (in fact, a retract) of the fundamental group of a closed 3–manifold. It is an immediate consequence of [17, Lemma 3.9] that a family of groups $\mathscr{G}$ is equationally noetherian if and only if the family of all subgroups of elements of $\mathscr{G}$ is equationally noetherian, so it suffices to consider only closed manifolds.

For each closed 3–manifold M, $\pi_1(M)$ has an index at most 2 subgroup belonging to $\mathscr{M}^{\pi}_{\mathrm{Ori}}$, so the result follows from Proposition 3.4. □

Each closed, orientable 3–manifold M has a unique Milnor–Kneser decomposition which induces a free product decomposition of $\pi_1(M)$:

$$\pi_1(M) \cong \pi_1(M_1) * \ldots * \pi_1(M_k) * F_n$$

where each $\pi_1(M_i) \in \mathscr{M}^{\pi}_{\mathrm{Irr}}$ (and F_n is a free group of rank n). The following is thus an immediate consequence of Theorem 3.7

c:irr EN then all EN **Corollary 3.12.** *If $\mathscr{M}^{\pi}_{\mathrm{Irr}}$ is equationally noetherian then so is $\mathscr{M}^{\pi}_{\mathrm{Ori}}$.*

Consider the case $\mathscr{M}^{\pi}_{\mathrm{LSF}}$ of large Seifert–fibered manifolds with orientable base orbifold. We prove some basic results, and also some results we need later. Let $\mathscr{M}^{\pi}_{\mathrm{Orb}}$ be the set of fundamental groups of orientable hyperbolic 2–orbifolds of finite-type. Each $\Gamma \in \mathscr{M}^{\pi}_{\mathrm{LSF}}$ admits a short exact sequence which is a central extension:

$$1 \to \mathbb{Z} \to \Gamma \to \Gamma_B \to 1,$$

for some $\Gamma_B \in \mathscr{M}^{\pi}_{\mathrm{Orb}}$ (the fundamental group of the base orbifold). From this, the following result follows quickly.

lem:central **Lemma 3.13.** *Let L be an $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit group. Then there is a central extension*

$$1 \to A \to L \to B \to 1,$$

where A is abelian and B is an $\mathscr{M}^{\pi}_{\mathrm{Orb}}$–limit group.

d:A-slender **Definition 3.14.** *A group G is $\mathscr{A}$–*slender *if every abelian subgroup of G is finitely generated.*

Our analysis of $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit groups is based on combining Lemma 3.13 with the following theorem of the third author.

thm:Liang **Theorem 3.15.** [25, Theorem 1.2] *Let B be an $\mathscr{M}^{\pi}_{\mathrm{Orb}}$–limit group. Then B is finitely presented and $\mathscr{A}$–slender.*

or:SF limit fp+A-slender **Corollary 3.16.** *Let L be an $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit group. Then L is finitely presented and $\mathscr{A}$–slender.*

Proof. Consider the central extension from Lemma 3.13. Since B is finitely presented by Theorem 3.15, A is finitely generated, for example by [18, Lemma 12.1]. It is now easy to see L is finitely presented. The fact that A is finitely generated and abelian subgroups of B are finitely generated implies that all abelian subgroups of L are finitely generated. □

l:abelian orb implies SF **Lemma 3.17.** *Let M be a large Seifert fibered 3–manifold with orientable base orbifold O. Let K be the kernel of the natural quotient map $\pi_1(M) \to \pi_1(O)$. Let $g,h \in \pi_1(M)$. If $[g,h] \in K$, then $[g,h] = 1 \in \pi_1(M)$.*

Proof. All abelian subgroups of $\pi_1(O)$ are cyclic, so the images of g and h in $\pi_1(O)$ generate a cyclic subgroup. A central extension of a cyclic group is abelian. □

The following result follows quickly from the fact that finite subgroups of elements of $\mathscr{M}^{\pi}_{\mathrm{Orb}}$ are cyclic.

lem:fin cyclic **Lemma 3.18.** *Let B be an $\mathscr{M}^{\pi}_{\mathrm{Orb}}$–limit group. Any finite subgroup of B is cyclic.*

The following result is now immediate from Lemmas 3.17 and 3.18.

lem:ab **Corollary 3.19.** *Suppose that L is an $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit group, and let $1 \to A \to L \to B \to 1$ be as in Lemma 3.13. Suppose that $E \le L$ has finite image in B. Then E is abelian.*

JSJ properties **Lemma 3.20.** *Let $\Gamma \in \mathscr{M}^{\pi}_{\mathrm{Gen}}$, and let Γ_e be an edge group of the geometric decomposition of Γ, with adjacent vertex groups Γ_v, Γ_w. The pre-images in Γ_e of the centers $Z(\Gamma_v)$ and $Z(\Gamma_w)$ intersect trivially.*

Proof. The only way that $Z(\Gamma_v)$ and $Z(\Gamma_w)$ can both be nontrivial is if the corresponding pieces are both Seifert–fibered. The centers of adjacent Seifert-fibered pieces intersect trivially, since otherwise the connecting torus is not part of the characteristic sub-manifold. □

t:geo EN **Theorem 3.21.** *$\mathscr{M}^{\pi}_{\mathrm{Geo}}$ is an equationally noetherian family of groups.*

Proof. It is a straightforward observation that any finite union of equationally noetherian families is equationally noetherian [17, Lemma 3.8]. Hence it suffices to deal with each geometry individually.

Suppose first that M supports one of the following geometries: $\mathbb{S}^3$, $\mathbb{S}^2 \times \mathbb{R}$, $\mathbb{E}^3$, NIL or SOL. Then (as, for example, explained in [18, §7 − 11]) $\pi_1(M)$ has a subgroup of index at most 240 which is either cyclic (finite or infinite), $\mathbb{Z}^3$, or some semi-direct product $\mathbb{Z} \rtimes \mathbb{Z}$. Each of these groups embeds in $\mathrm{SL}_3(\mathbb{Z})$. Therefore, the union of the set of fundamental groups of manifolds supporting these five geometries in equationally noetherian by Proposition 3.4 and Lemma 3.1.

The fundamental group of any hyperbolic 3–manifold is (isomorphic to) a subgroup of $\mathrm{PSL}(2,\mathbb{C})$ (and hence also contained in $\mathrm{SL}(3,\mathbb{C})$ – see [18, Lemma 12.9]). Thus, the family of fundamental groups of hyperbolic 3–manifolds is equationally noetherian by Lemma 3.1.

The two remaining geometries correspond to *large Seifert fibered manifolds*. Since we can pass to an index 2 subgroup to lie in $\mathscr{M}^{\pi}_{\mathrm{LSF}}$, by Proposition 3.4 it suffices to consider $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit groups. If L is an $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit group then L is finitely presented by 3.16. It follows that the sequence of maps defining L ω–factors through the limit (see Lemma 3.2). Thus the family of fundamental groups of large Seifert-fibered manifolds is equationally noetherian by Definition 2.3. □

Summarizing the above discussion and the previous results in the section, we have the following result from Section 2.

Theorem 2.4. *If $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is equationally noetherian then so is $\mathscr{M}^{\pi}$.*

Proof. By Lemma 3.11 and Corollary 3.12 it is enough to prove that $\mathscr{M}^{\pi}_{\mathrm{Irr}}$ is equationally noetherian. By Theorem 3.21 $\mathscr{M}^{\pi}_{\mathrm{Geo}}$ is equationally noetherian, so it is enough to consider closed, orientable, irreducible, non-geometric 3–manifolds. By Lemma 3.8 any such 3–manifold has a degree at most 4 cover which lies in $\mathscr{M}_{\mathrm{Gen}}$, so by Proposition 3.4 in suffices to prove that $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is equationally noetherian. □

4. GENERALIZED GEOMETRIC DECOMPOSITIONS

s:geo decomp

In order to prove Theorem A, it remains to prove Theorems 2.5, 2.6 and 2.7. In this section, we provide the setting and some basic definitions, and prove Theorems 2.5 and 2.6.

Let $G = \langle S \rangle$ be a finitely generated group, and let $(\phi_i \colon G \to \Gamma_i)$ be a sequence from $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$. Let $L = G/\mathrm{Ker}^{\omega}(\phi_i)$ be the $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group associated to (ϕ_i), with associated limit map $\phi_\infty \colon G \twoheadrightarrow L$ (see Definition 2.2).

Definition 4.1. *Suppose that $G, (\phi_i)$ and L are as above, and that $g \in L$. An* ω–approximation *to g is a sequence (g_i) so $g_i \in \Gamma_i$ and there is an element $\widetilde{g} \in G$ so that $\phi_\infty(\widetilde{g}) = g$ and for each i we have $\phi_i(\widetilde{g}) = g_i$.*

Similarly, if $F \subset L$ is a finite set then an ω–approximation *to F is a sequence of (ordered) tuples (F_i) from Γ_i so that for some lift $\widetilde{F}$ of F to G (so ϕ_∞ restricts to an ordered bijection from $\widetilde{F}$ to F) we have $\phi_i(\widetilde{F}) = F_i$ (again, as an ordered bijection).*

The following lemmas are immediate.

Lemma 4.2. *Suppose that $g \in L$ and that (g_i) and (g'_i) are ω–approximations to g. Then ω–almost surely $g_i = g'_i$. Similarly, if $F \subset L$ is finite and (F_i), (F'_i) are ω–approximations to F then ω–almost surely $F_i = F'_i$.*

l:appro relation **Lemma 4.3.** *Suppose $g^1, \ldots, g^s \in L$ and for each $1 \le j \le s$, (g^j_i) is an ω–approximation to g^j. For any word $w(x_1, \ldots, x_s)$ in $\{x_i^{\pm}\}$ we have $w(g^1, \ldots, g^s) = 1$ in L if and only if for ω–almost all i $w(g^1_i, \ldots, g^s_i) = 1$ in Γ_i.*

We are only concerned with properties of ω–approximations that hold ω–almost surely for terms in the sequence, and so it is always irrelevant which ω–approximation we choose for an element (or finite subset) of L. We use this observation frequently without mention.

Let S be a finite generating set for G, (S_i) an ω–approximation to S, and T_i the tree dual to the geometric decomposition of Γ_i. Let

$$\|\phi_i\|_{T_i} = \inf_{t \in T_i} \max_{s \in S_i} d_{T_i}(t, s.t).$$

d:Tdiv **Definition 4.4.** *The sequence (ϕ_i) is* $\mathscr{T}$–divergent *if $\lim^{\omega} \|\phi_i\|_{T_i} = \infty$.*

We are now ready to prove the following theorem from Section 2.

Theorem 2.5. *Suppose that for every finitely generated group G, every sequence from* $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$ *which is not $\mathscr{T}$–divergent ω–factors through the limit. Then for every finitely generated group G* every *sequence from* $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$ *ω–factors through the limit.*

Proof. The action of any $\Gamma \in \mathscr{M}^{\pi}_{\mathrm{Gen}}$ on its geometric tree is 2–acylindrical (see Lemma 3.9). In particular, $\mathscr{M}^{\pi}_{\mathrm{Gen}}$ is (in the terminology of [17]) a uniformly acylindrical family of groups (see, for example, [17, p.7142]). Moreover, any sequence from $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$ which is not $\mathscr{T}$–divergent is 'non-divergent' in the sense of [17]. Thus, Theorem 2.5 is an immediate consequence of [17, Theorem B]. □

As explained in Section 2, by Theorem 2.5 we may assume the sequence (ϕ_i) is not $\mathscr{T}$–divergent.

def:stably para

Definition 4.5. *Let $g \in L$, and let (g_i) be an ω–approximation to g. Then g is* stably parabolic with respect to (ϕ_i) *if ω–almost surely g_i fixes an edge in T_i. A subgroup $H \leq L$ is* stably parabolic with respect to (ϕ_i) *if for any finite $F \subset H$ and any ω–approximation (F_i) of F ω–almost surely there is an edge in T_i fixed by each $f \in F_i$. The set of stably parabolic subgroups of L with respect to (ϕ_i) is $\mathscr{H}_{L,(\phi_i)}$, or just $\mathscr{H}_L$ when (ϕ_i) is implied.*

Since each subgroup of Γ_i which stabilizes an edge in T_i is isomorphic to $\mathbb{Z}^2$, the following result is immediate from Lemma 4.3.

l:sp abelian

Lemma 4.6. *Any stably parabolic subgroup of L is abelian.*

def:sub_geom

Definition 4.7. *A subgroup H of L is* ω–geometric *if for any finite subset $F \subset H$ and any ω–approximation (F_i) of F, ω–almost surely there is a vertex v_i of T_i so that F_i fixes v_i.*

Terminology 4.8. *Let H be an ω–geometric subgroup of L and v_i be as in Definition 4.7. Then either*

(1) ω–almost surely v_i is of hyperbolic type; or
(2) ω–almost surely v_i is of LSF-type.

In the first case H is hyperbolic-type *and in the second H is* LSF-type.

def:stable center

Definition 4.9. *Let H be an ω–geometric subgroup of L, with associated sequence v_i of vertices of T_i, and let Γ_{i,v_i} be the associated vertex group of Γ_i. The* stable center *of H, denoted $Z^{\omega}(H)$, is the set of $g \in H$ so that for any ω–approximation (g_i) of g, ω–almost surely $g_i \in Z(\Gamma_{i,v_i})$*

Observe that if H is a hyperbolic-type ω–geometric subgroup of L then $Z^{\omega}(H) = \{1\}$.

As we see in Definition 4.15 below, L admits a splitting arising from a limiting action on the trees T_i. This geometric decomposition is the natural splitting associated to L, but for technical reasons in Section 9 we need a more general class of splittings of L.

def:ggd

Definition 4.10. *A minimal abelian splitting $\mathbb{G}$ of L is a* generalized geometric decomposition (or GGD) with respect to (ϕ_i) *if*

ggd:vertices
ggd:edges

(1) Each vertex group of $\mathbb{G}$ is ω–geometric
(2) Let v and w be adjacent vertices of $\mathbb{G}$ with sequences (v_i) and (w_i) as in the definition of ω–geometric. Then ω–almost surely v_i and w_i are adjacent in T_i;
(3) $\mathbb{G}$ is an $\left(\mathscr{H}_{L,(\phi_i)}, \mathscr{H}_{L,(\phi_i)}\right)$–splitting.
(4) For each LSF–type vertex group $\mathbb{G}_v$ of $\mathbb{G}$, $Z^{\omega}(\mathbb{G}_v)$ is contained in each edge group of $\mathbb{G}$ adjacent to $\mathbb{G}_v$.

For the rest of this subsection, fix the above notation, and let $\mathbb{G}$ be a GGD of L with respect to (ϕ_i). The sequences (v_i) and (w_i) in Definition 4.10 are implicitly fixed.

perties of stable center

Lemma 4.11. *If $\mathbb{G}_v$ is an LSF-type vertex group of $\mathbb{G}$ then $Z^{\omega}(\mathbb{G}_v)$ is central in $\mathbb{G}_v$.*

Proof. Let $g \in Z^{\omega}(\mathbb{G}_v)$ and $h \in \mathbb{G}_v$. If (g_i) and (h_i) are ω–approximations to g and h, respectively then ω–almost surely there is an LSF-type vertex group Γ_{i,v_i} of Γ_i so that $g_i \in Z(\Gamma_{i,v_i})$ and $h_i \in \Gamma_{i,v_i}$. Therefore ω–almost surely $[g_i,h_i] = 1$ in Γ_i. By Lemma 4.3 $[g,h] = 1$ in L. □

Let V be a vertex group of $\mathbb{G}$. By Lemma 4.11, $Z^{\omega}(V) \trianglelefteq V$, and we define $\overline{V} = V/Z^{\omega}(V)$. Let $\pi : V \to \overline{V}$ be the quotient map.

l:preimage abelian

Lemma 4.12. *Let V be an LSF type vertex group of $\mathbb{G}$. If H is an abelian subgroup of $\overline{V}$ then $\pi^{-1}(H)$ is an abelian subgroup of V.*

Proof. Let $\bar{g},\bar{h} \in H$, let $g \in \pi^{-1}(\bar{g})$ and $h \in \pi^{-1}(\bar{h})$, and let (g_i) and (h_i) be ω–approximations to g and h, respectively. Since $[\bar{g},\bar{h}] = 1$, we have $[g,h] \in Z^{\omega}(V)$, and so ω–almost surely $[g_i,h_i] \in Z(\Gamma_{v_i})$. Since $Z(\Gamma_{v_i})$ is the kernel of $p_i : \Gamma_{v_i} \twoheadrightarrow \pi_1(O_i)$, where O_i is the base orbifold of the LSF manifold corresponding to Γ_{v_i}, by Lemma 3.17, we have $[g_i,h_i] = 1$ ω–almost surely. By Lemma 4.3 $[g,h] = 1$, as required. □

Let V, $\overline{V}$ be as above, let $\mathscr{E}$ be the edge groups of $\mathbb{G}$ adjacent to V, and $\overline{\mathscr{E}}$ their image in $\overline{V}$. By Corollary B.4 $\overline{V}$ admits a graph of groups decomposition $\overline{\mathbb{L}}$ rel $\overline{\mathscr{E}}$ so that all edge groups of $\overline{\mathbb{L}}$ have cardinality at most 4, and no vertex group of $\overline{\mathbb{L}}$ splits rel $\overline{\mathscr{E}}$ over a subgroup of size at most 4. We call $\overline{\mathbb{L}}$ the *relative* 4*–Linnell decomposition of $\overline{V}$ rel $\overline{\mathscr{E}}$*. Notice that by Corollary 3.19 the preimages in V of edges in $\overline{\mathbb{L}}$ are abelian.

def:refined ggd

Definition 4.13. *Let V be an LSF–type vertex group of $\mathbb{G}$. Let $\overline{V} = V/Z^{\omega}(V)$. Let $\overline{\mathbb{L}}$ be the relative* 4*–Linnell decomposition of $\overline{V}$ relative to (images of) edge groups of $\mathbb{G}$ adjacent to V, and $\mathbb{L}$ the induced abelian splitting of V relative to adjacent edge groups. The* canonical refinement of $\mathbb{G}$*, denoted $\mathbb{G}^r$, is the refinement of $\mathbb{G}$ obtained by replacing each LSF-type vertex group V by $\mathbb{L}$.*

def:good generating set

Definition 4.14. *Suppose V is a vertex group of $\mathbb{G}^r$. A* good relative generating set *for V is a finite set A so that*

(1) A together with the adjacent edge groups generates V; and
(2) For each edge group E adjacent to V, there exists $a \in A \cap E \setminus Z^{\omega}(V)$.

4.1. The geometric decomposition.

Let $(\phi_i\colon\ G \to \Gamma_i)$ be a sequence from $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$ which is not $\mathscr{T}$–divergent and let T_i be the minimal $\phi_i(G)$–invariant subtree of the tree dual to the geometric splitting of Γ_i. Let o_i be a point in T_i so $\max_{s\in S} d_{T_i}(o_i, \phi_i(s)o_i) = \|\phi_i\|_{T_i}$.

Let $T_\infty = \lim^{\omega}(T_i, o_i)$. Since each T_i is a tree in which edges have length one, T_∞ is a simplicial tree. Moreover, the actions of G on T_i induce a 2–acylindrical action of L on T_∞; see [17, Proposition 6.1].

def:gd

Definition 4.15. *Let G, S, (ϕ_i), T_i, L and T_∞ be as above. The* geometric tree *of L, denoted T_{geom}, is the minimal L–invariant subtree of T_∞.*

The geometric decomposition *of L is the splitting dual to T_{geom}.*

It is straightforward to check that the geometric decomposition of L is in fact a GGD in the sense of Definition 4.10, which we record in the following lemma.

l:gd is ggd

Lemma 4.16. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a sequence $(\phi_i\colon G \to \Gamma_i)$ (so $L = G/\mathrm{Ker}^{\omega}(\phi_i)$). If (ϕ_i) is not $\mathscr{T}$–divergent then the geometric splitting of L is a GGD with respect to (ϕ_i).*

We now prove the following result from Section 2.

Theorem 2.6. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a non–$\mathscr{T}$–divergent sequence (ϕ_i) and suppose the edge groups of the geometric decomposition of L (wrt (ϕ_i)) are finitely generated. Then (ϕ_i) ω–factors through the limit.*

Proof. Let $\mathbb{G}$ be the geometric decomposition of L. The edge groups of $\mathbb{G}$ are finitely generated, so the vertex groups are also and hence each vertex group of $\mathbb{G}$ is an $\mathscr{M}^{\pi}_{\mathrm{Geo}}$–limit group. By Theorem 3.21, the defining sequence for each vertex group ω–factors through the limit. The edge groups of $\mathbb{G}$ are finitely generated and abelian, so L is finitely presented relative to the vertex groups of $\mathbb{G}$, and by Lemma 3.2 the defining sequence for L ω–factors through the limit, as required. □

Given Theorem 2.6, our goal is to prove that the edge groups of $\mathbb{G}$ are finitely generated. It is tempting to think that this follows from the fact that they are abelian edge groups in an acylindrical splitting of a finitely generated group. Moreover, we also know that some (and we believe all) vertex groups of this splitting have the property that their finitely generated subgroups are $\mathscr{A}$–slender. The next example shows that this is not enough.

ex:Hao HNN

Example 4.17. *Let $A = \langle \{a_i\}_{i\in\mathbb{N}}\rangle$ and $B = \langle \{b_i\}_{i\in\mathbb{N}}\rangle$ be free abelian groups of countably infinite rank, and let $H = A * B$. Define:*

- $A_e = \langle \{a_{2k}\}_{k\in\mathbb{N}}\rangle$ $\quad A_o = \langle \{a_{2k-1}\}_{k\in\mathbb{N}}\rangle$;
- $B_e = \langle \{b_{2k}\}_{k\in\mathbb{N}}\rangle$ $\quad B_o = \langle \{b_{2k+1}\}_{k\in\mathbb{N}}\rangle$.

