Geometry & Topology 11 (2007) 643–666 **643**

Makanin–Razborov diagrams for limit groups

EMINA ALIBEGOVIĆ

We give a description of $\mathrm{Hom}(G, L)$, where L is a limit group (fully residually free group). We construct a finite diagram of groups, the Makanin–Razborov diagram, that gives a convenient representation of all such homomorphisms.

20F65; 57M07, 20F67

1 Introduction

The subject of this paper has the roots in the following problem: given a finitely generated group G, describe the set of all homomorphisms $G \to F$ to a fixed free group F.

When G is the fundamental group of a closed surface, say of genus g, the answer has been known for a while (see for example Zieschang [17]). Denote by $q\colon G \to F_g$ an epimorphism to the free group of rank g (defined by inclusion of the boundary to the handlebody). Then every $f\colon G \to F$ factors as $f = \phi \circ q \circ \alpha$, for some automorphism $\alpha\colon G \to G$ and some $\phi\colon F_g \to F$. Thus $\mathrm{Hom}(G, F)$ is "parametrized" by the product of the Teichmüller modular group of G and the "affine space" F^g. This theorem of Stallings was generalized to arbitrary finitely generated groups G by Sela in [16] and Kharlampovich and Myasnikov in [13].

All homomorphisms $G \to F$ are encoded into a finite diagram of groups, called the *Makanin–Razborov diagram*. Each group in this diagram has a finite number of directed edges issuing from it. This number will be zero if the group in question is a free group. Each edge represents a quotient map, and all quotients are proper; see Figure 1.

Every homomorphism $h\colon G \to F$ can be written as

$$h = h_0 \circ \vartheta_k \circ \varphi_{k-1} \circ \ldots \circ \vartheta_1 \circ \varphi_0,$$

where $\varphi_i \in \mathrm{Mod}(L_i)$, ϑ_i are the quotient maps and $h_0\colon L_{t_i} \to F$, where L_{t_i} is a free group. We say that h factors through a branch of the M–R diagram. In addition, all groups in the M–R diagram naturally belong to the class of limit groups. Limit groups, also known as fully residually free groups, have been studied by Sela [16],

Published: 10 May 2007 DOI: 10.2140/gt.2007.11.643

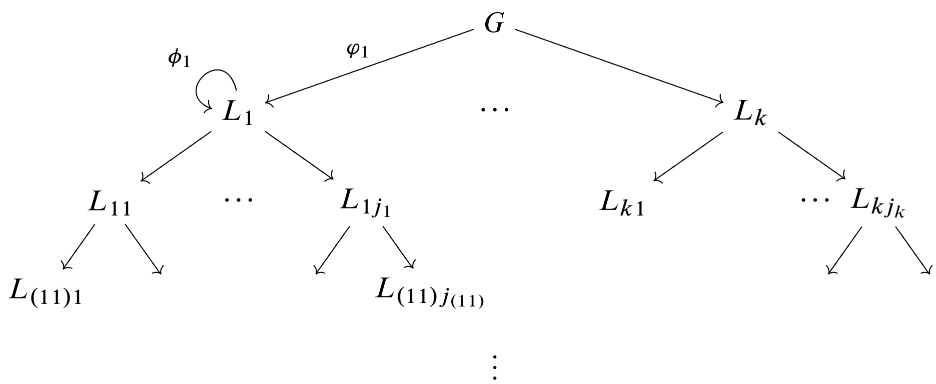

Figure 1: M-R diagram

Kharlampovich and Myasnikov [12; 13] and Chiswell [8]. The structure of limit groups can be easily described. These groups can be built inductively: level 0 limit groups are finitely generated free groups, finitely generated free abelian groups and surface groups. A level n limit group is obtained by taking a finite number of free products or amalgamated free products or HNN extensions of level $n-1$ limit groups along their maximal cyclic subgroups. We will talk more about this decomposition in Section 2.