Define $\phi : A_o \to A_e$ by $\phi(a_{2k-1}) = a_{2k}$ and $\psi : B_e \to B_o$ by $\psi(b_{2k}) = b_{2k+1}$. Then ϕ and ψ are both isomorphisms. Considering ϕ and ψ to be isomorphisms between subgroups of H, let K be the the double HNN extension of H via maps ϕ and ψ, and observe K is not finitely generated.

The group K naturally contains A and B as subgroups. Let $\theta : A \to B$ be the isomorphism $\theta(a_i) = b_i$, and G the HNN extension of K with map θ. Then G is generated by the stable letters of the HNN extensions and a_1.

G is finitely generated and the defining HNN extension of G is a 2–acylindrical splitting with abelian edge groups, but the vertex group and the edge groups of this splitting are not finitely generated. The vertex group also has the property that finitely generated subgroups are $\mathscr{A}$–slender (recall Definition 3.14).

4.2. **Assumptions.** It remains to prove Theorem 2.7. If G is the group used to defined the limit group L, then we can replace G by a finitely presented group. For the remainder of the paper we make the following assumptions, which are key for the proof of Theorem 2.7.

ass:fixed things

Standing Assumption 4.18. *Let G be a finitely presented group with finite generating set S. Let $(\phi_i\colon G \to \Gamma_i = \pi_1(M_i))$ be a sequence from $\mathrm{Hom}(G, \mathscr{M}^{\pi}_{\mathrm{Gen}})$ and let T_i be the minimal $\phi_i(G)$–invariant subtree of the geometric tree associated to M_i. Suppose (ϕ_i) is not $\mathscr{T}$–divergent. Let $L = G/\mathrm{Ker}^{\omega}(\phi_i)$ be the $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group, and $\phi_\infty\colon G \twoheadrightarrow L$ the limit map. Let $\mathbb{G}$ be a GGD of L wrt (ϕ_i) and $T_{\mathbb{G}}$ the Bass–Serre tree. Let $\mathbb{G}^r$ be the canonical refinement of $\mathbb{G}$ (Definition 4.13). Fix a vertex group V of $\mathbb{G}^r$ associated to the vertex v, and a good relative generating set A of V (Definition 4.14). Let (A_i) be an ω–approximation to A. Denote by $\mathscr{H}_V$ the set of stably parabolic subgroups of V.*

s:collapsing Divergent

5. THE COLLAPSED SPACE

Suppose $\Delta \in \mathscr{M}^{\pi}_{\mathrm{Gen}}$, and consider the geometric decomposition of Δ. There are two kinds of vertex groups: (i) Hyperbolic vertex groups which act on $\mathbb{H}^3$; and (ii) Large Seifert-fibered vertex groups which (through projection to the base orbifold and a choice of hyperbolic structure on this base orbifold) act on $\mathbb{H}^2$. Each of $\mathbb{H}^3$ and $\mathbb{H}^2$ are δ–hyperbolic. However, these actions do not have properties that are well suited for the arguments in this paper.

In this section we build the *collapsed space*, a variation on Farb's construction of the "Electric" space from [15]. The collapsed space is geodesic and δ–hyperbolic for a uniform δ (see Proposition 5.7 and Theorem 5.8). The key advantage of the collapsed space over $\mathbb{H}^2$ and $\mathbb{H}^3$ is Lemma 5.15. As well as these results, we prove a collection of basic structural results which are used in later sections.

ss:collapsed geometric

5.1. **The collapsed space for Kleinian and Fuchsian groups.** Let Γ be a finitely generated Fuchsian or torsion-free Kleinian group. In case Γ is Kleinian, let $\mathbb{H}^* = \mathbb{H}^3$, and in case Γ is Fuchsian let $\mathbb{H}^* = \mathbb{H}^2$. Fix $\delta_{\mathbb{H}}$ so that all geodesic triangles in $\mathbb{H}^2$ and $\mathbb{H}^3$ are $\delta_{\mathbb{H}}$–thin. Let ε be the 3–dimensional (or 2–dimensional) Margulis constant. (see [6] for more information). Define the *ε–thin part* of $\mathbb{H}^*$ with respect to Γ:

$$\mathscr{U}_\varepsilon(\Gamma) = \{x \in \mathbb{H}^* \mid d(x, \gamma x) \leq \varepsilon \text{ for some } \gamma \neq 1 \in \Gamma\}.$$

The set $\mathscr{U}_\varepsilon(\Gamma)$ is a disjoint union of a collection of horoballs, balls and neighborhoods of geodesics (see, for example, [6, Theorem D.3.3], though we work in the universal cover). We refer to the neighborhoods of geodesics in the thin part of $\mathbb{H}^*$ as *tubes*. The set $\mathscr{U}^K(\Gamma)$ below can be found by shrinking the components of $\mathscr{U}_\varepsilon(\Gamma)$ until they are distance K apart, and then keeping only the tubes and balls which remain that are of radius at least K.

p:props of U^K

Proposition 5.1. *For any $K > 0$ there exists $D > 0$ and a Γ–invariant open set $\mathscr{U}^K(\Gamma) \subseteq \mathscr{U}_\varepsilon(\Gamma)$ so that $\mathscr{U}^K(\Gamma)$ satisfies the following:*

(1) The components of $\mathscr{U}^K(\Gamma)$ are tubes, horoballs and balls;
(2) Distinct components of $\mathscr{U}^K(\Gamma)$ are at least K apart;
eq:lower bound
(3) Suppose $x \in \partial U$ for some connected component of U of $\mathscr{U}^K(\Gamma)$. Then every $g \in \Gamma$ moves x by at least D;
eq:large radius
(4) Tubes and balls in $\mathscr{U}^K(\Gamma)$ have radius at least K; and
(5) Any horoball in $\mathscr{U}_\varepsilon(\Gamma)$, and every tube and ball of radius at least $2K$ is contained in the K–neighborhood of $\mathscr{U}^K(\Gamma)$.

Given points x, y in a tube U around a geodesic γ in $\mathbb{H}^*$, let $\pi_\gamma \colon U \to \gamma$ be the closest point projection and let $d_\gamma(x,y) = d_{\mathbb{H}^*}(\pi_\gamma(x), \pi_\gamma(y))$. Suppose $K > 1$ (a specific value for K is fixed later). Define an "electric" distance d on $\mathbb{H}^*$ by

$$d_{el}(x,y) = \begin{cases} 0 & \text{if } x,y \text{ lie in the same horoball or ball of } \mathscr{U}^K(\Gamma); \\ d_\gamma(x,y) & \text{if } x,y \text{ both lie in a tube in } \mathscr{U}^K(\Gamma); \text{ and} \\ d_{\mathbb{H}^*}(x,y) & \text{otherwise.} \end{cases}$$

Define a pseudo-metric $\widehat{d}$ on $\mathbb{H}^*$ by setting

$$\widehat{d}(x,y) = \inf_{\{x_i\}_{i=0}^n} \left\{ \sum_{i=1}^n d_{el}(x_i, x_{i-1}) \right\},$$

where the infimum is taken over all finite sets $x_0 = x, \ldots, x_{n-1}, x_n = y$.

Denote by $(\mathscr{C}^K(\Gamma), d)$ the induced metric space of $(\mathbb{H}^*, \widehat{d})$, and let $q_\Gamma^K\colon \mathbb{H}^* \to \mathscr{C}^K(\Gamma)$ be the canonical quotient map. Observe $(\mathscr{C}^K(\Gamma), d)$ is a length space. Denote the length of a path p in $\mathscr{C}^K(\Gamma)$ by $\widehat{\ell}(p)$.

Definition 5.2. *The* collapsing locus *of $\mathscr{C}^K(\Gamma)$ is the collection of points $x \in \mathscr{C}^K(\Gamma)$ so that $\left|(q_\Gamma^K)^{-1}(x)\right| > 1$.*

Let γ be a path from x to y in $\mathscr{C}^K(\Gamma)$ so that the intersection of γ with each component of the collapsing locus is connected. Let the components of the intersection of γ with the collapsing locus of $\mathscr{C}^K(\Gamma)$ be $\mathbf{C} = \{C_1, \ldots, C_k\}$. The *collapsed set associated to* γ is $\left\{U_i = (q_\Gamma^K)^{-1}(C_i)\right\}$, a disjoint collection of horoballs, balls and tubes. If C_i does not contain x or y then γ *penetrates* U_i. A component of the collapsed set is a horoball, a ball, or a tube around a geodesic segment. Let $\{\gamma_0, \ldots, \gamma_k\}$ be the components of $\gamma \smallsetminus \mathbf{C}$ (γ_0 and γ_k may be empty). Let $\widetilde{\gamma}_i$ be the closure of $(q_\Gamma^K)^{-1}(\gamma_i)$ in $\mathbb{H}^*$. The *entry point* of γ into U_i is $e(\gamma, U_i) = \widetilde{\gamma}_{i-1} \cap U_i$ and the *exit point* out of U_i is $o(\gamma, U_i) = \widetilde{\gamma}_i \cap U_i$. Let $\widetilde{\gamma}_i^p$ be the geodesic between $e(\gamma, U_i)$ and $o(\gamma, U_i)$. The *lift* of γ is $\{\widetilde{\gamma}_0, \widetilde{\gamma}_1^p, \widetilde{\gamma}_1, \ldots, \widetilde{\gamma}_k^p, \widetilde{\gamma}_k\}$. If each $\widetilde{\gamma}_i$ is geodesic then γ is a *local collapsed geodesic*.

A local collapsed geodesic γ between $x, y \in \mathscr{C}^K(\Gamma)$ is an *almost geodesic* if $\widehat{\ell}(\gamma) \leq \widehat{d}(x,y) + 1$. Any path $[x,y]$ can be replaced by a local collapsed geodesic connecting x, y without increasing its length, so an almost geodesic always exists between any two points in $\mathscr{C}^K(\Gamma)$.

Notation 5.3. *Suppose (X, d_X) is a metric space, and $A, B \subset X$ are closed. If $x \in X$ then $\pi_A(x) = \{y \in A \mid d_X(x,y) = d_X(x,A)\}$ and $\pi_A(B) = \bigcup_{b \in B} \pi_A(b)$.*

Here is a standard fact about hyperbolic spaces.

l:visual size **Lemma 5.4.** *Let (X, d_X) be a ν–hyperbolic metric space and suppose that $W_1, W_2 \subset X$ are convex and that $d(W_1, W_2) > 6\nu$. There exists $x \in X$ with $d_X(x, W_2) \leq 3\nu$ so that $\pi_{W_1}(W_2) \subseteq \pi_{W_1}(B(x, 2\nu))$.*

Bounded projection **Lemma 5.5.** *Suppose that U_1 and U_2 are convex subsets of $\mathbb{H}^*$ that are at distance $t \geq 6\delta_{\mathbb{H}}$ apart. Then* $\mathrm{Diam}\,(\pi_{U_1}(U_2)) \leq 4\delta_{\mathbb{H}} e^{-t+5\delta_{\mathbb{H}}}$.

Proof. This follows from Lemma 5.4 and the fact that for any t' the length of paths outside of $N(U_2, t')$ decreases by a factor of at least $e^{-t'}$ under projection to U_2. □

The following lemma is similar to [15, Lemma 4.5].

l:track geodesics **Lemma 5.6.** *There an absolute constant l with the following property: Let $\widetilde{\gamma}$ be the geodesic between $\widetilde{x}, \widetilde{y} \in \mathbb{H}^*$ and let $\gamma = q_\Gamma^K(\widetilde{\gamma})$. Let β be any almost geodesic in $\mathscr{C}^K(\Gamma)$ between $x = q_\Gamma^K(\widetilde{x})$ and $y = q_\Gamma^K(\widetilde{y})$. Any subsegment of β which lies outside of $N(\gamma, l)$ has length at most $4l + 2$. In particular, any almost geodesic from x to y stays completely inside $N(\gamma, 3l+1)$.*

Proof. We claim any l satisfying the following conditions suffices:

(1) $l \geq 6\delta_{\mathbb{H}}$; and
(2) $[1 - (4\delta_{\mathbb{H}} + 1)e^{-l+5\delta_{\mathbb{H}}}] \geq 1/2 > 0$.

Let β' be a subsegment of β lying completely outside $N(\gamma, l)$, and let z and w be the first point and the last point of β', respectively. Since β' lies outside $N(\gamma, l)$, all collapsed sets penetrated by β' are at least l away from $\widetilde{\gamma}$. Let $\{U_1, \ldots, U_k\}$ be the collapsed set for β', and let

$$\widetilde{\beta}_0, \widetilde{\beta}_1^p, \widetilde{\beta}_1, \ldots, \widetilde{\beta}_k^p, \widetilde{\beta}_k$$

be the lift of β' to $\mathbb{H}^*$. Then $q_\Gamma^K(\widetilde{\beta}_i)$ is subsegment of β' and hence $\widetilde{\beta}_i$ is at least l away from $\widetilde{\gamma}$. Since U_i is at least l away from $\widetilde{\gamma}$ and $l \geq 6\delta_{\mathbb{H}}$, by Lemma 5.5 the projection of $\widetilde{\beta}_i$ to $\widetilde{\gamma}$ has size at most $4\delta_{\mathbb{H}} e^{-l+5\delta_{\mathbb{H}}}$. Let $l' = l - 5\delta_{\mathbb{H}}$.

Let $\widetilde{z}$ and $\widetilde{w}$ be the q_Γ^K–pre-image of z and w, respectively, and let $\widetilde{z}'$ and $\widetilde{w}'$ denote the images of $\widetilde{z}$ and $\widetilde{w}$ under orthogonal projection onto $\widetilde{\gamma}$, respectively. Let $z' = q_\Gamma^K(\widetilde{z}')$ and $w' = q_\Gamma^K(\widetilde{w}')$. Then

$$\begin{aligned}\widehat{\ell}(\beta') &\leq d(z,z') + d(z',w') + d(w',w) + 1 \\ &\leq l + d(\widetilde{z}',\widetilde{w}') + l + 1 \\ &\leq l + \widehat{\ell}(\beta')\cdot e^{-l'} + 4\delta_{\mathbb{H}}\cdot k \cdot e^{-l'} + l + 1\end{aligned}$$

Components are at least K apart (in $\mathbb{H}^*$–distance) with $K > 1$, so $k+1 \leq \widehat{\ell}(\beta')$, and

$$\widehat{\ell}(\beta') \leq 2l + \widehat{\ell}(\beta')\cdot e^{-l'} + 4\delta_{\mathbb{H}}\cdot(\widehat{\ell}(\beta') - 1)\cdot e^{-l'} + 1$$

Therefore

$$\widehat{\ell}(\beta')\cdot[1 - (4\delta_{\mathbb{H}} + 1)e^{-l'}] \leq 2l - 4\delta_{\mathbb{H}}\cdot e^{-l'} + 1$$

We chose l so that $[1 - (4\delta_{\mathbb{H}}+1)e^{-l'}] \geq 1/2 > 0$, so

$$\widehat{\ell}(\beta') \leq 2(2l - 4\delta_{\mathbb{H}}\cdot e^{-l'} + 1) \leq 4l + 2,$$

as required. □

p: geodesic

Proposition 5.7. *If $K \geq 8l$ then $(\mathscr{C}(\Gamma), d)$ is a geodesic metric space.*

Proof. Since $(\mathscr{C}(\Gamma), d)$ is a length space, for any $x, y \in \mathscr{C}(\Gamma)$ there is a sequence of almost geodesics p_i from x to y such that $\widehat{\ell}(p_i) \to d(x,y)$. Choose lifts $\widetilde{x}$ and $\widetilde{y}$ of x and y to $\mathbb{H}^*$, let $\widetilde{\gamma}$ denote the $\mathbb{H}^*$–geodesic from $\widetilde{x}$ to $\widetilde{y}$ and let $\gamma = q_\Gamma(\widetilde{\gamma})$. By Lemmas 5.6, each of the p_i lies within $3l+1$ of γ. Therefore, there is a finite collection of components of the collapsing locus which any of the p_i intersect, and passing to a subsequence we may assume that the p_i intersect exactly the same collection of components of the collapsing locus, say $C_1, \ldots, C_n$.

The space $\mathscr{C}(\Gamma)$ is locally compact except at points in the collapsing locus corresponding to the q_Γ–image of horoballs in $\mathscr{U}^K(\Gamma)$. Suppose that C_i is a component corresponding to a horoball U_i. Then an almost geodesic travelling through C_i is locally the concatenation of two paths, each of which lift to paths in $\mathbb{H}^*$ which travel almost as quickly as possible towards U_i. In particular, if a_i is a point on p_i at distance $4l$ from C_i, then $d(a_i, C_i) > 3l+1$. The path from a_i to γ of length at most $3l+1$ cannot intersect C_i, and since $K \geq 8l$ it cannot intersect any other component of the collapsing locus either. It follows that such a_i lift to a bounded region of $\mathbb{H}^*$, and hence all of p_i lifts to a bounded region of $\mathbb{H}^*$. Since $\mathbb{H}^*$ is locally compact, (p_i) sub-converges to a geodesic p from x to y. □

The hyperbolicity (along with calculation of constants) of $(\mathscr{C}^K(\Gamma), d)$ follows from Lemma 5.6 in the same way that [15, Proposition 4.6] follows from [15, Lemma 4.5]. In the following, l is as in Lemma 5.6.

t:hyperbolicity

Theorem 5.8. *$(\mathscr{C}^K(\Gamma), d)$ has $(6l + \delta_{\mathbb{H}} + 2)$–slim triangles. Therefore, $(\mathscr{C}^K(\Gamma), d)$, is δ–hyperbolic, where $\delta = 4\,(6l + \delta_{\mathbb{H}} + 2)$.*

not:Fix K

Notation 5.9. *Fix $K = 40\delta$, and let l be the constant from Lemma 5.6.*

We write $\mathscr{C}(\Gamma) = \mathscr{C}^K(\Gamma)$, and drop the superscript K from other constructions also. Therefore, for example, we have the map $q_\Gamma \colon X \to \mathscr{C}(\Gamma)$, and when Γ is assumed we just write $q \colon X \to \mathscr{C}$, etc.

l:Hausdorff distance

Lemma 5.10. *Let $x,y \in \mathscr{C}(\Gamma)$ and let $\gamma_2 = [x,y]$ be a $\mathscr{C}(\Gamma)$–geodesic. Let $\widetilde{x},\widetilde{y} \in \mathbb{H}^*$ be lifts of x,y, respectively. Let $\widetilde{\gamma}_1 = [\widetilde{x},\widetilde{y}]$ be the $\mathbb{H}^*$–geodesic and $\gamma_1 = q_\Gamma(\widetilde{\gamma}_1)$. The Hausdorff distance between γ_1 and γ_2 is at most 9δ.*

Proof. By Lemma 5.6 $\gamma_2 \subseteq N(\gamma_1, 3l+1)$, and $3l+1 \leq 9\delta$. So it remains to prove that $\gamma_1 \in N(\gamma_2, 9\delta)$.

First note that any segment of γ_1 which lies outside the $4l$–neighborhood of the collapsing locus is a geodesic. For, if not the geodesic must intersect the collapsing locus. But the geodesic cannot cannot intersect the collapsing locus, since by Lemma 5.6 the geodesic remains within distance $3l+1$ of γ_1.

Given a convex set $W \in \mathbb{H}^*$ and a geodesic γ, the distance from $\gamma(t)$ to W is a strictly convex function, so for any r there are at most two points on γ lying at distance exactly r from W.

Now, let $C_1, \ldots, C_k$ denote the components of the collapsing locus that lie within 6δ of γ_2, and suppose that $\gamma_1 = p_1 \cdot p_2 \cdots p_j$ (where $j \in \{2k-1, 2k, 2k+1\}$) is so that the endpoints of each p_i lie at distance exactly 12δ from C_i. Thus, the p_i alternate between paths which lie outside the 12δ–neighborhood of the collapsing locus, and within 12δ of some C_l. Suppose that p_i lies within 12δ of C_l, and let q be a $\mathscr{C}(\Gamma)$-geodesic with the same endpoints as p_i. By lifting p_i to $\mathbb{H}^*$ and considering projection to $q_\Gamma^{-1}(C_l)$, we see that p_i is within $2\delta_{\mathbb{H}} \leq \delta$ of the projection of the endpoints of p_i to C_l together with (possibly) a geodesic through C_l. By considering the projection of the endpoints to C_l, and slim quadrilaterals, it is straightforward to check that the Hausdorff distance between q and this union of projections to C_l and a geodesic through C_l is at most δ (since q intersect C_l or $p_i = q$). Thus, the Hausdorff distance between p_i and q is at most 2δ. Thus, if we replace each of the p_i which lies in the 12δ–neighborhood of some C_l by a $\mathscr{C}(\Gamma)$–geodesic, then we obtain a concatenation of $\mathscr{C}(\Gamma)$–geodesics which lies at Hausdorff distance at most 3δ from γ_2.

It is straightforward to check (using projections to the collapsing locus) that the Gromov product at each point of concatenation is at most 2δ, and that each of the paths in this concatenation, except possibly the first and last, have length greater than 12δ (since $K \geq 40\delta > 24\delta$). Therefore, by [1, Lemma 4.9] (with $l = 2\delta$), the Hausdorff distance between this concatenation and γ_2 is at most 7δ, and the result follows. □

The following is similar to [15, Lemma 4.8].

l:generalized Farb's 4.8

Lemma 5.11. *There exists D_1 satisfying the following: Let $x,y \in \mathscr{C}(\Gamma)$ and let $\gamma_2 = [x,y]$ be $\mathscr{C}(\Gamma)$–geodesic and $\widetilde{\gamma}_2$ be its lift. Let $\widetilde{x},\widetilde{y} \in \mathbb{H}^*$ be lifts of x,y, respectively. Let $\widetilde{\gamma}_1 = [\widetilde{x},\widetilde{y}]$ and $\gamma_1 = q_\Gamma(\widetilde{\gamma}_1)$. If precisely one of $\{\gamma_1,\gamma_2\}$ penetrates a component U of the collapsed set then $d_{\mathbb{H}^*}(e(\gamma_i,U), o(\gamma_i,U)) \leq D_1$.*

Proof. Suppose γ_2 penetrates U and γ_1 does not. Let $\widetilde{z}$ be the point on $\widetilde{\gamma}_2$ between $\widetilde{x}$ and $e(\gamma_2,U)$ and $10l$ away from U, or $\widetilde{z} = \widetilde{x}$ if no such point exists. Let $\widetilde{w}$ be the point on $\widetilde{\gamma}_2$ between $\widetilde{y}$ and $o(\gamma_2,U)$ and $10l$ away from U, or $\widetilde{w} = \widetilde{y}$ if no such point exists. Let $z = q_\Gamma(\widetilde{z})$ and $w = q_\Gamma(\widetilde{w})$. By Lemma 5.6, there exist z', w' on γ_1 with $d(z,z'), d(w,w') \leq 5l$. Let $\widetilde{z}'$ and $\widetilde{w}'$ be the lifts of z' and w', respectively. By construction $[z',z]$ and $[w',w]$ are disjoint from the collapsing locus. Hence $d_{\mathbb{H}^*}(\widetilde{z},\widetilde{z}'), d_{\mathbb{H}^*}(\widetilde{w},\widetilde{w}') \leq 5l$. Since these paths are in $\mathbb{H}^*$, $[\widetilde{z}',\widetilde{w}'] \subseteq N_{5l}([\widetilde{z},\widetilde{w}])$. Thus, $q_\Gamma([\widetilde{z}',\widetilde{w}']) \subseteq N_{5l}([z,w])$, since $[z,w] = q_\Gamma([\widetilde{z},\widetilde{w}])$. By

Lemma 5.6, the subsegment β of γ_1 between z' and w' lies in $N_{3l+1}(q_\Gamma([\widetilde{z}', \widetilde{w}']))$, so $\beta \subseteq N_{8l+1}([z,w]) \subseteq N_{18l+1}(q_\Gamma(U))$. Components of the collapsing locus are $40l$–separated, so β does not penetrate any component of the collapsed set except U. On the other hand, since β is a subsegment of γ_1, it does not penetrate U. Therefore the lift of β is $[\widetilde{z}', \widetilde{w}']$ and $[\widetilde{z}', \widetilde{w}']$ is disjoint from U, so the projection of $[\widetilde{z}', \widetilde{w}']$ onto U is bounded by a uniform constant D. Since $[z,z']$ and $[w,w']$ do not penetrate any component of the collapsing locus, the projections of $[\widetilde{z}, \widetilde{z}']$ and $[\widetilde{w}, \widetilde{w}']$ onto U are also bounded by D. By choice of $\widetilde{z}$ and $\widetilde{w}$, $d_{\mathbb{H}^*}(\pi_U(\widetilde{z}), e(\gamma_2, U))$ and $d_{\mathbb{H}^*}(\pi_U(\widetilde{w}), o(\gamma_2, U))$ are bounded by a uniform constant C_2. The above distance bounds imply $d_{\mathbb{H}^*}(o(\gamma_2,U), e(\gamma_2,U)) \leq 3C_1 + 2C_2$.

The case where γ_1 penetrates U and γ_2 does not is similar, the only difference being that instead of using Lemma 5.6, we use Lemma 5.10. As a result, the constants are different but they are again uniform constants. □

The following is similar to [15, Lemma 4.9].

l:generalized Farb's 4.9 **Lemma 5.12.** *There exists D_2 satisfying: Let $x, y, \widetilde{x}, \widetilde{y}, \gamma_1$ and γ_2 be as in Lemma 5.11. If both γ_1 and γ_2 penetrate a component U of the collapsed set then $d_{\mathbb{H}^*}(e(\gamma_1,U), e(\gamma_2,U)) \leq D_2$. The same is true for exit points.*

Proof. Let $\widetilde{z}$ be the point on $\widetilde{\gamma}_2$ between $\widetilde{x}$ and $e(\gamma_2,U)$ and $10l$ away from U, or $\widetilde{z} = \widetilde{x}$ if no such point exists. The distance between $\widetilde{z}$ and $e(\gamma_2,U)$ is bounded by a constant depending only on l. By Lemma 5.6, there exists w on γ_1, so that the collapsed distance between w and $z = q_\Gamma(\widetilde{z})$ is at most $5l$. There is no other connected component of the collapsing locus in the $20l$–neighborhood of $q_\Gamma(U)$, so the geodesic $[z,w]$ does not penetrate any component of the collapsing locus. Let $\widetilde{w} = q_\Gamma^{-1}(w)$. The geodesic $[\widetilde{z}, \widetilde{w}]$ has length at most $5l$. The distance between $\widetilde{w}$ and $e(\gamma_2,U)$ is bounded by a constant depending only on l. By the triangle inequality, the distance between $e(\gamma_1,U)$ and $e(\gamma_2,U)$ is bounded by a constant depending only on l. □

Definition 5.13. *Let $S \subseteq \mathbb{H}^*$, and let γ be a geodesic in $\mathbb{H}^*$ so $\gamma \cap S = \emptyset$. Let T be the set of points $s \in S$ so that there exists some t for which $[\gamma(t), s] \cap S = s$. The* visual size *of S with respect to γ is the diameter of T. The* visual size *of the S is the supremum of the visual size of S with respect to γ, where the supremum is taken over all geodesics γ disjoint from S.*

The following is similar to [15, Lemma 4.4].

bounded visual size **Lemma 5.14.** *There exists D_3 such that for each connected component C of the collapsing locus, the visual size of $q_\Gamma^{-1}(C)$ is less than D_3.*

Proof. The case where $U = q_\Gamma^{-1}(C)$ is a horoball is proved in [15, Lemma 4.4]. The other cases are similar. In the last step, use that fact that (by Proposition 5.1.(4)) the radius of U is at least K, which is much bigger than the hyperbolicity constant of $\mathbb{H}^*$. □

equal trans len for para **Lemma 5.15.** *There exist uniform constants B_1 and B_2 satisfying: Let $g \in \Gamma \smallsetminus \{1\}$ and $x \in \mathscr{C}(\Gamma)$. If g is parabolic, let γ be the component of the collapsing locus fixed by g. Otherwise let γ be the q_Γ-image of the minimal invariant set of g in $\mathbb{H}^*$. Then $d(x, gx) - 2d(x,\gamma) \geq B_1$ and if g is elliptic or parabolic then $d(x,gx) - 2d(x,\gamma) \leq B_2$.*

Proof. Let $\widetilde{x} \in q_\Gamma^{-1}(x)$. Let $U_1 = q_\Gamma^{-1}(\gamma)$ and $\widetilde{y} = \pi_{U_1}(\widetilde{x})$. Let U_2 be a horoball (or ball or tube) with the same center (or axis) as U_1 so g moves points on ∂U_2 by $5\delta + l_g$, where l_g is the translation length of g on $\mathbb{H}^*$. Let $\widetilde{w} = \pi_{U_2}(\widetilde{x})$. We claim that $d(\widetilde{w}, U_1) \leq D_1$ for some absolute constant D_1. The claim is trivial if $U_2 \subset U_1$, so suppose $U_1 \subset U_2$.

First suppose that γ is a component of the collapsing locus. In this case by Proposition 5.1.(3) $d(g\widetilde{y},\widetilde{y}) \geq \max\{D, l_g\}$, so the length of $[\widetilde{w},\widetilde{y}]$ is bounded above by an absolute constant.

Now suppose γ is *not* a component of the collapsing locus. In this case, g is either loxodromic and l_g is bounded from below by an absolute positive constant, or g is elliptic and the angle of rotation is bounded below by an absolute positive constant, so U_2 has uniformly bounded size (tube radius).

The claim gives an absolute upper bound on $d(w,\gamma)$, where $w = q_\Gamma(\widetilde{w})$. By hyperbolicity of $\mathbb{H}^*$, up to an error of $4\delta_{\mathbb{H}}$, $[\widetilde{x}, g\cdot\widetilde{x}]$ tracks $[\widetilde{x},\widetilde{w}] \cup [\widetilde{w}, g\cdot\widetilde{w}] \cup [g\widetilde{w}, g\cdot\widetilde{x}]$. By Lemma 5.6, $[x,gx]$, $[x,w]$ and $[gx,gw]$ are in the δ-neighborhoods of $q_\Gamma([\widetilde{x}, g\cdot\widetilde{x}])$, $q_\Gamma([\widetilde{x},\widetilde{w}])$ and $q_\Gamma([g\cdot\widetilde{x}, g\cdot\widetilde{w}])$, respectively. Hence, up to an error depending on δ, $[x, g\cdot x]$ tracks $[x,w] \cup [w, g\cdot w] \cup [g\cdot w, g\cdot x]$. Thus, there is a constant L_1 depending only on δ so that $d(x, g\cdot x) > 2d(x,w) - L_1$. If g is elliptic or parabolic, then $d(w, g\cdot w) \leq 5\delta$ and hence $|d(x, g\cdot x) = 2d(x,w)|$ is bounded by a constant depending only on δ. Combining this with the fact that $d(w,\gamma)$ has an absolute upper bound, the lemma follows. $\square$

least shift **Lemma 5.16.** *Suppose $g \in \Gamma$ acts loxodromically on $\mathbb{H}^*$ with axis $\widetilde{\gamma}$, and let $\gamma = q(\widetilde{\gamma})$. If γ comes within 5δ of the collapsing locus but is not contained in the collapsing locus then g moves every point of γ by at least 30δ.*

Proof. By assumption, there exists x on γ within 5δ of a component C of the collapsing locus and γ is not contained in C. Hence C and $g\cdot C$ are distinct components of the collapsing locus, which are separated by $K = 40\delta$. Hence x and gx are separated by a distance at least $40\delta - 2(5\delta) = 30\delta$. So the translation length of g on $\widetilde{\gamma}$ is at least 30δ. Let y be a point on γ. Suppose $d(y,gy) < 30\delta$. Since the translation length of g on $\widetilde{\gamma}$ is at least 30δ, $[y,gy]$ intersects some connected component C' of the collapsing locus and we have $d(y,C') + d(C',gy) \leq 30\delta$. Hence $d(gy,gC') + d(C',gy) \leq 30\delta$, so $d(C',gC') \leq 30\delta$. Thus $gC' = C'$, so C' is a line and $\gamma = C'$. This is a contradiction. $\square$

s:non-divergent

6. Sequences which are not $\mathscr{C}$–divergent

Throughout this section make Standing Assumption 4.18. We apply the results in the previous section to the vertex groups of the geometric decomposition of a group in $\mathscr{M}^{\pi}_{\mathrm{Gen}}$. If Γ_v is such a vertex group, and it is hyperbolic, then it admits a hyperbolic structure, unique by Mostow–Prasad Rigidity, and this exhibits Γ_v as a Kleinian group. If Γ_v is an LSF-type vertex group, let $\overline{\Gamma_v}$ be the quotient (hyperbolic) 2–orbifold group, and fix a hyperbolic structure on the orbifold, witnessing $\overline{\Gamma_v}$ as a Fuchsian group. There are many such hyperbolic structures, but any suffices for our purposes.

The purpose of this section is to prove the following result, required for the proof of Theorem 2.7. The proof of Theorem 6.2 is contingent on the technical result Theorem 6.7, proved in Appendix A.

Recall from Assumption 4.18 that A_i is the ω–approximation to the good relative generating set A of the vertex group V of the refined GGD $\mathbb{G}^r$ of L. By Definition 4.10 ω–almost surely there are vertices v_i of T_i so A_i fixes v_i.

d:P-divergent **Definition 6.1** ($\mathscr{C}$–divergent)**.** *Let Γ_{v_i} be the vertex group of Γ_i associated to v_i and let $\mathscr{C}(\Gamma_{v_i})$ be the associated collapsed space. We set*

$$\|\phi_i\|_{\mathscr{C},v} = \inf_{x\in\mathscr{C}(\Gamma_{v_i})} \max_{s\in A_i} d_i\,(s\cdot x, x)\,,$$

which is defined ω–almost surely. The sequence (ϕ_i) is $\mathscr{C}$–divergent with respect to $\mathbb{G}$ *if $\lim^{\omega} \|\phi_i\|_{\mathscr{C},v} = \infty$ for some vertex v of $\mathbb{G}^r$.*

thmt@@nondivtheorem@data

Theorem 6.2. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a non-$\mathscr{T}$–divergent sequence (ϕ_i). If there is a GGD $\mathbb{G}$ of L with respect to which (ϕ_i) is not $\mathscr{C}$–divergent then all stably parabolic subgroups of L are finitely generated.*

When (ϕ_i) is not $\mathscr{C}$–divergent with respect to $\mathbb{G}$, we obtain the following crucial information about the vertex group V of $\mathbb{G}^r$:

p:cut point splitting

Proposition 6.3. *Suppose (ϕ_i) is not $\mathscr{C}$–divergent with respect to $\mathbb{G}$. Then V admits a splitting $\mathbb{D}$ so:*

eq:star

eq:rank 2

eq:non-central

(1) The underlying graph of $\mathbb{D}$ is a star, i.e. a bipartite graph where one type of vertices has a single vertex.
(2) $Z^{\omega}(V)$ is contained in each edge group of $\mathbb{D}$.
(3) For each the edge group H of $\mathbb{D}$, $H/Z^{\omega}(V)$ has rank at most two.
(4) The non-central vertices of $\mathbb{D}$ are in bijection with the edges in $\mathbb{G}^r$ adjacent to V. This bijection induces isomorphisms of associated groups.

Furthermore, if V is of hyperbolic type, and $P \in \mathscr{H}_V$ is not conjugate into an edge group of $\mathbb{G}^r$ adjacent to V, then P is finitely generated.

We prove Proposition 6.3 after we use it to prove Theorem 6.2.

6.1. **The proof of Theorem 6.2.** As a consequence of Proposition 6.3, each vertex group of $\mathbb{G}$ also enjoys the properties (1)–(4) from the conclusion of Proposition 6.3 enjoyed by V:

p:cut point splitting-c

Corollary 6.4. *Suppose (ϕ_i) is not $\mathscr{C}$–divergent with respect to $\mathbb{G}$. Then any vertex group $\mathbb{G}_v$ of $\mathbb{G}$ admits a splitting $\mathbb{D}(v)$ so that properties of Proposition 6.3 hold when $\mathbb{D}$ and $\mathbb{G}^r$ are replaced by $\mathbb{D}(v)$ and $\mathbb{G}$, respectively.*

Proof. Let $\mathbb{R}_v$ be the sub-splitting of $\mathbb{G}^r$ corresponding to $\mathbb{G}_v$. Let $V_1,\ldots,V_t$ be the vertex groups of $\mathbb{R}_v$ and let $\mathbb{D}_1,\ldots,\mathbb{D}_t$ be the corresponding splittings given by Proposition 6.3, respectively. Refined $\mathbb{R}_v$ by $\mathbb{D}_1,\ldots,\mathbb{D}_t$ and collapse the sub-graph of group of this refinement spanned by the central vertex groups of $\mathbb{D}_1,\ldots,\mathbb{D}_t$. The resulting splitting is the required splitting $\mathbb{D}(v)$. □

thmt@@defedgetwist@data

Definition 6.5. *A graph of groups $\mathbb{E}$ is* edge-twisted *if:*

The underlying graph of $\mathbb{E}$ is bipartite with colors A and B. Type A vertices have valence 2, and abelian vertex groups (thus the edge groups of $\mathbb{E}$ are also abelian). Let W be a Type A vertex group of $\mathbb{E}$ and let E_1 and E_2 be the images in W of the adjacent edge groups. There are subgroups $K_j \leq E_j$ (for $j=1,2$) so that

(1) $K_1 \cap K_2 = \{1\}$; and
(2) For $j=1,2$, the group E_j/K_j is finitely generated.

Refine $\mathbb{G}$ by replacing each vertex group $\mathbb{G}_v$ by the splitting given by Corollary 6.4. Collapse the edges in the resulting splitting corresponding to the edges in $\mathbb{G}$, and denote the resulting splitting of L by $\mathbb{K}$.

t:tree output

Proposition 6.6. *$\mathbb{K}$ is an edge-twisted splitting.*

Proof. The vertices of $\mathbb{K}$ corresponding to the central vertices of the splittings given by Corollary 6.4 are of Type B and the vertices corresponding to the non-central vertices are of Type A. That Type A vertex groups are abelian follows from the last property of Corollary 6.4 and the fact that edge groups of $\mathbb{G}^r$ are abelian. Type A vertex groups clearly have valence 2. By Lemma 3.20 and Definition 4.10(2) $Z^{\omega}(W_1) \cap Z^{\omega}(W_2) = \{1\}$ for adjacent vertex groups W_1, W_2 of $\mathbb{G}$. For the K_i from Definition 6.5 we choose the stable center of the vertex groups of $\mathbb{G}$. With this choice, the first condition of Definition 6.5 is satisfied. The last condition of Definition 6.5 is satisfied because of the third property of Corollary 6.4. ☐

thmt@@edgetwistfg@data

The following Theorem 6.7 is proved in Appendix A.

thmt@@edgetwistfg

Theorem 6.7. *Let $\mathbb{E}$ be a finite edge-twisted graph of groups so that $\pi_1(\mathbb{E})$ is finitely generated. The Type B vertex groups of $\mathbb{E}$ are finitely generated.*

Proof of Theorem 6.2. Let $\mathbb{K}$ be as in Proposition 6.6. By Proposition 6.6 and Theorem 6.7, Type B vertex groups of $\mathbb{K}$ are finitely generated. Thus, each Type B vertex group which is a subgroup of an LSF-type vertex groups of $\mathbb{G}$ is an $\mathscr{M}^{\pi}_{\mathrm{LSF}}$–limit group. By Corollary 3.16, all abelian subgroups of these vertex groups are finitely generated, so edge groups adjacent to these vertex groups are finitely generated. The other edge groups are adjacent to a hyperbolic type vertex group of $\mathbb{K}$ and are finitely generated by the third property of Corollary 6.4, so all vertex groups of $\mathbb{K}$ are finitely generated also. Hence all vertex groups of $\mathbb{G}$ are finitely generated. A stably parabolic subgroup is finitely generated by Corollary 3.16 if it is in an LSF-type vertex group or by Corollary 6.4 if it is in a hyperbolic type vertex group. ☐

The rest of the section is devoted to the proof of Proposition 6.3.

6.2. **The cut point tree.** The defining sequence for L is not $\mathscr{C}$–divergent with respect to $\mathbb{G}$. We consider how V maps into collapsed spaces of the vertex groups of the geometric decompositions of the Γ_i.

Now, ω–almost surely $A_i \subset \Gamma_{v_i}$. For ω–almost every i, fix $\mathfrak{o}_i^v \in \mathscr{C}(\Gamma_{v_i})$ satisfying

$$\max_{g \in A_i} \{d_i(g.\mathfrak{o}_i^v, \mathfrak{o}_i^v)\} \leq \|\phi_i\|_{\mathscr{C},v} + \frac{1}{i}.$$

By Definition 4.14 V is generated by A and the adjacent edge groups. Let S_v be a (possibly infinite) generating set for V consisting of A and elements from the adjacent edges groups. Choose a lift $\widetilde{S_v} \subset G$ of S_v.

Fix $s \in S_v$, and denote by $\widetilde{s}$ the corresponding element in $\widetilde{S_v}$. For each i consider a geodesic $\gamma_{i,s}$ in $\mathscr{C}(\Gamma_{v_i})$ between $\mathfrak{o}_i^v$ and $\phi_i(\widetilde{s}).\mathfrak{o}_i^v$. Since (ϕ_i) is not $\mathscr{C}$–divergent, by Lemma 5.15 the ω-limit of the length of $\gamma_{i,s}$ is finite; denote this limiting length by l_s. Let $P_{i,s}$ denote the set of points on $\gamma_{i,s}$ which map to parabolic collapsed points in $\mathscr{C}(\Gamma_{v_i})$. Since $K \geq 1$, distinct parabolic collapsed points lie at least distance 1 from each other. Hence the size of $P_{i,s}$ is ω–almost surely bounded independent of i. Let n_s be the ω–limit of the size of $P_{i,s}$. The positions of these points on $\gamma_{i,s}$ ω–converge to n_s distinct points on a segment of length l_s. Let γ_s be a segment of length s marked with n_s points whose positions correspond to the ω–limit of the n_s marked points on the $\gamma_{i,s}$.

Let $\mathrm{Cay}(V, S_v)$ be the (possibly locally infinite) Cayley graph of V with respect to S_v. For $g \in V$ and $s \in S_v$ let the length of the edge $[g, gs]$ be l_s and mark the edge with n_s *marked* points. The locations of these marked points are chosen so there is an isometry between $[g, gs]$ and γ_s taking marked points to marked points.

Now define a partial map f_i from $\mathrm{Cay}(V, S_v)$ to $\mathscr{C}(\Gamma_{v_i})$ as follows:

For each $s \in S_v$, let f_i be a continuous map from $[1, s]$ to $\gamma_{i,s}$. By our choice of length for the edges and the number of the marked points on them, for ω–almost all i we can choose f_i so that:

(1) f_i maps the marked points of $[1,s]$ bijectively to the parabolic collapsed points on $\gamma_{i,s}$; and
(2) f_i is bi-Lipschitz on $[1,s]$ with constant 2.

Extend f_i to other edges of $\mathrm{Cay}(V,S_v)$ equivariantly whenever possible: For each $g \in V$, choose an ω–approximation (g_i) of g. For each $g \in V$ so $g_i \in \Gamma_{v_i}$, extend f_i to $[g,gs]$ equivariantly.

Since f_i depends on the choice of (g_i) of $g \in V$ and there is no canonical choice of (g_i), f_i is not well defined on the vertices of $\mathrm{Cay}(V,S_v)$. So formally, each f_i is a map defined on the interior of some edges of $\mathrm{Cay}(V,S_v)$. However, we have the following:

Lemma 6.8. *Let B be a finite set of edges in* $\mathrm{Cay}(V,S_v)$*. For ω–almost all i, f_i is defined on B and can be continuously extended to the closure of B.*

Define an equivalence relation on the set of marked points of $\mathrm{Cay}(V,S_v)$ by setting marked points x and y to be equivalent if ω–almost surely $f_i(x) = f_i(y)$. Let Q be the quotient of $\mathrm{Cay}(V,S_v)$ obtained by identifying equivalent marked points.