We are interested in studying $\mathrm{Hom}(G, L)$, for G a f.g. group and L an arbitrary limit group, and constructing M–R diagrams for this case. The problem that occurs here lies in the fact that a homomorphism $h\colon G \to L$ might not factor through a complete branch of the M–R diagram. That is, the final homomorphism h_0 in the above representation of h might be an embedding of some limit group into L. We need to ensure that there are only finitely many such embeddings, up to some equivalence relation. Further, under this equivalence relation the homomorphism will change, albeit in a controlled manner. For this reason the homomorphisms that factor through the M–R diagram are the shortest homomorphisms, rather than all homomorphisms. The following theorem is the objective of this work:

Main Theorem *Let G be a f.g. freely indecomposable group and L a freely indecomposable limit group. There exist finitely many proper quotients $G_1, \dots, G_r$ of G so that for every homomorphism $f\colon G \to L$ with $d_{[f]} \gg 0$, an element of the equivalence class $\sim$ of f factors through some G_i. Furthermore, there are only finitely many homomorphisms with uniformly bounded $d_{[f]}$ and nonabelian image, up to conjugacy.*

We will define all the terms used in the statement of this theorem in the next section, and the proof will be given in Section 3.

Acknowledgments We thank the referee, Daniel Groves and Mladen Bestvina for their helpful comments and suggestions.

2 Background

We list some of the properties of limit groups that are often used, without proofs, and we refer reader to Sela [16] and Bestvina and Feighn [6].

Lemma 2.1 *Let L be a limit group.*

(L0) *L is torsion free.*

(L1) *L is finitely presented, in fact coherent.*

(L2) *Every f.g. subgroup of L is a limit group.*

(L3) *Every abelian subgroup of L is finitely generated and free, and there is a uniform bound on its rank.*

(L4) *Every abelian subgroup is contained in a unique maximal abelian subgroup.*

(L5) *Every maximal abelian subgroup of L is malnormal.*

In [1] we defined a class of groups that contains limit groups, and for this class we found δ–hyperbolic spaces on which they act freely, by isometries. This class of groups was defined as follows:

Definition 2.2 A torsion-free, f.g. group G is a depth 0 $\mathcal{C}$–group if it is either an f.g. free group, or an f.g. free abelian group or the fundamental group of a closed hyperbolic surface. A torsion-free f.g. group G is a $\mathcal{C}$–group of depth $\leq n$ if it has a graph of groups decomposition with three types of vertices: abelian, surface or depth $\leq (n-1)$, cyclic edge stabilizers and the following holds:

- Every edge is adjacent to at most one abelian vertex v. Further, G_v, the stabilizer of v, is a maximal abelian subgroup of G.
- Each surface vertex group is the fundamental group of a surface with boundary, and to each boundary component corresponds an edge of this decomposition. Each edge group is conjugate to a boundary component.
- The stabilizer of a depth $\leq (n-1)$ vertex v, G_v, is $\mathcal{C}$–group of depth $\leq (n-1)$. The images in G_v of incident edge groups are distinct maximal abelian subgroups of G_v (ie cyclic subgroups generated by distinct, primitive, hyperbolic elements of G_v).

We say that the depth of a C–group G is the smallest n for which G is of depth $\leq n$.

That limit groups belong to the class $\mathcal{C}$ follows from Theorem 3.2 and Theorem 4.1 in [16]. In fact, the decomposition of a limit group from this definition coincides with cyclic JSJ decomposition defined in [16].

Let L be a depth n limit group. Let Δ_L be a graph of groups decomposition of L given in the above definition, call T_L the underlying graph, and let T be the tree so that $T/L = T_L$. In [1] we showed that we can find a graph of spaces X/L corresponding to Δ_L so that its universal cover X is δ–hyperbolic. In order to establish necessary notation we give some properties of this covering space. X/L is quasiisometric to the wedge of k rays $[0,\infty)$ joined at 0. We lift k rays that correspond to $\partial X/L$ to rays $r_i\colon [0,\infty) \to X$, $i = 1,\dots,k$, and we let h_i be the horofunction corresponding to r_i. The stabilizer, $L_i < L$, of the limit point $r_i(\infty)$ of r_i preserves h_i. Denote by $B_i(\rho)$ the horoballs $h_i^{-1}(-\infty,\rho) \subset X$. For sufficiently small ρ the intersection $\gamma B_i(\rho) \cap B_j(\rho)$ is empty unless $i = j$ and $\gamma \in L_i$. Let

$$LB(\rho) = \bigcup_{i,\gamma} \gamma B_i(\rho),$$

$i = 1,\dots,k$, $\gamma \in L$. Let $X(\rho) = X \setminus LB(\rho)$. $X(\rho)/L$ is compact for all $\rho \in (-\infty,\infty)$.