This construction is clearly equivariant, so the V–action on $\mathrm{Cay}(V,S_v)$ descends to an V–action on Q by isometries. We continue to refer to images in Q of the marked points in $\mathrm{Cay}(V,S_v)$ as marked points of Q. The following lemma is straightforward.

Lemma 6.9. *The equivalence relation on the marked points does not depend on the choices of (g_i).*

Since the marked points in $\mathrm{Cay}(V,S_v)$ are discrete in each edge, we may consider Q to be a metric graph in the obvious way.

By construction, each point in the complement of the marked points in Q has a unique pre-image in $\mathrm{Cay}(V,S_v)$. Hence f_i induces a partial map from the interiors of the edges of Q to $\mathscr{C}(\Gamma_{v_i})$, which we denote by f_i^Q. Although f_i^Q is only partially defined, we have the following:

aps to collapsing spaces

Lemma 6.10. *Suppose B is a finite collection of edges in Q. For ω–almost all i the map f_i^Q well-defined on B. For any finite path γ in Q, the map f_i^Q is ω–almost surely well-defined on γ. Moreover f_i^Q is ω–*equivariant *in the sense that for any $x \in B$ and any $g \in V$ ω–almost surely $f_i^Q(g.x) = g_i.f_i^Q(x)$.*

Let $\mathscr{C}_{\infty,v}$ be the ω–limit of the collapsed spaces $\mathscr{C}(\Gamma_{v_i})$. The sequence $\{f_i^Q\}$ of partially defined maps induces a map $f_\infty \colon Q \to \mathscr{C}_{\infty,v}$.

l:1-1 map

Lemma 6.11. *The map $f_\infty \colon Q \to \mathscr{C}_{\infty,v}$ is well-defined, continuous and V–equivariant. The map f_∞ restricts to an injective map from the marked points of Q to the set of limits of parabolic collapsed points in $\mathscr{C}_{\infty,v}$. Moreover, for any $p \in \mathscr{C}_{\infty,v}$ which is a limit of parabolic collapsed points, $f_\infty^{-1}(p)$ is either a single marked point in Q or is empty.*

Definition 6.12. *Suppose $H \in \mathscr{H}_V$. If $g \in H$ and (g_i) is an ω–approximation of g, then ω–almost surely g_i is contained in Γ_{v_i} and fixes a parabolic collapsed point $\xi_i(g)$ in $\mathscr{C}(\Gamma_{v_i})$.*

An element $g \in H$ has Property $\mathscr{B}$ *if for some $x \in Q$ there is a path τ in Q between x and $g \cdot x$ so ω–almost surely $\xi_i(g) \notin f_i^Q(\tau)$.*

The point $\xi_i(g)$ ω–almost surely depends only on g and not on the choices of (g_i).

:bounded parabolic image

Lemma 6.13. *Let $H \in \mathscr{H}_V$ be such that every element of H has Property $\mathscr{B}$. Then $H/Z^\omega(V)$ is a free abelian group of rank at most two.*

Proof. Let $g \in H$ and fix an ω–approximation (g_i) of g. By the definition of Property $\mathscr{B}$, there is a path τ in Q from x to $g.x$ such that ω–almost surely $f_i^Q(\tau)$ does not intersect the parabolic collapsed point $\xi_i(g)$ in $\mathscr{C}(\Gamma_{v_i})$. Let $x_i = f_i^Q(x)$. By Lemma 6.10 ω–almost surely $f_i^Q(\tau)$ is a path τ_i in $\mathscr{C}(\Gamma_{v_i})$ connecting x_i and $g_i.x_i$. This path τ_i can be lifted to a path $\widetilde{\tau_i}$ in $\mathbb{H}^*$ connecting $\widetilde{x_i}$ and $g_i.\widetilde{x_i}$, where is $\widetilde{x_i}$ is the pre-image of x_i under the natural projection from $\mathbb{H}^*$ to $\mathscr{C}(\Gamma_{v_i})$. Let $U_i \subset \mathbb{H}^*$ be the maximal horoball in $\mathbb{H}^*$ that is collapsed to make the collapsed point ξ_i. Let $\widetilde{y_i}$ be the closest point projection of $\widetilde{x_i}$ to U_i. Then $g_i.\widetilde{y_i}$ is the closest point projection of $g_i.\widetilde{x_i}$ to A_i. Moreover, the path $\widetilde{\tau_i}$ projects to a path in the boundary of U_i connecting $\widetilde{y_i}$ and $g_i.\widetilde{y_i}$. The length of τ_i is ω–almost surely bounded. Therefore, ω–almost surely there is a bound on the number of components of the pre-image of the collapsing locus which $\widetilde{\tau_i}$ intersects. Then ω–almost surely $\widetilde{\tau_i}$ does not intersect U_i. If C is a component of the pre-image of the collapsing locus other than U_i then by Lemma 5.5 the projection of C to U_i has uniformly bounded diameter. Thus there exists N so that ω–almost surely the projection of $\widetilde{\tau_i}$ to U_i has length at most N. Hence, the translation length of g_i on U_i is ω–almost surely bounded.

Let P_i be the subgroup of Γ_{v_i} fixing $\xi_i(g)$. We may identify each P_i with $\mathbb{Z}$ or $\mathbb{Z}^2$ via a choice of short generators. Suppose that $P_i \cong \mathbb{Z}^2$. Using this identification, we have that ω–almost surely $\phi_i(g_i) \in \mathbb{Z}^2$. The upper bound on the translation length of g_i on U_i imply that there exist $a, b \in \mathbb{Z}$, independent of i, so that $\phi_i(g_i) = (a,b)$ ω–almost surely. Note that (a,b) does not depend on the choice of (g_i).

By Proposition 5.1.(3), as long as $g \notin Z^\omega(V)$ there is a lower bound, independent of i, on the translation length of g_i on U_i. Defining $\phi'(g) = (a,b)$ for each $g \in H$ induces a homomorphism $H \to \mathbb{Z}^2$, and the lower bound on translation length for $g \notin Z^\omega(V)$ implies that this homomorphism has kernel $Z^\omega(V)$. The case where $P_i \cong \mathbb{Z}$ is entirely similar. The proof of the lemma is complete. □

l:not mark point

Lemma 6.14. *Let $H \in \mathscr{H}_V$. Then either $H/Z^\omega(V)$ has rank at most two or H is the stabilizer of a marked point in Q.*

Proof. Let P_i and ξ_i be as above. By Lemma 5.15 and the fact the translation length of ϕ_i with respect to S_v and $\mathfrak{o}_i^v$ is bounded, the distance between ξ_i and $\mathfrak{o}_i^v$ is ω–almost surely bounded independent of i. As a result, $\{\xi_i\}$ defines a point ξ in $\mathscr{C}_{\infty,v}$. The stabilizer of ξ is H. If ξ is the image of a marked point under the map f_∞, then by Lemma 6.11 H is the stabilizer of a marked point of Q. Now suppose ξ is *not* the image of any marked point under the map f_∞. Then ω–almost surely the f_i^Q–image of any segment in Q does not contain ξ_i, so for any $h \in H$, ω–almost surely $f_i^Q(1, h.1)$ does not contain ξ_i. Hence every element of H has Property $\mathscr{B}$. Therefore, by Lemma 6.13, $H/Z^\omega(V)$ has rank at most two. □

By Lemma 6.14, to further understand the elements of $\mathscr{H}_V$, we study the case when they are stabilizers of marked points in Q.

l:not cut point

Lemma 6.15. *Let $H \in \mathscr{H}_V$. Suppose H is the stabilizer of a marked point $z \in Q$ which is not a cut point. Then $H/Z^\omega(V)$ has rank at most two .*

Proof. By Lemma 6.13, it suffices to show that every element of H has Property $\mathscr{B}$. Let $g \in H$ and $x \in Q \smallsetminus \{z\}$. Since z is not a cut point, there is a path τ in $Q \smallsetminus \{z\}$ from x to $g \cdot x$. By Lemma 6.11, $f_\infty(z) \notin f_\infty(\tau)$. Note that $\xi = f_\infty(z)$ is the fixed point of H in $\mathscr{C}_{\infty,v}$. Hence, ω–almost surely $f_i^Q(\tau)$ does not intersect the parabolic collapsed point ξ_i associated to H, so each $g \in H$ has Property $\mathscr{B}$ as required. □

By Lemma 6.15, to understand stabilizers of marked points we may assume that H is the stabilizer of a marked point z which is a cut point of Q. Let $M \subset Q$ be the collection of all marked points that are cut points. Consider the *cut-point tree* T_{cut} associated to (Q,M). The vertices of T_{cut} are:

(1) Points in M; and
(2) The maximum connected subsets of Q not separated by an element of M,

Vertices of the first type are cut points, and vertices of the second type are *blocks*. Edges of T_{cut} correspond to the inclusion of a cut point in a block. Clearly, V acts on T_{cut}.

l:cut point tree edge

Lemma 6.16. *Let $z \in M$ and let B be a block containing z. Let $e = [z \subset B]$ be the corresponding edge of T_{cut}, and let H be the stabilizer in V of e. Then $H/Z^{\omega}(V)$ has rank at most two.*

Proof. Denote the stabilizer of z by H'. By Lemma 6.11, H' fixes $\xi = f_\infty(z) \in \mathscr{C}_{\infty,v}$, a limit point of parabolic points $\xi_i \in \mathscr{C}(\Gamma_{v_i})$. In particular, $H' \in \mathscr{H}_V$, hence $H \in \mathscr{H}_V$. By Lemma 6.13, it suffices to show that H consists entirely of elements satisfying Property $\mathscr{B}$. Let $g \in H$ and $x \in B$. Since g fixes B as a set, we know $g \cdot x \in B$. By definition of blocks, $B \smallsetminus \{z\}$ is connected. Hence there is a path τ connecting x and $g \cdot x$ and τ does not intersect z. Thus ω–almost surely $f_i^Q(\tau)$ does not intersect ξ_i. Therefore, g has Property $\mathscr{B}$. □

We are now ready to prove Proposition 6.3. As we showed above, this is enough to finally complete the proof of Theorem 6.2

Proof of Proposition 6.3. Let $\mathbb{D}_0$ be the splitting of V dual to the action of V on T_{cut}. Note that $Z^{\omega}(V)$ acts on T_{cut} trivially, hence $Z^{\omega}(V)$ is contained in all the edge groups of $\mathbb{D}_0$. The underlying graph of $\mathbb{D}_0$ is bipartite, with one type of vertices corresponding to cut point vertices in T_{cut} and the other type corresponding to block vertices. For each cut point vertex of $\mathbb{D}_0$, fold all the adjacent edges together to one edge, and denote the resulting splitting by $\mathbb{D}'$. Now, each cut point vertex of $\mathbb{D}'$ has one adjacent edge and $\mathbb{D}'$ is still connected, so the underlying graph of $\mathbb{D}'$ is a star.

Collapse the cut point vertices of $\mathbb{D}'$ whose stabilizers are not conjugate into edge groups of $\mathbb{G}^r$ adjacent to V. Add vertices to the resulting graph of groups corresponding to edge groups adjacent to V not corresponding to vertices of $\mathbb{D}'$ (with identical edge and vertex groups). It follows that the resulting graph of groups $\mathbb{D}$ satisfies (4). Note that since $\mathbb{D}'$ satisfies (1) and (2), so does $\mathbb{D}$. Statement (3) follows from Lemmas 6.14, 6.15 and 6.16.

Suppose V is of hyperbolic type and that $P \in \mathscr{H}_V$ is not conjugate into an adjacent edge group. Then $Z^{\omega}(V) = \{1\}$. If P does not stabilize a cut point vertex of T_{cut} then P is finitely generated by Lemma 6.14 and Lemma 6.15. If P does stabilize a cut point vertex of T_{cut}, then P is in a valence-one vertex group of $\mathbb{D}'$ whose adjacent edge group is finitely generated by Lemma 6.16. Refining $\mathbb{G}^r$ by $\mathbb{D}'$ yields a splitting $\mathbb{K}'$ of L. Since P is not contained in the conjugate of any edge group of $\mathbb{G}^r$, the corresponding vertex group in $\mathbb{K}'$ containing P is still valence-one with the same adjacent edge group, so it is finitely generated since L is. In all cases, P is finitely generated and the proof of Proposition 6.3 is complete. □

s:R-trees

7. LIMITS AND $\mathbb{R}$–TREES

Throughout this section make Standing Assumption 4.18. Suppose further that (ϕ_i) is $\mathscr{C}$–divergent with respect to $\mathbb{G}$ (see Definition 6.1) and that the vertex v of $\mathbb{G}^r$ associated to V is one for which $\lim^{\omega} \|\phi_i\|_{\mathscr{C},v} = \infty$. The space $\mathscr{C}_{\infty,v}$ is defined in Definition 7.4 below.

The goal of this section is to prove the following theorem, which is the only result from this section needed in future sections.

t:Rtree summary

Theorem 7.1.

eq:rtree

(1) $\mathscr{C}_{\infty,v}$ is an $\mathbb{R}$–tree equipped with a nontrivial isometric V–action.

eq:para unique fp

(2) Let $H \in \mathscr{H}_V$. There exists $x_H \in \mathscr{C}_{\infty,v}$, fixed by H. Each nontrivial element of $H \smallsetminus Z^{\omega}(V)$ fixes only the point x_H.

eq:ell unique fp

(3) If $g \in V$ and (g_i) is an ω–approximation to g so ω–almost surely g_i is (nontrivial and) elliptic then g fixes a unique point in $\mathscr{C}_{\infty,v}$.

eq:triv trip

(4) The $\overline{V}$–action on $\mathscr{C}_{\infty,v}$ has trivial tripod stabilizers.

eq:abel seg

(5) The $\overline{V}$–action on $\mathscr{C}_{\infty,v}$ has abelian segment stabilizers.

The above theorem, together with the Rips machine, implies that the $\overline{V}$–action on $\mathscr{C}_{\infty,v}$ admits a graph of actions decomposition with simplicial, axial and Seifert type vertex actions, which induces an abelian splitting of $\overline{V}$. The splitting of V (rel $\mathscr{H}_V$) induced by this splitting also has abelian edge groups. [1]

The remainder of this section is dedicated to the proof of Theorem 7.1. Recall the definitions from Section 4, particularly Definitions 4.10 and 4.14.

Let A be the good relative generating set for V from Assumption 4.18 and let (A_i) be the ω–approximation of A. Let $v_i \in T_i$ be the sequence of vertices associated to v. Let $\Gamma_{v_i} \subset \Gamma_i$ be the stabilizer of v_i. Then since $\mathbb{G}$ is a GGD of L with respect to (ϕ_i) ω–almost surely $A_i \subset \Gamma_{v_i}$. Let $\mathscr{C}_i = \mathscr{C}(\Gamma_{v_i})$. For ω–almost every i, fix $\mathfrak{o}_i \in \mathscr{C}_i$ so that

$$\max_{g \in A_i} \{ d_i(g.\mathfrak{o}_i, \mathfrak{o}_i) \} \leq \|\phi_i\|_{\mathscr{C},v} + \frac{1}{i}.$$

Definition 7.2. *Let $\mathscr{C}^0_{\infty,v}$ be the ultra-limit (with respect to ω) of the sequence $\left(\mathscr{C}_i, \frac{1}{\|\phi_i\|_{\mathscr{C},v}} d_i, \mathfrak{o}_i\right)$.*

p:Limit R-tree

Proposition 7.3 (Limiting $\mathbb{R}$–tree)**.** *The space $\mathscr{C}^0_{\infty,v}$ is a (pointed) $\mathbb{R}$–tree equipped with a non-trivial isometric V–action.*

Proof. Most parts of this proposition are standard facts in the theory of ultra-limits and $\mathbb{R}$–trees (see, for example, [17, §4] and the references therein). The only new thing here is to show that $\{\phi_i\}$ induces an isometric action of V on $\mathscr{C}^0_{\infty,v}$. This does not follow from the standard theory since A may not generate V. However, V is generated by $A \cup \mathbb{E}(V)$, where $\mathbb{E}(V)$ are the edge groups of $\mathbb{G}^r$ adjacent to V.

Let $\mathfrak{o} = [\{\mathfrak{o}_i\}] \in \mathscr{C}^0_{\infty,v}$ be the basepoint of $\mathscr{C}^0_{\infty,v}$. By the definition of $\|\phi_i\|_{\mathscr{C},v}$, for all $g \in A_{v,i}$

eq:nondiv

$$(\dagger) \qquad \lim\nolimits^{\omega} \frac{1}{\|\phi_i\|_{\mathscr{C},v}} d_i(\mathfrak{o}_i, g \cdot \mathfrak{o}_i) < \infty.$$

Let $H \in \mathbb{E}(V)$, and suppose $h \in H$ and $g \in A \cap H$ (note that such a g exists by Definition 4.14). Let (h_i) and (g_i) be ω–approximations of h and g, respectively. Then ω–almost surely g_i and h_i are parabolic fixing the same point in $\mathscr{C}_i$. By Lemma 5.15, there exists some constant D so that ω–almost surely $|d_i(\mathfrak{o}_i, h_i\mathfrak{o}_i) - d_i(\mathfrak{o}_i, g_i\mathfrak{o}_i)| < D$. Therefore, by $(\dagger)$

$$\lim\nolimits^{\omega} \frac{1}{\|\phi_i\|_{\mathscr{C},v}} d_i(\mathfrak{o}_i, h_i \cdot \mathfrak{o}_i) < \infty,$$

as required. □

[1]See [19] for the definitions of graphs of actions, simplicial vertex actions, axial vertex actions and vertex actions of Seifert type, as well as a statement of the Rips machine (Theorem 5.1 in that paper).

d:Pinfty **Definition 7.4.** *Let $\mathscr{C}_{\infty,v}$ be the minimal V–invariant subtree of $\mathscr{C}^0_{\infty,v}$.*

Theorem 7.1.(1) follows immediately from Proposition 7.3.

Recall that $\overline{V} = V/Z^\omega(V)$, where $Z^\omega(V)$ is given by Definition 4.9. Each element of $Z^\omega(V)$ ω–almost surely acts trivially on $\mathscr{C}_i$. Thus we have

Corollary 7.5. *The V–action descends to a non-trivial isometric $\overline{V}$–action on $\mathscr{C}^0_{\infty,v}$, with minimal $\overline{V}$–invariant subtree $\mathscr{C}_{\infty,v}$.*

Theorem 7.1.(2) and (3) follow from the next lemma.

l:fixed for ell/par **Lemma 7.6.** *Let $g \in V$ and let (g_i) be an ω–approximation of g. Suppose that ω–almost surely g_i is parabolic (or nontrivial elliptic). Then g fixes a unique point in $\mathscr{C}_{\infty,v}$. If $g^1, g^2 \in V$, (g_i^1) and (g_i^2) are ω–approximations of g^1 and g^2, respectively, and ω–almost surely g_i^1 and g_i^2 are parabolic with the same fixed point then g^1 and g^2 fix the same point in $\mathscr{C}_{\infty,v}$.*

Proof. Let C_i be the component of the collapsing locus fixed by g_i (or the projection in $\mathscr{C}_i$ of the fixed point of the elliptic g_i in $\mathbb{H}^*$). Consider the geodesic triangle with vertices $\mathfrak{o}_i, g_i \cdot \mathfrak{o}_i$ and C_i. Let x_i be the point on $[\mathfrak{o}_i, C_i]$ with $d_i(x_i, C_i) = (\mathfrak{o}_i \mid g \cdot \mathfrak{o}_i)_{C_i}$. Note that $2d_i(x_i, \mathfrak{o}_i) \leq d_i(\mathfrak{o}_i, g_i \cdot \mathfrak{o}_i) + 4\delta$. Hence

$$\lim^\omega \frac{1}{\|\phi_i\|_{\mathscr{C},v}} d_i(\mathfrak{o}_i, x_i) < \infty$$

Hence $\{x_i\}$ defines a point in $\mathscr{C}^0_{\infty,v}$. By hyperbolicity of $\mathscr{C}_i$, $d_i(x_i, g_i \cdot x_i) \leq \delta$. Hence g fixes x. Therefore g fixes a point in $\mathscr{C}_{\infty,v}$. We now show that g fixes exactly one point in $\mathscr{C}_{\infty,v}$. Suppose g fixes $x \in \mathscr{C}_{\infty,v}$ and $(x_i \in \mathscr{C}_i)$ represents x. By Lemma 5.15, $|2d_i(x_i, C_i) - d_i(x_i, g_i x_i)|$ is ω–almost surely bounded independent of i. Since g fixes x,

$$\lim^\omega d_i(x_i, g_i x_i)/\|\phi_i\|_{\mathscr{C},v} = 0$$

so $\lim^\omega d_i(x_i, C_i)/\|\phi_i\|_{\mathscr{C},v} = 0$, and $x = [\{C_i\}] \in \mathscr{C}_{\infty,v}$.

The second assertion follows immediately from the above argument and the assumption that ω–almost surely g_i^1 and g_i^2 fix the same point. □

c:arc stab loxo **Corollary 7.7.** *Suppose $g \in V \smallsetminus Z^\omega(V)$ fixes a non-degenerate arc in $\mathscr{C}_{\infty,v}$. If (g_i) is an ω–approximation of g then ω–almost surely g_i is loxodromic.*

Theorem 7.1.(4) follows from Lemma 7.6 and the following lemma.

l:fixed for loxo **Lemma 7.8.** *Let $g \in V$ and (g_i) be an ω–approximation of g. Suppose that ω–almost surely g_i is loxodromic. The fixed point set of g in $\mathscr{C}_{\infty,v}$ is either empty or a single geodesic. In case it is a single geodesic, it is a limit of images in the collapsed space of geodesic axes in $\mathbb{H}^*$.*

Proof. Suppose the fixed point set of g is not empty. Let $x = [\{x_i\}] \in \mathscr{C}_{\infty,v}$ be a point fixed by g. Let $\mathbb{H}^*$ be the domain of the map $q_i = q_{\Gamma_{v_i}}$. Denote the axis of g_i in $\mathbb{H}^*$ by $\widetilde{\gamma}_i$. Let $\gamma_i = q_i(\widetilde{\gamma}_i)$. Let $\widetilde{x}_i \in \mathbb{H}^*$ be a point in $q_i^{-1}(x_i)$. By Lemma 5.15, there is a constant D so that ω–almost surely $d_i(x_i, g_i x_i) \geq 2d_i(x_i, \gamma_i) - D$. Since x is fixed by g, we have $\lim^\omega \frac{d_i(x_i, g_i x_i)}{\|\phi_i\|_{\mathscr{C},v}} = 0$, and so $\lim^\omega \frac{d_i(x_i, \gamma_i)}{\|\phi_i\|_{\mathscr{C},v}} = 0$. Therefore x is on the limit of $\{\gamma_i\}$, which is a geodesic in $\mathscr{C}_{\infty,v}$. □

The goal of the rest of this section is to prove Theorem 7.1.(5). To that end, fix a non-trivial segment $I = [a, b]$ in $\mathscr{C}_{\infty,v}$, and suppose that $\overline{g}, \overline{h} \in \overline{V}$ stabilize I. Fix lifts $g, h \in V$ of $\overline{g}, \overline{h}$ and let (g_i) and (h_i) be ω–approximations of g and h, respectively. We know that

ω–almost surely g_i and h_i both lie in the same $\Gamma_{v,i}$, and our goal is to show that ω–almost surely $[g_i,h_i] \in Z(\Gamma_{v,i})$. Recall that if $\Gamma_{v,i}$ is of hyperbolic type then $Z(\Gamma_{v,i}) = \{1\}$, and that in either case $Z(\Gamma_{v,i})$ is the kernel of the action of $\Gamma_{v,i}$ on the associated hyperbolic space $\mathbb{H}^*$.

Observe that $[g_i,h_i]$ is an ω–approximation of $[g,h]$. We may suppose $[g_i,h_i]$ is ω–almost surely non-trivial, or else there is nothing to prove. Because $[g,h]$ stabilizes the non-trivial segment I in $\mathscr{C}_{\infty,v}$, by Corollary 7.7 ω–almost surely $[g_i,h_i]$ corresponds to a loxodromic isometry of $\mathbb{H}^*$. Let $\widetilde{\gamma}_i$ be the invariant geodesic for $[g_i,h_i]$ in $\mathbb{H}^*$ and let $\gamma_i = q_i(\widetilde{\gamma}_i)$. By Lemma 7.8, I is a subsegment of the limit of γ_i. As a result, there are a_i, b_i on γ_i such that $a = [a_i]$ and $b = [b_i]$. Let I_i be a geodesic segment between a_i and b_i.

That segment stabilizers are abelian is proved in the following three lemmas.

nbhd of collapsing locus

Lemma 7.9. *Suppose that ω–almost surely γ_i is not entirely contained in the collapsing locus but that it comes within 5δ of the collapsing locus. Then ω–almost surely $[g_i,h_i] \in Z(\Gamma_{v,i})$.*

Proof. Since a_i and b_i are on γ_i, by Lemma 5.6 and the choice of δ, $I_i \subset N_\delta(\gamma_i)$. Thus for $w_i \in I_i$ satisfying $d_i(a_i,w_i), d_i(b_i,w_i) \geq \frac{1}{10} d_i(a_i,b_i)$, we have $d_i([g_i,h_i]\cdot w_i, w_i) \leq 16\delta$ (See [29, Lemma 5.7]). Hence $[g_i,h_i]$ moves points in the middle part of γ_i by less than 30δ, contradicting Lemma 5.16. Hence $[g_i,h_i]$ acts trivially on $\mathbb{H}^*$ ω–almost surely, so ω–almost surely $[g_i,h_i] \in Z(\gamma_{v,i})$, as required. □

nt from collapsing locus

Lemma 7.10. *Suppose that ω–almost surely γ_i does not come within 5δ of the collapsing locus. Then ω–almost surely $[g_i,h_i] \in Z(\Gamma_{v,i})$.*

Proof. We consider an ω–large set of indices i so that g_i and h_i correspond to loxodromic isometries in $\mathbb{H}^*$, as above, and implicitly concentrate only on such indices.

Suppose $[a_i, g_i\cdot a_i]$ penetrates a component W of the collapsing locus. We claim the distance between the entry and exit points of $[a_i, g_i\cdot a_i]$ in W is bounded independent of i. To that end, let $\widetilde{a}_i = q_i^{-1}(a_i)$ and $\widetilde{b}_i = q_i^{-1}(b_i)$. If $q_i([\widetilde{a}_i, g_i\cdot\widetilde{a}_i])$ does not penetrate W, the claim follows from Lemma 5.11. Suppose then that $q_i([\widetilde{a}_i, g_i\cdot\widetilde{a}_i])$ does penetrate W and let a_i' and u_i' be the points of entry and exit of $q_i([\widetilde{a}_i, g_i\cdot\widetilde{a}_i])$ into W, respectively. By Lemma 5.12, it suffices to show $d_{\mathbb{H}^*}(a_i', u_i')$ is bounded independent of i. Let $\pi_{\widetilde{W}}$ be the closest-point projection onto $\widetilde{W} = q_i^{-1}(W)$. By Lemma 5.14, both $d_{\mathbb{H}^*}(a_i', \pi_{\widetilde{W}}(\widetilde{a}_i))$ and $d_{\mathbb{H}}(u_i', \pi_{\widetilde{W}}(g_i\cdot\widetilde{a}_i))$ are bounded independent of i. Suppose $q_i([\widetilde{b}_i, g_i\cdot\widetilde{b}_i])$ penetrates W. Then by Lemma 5.6 $[b_i, g_i\cdot b_i]\cap N_\delta(W) \neq \emptyset$. Since $d(a_i,b_i)$ is much bigger than both $d(b_i, g_i\cdot b_i)$ and $d(a_i, g_i\cdot a_i)$ for large i, by δ–hyperbolicity, $[a_i,b_i]\cap N_{2\delta}(W) \neq \emptyset$. By Lemma 5.6, the subsegment of γ_i between a_i and b_i lies entirely within $N_\delta([a_i,b_i])$. As a result, $\gamma_i \cap N_{3\delta}(W) \neq \emptyset$, contradicting the assumptions of the lemma. Hence $q_i([\widetilde{b}_i, g_i\cdot\widetilde{b}_i])$ does not penetrate W. It follows that the projections of $[\widetilde{a}_i,\widetilde{b}_i]$, $[\widetilde{b}_i, g_i\cdot\widetilde{b}_i]$ and $g_i\cdot[\widetilde{a}_i,\widetilde{b}_i]$ to W have bounded length as none of them intersect W. The above bounds together show that $d_{\mathbb{H}^*}(a_i', u_i')$ is ω–almost surely bounded independent of i.

The number of components penetrated by $[a_i, g_i\cdot a_i]$ is at most $d(a_i, g_i\cdot a_i)$. Hence $d_{\mathbb{H}^*}(\widetilde{a}_i, g_i\cdot\widetilde{a}_i)$ is bounded by $d(a_i, g_i\cdot a_i)D$, where D is independent of i. As a result, the geodesic rectangle $[\widetilde{a}_i, \widetilde{b}_i, g_i\cdot\widetilde{a}_i, g_i\cdot\widetilde{b}_i]$ is arbitrarily thin in the middle of $[\widetilde{a}_i,\widetilde{b}_i]$. Similarly, the same is true for the geodesic rectangle $[\widetilde{a}_i, \widetilde{b}_i, h_i\cdot\widetilde{a}_i, h_i\cdot\widetilde{b}_i]$. Hence ω–almost surely, $[g_i,h_i]$ moves a point in the middle of $[\widetilde{a}_i,\widetilde{b}_i]$ by arbitrarily small amount. If $[g_i,h_i] \notin Z(\Gamma_{v_i})$, then by the definition of collapsed spaces γ_i is in the collapsing locus ω–almost surely. This contradicts the assumption of the lemma. Hence $[g_i,h_i] \in Z(\Gamma_{v_i})$, as required. □

ined in collapsing locus

Lemma 7.11. *If ω–almost surely γ_i is contained in the 5δ–neighborhood of the collapsing locus then ω–almost surely $[g_i,h_i] \in Z(\Gamma_{v,i})$.*

Proof. We first show that γ_i is entirely contained in the collapsing locus of $\mathscr{C}_i$. By the choice of K, if γ_i is entirely contained in the 5δ–neighborhood of the collapsing locus, then it is entirely contained in the 5δ–neighborhood of some connected component C of the collapsing locus. So $q_i^{-1}(C)$ is a tube around some geodesic β in $\mathbb{H}^*$. The q_i–preimage of the 5δ neighborhood of C is the D' neighborhood of β for some constant D'. The only bi-infinite geodesic contained in this neighborhood is β, so $\widetilde{\gamma}_i = \beta$. In particular, $\widetilde{\gamma}_i$ is contained in $q_i^{-1}(C)$, so γ_i is contained in C.

As above, choose subsegments of γ_i denoted $[a_i,b_i]$. Let C_i be the component of the collapsing locus containing γ_i. Suppose $g_i \cdot C_i \neq C_i$. Then $d_{\mathscr{C}_i}(C_i, g_i \cdot C_i) \geq K$. For ω–almost all i, $[a_i,b_i]$ is much longer than both $[a_i, g_i \cdot a_i]$ and $[b_i, g_i \cdot b_i]$. By δ–hyperbolicity of $\mathscr{C}_i$, the middle part of $[a_i,b_i]$ is in a 2δ–neighborhood of $g[a_i,b_i]$. But $K \geq 2\delta$, so C_i intersects the K–neighborhood of $g_i \cdot C_i$, a contradiction. Thus, $g_i \cdot C_i = C_i$, and similarly $h_i \cdot C_i = C_i$. Hence g_i and h_i act as translations along $\widetilde{\gamma}_i$, and ω–almost surely $[g_i,h_i] = 1$, proving the lemma. □

The previous three lemmas finish the proof of Theorem 7.1.(5).

s:JSJ

8. JSJ-DECOMPOSITIONS AND MODULAR AUTOMORPHISMS

We continue to make Standing Assumption 4.18 and use the notation from there.

8.1. **Summary.** In the last section we showed that if v is the vertex associated to V and $\lim^{\omega} \|\phi_i\|_{\mathscr{C},v} = \infty$ then V admits a nontrivial action on an $\mathbb{R}$-tree. By the Rips machine (and the results in the last section) this induces a graph of groups decomposition of V. Associated to any such decomposition is a group of *modular automorphisms* of V (see Definition 8.11). For a group H equipped with a splitting $\mathbb{A}$ and a family of subgroups $\mathscr{H}$, the group of modular automorphisms of H rel $(\mathbb{A},\mathscr{H})$ is denoted $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}}(H)$.

To apply the shortening argument in the next section, we need to use modular automorphisms of V without a priori knowing the particular splitting of V. For this we need a *JSJ–decomposition* of V, which is a graph of groups decomposition which "sees" all possible modular automorphisms of V. The JSJ–decomposition we use was essentially constructed by Guirardel-Levitt in [21]. Our situation does not quite satisfy the assumptions used in the construction in [21], but the same arguments work with only minor changes. These changes are outlined in Appendix B.

Recall $Z^{\omega}(V)$ is the stable center of V and $\overline{V} = V/Z^{\omega}(V)$. Let $\pi_v \colon V \to \overline{V}$ denote the quotient map. If V is a hyperbolic-type vertex group, $Z^{\omega}(V) = \{1\}$ and $\overline{V} = V$ in this case. Let $\mathscr{A}$ be the family of all virtually abelian subgroups of $\overline{V}$, and $\mathscr{H} = \mathscr{H}_{\overline{V}}$ the family of all stably parabolic subgroups of $\overline{V}$ (recall Definition 4.5). When working with a subgroup $H \leq \overline{V}$ or $H \leq V$, we abuse notation and write $\mathscr{A}$ and $\mathscr{H}$ for the families of virtually abelian and stably parabolic subgroups of H, respectively.

We construct a JSJ–decomposition for $\overline{V}$ in Appendix B. The next theorem, which is the main result of this sections, shows that this decomposition has the properties we need in Section 9.

thm:decomps

Theorem 8.1.

(1) $\overline{V}$ has an $(\mathscr{A},\mathscr{H})$–splitting $\overline{\mathbb{J}}(v)$ so that for any $(\mathscr{A},\mathscr{H})$–splitting $\mathbb{A}$ of $\overline{V}$ we have $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}}(\overline{V}) \leq \mathrm{Mod}_{\overline{\mathbb{J}}(v)}^{\mathscr{H}}(\overline{V})$.

(2) Any element of $\mathrm{Mod}^{\mathscr{H}}_{\mathbb{J}(v)}(\overline{V})$ *lifts to an automorphism of* V *which acts as the identity on* $Z^{\omega}(V)$ *and by conjugation on each stably parabolic subgroup.*

def:J^v splitting **Definition 8.2.** *Let* $\mathbb{J}(v)$ *be the lift of* $\overline{\mathbb{J}}(v)$ *to* V*.*

If $\overline{V}$ is finitely generated and K–CSA, the existence of $\overline{\mathbb{J}}(v)$ follows from [21, Theorem 9.14]. In our case $\overline{V}$ is not necessarily K–CSA since $\overline{V}$ can have arbitrarily large finite subgroups. In the next section we show that $\overline{V}$ is *weakly K–CSA* (see Definition 8.4 below), which means that it shares enough of the properties of K–CSA groups that the proof of [21, Theorem 9.14] holds with only minor changes. We explain these changes in Appendix B, see Theorem B.5.

8.2. **Weakly K–CSA groups.**

Definition 8.3. *A group is* K–virtually abelian *if it has an abelian subgroup of index at most* K*.*

thmt@@defwcsa@data

thmt@@defwcsa **Definition 8.4.** *Fix* $K \geq 1$*. A group is* weakly K–CSA *if (i) any element* g *of order greater than* K *is contained in a unique maximal virtually abelian subgroup* $M(g)$*, so that* $M(g)$ *is* K*–virtually abelian and equal to its normalizer, and (ii) every two infinite, virtually abelian subgroups* A *and* B *with* $\langle A,B\rangle$ *not virtually abelian satisfy* $|A\cap B| \leq K$*.*

Note that if a group H is weakly K–CSA and all finite subgroups of H have order $\leq K$, then H is K–CSA.

lem:CSA1 **Lemma 8.5.** *Let* A *and* B *be virtually abelian subgroups of* $\overline{V}$*. Either* $|A\cap B| \leq 2$ *or* $\langle A,B\rangle$ *is* 2*–virtually abelian.*

Proof. If all finitely generated subgroups of a countable group C are 2–virtually abelian, then C is 2–virtually abelian by [21, Lemma 9.6]. Thus we may assume that A and B are finitely generated. Let A_i and B_i be subgroups of Γ_i generated by ω–approximations of fixed finite generating sets of A and B, respectively. There exist vertex groups V_i in the geometric decomposition of Γ_i so that ω–almost surely $A_i, B_i \leq V_i$. Let $\overline{A}_i$ and $\overline{B}_i$ be the images of A_i and B_i in $V_i/Z(V_i)$, respectively. Then $\overline{A}_i$ and $\overline{B_i}$ are either both subgroups of fundamental groups of orientable finite volume hyperbolic 3–manifolds or both subgroups of fundamental groups of orientable hyperbolic 2–orbifolds. Also, A and B are limit groups over the families $\{\overline{A}_i\}$ and $\{\overline{B}_i\}$, respectively.

Assume $|A| \geq 3$ and $|B| \geq 3$, otherwise the result is trivial. Since A and B are finitely generated and virtually abelian, $\overline{A}_i$ and $\overline{B}_i$ are ω–almost surely virtually abelian. Hence each $\overline{A}_i$ is contained in a unique maximal virtually abelian subgroup M_i which is either cyclic or infinite dihedral. Similarly, each $\overline{B}_i$ is contained in a unique maximal virtually abelian subgroup N_i. Either $|M_i \cap N_i| \leq 2$, or $M_i = N_i$. If the first case happens ω–almost surely, then $|A\cap B| \leq 2$, and if the second happens ω–almost surely then $\langle A,B\rangle$ is 2–virtually abelian. □

The following lemma follows easily from the proof of Lemma 8.5.

lem:dihedraltype **Lemma 8.6.** *Let* $H \leq \overline{V}$ *be virtually abelian but not abelian. Then* H *has an abelian subgroup* H^+ *of index two and an element* $\tau \in H \setminus H^+$ *so that for all* $g \in H^+$*,* $\tau^{-1} g \tau = g^{-1}$*.*

lem:CSA2 **Lemma 8.7.** *Let* $g \in \overline{V}$ *be an element of order at least 3. Let*

$$M(g) = \left\langle \{h \in \overline{V} \mid \langle g,h\rangle \text{ is virtually abelian}\} \right\rangle$$

Then $M(g)$ *is* 2*–virtually abelian and* $M(g)$ *is the unique, maximal virtually abelian subgroup of* $\overline{V}$ *containing* g*.*

Proof. By induction and Lemma 8.5, all finitely generated subgroups of $M(g)$ are 2–virtually abelian so $M(g)$ is 2–virtually abelian by [21, Lemma 9.6]. The maximality and uniqueness of $M(g)$ follow. □

lem:CSA3 **Lemma 8.8.** *Let $g \in \overline{V}$ have order at least 3. Then $M(g)$ is equal to its normalizer.*

Proof. Suppose $x \in \overline{V}$ and $x^{-1}M(g)x = M(g)$. Let V_i be as in the proof of Lemma 8.5. Let g_i and x_i be ω–approximations of g and x respectively, and let $\overline{g}_i$ and $\overline{x}_i$ be the images of g_i and x_i in $V_i/Z(V_i)$ respectively.

Then ω–almost surely, $\langle \overline{g}_i, \overline{x}_i^{-1}\overline{g}_i\overline{x}_i \rangle$ is virtually abelian, so it is contained in a maximal virtually abelian subgroup M_i, which is either cyclic or infinite dihedral. Since $\overline{g}_i$ and $\overline{x}_i^{-1}\overline{g}_i\overline{x}_i$ both have order at least 3, they commute. Since M_i is a maximal virtually abelian subgroup of a hyperbolic 3–manifold group or a hyperbolic 2–orbifold group, ω–almost surely $\overline{x}_i \in M_i$, so $x \in M(g)$. □

Lemmas 8.5, 8.7, and 8.8 together imply the following.

prop:CSA **Proposition 8.9.** *$\overline{V}$ is weakly 2–CSA.*

Using Proposition 8.9, Theorem B.5 applies to $\overline{V}$. Hence $\overline{V}$ has an $(\mathscr{A}, \mathscr{H})$–JSJ decomposition, and we let $\overline{\mathbb{J}}(v)$ be the tree of cylinders of this splitting. By Theorem B.5, $\overline{\mathbb{J}}(v)$ is compatible with every $(\mathscr{A}, \mathscr{H})$–splitting of $\overline{V}$.

8.3. **Modular automorphisms.** We define the group of modular automorphisms relative to a particular graph of groups decomposition and a fixed family of subgroups. These definitions are standard except that some care has to be taken to ensure we are able to lift the automorphisms from $\overline{V}$ to V.

Suppose $H = A *_C B$ and $c \in Z_H(C)$. The *Dehn twist by c* is the automorphism of H defined by fixing each $a \in A$ and mapping each $b \in B$ to cbc^{-1}. If $H = A*_C$ and $c \in Z_H(C)$ then the *Dehn twist by c* is the automorphism defined by fixing each $a \in A$ and mapping the stable letter t to tc. If $\mathbb{A}$ is a graph of groups decomposition and e is an edge of $\mathbb{A}$, a *Dehn twist over e* is a Dehn twist in the one edge splitting corresponding to collapsing all edges of $\mathbb{A}$ other than e.

Now suppose H is finitely generated relative to a family of subgroups $\mathscr{H}_0$. Let $\mathbb{A}$ be a splitting of H rel $\mathscr{H}_0$ and let A be an abelian vertex group of $\mathbb{A}$. Let P_A be the minimal direct factor of A which contains the edge groups of $\mathbb{A}$ adjacent to A, the elements of $\mathscr{H}_0$ contained in A, and the torsion of A. The subgroup P_A can be characterized as the intersection of the kernels of all homomorphisms $A \to \mathbb{Z}$ which are identically 0 on the adjacent edge groups and on elements of $\mathscr{H}_0$. We call P_A the *$(\mathbb{A}, \mathscr{H}_0)$–envelope of A*. Write $A = A_0 \oplus P_A$, and note that since A is finitely generated relative to adjacent edge groups and $\mathscr{H}_0$, A_0 is a finitely generated free abelian group. Thus the group of automorphisms of A fixing P_A can be identified with $GL_n(\mathbb{Z})$ for some $n \geq 1$. In particular, this group is generated by transvections: automorphisms sending g to gh for some basis elements g and h of A_0 and fixing all other elements of the basis of A_0, and all elements of P_A.

A *dihedral type* vertex group of $\mathbb{A}$ is a vertex group D which has an abelian subgroup D^+ of index two and $\tau \in D \setminus D^+$ so that for all $g \in D^+$, $\tau^{-1} g \tau = g^{-1}$. The $(\mathbb{A}, \mathscr{H}_0)$–envelope P_{D^+} is defined exactly as for abelian vertex groups. Write $D^+ = D_0 \oplus P_{D^+}$, and suppose g and h are distinct elements of a basis for D_0. A *squared Nielsen transformation rel $(\mathbb{A}, \mathscr{H}_0)$* is an automorphism of D fixing τ and extending the automorphism of D^+ sending g to gh^2 and fixing P_{D^+} and all other basis elements of D_0.