These properties in fact constitute the definition of *relatively hyperbolic* groups given by Gromov in [10]. Hence limit groups are hyperbolic relative the collection of the representatives of conjugacy classes of their maximal noncyclic abelian subgroups (see Alibegović [1] and Dahmani [9]). We will call the subgroups L_i *parabolic subgroups*.

The *modular group* $\mathrm{Mod}(L)$ associated to the decomposition Δ_L of L is the subgroup of $\mathrm{Aut}(L)$ generated by

- inner automorphisms of L,
- Dehn twists in the centralizers of edge groups,
- automorphisms induced by automorphisms of an abelian vertex group that are identity on its *peripheral subgroup* (subgroup generated by the incident edge groups) and all other vertex groups, and
- automorphisms induced by homeomorphisms of surfaces underlying surface vertex groups that fix all boundary components.

Let G be an f.g. group with a finite generating set S. We consider a sequence of homomorphisms $f_i\colon G \to L$. Each of the given homomorphisms induces an action of G on the space X, and for each we define

$$d_i = \inf\{\sup\{d(x, f_i(g)x) : g \in S\} :\ x \in X\}.$$

If this infimum is attained at a point x_i then x_i is called a centrally located point for the action.

In order to determine if we can extract a subsequence of actions that converges, we need to see whether each of these actions has a centrally located point and how the sequence d_i behaves; see Bestvina [4, Theorem 3.3, Proposition 3.8].

The following lemma has been proved by Bestvina [3] for hyperbolic spaces $\mathbb{H}^n$ and by Paulin [15] for Gromov hyperbolic spaces.

Lemma 2.3 *For every $f\colon G \to L$ whose image is not an abelian group and for which an induced action on X is nonelementary (no point at infinity is fixed by the whole group), there exists a centrally located point.*

Proof Suppose $d_{[f]} = 0$. Since the group L contains no elliptic elements, this implies that for every $g \in S$, $f(g)$ is a parabolic element. Furthermore, $f(G)$ is an abelian subgroup of L. Note that a partial converse holds: if $f(G)$ is an abelian group and all generators are mapped into parabolic elements then $d_{[f]} = 0$ and centrally located point does not exist. So we need to show that if $d_{[f]} > 0$ then a centrally located point exists.

We would like to show that a map $F\colon X \to \mathbb{R}_+$ defined by

$$F(x) = \sup_{g \in S} d(x, f(g)x)$$

is a proper map away from the horoballs and consequently attains its infimum. Suppose it is not, that is there is a sequence $\{x_n\}$ in X not contained in a compact set, such that $\{F(x_n)\}$ is bounded. Recall from the definition of X that there is ρ such that $B_i(\rho) \cap \gamma B_j(\rho) = \varnothing$, unless $i = j$ and $\gamma \in L_i$ and $X(\rho)/L$ is compact, where $X(\rho) = X \setminus \bigcup_{i,\gamma} \gamma B_i(\rho)$. We first note that our sequence has to stay within bounded distance from $X(\rho)$, for otherwise $F(x_n) \to \infty$ since $d_{[f]} > 0$. Namely, suppose $d(x_n, X(\rho)) \longrightarrow \infty$ as $n \to \infty$. We also may assume that $x_n \in B_i(\rho)$. Since $d_{[f]} > 0$ not every $f(g), g \in S$, is contained in L_i. Let us assume that $h \in S$ is such that $f(h) \notin L_i$. Then $f(h)B_i(\rho) \cap B_i(\rho) = \varnothing$ implies that $d(x_n, f(h)x_n) \geq 2d(x_n, X(\rho))$, hence $F(x_n) \to \infty$ as $n \to \infty$. Contradiction. After passing to a subsequence, $\{x_n\}$ will converge to a point $x \in \partial X$. Since $d(x_n, f(g)x_n)$ is bounded for every $g \in S$, we conclude that the sequence $\{f(g)x_n\}$ also converges to x. Hence, x is fixed by $f(G)$, which is a contradiction. □

Proposition 2.4 *There are only finitely many homomorphisms $f\colon G \to L$ with (uniformly) bounded $d_{[f]}$ and nonabelian image, up to conjugation.*

Proof Suppose this is not true: there are infinitely many homomorphisms $f_i\colon G \to L$ with d_i bounded by $D > 0$. We assume that no $d_i = 0$, since otherwise f_i has an abelian image, by Lemma 2.3. By the same lemma we know that a centrally located point x_i exists for the action given by f_i.