Let $\mathrm{Aut}_s(D)$ be the subgroup of the automorphism group of D generated by squared Nielsen transformations rel $(\mathbb{A},\mathscr{H}_0)$. In previous applications of the shortening argument, one is allowed to pick α from the whole automorphism group of D. In that case Lemma 8.10 below follows from the Euclidean algorithm. Our situation, where fewer shortening automorphisms are allowed, requires only slightly more care.

l:snt is enough

Lemma 8.10. *Suppose D acts on the real line T with indiscrete orbit and $g_1,\dots,g_k \in D$ project to form a basis of D^+/P_{D^+}. There exists a sequence $\alpha_i \in \mathrm{Aut}_s(D)$ so the translation length of $\alpha_i(g_j)$ goes to zero as $i \to \infty$.*

Suppose $\mathbb{A}$ is a graph of groups decomposition of a group H and v_0 is a vertex of $\mathbb{A}$. Any element of g can be represented (non-uniquely) by $[a_0,e_1,a_1,...,e_n,a_n]$ where $e_1,...,e_n$ is an edge path in $\mathbb{A}$ from v_0 to v_0, $a_0 \in \mathbb{A}_{v_0}$, and each $a_i \in \mathbb{A}_{v_i}$ where v_i is the terminal vertex of e_i for $1 \le i \le n$. Suppose v is a vertex of $\mathbb{A}$ and $\sigma \in \mathrm{Aut}(\mathbb{A}_v)$ acts by conjugation on each adjacent edge group. Then σ can be extended to an automorphism of G as in [48, Definition 4.13] as follows: for each adjacent edge e there is $\gamma_e \in \mathbb{A}_v$ so that $\sigma(h) = \gamma_e h \gamma_e^{-1}$ for all h in the image of $\mathbb{A}_e$ in $\mathbb{A}_v$. Then σ extends to $\overline{\sigma} \in \mathrm{Aut}(H)$ by defining

$$\overline{\sigma}([a_0,e_1,a_1,...,e_n,a_n]) = [\overline{a}_0,e_1,\overline{a}_1,...,e_n,\overline{a}_n]$$

where

$$\overline{a_i} = \begin{cases} a_i & \text{if } a_i \notin \mathbb{A}_v \\ \gamma_{e_i}^{-1}\sigma(a_i)\gamma_{e_{i+1}} & \text{if } a_i \in \mathbb{A}_v. \end{cases}$$

We call $\overline{\sigma}$ the *natural extension* of σ. This natural extension is not unique: it depends on the choice of γ_e.

In the splittings we are interested in, some vertex groups are identified with finite extensions of fundamental groups of hyperbolic orbifolds. These are called *QH–vertex groups*.

defn:modaut

Definition 8.11. *Let $\mathbb{A}$ be a reduced $(\mathscr{A},\mathscr{H})$ splitting of a group H. The* modular automorphism group of H rel $(\mathbb{A},\mathscr{H})$*, denoted* $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}}(H)$*, is the subgroup of* $\mathrm{Aut}(H)$ *generated by:*

(1) Inner automorphisms.
(2) Dehn twists over edges of $\mathbb{A}$ with abelian edge groups.
(3) Natural extensions of automorphisms of QH-vertex groups of $\mathbb{A}$ induced by homeomorphisms of the underlying orbifold fixing the boundary and cone points which act by conjugation on elements of $\mathscr{H}$.
(4) Natural extensions of automorphisms of maximal abelian vertex groups of $\mathbb{A}$ fixing the $(\mathbb{A},\mathscr{H}_0)$–envelope.
(5) Natural extensions of squared Nielsen transformations rel $(\mathbb{A},\mathscr{H}_0)$ of maximal dihedral-type vertex groups.
(6) Natural extensions of automorphisms of vertex groups U where U is a central extension $1 \to Z \to U \to B \to 1$, Z is contained in all edge groups adjacent to U, B is a QH or maximal dihedral-type group and the automorphism fixes Z and projects to a type 3 or type 5 automorphism of B.

We only need automorphisms of type 1-5 when working with the quotient $\overline{V}$, the type 6 automorphisms are used when we lift the splitting of $\overline{V}$ to a splitting of V.

If D is a dihedral type vertex group then squared Nielsen transformations fix P_{D^+} and act by conjugation on each subgroup U with $U \cap D^+ \subseteq P_{D^+}$.

The following is a straightforward consequence of the definition.

l:conjugate parabolic

Lemma 8.12. *If $\sigma \in \mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}}(H)$ and $P \in \mathscr{H}$, then $\sigma(P)$ is conjugate to P.*

Dihedral type groups do not arise as subgroups of hyperbolic type vertex groups. For LSF-type vertex groups, they arise as subgroups which are locally limit groups over the infinite dihedral group D_∞.

By Proposition 8.9 and Theorem B.5, $\overline{V}$ has a graph of groups decomposition $\overline{\mathbb{J}}(v)$ which is compatible with every $(\mathscr{A},\mathscr{H})$–splitting of $\overline{V}$. That is, for any $(\mathscr{A},\mathscr{H})$–splitting $\mathbb{A}$ there exists an $(\mathscr{A},\mathscr{H})$–splitting $\mathbb{B}$ together with collapse maps $\mathbb{B}\to\mathbb{A}$ and $\mathbb{B}\to\overline{\mathbb{J}}(v)$.

Using the properties of $\overline{\mathbb{J}}(v)$ from Corollary B.6, the proof of the following can be completed as in [17, Theorem 5.23]. There are some additional complications in [17, Theorem 5.23] which do not arise here, but the differences between the proofs are almost all notational. Similar arguments also appear in the proof of [48, Proposition 4.17].

lem:p1 **Lemma 8.13.** *For any $(\mathscr{A},\mathscr{H})$–splitting $\mathbb{A}$ of $\overline{V}$, $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}}(\overline{V}) \le \mathrm{Mod}_{\overline{\mathbb{J}}(v)}^{\mathscr{H}}(\overline{V})$.*

Next we show modular automorphisms of $\overline{V}$ lift to V. By Lemma 4.12, the pre-image in V of any abelian subgroup of $\overline{V}$ is abelian. Recall that $\mathbb{J}(v)$ is the lift of $\overline{\mathbb{J}}(v)$ to V.

ting the shortening auto **Lemma 8.14.** *Let $\alpha\in\mathrm{Mod}_{\overline{\mathbb{J}}(v)}^{\mathscr{H}}(\overline{V})$. There exists $\widetilde{\alpha}\in\mathrm{Mod}_{\mathbb{J}(v)}^{\mathscr{H}}(V)$ so that for all $g\in V$ $\pi_v(\widetilde{\alpha}(g))=\alpha(\pi_v(g))$. Moreover, $\widetilde{\alpha}$ acts trivially on the stable center and by conjugation on each stably parabolic subgroup.*

Proof. Suppose α is a Dehn twist by c over an edge e of $\overline{\mathbb{J}}(v)$. Then V has a one-edge splitting $\mathbb{B}$ with edge f such that $\mathbb{B}_f=\pi_v^{-1}(\overline{\mathbb{J}}(v)_e)$, and $\mathbb{B}_f$ is abelian, by Lemma 4.12. Let $\widetilde{c}\in\pi_v^{-1}(c)$. Define $\widetilde{\alpha}$ to be the Dehn twist by $\widetilde{c}$ over f by the element.

Type (3) and (4) automorphisms are generated by Dehn twists and Nielsen transformations respectively, so these can also be lifted to V.

Suppose now that D is a dihedral type vertex group of $\overline{\mathbb{J}}(v)$, let $\widetilde{D}$ be the pre-image of D in V, and let $\widetilde{D}^+$ be the pre-image of D^+. By Lemma 4.12 $\widetilde{D}^+$ is abelian. Let t be a pre-image of τ. For each $x\in\widetilde{D}^+$, let $c_x=t^{-1}xtx$. A straightforward calculation shows that the map $x\to c_x$ is a homomorphism.

Suppose g,h are elements of D^+ with pre-images $\widetilde{g}$ and $\widetilde{h}$ in $\widetilde{D}^+$, and α is a squared Nielsen transformation of D of the form $g\to gh^2$. Let $\widetilde{\alpha}$ be the map which is the identity on t, the stable center of $\widetilde{D}$, P_{D+} and on the pre-image of all basis elements of D^+/P_{D+} except $\widetilde{g}$ and which sends $\widetilde{g}\to\widetilde{g}\widetilde{h}^2c_{\widetilde{h}^{-1}}$. A calculation shows that $c_{\widetilde{\alpha}(\widetilde{g})}=c_{\widetilde{g}}$, so $\widetilde{\alpha}$ is an automorphism of $\widetilde{D}$ inducing α on D. For any element $d\in\widetilde{D}$, $\widetilde{\alpha}(td)$ is conjugate to td. □

Theorem 8.1 follows immediately from Lemmas 8.13 and 8.14.

s:res-short

9. Resolutions and factoring

In this section we complete the proof of Theorem 2.7. Throughout this section we continue to make Standing Assumption 4.18 and use the notation from there.

9.1. **Shortening quotients.** We first use the decompositions $\mathbb{J}(v)$ from Definition 8.2 to refine the decomposition $\mathbb{G}^r$ of L.

def:J splitting **Definition 9.1.** *Since the edge groups of $\mathbb{G}^r$ are elliptic in each $\mathbb{J}(v)$, we can refine $\mathbb{G}^r$ by replacing each vertex v by the splitting $\mathbb{J}(v)$ (see* [21, Lemma 4.12]*). Denote the resulting splitting of L by $\mathbb{J}$.*

As in [48] (see also [17]), we create a sequence of graphs of groups approximating $\mathbb{J}$ so that the ϕ_i factor through the terms in the approximating graphs of groups. This is essentially the same as [48, Lemma 7.1], and the proof from there works in our situation without change. The idea comes from [38].

lem:ugly hack **Lemma 9.2.** *Let $\mathbb{J}$ be the splitting of L from Definition 9.1. There exists a sequence of finitely presented groups $G = W_0, W_1, \ldots$ and epimorphisms $f_i \colon W_i \to W_{i+1}$ and $h_i \colon W_i \to L$ for $i \geq 0$ so that:*

(1) $\phi_\infty = h_0$;
(2) for all $i \geq 1$ we have $h_i = h_{i+1} \circ f_i$;
eq:direct limit *(3) L is the direct limit of the sequence $G \twoheadrightarrow W_1 \twoheadrightarrow \ldots$. Equivalently,*

$$\mathrm{Ker}^{\omega}(\phi_i) = \bigcup_{k=1}^{\infty} \mathrm{Ker}(f_{k-1} \circ \ldots \circ f_0)$$

(4) Each W_i has a graph of groups decomposition $\mathbb{A}_i$ whose underlying graph is isomorphic to the underlying graph of $\mathbb{J}$.
(5) If V is a vertex group of $\mathbb{J}$ and V_i is the corresponding vertex group of $\mathbb{A}_i$, then $h_i(V_i) \subseteq V$. Furthermore,
   (a) $V = \bigcup\limits_{i=1}^{\infty} h_i(V_i)$
   (b) V_i has a central subgroup Z_i which is contained in all edge groups adjacent to V_i such that h_i maps Z_i injectively into $Z^{\omega}(V)$ and $Z^{\omega}(V) = \bigcup\limits_{i=1}^{\infty} h_i(Z_i)$
   (c) If $\overline{V}$ is QH–vertex group, then h_i induces an isomorphism from V_i/Z_i to $\overline{V}$.
   (d) If $\overline{V}$ is a virtually abelian vertex group, then h_i induces an injective map from V_i/Z_i to $\overline{V}$.
(6) If E is an edge group of $\mathbb{J}$ and E_i is the corresponding edge group of $\mathbb{A}_i$, then h_i maps E_i injectively into E and $E = \bigcup_{i=1}^{\infty} h_i(E_i)$.
(7) The edge groups and vertex groups of $\mathbb{A}_i$ are finitely generated.

Fix a vertex v of $\mathbb{G}^r$. There is a collapse map $\mathbb{J} \to \mathbb{G}^r$, such that the pre-image of v is $\mathbb{J}(v)$. For each $\mathbb{A}_i$, the underlying graph is isomorphic to $\mathbb{J}$. Let $\mathbb{A}_{v,i}$ denote the sub-(graph of groups) corresponding to $\mathbb{J}(v)$. The *generalized geometric splitting of W_i corresponding to* $\mathbb{G}^r$ is the splitting of W_i obtained from $\mathbb{A}_i$ by collapsing each $\mathbb{A}_{v,i}$. Let $W_{v,i}$ be the fundamental group of $\mathbb{A}_{v,i}$ (as a subgroup of W_i). It follows from the construction that $f_i|_{W_{v,i}}$ maps into (though possibly not *onto*) $W_{v,i+1}$ and V is the direct limit of the $W_{v,i}$.

We use modular automorphisms associated to these graphs of groups to "shorten" ϕ_i relative to the vertex v, as we now explain.

Let $\xi_i \colon G \to W_i$ be the natural map, that is $\xi_i = f_{i-1} \circ \ldots \circ f_0$. Since W_j is finitely presented, for fixed j, $\mathrm{Ker}(\xi_j) \subseteq \mathrm{Ker}(\phi_i)$ for an ω large set of i. Hence after passing to a subsequence of ϕ_i and re-indexing, we may assume $\mathrm{Ker}(\xi_j) \subseteq \mathrm{Ker}(\phi_i)$ for all j and all $i \geq j$, so for all j and all $i \geq j$ the map ϕ_i factors through ξ_j, so there is $\lambda_i^j \colon W_j \to \Gamma_i$ so

$$\phi_i = \lambda_i^j \circ \xi_j.$$

Associated to v is a sequence $\{v_i\}$ of vertices in the geometric trees of $\{\Gamma_i\}$, and $g \in V$ if and only if for some ω–approximation (g_i) to g ω–almost surely $g_i \in \Gamma_{v_i}$. Since each $W_{v,j}$ is finitely generated, after again passing to a subsequence of $\{\phi_i\}$ and hence $\{\lambda_i^j\}$ and re-indexing we can assume λ_i^j is in fact a morphism of graphs of groups from the generalized geometric splitting of W_j to the geometric splitting of Γ_i. Let $\lambda_i = \lambda_i^i$. Summarizing, we have

lambda_i **Lemma 9.3.**

(1) $\phi_i = \lambda_i \circ \xi_i$;

(2) λ_i is a morphism of graphs of groups from the generalized geometric splitting of W_i corresponding to $\mathbb{G}^r$ to the geometric splitting of Γ_i.

Let $\widetilde{A}_v$ be a lift to G of the good generating set A_v of V and for each i let $\widetilde{A_{v,i}} = \xi_i(\widetilde{A}_v)$. Lemma 9.3 allows us make the following definition.

d:stretch factor

Definition 9.4.

$$\|\lambda_i\|_{\mathscr{C},v} = \inf_{x\in\mathscr{C}(\Gamma_{v_i})} \max_{a\in\widetilde{A_{v,i}}} d_i(\lambda_i(a)\cdot x, x),$$

For ω–almost every i, choose a point $\mathfrak{o}_i \in \mathscr{C}(\Gamma_{v_i})$ which satisfies

$$\max_{a\in\widetilde{A_{v,i}}} \{d_i(\lambda_i(a)\cdot \mathfrak{o}_i, \mathfrak{o}_i)\} \leq \|\lambda_i\|_{\mathscr{C},v} + \frac{1}{i}.$$

Define

$$\|\lambda_i\|_{\mathscr{C},v,\mathfrak{o}} = \max_{a\in\widetilde{A_{v,i}}} \{d_i(\lambda_i(a)\cdot \mathfrak{o}_i, \mathfrak{o}_i)\}$$

Let $\mathscr{H}_i$ denote the family of subgroups H of W_i such that H is elliptic in $\mathbb{A}_i$ and $h_i(H)$ is a stably parabolic subgroup of L. For a vertex v of $\mathbb{G}^r$ let $\mathscr{H}_{v,i}$ denote the family of subgroups of $W_{v,i}$ which belong to $\mathscr{H}_i$. Note that any edge group of the generalized geometric splitting of W_i corresponding to $\mathbb{G}^r$ which belongs to $W_{v,i}$ maps to a stably parabolic subgroup of L and hence belongs to $\mathscr{H}_{v,i}$. By Lemma 8.12 $\mathrm{Mod}_{\mathbb{A}_{v,i}}^{\mathscr{H}_{v,i}}(W_{v,i})$ acts on these subgroups by conjugation, and hence can be extended to automorphisms of W_i. Let $\mathrm{Mod}_v(W_i)$ be the subgroup of $\mathrm{Aut}(W_i)$ generated by these extensions. Denote by $\mathrm{Mod}(W_i)$ the subgroup of $\mathrm{Aut}(W_i)$ generated by all subgroups $\mathrm{Mod}_v(W_i)$ over all vertices v of $\mathbb{G}^r$. Note that $\mathrm{Mod}(W_i)$ acts on subgroups in $\mathscr{H}_i$ by conjugation. We define an equivalence relation on $\mathrm{Hom}(W_i,\Gamma_i)$ by setting $\lambda \sim \lambda'$ if there is an $\alpha \in \mathrm{Mod}(W_i)$ so $\lambda' = \lambda \circ \alpha$.

l:shortest

Lemma 9.5. *For each λ_i, there exists $\widehat{\lambda}_i \sim \lambda_i$ so that for any vertex v of $\mathbb{G}^r$ and any $\lambda' \sim \lambda_i$ $\|\widehat{\lambda}_i\|_{\mathscr{C},v,\mathfrak{o}} \leq \|\lambda'\|_{\mathscr{C},v,\mathfrak{o}}$. We call $\widehat{\lambda}_i$ a* shortest element *in the $\sim$-class of λ_i.*

Proof. The fact that the minimum of $\|\cdot\|_{\mathscr{C},v,\mathfrak{o}}$ within a $\sim$-class is realized follows directly from the fact that the $\lambda_i(W_i)$-orbit of $\mathfrak{o}_i$ in $\mathscr{C}(\Gamma_{v_i})$ is discrete. "Shortening" with respect to a different vertex v' does not change $\|\cdot\|_{\mathscr{C},v,\mathfrak{o}}$ as we are changing $\lambda_i|_{W_{v,i}}$ by conjugation, so we can realize the minimal of $\|\cdot\|_{\mathscr{C},v,\mathfrak{o}}$ for all vertices v simultaneously. □

Replacing each λ_i with a shortest element $\widehat{\lambda}_i$ given in Lemma 9.5 gives a new sequence $(\eta_i\colon G \to \Gamma_i)$, where $\eta_i\colon = \widehat{\lambda}_i \circ \xi_i$. This produces a new sequence $\eta_i \in \mathrm{Hom}(G, \mathscr{M}_{\mathrm{Gen}}^{\pi})$.

l:sq not tdiv

Lemma 9.6. *η_i is not $\mathscr{T}$–divergent.*

Proof. Let $g \in G$. Represent $\phi_\infty(g)$ as a reduced $\mathbb{G}^r$–loop $[a_0, e_1, ..., e_n, a_n]$ based at v. For ω–almost every i, $\xi_i(g)$ is a lift of $\phi_\infty(g)$ in W_i and $\xi_i(g) = [\widetilde{a}_0, e_1, ..., e_n, \widetilde{a}_n]$, where $\widetilde{a}_k$ is a lift of a_k in the corresponding vertex group of the geometric splitting of W_i. By construction, the factoring map λ_i is a morphism from the geometric splitting of W_i to the geometric splitting of Γ_i. Combining this with the definition of the α_i, $\widehat{\lambda}_i = \lambda_i \circ \alpha_i$ is also a morphism from the geometric splitting of W_i to the geometric splitting of Γ_i. Hence $\eta_i(g) = \widehat{\lambda}_i(\xi_i(g))$ moves $v_i \in T_i$ by a distance not bigger than the amount by which $\phi_i(g)$ moves v_i. Therefore, $\eta_i(g)$ moves a point in T_i independent of g by an amount independent of i, so η_i is not $\mathscr{T}$–divergent. □

d:shortening quotient

Definition 9.7. *A* shortening quotient of *L is a group of the form $S = G/\mathrm{Ker}^{\omega}(\eta_i)$.*

By Lemma 9.2.(3) and construction we have

$$\mathrm{Ker}^{\omega}(\phi_i) = \mathrm{Ker}^{\omega}(\xi_i) \subseteq \mathrm{Ker}^{\omega}(\eta_i),$$

so there is a natural quotient map $\pi\colon L \to S$.

Recall $\mathscr{H}_{L,(\phi_i)}$ is the family of stably parabolic subgroups of L with respect to (ϕ_i). This family includes all subgroups fixing edges in T_∞.

lem:shortq

Lemma 9.8. *Each $P \in \mathscr{H}_{L,(\phi_i)}$ maps injectively into an element of $\mathscr{H}_{S,(\eta_i)}$.*

Proof. If ω–almost surely ϕ_i maps $g \in G$ to a non-trivial element in a parabolic subgroup of Γ_i, then $\phi_i(g)$ and $\eta_i(g)$ are conjugate ω–almost surely, and hence $\eta_\infty(g)$ is a non-trivial element in a stably parabolic subgroup of S. It is easy to see that π takes stably parabolic *subgroups* of L injectively to stably parabolic subgroups of S. □

We first consider the case where the quotient map π is injective, which implies that $L \cong S$. Hence we can consider (η_i) as a defining sequence for L in this case. The following lemma is the reason why we consider GGDs instead of geometric decompositions. It follows directly from the construction of η_i from ϕ_i.

lem:still ggd

Lemma 9.9. *If $\pi\colon L \to S$ is injective, then $\mathbb{G}^r$ is the canonical refinement of a GGD with respect to (η_i).*

If $\pi\colon L \to S$ is injective then for a vertex v of $\mathbb{G}^r$ we define

$$\|\eta_i\|_{\mathscr{C},v} = \inf_{x \in \mathscr{C}(\Gamma_{v_i})} \max_{s \in \widetilde{A}_v} d_i(\eta_i(s)x, x).$$

t:shortening argument

Theorem 9.10. *If π is injective then for every vertex v of $\mathbb{G}^r$ $\|\eta_i\|_{\mathscr{C},v}$ does not diverge.*

Proof. The proof is based on Sela's shortening argument. We give a sketch. For sake of contradiction, suppose $\|\eta_i\|_{\mathscr{C},v}$ diverges for some vertex v of $\mathbb{G}^r$, and let $\mathfrak{o}_i \in \mathscr{C}(\Gamma_{v_i})$ be as in Definition 9.4. Suppose v is a hyperbolic type vertex. By Theorem 7.1(1), (η_i) induces an action of $\overline{V}$ on the $\mathbb{R}$–tree $\mathscr{C}_{v,\infty}$ of $\mathscr{C}(\Gamma_{v_i})$. Let T be the minimal V-invariant subtree of $\mathscr{C}_{v,\infty}$, and $\mathfrak{o} \in \mathscr{C}_{v,\infty}$ the point defined by $\{\mathfrak{o}_i\}$.

claim:intersect

Claim 1. *For some $s \in \widetilde{A}$, the geodesic segment $[\mathfrak{o}, s\cdot\mathfrak{o}]$ intersects T in a non-degenerate segment.*

Proof of Claim 1. If $\mathfrak{o} \in T$, this follows by construction, as the action of $\langle\widetilde{A}\rangle$ on T is non-trivial. Now suppose $\mathfrak{o} \notin T$. Let o be the projection of $\mathfrak{o}$ in the closure (in $\mathscr{C}_{v,\infty}$) of T. Observe $\langle\widetilde{A}\rangle$ does not fix o. Let $s \in \widetilde{A}$ such that $s\cdot o \neq o$. Then $[s\cdot o, o]$ has non-degenerate intersection with T and hence $[s\cdot\mathfrak{o}, \mathfrak{o}] = [s\cdot\mathfrak{o}, s\cdot o] \cup [s\cdot o, o] \cup [o, \mathfrak{o}]$ also does. □

It follows from Theorem 7.1(4) and (5) and a standard argument that the action of $\overline{V}$ on T is super-stable. Since $\overline{V}$ is freely indecomposable, by [19, Theorem 5.1], there is a graph of actions decomposition of the $\overline{V}$-tree T, whose vertex actions are either simplicial, axial, or Seifert type. Let $\mathbb{A}$ be the splitting of $\overline{V}$ induced by this decomposition. Given a generator $s \in A$, consider the path $[\mathfrak{o}, s\cdot\mathfrak{o}]$. Choosing a generator for which this path is the longest, it has a non-trivial intersection with at least one piece in the graph of actions decomposition. Note that stably parabolic subgroups of V fix points in T by Lemma 7.6. Let $\mathscr{H}_{\overline{V}}$ be the collection of images of stably parabolic subgroups of V in $\overline{V}$. For the axial and Seifert type pieces, we can directly find a modular automorphism in $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}_{\overline{V}}}(\overline{V})$ which shortens the segment of $[\mathfrak{o}, s\cdot\mathfrak{o}]$ in this piece without changing the lengths of the segments in any other pieces and without changing the lengths of any other generators. In the axial case, the shortening automorphisms in $\mathrm{Mod}_{\mathbb{A}}^{\mathscr{H}_{\overline{V}}}(\overline{V})$ are more restrictive than those used in other applications of the shortening argument. However, Lemma 8.10 ensures these automorphisms suffice. By Theorem 8.1 this modular automorphism belongs to $\mathrm{Mod}_{\mathbb{J}(v)}^{\mathscr{H}_{\overline{V}}}(\overline{V})$ and by Lemma 8.14 it can be lifted to an automorphism in $\mathrm{Mod}_{\mathbb{J}(v)}^{\mathscr{H}_V}(V)$, where $\mathscr{H}_V$ is the collection of stably parabolic subgroups of V. By construction ω–almost surely this automorphism can be lifted to an element of $\mathrm{Mod}_v(W_i)$. Since the action of $W_{v,i}$ on $\mathscr{C}(\Gamma_{v_i})$ approximates the action of V on T, this automorphism ω–almost surely shortens $\widehat{\lambda}_i$, a contradiction. In case $[\mathfrak{o}, s\cdot\mathfrak{o}]$ intersects a simplicial piece, we cannot find a automorphism which shortens the action of V on T. However, by Theorem 7.1.(2) and (3) any non-trivial element in a segment stabilizer is stably loxodromic. Thus, as above, we can find an automorphism which shortens the action of $W_{v,i}$ on $\mathscr{C}(\Gamma_{v_i})$, so we again reach a contradiction. □

9.2. **The descending chain condition.** The proof of the following result is inspired by Sela's proof of [39, Theorem 1.12].

dcc **Theorem 9.11.** *There is no infinite sequence of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit groups*

$$L_1 \overset{\alpha_1}{\twoheadrightarrow} L_2 \overset{\alpha_2}{\twoheadrightarrow} \dots$$

such that each α_i is a proper quotient map.

Proof. Towards a contradiction, suppose there is an infinite descending sequence of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit groups as above. Choose a sequence of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit groups $R_1 \overset{\beta_1}{\twoheadrightarrow} R_2 \dots$ with each $R_n = F_k/\mathrm{Ker}^{\omega}(\phi^n_i)$ and $\phi^n_\infty \circ \beta_n = \phi^{n+1}_\infty$ for all n, and further assume that, for each $n \geq 1$, R_{n+1} is chosen such that if $R_n \twoheadrightarrow S \twoheadrightarrow \dots$ is any infinite descending sequence of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit groups with $S = F_k/\mathrm{Ker}^{\omega}(\rho_i)$, then

$$|\,\mathrm{Ker}(\rho_\infty) \cap B_n| \leq |\,\mathrm{Ker}(\phi^{n+1}_\infty) \cap B_n|$$

where B_n is the ball of radius n in F_k with respect to the word metric.

Choose a diagonal sequence $(\psi_n = \phi^n_{j_n})$, where j_n is so that $\mathrm{Ker}(\psi_n) \cap B_n = \mathrm{Ker}(\phi^n_\infty) \cap B_n$ and for some $g \in \mathrm{Ker}(\phi^{n+1}_\infty)$, $g \notin \mathrm{Ker}(\psi_n)$. Let $R_\infty = G/\mathrm{Ker}^{\omega}(\psi_n)$; by construction, $\mathrm{Ker}(\psi_\infty) = \bigcup_{j=1}^{\infty} \mathrm{Ker}(\phi^j_\infty)$, so R_∞ is also the direct limit of the sequence $F_n \twoheadrightarrow R_1 \twoheadrightarrow R_2 \twoheadrightarrow \dots$.

We claim every descending sequence of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit groups

$$R_\infty \twoheadrightarrow L_1 \twoheadrightarrow L_2 \dots$$

with proper quotient maps is finite. Indeed, if some element $g \in B_n$ maps trivially to L_1 but not to R_∞, then g maps non-trivially to R_{n+1}. But then there is an infinite descending sequence $R_n \twoheadrightarrow L_1 \twoheadrightarrow \dots$, so this contradicts our choice of R_{n+1}.

Since there is no infinite descending chain starting at R_∞, any sequence of proper $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit quotients of R_∞ is finite. Consider the case where the defining sequence of homomorphisms for R_∞ is $\mathscr{T}$–divergent. In this case, by [17, Lemmas 6.4 and 6.6] R_∞ has a proper shortening quotient S_1. Since R_∞ does not admit an infinite descending sequence of proper $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit quotients, repeating this argument finitely many times we obtain a sequence of proper quotients

$$R_\infty \twoheadrightarrow S_1 \dots \twoheadrightarrow S_q$$

where the defining sequence of homomorphisms for S_q is not $\mathscr{T}$–divergent. If the defining sequence of homomorphisms for S_q is $\mathscr{C}$–divergent then by applying the shortening argument above finitely many times we obtain a sequence of quotient maps

$$S_q \twoheadrightarrow U_1 \ldots \twoheadrightarrow U_s$$

of $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit quotients, each of whose defining sequence of homomorphisms is not $\mathscr{T}$–divergent (by Lemma 9.6) and terminating in an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group U_s which has a GGD $\mathbb{G}$ so that the defining sequence of homomorphisms for U_s is not $\mathscr{C}$–divergent with respect to $\mathbb{G}$ (a sequence of proper quotients terminates by the construction of R_∞, and once a shortening quotient is not a proper quotient, the defining sequence of the quotient is not $\mathscr{C}$–divergent by Theorem 9.10). Each map $U_i \twoheadrightarrow U_{i+1}$ maps the stably parabolic subgroups of U_i injectively into the stably parabolic subgroups of U_{i+1}, by Lemma 9.8. By Theorem 6.2 and the construction of U_s, all of the stably parabolic subgroups of U_s are finitely generated. It follows that all of the stably parabolic subgroups of S_q are finitely generated. It now follows from Theorem 2.6 that the defining sequence of homomorphisms for S_q ω–factors through the limit. We now apply the arguments of [17, §6]. Since R_∞ does not admit an infinite descending sequence of proper $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit quotients, and hence neither do any of the S_i, the hypotheses of [17, Lemmas 6.5 and 6.6] are satisfied for R_∞ and each of the S_i, and so these lemmas may be applied inductively from the bottom of the sequence to prove that the defining sequence for R_∞ ω–factors through the limit. Then, from the construction of R_∞, there exists i such that $\mathrm{Ker}(\psi_n) \subseteq \mathrm{Ker}(\psi_\infty) \subseteq \mathrm{Ker}(\phi_\infty^{n+1})$, contradicting our construction of ψ_n. □

We are now ready to prove the following theorem from Section 2. Except for the results proved in the appendices, this completes the proof of Theorem A, and hence the paper.

Theorem 2.7. *Let L be an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group defined by a non-$\mathscr{T}$–divergent sequence (ϕ_i). All stably parabolic subgroups of L are finitely generated.*

Proof. If L is an $\mathscr{M}^{\pi}_{\mathrm{Gen}}$–limit group whose defining sequence is not $\mathscr{T}$–divergent then L has a shortening quotient $L \twoheadrightarrow S_0$ as defined in Definition 9.7. If S_0 is a $\mathscr{C}$–divergent limit group, then by Theorem 9.10, $L \twoheadrightarrow S_0$ is a proper quotient map. Repeat the above to build a sequence of shortening quotient $S_0 \twoheadrightarrow S_1 \twoheadrightarrow \cdots$. By Theorem 9.11, this sequence of shortening quotients eventually terminates. Denote the last shortening quotient by S. By Theorem 9.10 there is a GGD with respect to which S is not $\mathscr{C}$–divergent. By Theorem 6.2 all stably parabolic subgroups of S are finitely generated. By Lemma 9.8 the stably parabolic subgroups of S_i are mapped injectively into stably parabolic subgroups of S_{i+1}, and by iterating this argument the stably parabolic subgroups of L are finitely generated. □

app:edge-twist

APPENDIX A. EDGE-TWISTED GRAPHS OF GROUPS

In this appendix we prove the following result from Section 6.

Theorem 6.7. *Let $\mathbb{E}$ be a finite edge-twisted graph of groups so that $\pi_1(\mathbb{E})$ is finitely generated. The Type B vertex groups of $\mathbb{E}$ are finitely generated.*

We repeat the definition of an edge-twisted splitting for convenience.

Definition 6.5. *A graph of groups $\mathbb{E}$ is* edge-twisted *if:*

The underlying graph of $\mathbb{E}$ is bipartite with colors A and B. Type A vertices have valence 2*, and abelian vertex groups (thus the edge groups of $\mathbb{E}$ are also abelian). Let W be a Type A vertex group of $\mathbb{E}$ and let E_1 and E_2 be the images in W of the adjacent edge groups. There are subgroups $K_j \leq E_j$ (for $j = 1,2$) so that*

(1) $K_1 \cap K_2 = \{1\}$*; and*
(2) For $j = 1,2$*, the group* E_j/K_j *is finitely generated.*

A.1. **Abelian vertex groups in graphs of groups.** We begin with some results about graphs of groups with abelian vertex groups.

W **Lemma A.1.** *Suppose* A *is an abelian group and* $K, L \leq A$ *satisfy* $K \cap L = \{1\}$*. Let* $N \leq A$ *be finitely generated. Then* $\langle N, K\rangle \cap L$ *is finitely generated.*

Proof. Let

$$N_0 = \{g \in N \mid \text{there exists } g' \in K, \text{such that } gg' \in L\}$$

Then N_0 is a subgroup of the finitely generated abelian group N, so N_0 is finitely generated. Define $\phi : N_0 \to \langle N, K\rangle \cap L$ by $\phi(g) = gg'$ where $g' \in K$ is so that $gg' \in L$. If $gg' \in L$ and $gg'' \in L$ for $g', g'' \in K$ then $g'(g'')^{-1} \in L$. But $g'(g'')^{-1} \in K$ and $K \cap L = \{1\}$. So $g' = g''$, and ϕ is well-defined. It is easy to check that ϕ is a homomorphism. For any $g_0 \in \langle N, K\rangle \cap L$, we have $g_0 = gg'$ for some $g \in N$ and $g' \in K$. Then $g_0 = \phi(g)$, so ϕ is surjective. Since $\langle N, K\rangle \cap L$ is the image of the finitely generated group N_0, it is finitely generated, as required. □

amalgam **Lemma A.2.** *Let* $G = M *_K A$*, where* A *is abelian. Let* $S_1 \subset M$*,* $S_2 \subset A$ *and let* S_3 *be a generating set of* $K \cap \langle S_2\rangle$*. Suppose* g *is a word in* $S_1 \cup S_2$*. Then*

(1) If $g \in M$*, then* g *can be written as a word in* $S_1 \cup S_3$*.*
eq:g=ma *(2) If* $g \in A$*, then* $g = ma$*, where* $m \in K \cap \langle S_1 \cup S_3\rangle$ *and* $a \in \langle S_2\rangle$*.*

In particular, if G *is finitely generated, then* M *is finitely generated.*

Proof. By the assumptions of the lemma, we have

$$g = r_1 s_1 \dots r_k s_k$$

where each $r_i \in \langle S_1\rangle$ (and hence in $\langle S_1 \cup S_3\rangle$) and each $s_i \in \langle S_2\rangle$. If some $i > 1$ we have $s_i \in K$ then $s_i \in \langle S_3\rangle$, so $r_i s_i r_{i+1} \in \langle S_1 \cup S_3\rangle$. We can repeat this reduction until no s_i lies in K. If $r_i \in K$ for some $i > 1$ then $r_i \in A$. Since A is abelian, we can write $r_{i-1}s_{i-1}r_i s_i$ as $r_{i-1}r_i s_{i-1}s_i$, and reduce the number of syllables in our description of g. After repeating finitely many times, this result is an expression:

e (1) $$g = p_1 q_1 \dots p_l q_l$$

where each $p_i \in \langle S_1 \cup S_3\rangle$, each $q_i \in \langle S_2\rangle$. Moreover, none of the p_i or q_i lies in K, except possibly that q_l is trivial, and we cannot control p_1 since the above reduction on s_i required $i > 1$.

Consider the case where $g \in M$. In this case, the above expression cannot contain any q_i, so $g = p_1 \in \langle S_1 \cup S_3\rangle$, as required.

Now suppose that $g \in A$. If $g \in K$ then $g \in M$ and we are in the first case. In this case, we can take $a = 1$ and $m = g$ in the conclusion of the lemma. Suppose then that $g \notin K$. In this case, we have $l = 1$, since otherwise we couldn't have $g \in A$. This proves (2).

The last conclusion is immediate. □

The next two results deal with the simplest cases of Theorem 6.7: those $\mathbb{E}$ with two edges. The general case follows quickly.

non-sep **Lemma A.3.** *Suppose* $\mathbb{E}$ *is an edge-twisted graph of groups so* $G = \pi_1(\mathbb{E})$ *is finitely generated. Suppose further that* $\mathbb{E}$ *has two edges and two vertices. Let* M *be the type B vertex group of* $\mathbb{E}$*. Then* M *is finitely generated.*

Proof. Let K and L be the edge groups of $\mathbb{E}$, corresponding to edges e_K and e_L respectively. Let A be the type A vertex group of $\mathbb{E}$. Choose the edge associated to e_K as a maximal tree, which allows us to consider K as a subgroup of both A and M. Consider L to be a subgroup of A, let t be the stable letter of the splitting, and let $L' := t^{-1}Lt \leq M$.

Let $K_0 \leq K$ and $L_0 \leq L$ be the subgroups of K and L guaranteed by Definition 6.5. Thus,

(1) $K_0 \cap L_0 = \{1\}$; and
(2) There are finite sets $T_K \subset K$ and $T_L \subset L$ so that $K = \langle T_K, K_0\rangle$ and $L = \langle T_L, L_0\rangle$.

Choose finite subsets $S_A \subset A$ containing $T_K \cup T_L$ and $S_M \subset M$ containing T_K so that $G = \langle S_A, S_M, t\rangle$. Let $N_0 = \langle S_A\rangle$. Since $K_0 \cap L_0 = \{1\}$, by Lemma A.1, $\langle N_0, K_0\rangle \cap L_0$ and $\langle N_0, L_0\rangle \cap K_0$ are finitely generated, by R_1 and R_2, respectively, say. Let $N_1 = \langle N_0, R_1, R_2\rangle$.

We claim $\langle N_1, K_0\rangle \cap L_0 \leq N_1$. Suppose $g \in \langle N_1, K_0\rangle \cap L_0$. Then $g = nkr_1r_2$ where $n \in N_0$, $k \in K_0$, $r_1 \in \langle R_1\rangle$ and $r_2 \in \langle R_2\rangle$. Since $g \in L_0$ and $\langle R_1\rangle \leq L_0$, we have $nkr_2 \in L_0$. But $k, r_2 \in K_0$, so $nkr_2 \in \langle N_0, K_0\rangle \cap L_0 = \langle R_1\rangle$. Thus, $g = (nkr_2)r_1 \in \langle R_1\rangle \leq N_1$. A similar argument shows $\langle N_1, L_0\rangle \cap K_0 \leq \langle R_2\rangle \leq N_1$.

We further claim that $\langle N_1, K\rangle \cap L \leq N_1$. Indeed, since $T_K \subset S_A \subset N_1$ we have $\langle N_1, K\rangle \cap L_0 = \langle N_1, K_0\rangle \cap L_0 \leq N_1$. Now suppose that $g \in \langle N_1, K\rangle \cap L$, and write $g = ng_0$ where $n \in \langle T_L\rangle \leq N_1$ and $g_0 \in L_0$. We have $g \in \langle N_1, K\rangle$ and $n \in N_1$, so $g_0 \in \langle N_1, K\rangle \cap L_0 \leq N_1$. Therefore, $g = ng_0 \in N_1$, as required. Similarly, we have $\langle N_1, L\rangle \cap K \subset N_1$.

Since N_1 is a finitely generated abelian group, $N_1 \cap K$ and $N_1 \cap L$ are both finitely generated. So

$$M' = \langle S_M, N_1 \cap K, t^{-1}(N_1 \cap L)t\rangle \tag{2}$$

is finitely generated. We claim that $M = M'$, which proves the lemma. We clearly have $M' \leq M$, so we have to prove the opposite inclusion.

To that end, let

$$N_2 = \langle N_1, M' \cap K, t(M' \cap L')t^{-1}\rangle \subset A. \tag{3}$$

(Recall that $L' = t^{-1}Lt \leq M$.)

We claim that

(1) $N_2 \cap K = \langle N_1 \cap K, M' \cap K\rangle$; and
(2) $N_2 \cap L = \langle N_1 \cap L, t(M' \cap L')t^{-1}\rangle$.

From the definition of N_2, it is clear that $\langle N_1 \cap K, M' \cap K\rangle \leq N_2 \cap K$. For the reverse inclusion, since $t(M' \cap L')t^{-1} \leq L$, we have

$$\langle N_1, t(M' \cap L')t^{-1}\rangle \cap K \leq \langle N_1, L\rangle \cap K \leq N_1 \cap K.$$

Since A is abelian, for any subgroup $J \leq K$ we have $\langle N, J\rangle \cap K = \langle N \cap K, J\rangle$. Therefore,

$$\begin{aligned} N_2 \cap K &= \langle N_1, t(M' \cap L')t^{-1}, M' \cap K\rangle \cap K \\ &= \langle \langle N_1, t(M' \cap L')t^{-1}\rangle \cap K, M' \cap K\rangle \\ &\leq \langle N_1 \cap K, M' \cap K\rangle, \end{aligned}$$

as required. The second equality of the claim is entirely similar.

As a result, we have $M' = \langle S_M, N_2 \cap K, t^{-1}(N_2 \cap L)t\rangle$. To finish the proof, we now show that $M \leq M'$. Choose sets $S_1 \supset S_M$ and $S_2 \supset S_A$ so that $\langle S_1\rangle = M'$ and $\langle S_2\rangle = N_2$. Let S_3 be a generating set for $K \cap N_2$. By the above $\langle S_3\rangle \leq M'$.

Consider the subgroup $\langle M, A\rangle \leq G$ (which is isomorphic to $M *_K A$). Since $S_M \subset S_1$ and $S_A \subset S_2$ we have $G = \langle S_1, S_2, t\rangle$.

Suppose that $g \in M$, and write $g = w_1 \dots w_j$ where each w_i is either t or t^{-1}, in $\langle S_1 \rangle$, or in $\langle S_2 \rangle$. Since $g \in M$, if there are any occurrences of t or t^{-1} then by Britton's Lemma there is a subword of one of the following two forms:

(1) twt^{-1} where $w \in \langle S_1 \cup S_2 \rangle \cap L'$; or
(2) $t^{-1}ut$ where $u \in \langle S_1 \cup S_2 \rangle \cap L$.

In the first case, $w \in L' \leq M$, so by Lemma A.2 $w \in \langle S_1 \cup S_3 \rangle$, and so $w \in M'$. Then $twt^{-1} \in t(M' \cap L')t^{-1} \leq N_2$, so twt^{-1} can be replaced by an element in $\langle S_2 \rangle$. In the second case, $u \in L \leq A$. By Lemma A.2, u can be written as $u = ma$ where $m \in K \cap \langle S_1 \cup S_3 \rangle$ and $a \in \langle S_2 \rangle$. So $m \in K \cap M' \leq N_2$ and $a \in N_2$, so $u \in N_2 \cap L$. Therefore, $t^{-1}ut \in t^{-1}(N_2 \cap L)t \leq M'$, so we can replace $t^{-1}ut$ by an element of $\langle S_1 \rangle$.

Repeating until there are no occurrences of t or t^{-1}, we obtain $g \in \langle S_1 \cup S_2 \rangle$. By Lemma A.2, $g \in \langle S_1 \cup S_3 \rangle = M'$, so $M = M'$ as required. □

The proof of the following lemma is very similar to that of Lemma A.3, and we omit it.

sep **Lemma A.4.** *Let $\mathbb{E}$ be an edge-twisted splitting with two edges, two type B vertices and one type A vertex so that $\pi_1(\mathbb{E})$ is finitely generated. Then the type B vertex groups are finitely generated.*

We now give the proof of the main result of this appendix.

Theorem 6.7. *Let $\mathbb{E}$ be a finite edge-twisted graph of groups so that $\pi_1(\mathbb{E})$ is finitely generated. The Type B vertex groups of $\mathbb{E}$ are finitely generated.*

Proof. We proceed by induction on the number of type A vertex groups of $\mathbb{E}$. If there are none, then $\mathbb{E}$ is a single type B vertex group, and the result is trivial. If there is a single type A vertex group, the result follows from Lemmas A.3 and A.4.

Suppose now that $\mathbb{E}$ has k type A vertices for $k > 1$ and that the result is true for any finite edge-twisted graph of groups with finitely generated fundamental group and at most $k-1$ type A vertices. Let A be a type A vertex group of $\mathbb{E}$. Let $\mathbb{E}_0$ be the graph of groups obtained from $\mathbb{E}$ by collapsing all of the edges that are not adjacent to A. It is clear that $\mathbb{E}_0$ is an edge-twisted graph of groups with a single type A vertex group, and so the by induction the type B vertex groups of $\mathbb{E}_0$ are finitely generated by the case $k = 1$. However, the type B vertex groups of $\mathbb{E}_0$ are the fundamental groups of (edge-twisted) sub-graphs of groups of $\mathbb{E}$, with fewer type A vertex groups than $\mathbb{E}$. Since these sub-graphs of groups have finitely generated fundamental group, the inductive hypothesis applies to them to prove that their type B vertex groups are finitely generated. But these are the type B vertex groups of $\mathbb{E}$, so the result is proved. □

app:decompositions APPENDIX B. RELATIVE LINNELL AND JSJ DECOMPOSITIONS

The purpose of this section is to prove Corollary B.4, about the existence of the relative C–Linnell decomposition which was used in Definition 4.13, and to prove Theorem B.5 which is used in Section 8. The only result other than these two which is used in the rest of the paper is Corollary B.6, which is also used in Section 8.

B.1. **Relative acylindrical accessibility.** We first construct the decomposition $\mathbb{L}$ from Theorem 8.1. We refer to this decomposition as a *relative Linnell decomposition*. The existence of such a decomposition is proved in [26] when G is finitely generated, and it can also be derived from Weidmann's version of acylindrical accessibility [47]. Here we modify Weidmann's argument to the case where G is finitely generated relative to a finite collection of subgroups $\mathscr{H}$ and all splittings considered are rel $\mathscr{H}$.

l:relative lemma

Lemma B.1. *Let T be a minimal G–tree. Let G_0 be a subgroup of G with no global fixed point in T and let T_0 be the minimal G_0–invariant subtree of T. Let H be a subgroup of G which fixes a point x such that $d(x, T_0) = k$. Let $E = \langle G_0, H\rangle$ and let T_E be the minimal E–invariant subtree of T. The number of edges in T_E/E is at most the number of edges in T_0/G_0 plus k.*

Proof. Let T' be the union of T_0 and the path from x to T_0. he E–orbit of T', denoted by T'', is connected since $gT' \cup T'$ is connected for each $g \in G_0 \cup H$, and T''/E has at most k more edges than T_0/G_0. On the other hand, $T_E \subset T''$, so the lemma follows. □

Definition B.2. *A G–action on a tree is* (k,C)–acylindrical *(for $k \geq 0$ and $C \geq 1$) if the stabilizer of all paths of length $\geq k+1$ has size at most C.*

Proposition B.3. *Fix $k \geq 0$ and $C \geq 1$. Suppose G is finitely generated relative to a finite collection of subgroups $\mathscr{H}$. There is a bound on the number of edges in minimal (k,C)–acylindrical splittings of G rel $\mathscr{H}$.*

Proof. Let $\mathscr{H} = \{H_1, ..., H_n\}$. It suffices to assume that each H_i is infinite. Let $S \subset G$ be a finite set such that $S \cup H_1 \cup \cdots \cup H_n$ generates G and $|S \cap H_i| \geq C+1$ for each i. Let $G_0 = \langle S\rangle$. By Weidmann's (k,C)–acylindrical-accessibility [47, Theorem 1], the number of edges of any (k,C)–acylindrical splitting of G_0 is bounded by a constant B. Let $\mathbb{A}$ be a (k,C)–acylindrical splitting of G rel $\mathscr{H}$, and let T be the corresponding Bass-Serre tree. First suppose that G_0 does not fix a point in T, and let T_0 be the minimal G_0 invariant subtree. For each i, let x_i be a fixed point of H_i and let y_i be the closest point in T_0 to x_i. Since $S \cap H_i$ fixes $[x_i, y_i]$ and $|S \cap H_i| \geq C+1$, we have $d_T(x_i, T_0) = d_T(x_i, y_i) \leq k$

Let G_i be the subgroup of G generated by G_0 and $H_1, \ldots, H_i$. Then $G_n = G$. Let T_i be the G_i–minimal invariant subtree. Clearly we have $T_i \subset T_{i+1}$. Each H_i fixes a point at distance at most k from T_{i-1}, so by Lemma B.1 T_i/G_i has at most k more edges than T_{i-1}/G_{i-1}. Then $T/G = T_n/G_n$ has at most $B + kn$ edges.

In case G_0 does fix a point in T the above argument shows that if $y \in \mathrm{Fix}(G_0)$ and $x_i \in \mathrm{Fix}(H_i)$ then for each i $d(y, x_i) \leq k$. Let $\mathscr{F}$ be the union of the geodesics $[y, x_i]$ and note $\mathscr{F}$ contains at most kn edges. Since the G–orbit of $\mathscr{F}$ covers T, we are done in this case also. □

cor:lindecomp

Corollary B.4. *Suppose G is finitely generated relative to a finite collection of subgroups $\mathscr{H}$. For all $C \geq 1$, G has a splitting rel $\mathscr{H}$ in which all edge groups have size $\leq C$ and no vertex group splits rel $\mathscr{H}$ over a subgroup of size $\leq C$.*

B.2. **JSJ-decompositions for relatively finitely generated groups with torsion.** The purpose of this section is to prove Theorem B.5 and Corollary B.6. We largely follow Guirardel–Levitt [21, 20], and explain the changes in our situation. We refer to [21] for terminology.

The JSJ decomposition we use is essentially [21, Theorem 9.14], but we need to weaken hypotheses. First, we need to accommodate relatively finitely generated groups. In addition, we allow groups that are not K–CSA in the sense of [21] but instead satisfy are weakly K–CSA as in Definition 8.4. We repeat the definition here for convenience.

Definition 8.4. *Fix $K \geq 1$. A group is* weakly K–CSA *if (i) any element g of order greater than K is contained in a unique maximal virtually abelian subgroup $M(g)$, so that $M(g)$ is*

K–virtually abelian and equal to its normalizer, and (ii) every two infinite, virtually abelian subgroups A and B with $\langle A,B\rangle$ not virtually abelian satisfy $|A\cap B|\leq K$.

If G is finitely generated and K–CSA the next result follows from [21, Theorem 9.14]. We briefly sketch how to modify their proof.

thm:JSJwithtorsion

Theorem B.5. *Let G be a weakly K–CSA group, let $\mathscr{H}$ be a finite collection of subgroups of G, and let $\mathscr{A}$ be the collection of virtually abelian subgroups of G. Suppose G is finitely generated rel $\mathscr{H}$ and G does not split over a subgroup of order $\leq 2K$ rel $\mathscr{H}$. The $(\mathscr{A},\mathscr{H})$–JSJ decomposition of G exists and all flexible vertex stabilizers are abelian or QH with fiber of size at most K. Its tree of cylinders is compatible with every $(\mathscr{A},\mathscr{H})$–tree.*

Proof. The proof of [21, Theorem 9.14] is based on [21, Sections 7 and 8]. In [21, Section 7] the finite generation assumption is never used. While it is used in [21, Section 8], there is a remark at the beginning of that section that it suffices to assume G is finitely generated rel $\mathscr{H}$ [21, p. 80]. Generalizing [21, Section 8] to the relatively finitely generated setting is straightforward and we omit the details.

Now we discuss replacing K–CSA groups with weakly K–CSA groups in the proof of [21, Theorem 9.14]. In [21, Section 7], they construct the tree of cylinders of a given $(\mathscr{A},\mathscr{H})$–tree T. In [21] all groups in $\mathscr{A}$ are infinite, whereas our $\mathscr{A}$ may contain finite groups. However, the construction works the same way in this case, as can be seen in [20]. We next discuss this construction and verify it has the properties we need even when $\mathscr{A}$ contains finite groups.

For $A,B\in\mathscr{A}$, define $A\sim B$ if $\langle A,B\rangle$ is virtually abelian. Since G is weakly K–CSA, this is an admissible equivalence relation on $\mathscr{A}$ (see [21, Lemma 9.13]), so for any $(\mathscr{A},\mathscr{H})$–tree T, we can form the tree of cylinders T_c (see [20]). If A and B are not equivalent, then by assumption $|A\cap B|\leq K$, so the action of G on T_c is $(2,K)$–acylindrical (See the proof of [21, Lemma 7.7] or [20, Proposition 5.13]). Now, T dominates T_c, and vertex stabilizers of T_c which are not elliptic in T are (maximal) non-(virtually cyclic) virtually abelian subgroups. That is, T smally dominates T_c. Then $\mathscr{A}$ contains all virtually cyclic and all finite subgroups of G, so by [21, Theorem 8.7] there exists an $(\mathscr{A},\mathscr{H})$–JSJ tree whose flexible vertices are either abelian or QH with fiber of size at most K.

If T_{JSJ} is the JSJ-tree its tree of cylinders is compatible with every $(\mathscr{A},\mathscr{H})$–tree. This follows from [21, Lemmas 7.14 and 7.15]. Note that [21, Lemma 7.15] assumes one-endedness. Here is how this is used: Suppose $S\to T$ collapses a single G–orbit of edges into the G–orbit of a vertex v of T. Let $H=\mathrm{Stab}_G(v)$. In case H is small in S the proof from [21] works as written. Assume that H is QH with fiber F, where $|F|\leq K$. Let S_v be the minimal subtree of S for the stabilizer H. Since G does not split over any subgroup of order $\leq 2K$, by [21, Lemma 5.16] all boundary components of the associated orbifold are used. Hence by [21, Lemma 5.18] the splitting of H corresponding to the action on S_v is dual to a family of geodesics on the orbifold.

The claim is that any cylinder of S containing an edge of S_v is contained in S_v. Suppose there are edges e and f with stabilizers $A=\mathrm{Stab}_G(e)$ and $B=\mathrm{Stab}_G(f)$ so that e is in S_v and f has exactly one endpoint in S_v. One-endedness would imply that B is infinite, and in this case the proof in [21] works as written. So assume that B is finite. Since B/F is a subgroup of the orbifold corresponding to the QH vertex v, it is in fact cyclic and contained (or conjugate into) a cone point subgroup of the orbifold (it does not come from a mirror since $|B/F|\geq 3$). By [21, Proposition 5.4], A/F is cyclic subgroup corresponding to a geodesic on the orbifold. Hence $\langle A,B\rangle$ is not virtually abelian, so these groups are not equivalent.

The rest of the proof of [21, Theorem 9.14] works verbatim in our situation, proving Theorem B.5. □

The next result follows because the JSJ is a tree of cylinders.

cor:extra JSJ props

Corollary B.6. *Let G be as in Theorem B.5, and let $\mathbb{J}$ be the $(\mathscr{A},\mathscr{H})$–JSJ decomposition of G. All virtually abelian vertex groups of $\mathbb{J}$ which are not virtually cyclic are maximal virtually abelian. Also, every edge connects a virtually abelian vertex group to a non-(virtually abelian) vertex group.*

REFERENCES

AGM_msqt [1] I. Agol, D. Groves, and J. F. Manning. An alternate proof of Wise's malnormal special quotient theorem. *Forum Math. Pi*, 4:e1, 54, 2016.

Agol-Liu [2] I. Agol and Y. Liu. Presentation length and Simon's conjecture. *J. Amer. Math. Soc.*, 25(1):151–187, 2012.

AFW [3] M. Aschenbrenner, S. Friedl, and H. Wilton. *3-manifold groups*. EMS Series of Lectures in Mathematics. European Mathematical Society (EMS), Zürich, 2015.

BMR99 [4] G. Baumslag, A. Myasnikov, and V. Remeslennikov. Algebraic geometry over groups. I. Algebraic sets and ideal theory. *J. Algebra*, 219(1):16–79, 1999.

BMR [5] G. Baumslag, A. Myasnikov, and V. Roman'kov. Two theorems about equationally Noetherian groups. *J. Algebra*, 194(2):654–664, 1997.

BP [6] R. Benedetti and C. Petronio. *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992.

Boileau-Boyer-Wang2008 [7] M. Boileau, S. Boyer, and S. Wang. Roots of torsion polynomials and dominations. In *The Zieschang Gedenkschrift*, volume 14 of *Geom. Topol. Monogr.*, pages 75–81. Geom. Topol. Publ., Coventry, 2008.

leau-Rubinstein-Wang2014 [8] M. Boileau, J. H. Rubinstein, and S. Wang. Finiteness of 3-manifolds associated with non-zero degree mappings. *Comment. Math. Helv.*, 89(1):33–68, 2014.

Boileau-Wang-JDG [9] M. Boileau and S. Wang. Non-zero degree maps and surface bundles over S^1. *J. Differential Geom.*, 43(4):789–806, 1996.

bridhaef:book [10] M. R. Bridson and A. Haefliger. *Metric Spaces of Non–Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer–Verlag, Berlin, 1999.

Calegari:mathscinet [11] D. Calegari. Review of "epimorphism sequences between hyperbolic 3-manifold groups". *Mathematical Reviews*, 2002.

CaoZhu [12] H.-D. Cao and X.-P. Zhu. A complete proof of the Poincaré and geometrization conjectures—application of the Hamilton-Perelman theory of the Ricci flow. *Asian J. Math.*, 10(2):165–492, 2006.

Derbez2003 [13] P. Derbez. Volume-convergent sequences of Haken 3-manifolds. *C. R. Math. Acad. Sci. Paris*, 336(10):833–838, 2003.

Derbez2007 [14] P. Derbez. Nonzero degree maps between closed orientable three-manifolds. *Trans. Amer. Math. Soc.*, 359(8):3887–3911, 2007.

farb:relhyp [15] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.

nzalez-Acuna-Ramirez2003 [16] F. González-Acuña and A. Ramírez. Epimorphisms of knot groups onto free products. *Topology*, 42(6):1205–1227, 2003.

GH [17] D. Groves and M. Hull. Homomorphisms to acylindrically hyperbolic groups I: Equationally noetherian groups and families. *Trans. Amer. Math. Soc.*, 372(10):7141–7190, 2019.

GMW [18] D. Groves, J. F. Manning, and H. Wilton. Recognizing geometric 3-manifold groups using the word problem. arxiv.org/abs/1210.2101, 2012.

Guirardel:Rtrees [19] V. Guirardel. Actions of finitely generated groups on $\mathbb{R}$-trees. *Ann. Inst. Fourier (Grenoble)*, 58(1):159–211, 2008.

GL11 [20] V. Guirardel and G. Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

GuiLev2 [21] V. Guirardel and G. Levitt. JSJ decompositions of groups. *Astérisque*, (395):vii+165, 2017.

t-legrand-Wang-Zieschang [22] C. Hayat-Legrand, S. Wang, and H. Zieschang. Any 3-manifold 1-dominates at most finitely many 3-manifolds of S^3-geometry. *Proc. Amer. Math. Soc.*, 130(10):3117–3123, 2002.

KleinerLott [23] B. Kleiner and J. Lott. Notes on Perelman's papers. *Geom. Topol.*, 12(5):2587–2855, 2008.

KreckLuck [24] M. Kreck and W. Lück. Topological rigidity for non-aspherical manifolds. *Pure Appl. Math. Q.*, 5(3, Special Issue: In honor of Friedrich Hirzebruch. Part 2):873–914, 2009.

Liang [25] H. Liang. Discrete representations of finitely generated groups into PSL(2,$\mathbb{R}$). Preprint, arXiv:2010.16008, 2020.

Linnell-Access [26] P. A. Linnell. On accessibility of groups. *J. Pure Appl. Algebra*, 30(1):39–46, 1983.

Liu [27] Y. Liu. Finiteness of nonzero degree maps between 3-manifolds. *J. Topol.*, 13(1):237–268, 2020.

MorganTian [28] J. Morgan and G. Tian. *The geometrization conjecture*, volume 5 of *Clay Mathematics Monographs*. American Mathematical Society, Providence, RI; Clay Mathematics Institute, Cambridge, MA, 2014.

paulin:rtrees [29] F. Paulin. Topologie de Gromov équivariante, structures hyperboliques et arbres réels. *Invent. Math.*, 94(1):53–80, 1988.

Per02 [30] G. Perelman. The entropy formula for the ricci flow and its geometric applications, 2002.

Per03a [31] G. Perelman. Finite extinction time for the solutions to the ricci flow on certain three-manifolds, 2003.

Per03b [32] G. Perelman. Ricci flow with surgery on three-manifolds, 2003.

Reid-Wang99 [33] A. W. Reid and S. Wang. Non-Haken 3-manifolds are not large with respect to mappings of non-zero degree. *Comm. Anal. Geom.*, 7(1):105–132, 1999.

RWZ [34] A. W. Reid, S. C. Wang, and Q. Zhou. Generalized Hopfian property, a minimal Haken manifold, and epimorphisms between 3-manifold groups. *Acta Math. Sin. (Engl. Ser.)*, 18(1):157–172, 2002.

Rong92 [35] Y. W. Rong. Degree one maps between geometric 3-manifolds. *Trans. Amer. Math. Soc.*, 332(1):411–436, 1992.

Scott:Coherence [36] G. P. Scott. Finitely generated 3-manifold groups are finitely presented. *J. London Math. Soc. (2)*, 6:437–440, 1973.

scott:geom [37] P. Scott. The geometries of 3-manifolds. *Bulletin of the London Mathematical Society*, 15:401–487, 1983.

sela:dio1 [38] Z. Sela. Diophantine geometry over groups. I. Makanin-Razborov diagrams. *Publ. Math. Inst. Hautes Études Sci.*, 93:31–105, 2001.

Sela:hyp [39] Z. Sela. Diophantine geometry over groups. VII. The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc. (3)*, 99(1):217–273, 2009.

Soma-JDG-98 [40] T. Soma. Non-zero degree maps to hyperbolic 3-manifolds. *J. Differential Geom.*, 49(3):517–546, 1998.

Soma2000 [41] T. Soma. Sequences of degree-one maps between geometric 3-manifolds. *Math. Ann.*, 316(4):733–742, 2000.

Soma [42] T. Soma. Epimorphism sequences between hyperbolic 3-manifold groups. *Proc. Amer. Math. Soc.*, 130(4):1221–1223, 2002.

thurston:notes [43] W. P. Thurston. Geometry and topology of three-manifolds. Princeton lecture notes available at `http://www.msri.org/publications/books/gt3m/`, 1980.

vdD-W [44] L. van den Dries and A. J. Wilkie. Gromov's theorem on groups of polynomial growth and elementary logic. *J. Algebra*, 89(2):349–374, 1984.

Wang:ICM [45] S. Wang. Non-zero degree maps between 3-manifolds. In *Proceedings of the International Congress of Mathematicians, Vol. II (Beijing, 2002)*, pages 457–468. Higher Ed. Press, Beijing, 2002.

Wang-Zhou2002 [46] S. Wang and Q. Zhou. Any 3-manifold 1-dominates at most finitely many geometric 3-manifolds. *Math. Ann.*, 322(3):525–535, 2002.

Weidmann-Access [47] R. Weidmann. On accessibility of finitely generated groups. *Q. J. Math.*, 63(1):211–225, 2012.

ReiWei [48] R. Weidmann and C. Reinfeldt. Makanin-Razborov diagrams for hyperbolic groups. *Ann. Math. Blaise Pascal*, 26(2):119–208, 2019.

WiltonZalesskii [49] H. Wilton and P. Zalesskii. Profinite properties of graph manifolds. *Geom. Dedicata*, 147(1):29–45, 2010.

DEPARTMENT OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT CHICAGO, 322 SCIENCE AND ENGINEERING OFFICES (M/C 249), 851 S. MORGAN ST., CHICAGO, IL 60607-7045

E-mail address: `groves@math.uic.edu`

DEPARTMENT OF MATHEMATICS & STATISTICS, UNIVERSITY OF NORTH CAROLINA AT GREENSBORO, 116 PETTY BUILDING, GREENSBORO, NC 27402-6170

E-mail address: `mbhull@uncg.edu`

SCHOOL OF MATHEMATICS, SUN YAT-SEN UNIVERSITY, GUANGZHOU, GUANGDONG, 510275, PEOPLE'S REPUBLIC OF CHINA

E-mail address: `liangh63@mail.sysu.edu.cn; lianghao1019@hotmail.com`