To simplify the notation, assume that X/L has only one cusp. Let $r\colon [0,\infty) \to X$ be a ray in X corresponding to this cusp, and let h be the horofunction corresponding to r. Also, let A be an abelian subgroup of L that stabilizes the horoball $h^{-1}(-\infty, 0)$. Choose ρ small enough so that $B(\rho) \cap \gamma B(\rho) = \varnothing$, for all $\gamma \in L \backslash A$. We first note that x_i can not be too deep inside the horoball. By "deep" we mean that the distance from x_i to the boundary of the horoball has to be smaller than D. If it is not, then the ball of radius D around x_i is not only completely contained within the horoball, but also contains $f_i(g)x_i$, for all $g \in S$, since $d_i \leq D$. This implies that $f_i(G)$ is abelian. Hence, the D–ball around x_i, call it B_i, has to intersect $X(\rho)$. We consider $X(\rho - D)$. The action of L on $X(\rho - D)$ is cocompact, hence we can find a compact set K whose translates cover $X(\rho - D)$; see Figure 2. We can find $l_i \in L$ for each x_i so that $l_i x_i \in K$. Since K is compact there exist $r > 0$ and $x \in K$ so that $B_r(x)$ contains the translates $l_i B_i$, for all i.

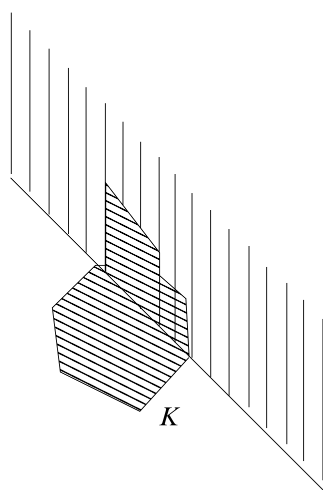

Figure 2: Fundamental domain for the action of L

Thus we have

$$\begin{aligned} d(x, l_i f_i(g) l_i^{-1} x) &= d(l_i^{-1} x, f_i(g) l_i^{-1} x) \\ &\leq d(l_i^{-1} x, x_i) + d(x_i, f_i(g) x_i) + d(f_i(g) x_i, f_i(g) l_i^{-1} x) \\ &\leq 2r + D \end{aligned}$$

and so $l_i f_i(g) l_i^{-1}$ moves a point x within a ball of radius $2r + D$, for all i and all $g \in S$. Since there are only finitely many translates of x within that ball we conclude that there can be only finitely many nonconjugate $l_i f_i l_i^{-1}$. □

We therefore conclude that if homomorphisms f_i belong to distinct conjugacy classes and have nonabelian images, the sequence of actions they induce contains a subsequence which converges to an action of G on an $\mathbb{R}$–tree T_∞ without a global fixed point. Let us call this limiting action ρ. In the following proposition we list some properties of this action. Before we do that let us define the kernel of the action of G on T_∞ as

$$K_\infty = \{g \in G : gy = y, \text{ for all } y \in T_\infty\}.$$

Let us also remind the reader what it means for an action of a group to be *stable* [4]. A subtree T of T_∞ is stable if all of its subtrees have the stabilizers equal to that of T. The action of a group on a tree T_∞ is stable if it is nontrivial, minimal and every nondegenerate tree in T_∞ has a stable subtree.

Proposition 2.5 *Let $L_\infty = G/K_\infty$.*

(1) *The stabilizer* $\mathrm{Stab}(I)$ *in L_∞ of every arc $I \subset T_\infty$ is abelian.*

(2) *The action of L_∞ on T_∞ is stable.*

Proof We first show (1). Assume that $H < G$ fixes an arc $I \subset T_\infty$ pointwise. Our claim is that H, when descended to L_∞, is abelian. Let Q be any finite subset of H. Following Paulin [15, Proposition 2.4] (see Bridson and Swarup [7] as well), one can show that there are points $y_i \in X$ so that for large enough i every element $f_i([s,t]), s, t \in Q$, moves y_i by no more than some small distance (eg 16δ). There are now two possibilities:

- $d(y_i, X(\rho)) \le D$, for large enough i and for some $D \in \mathbb{R}$. The number of elements of L that move any point of $X(\rho - D)$ within a ball of prescribed radius is uniformly bounded since the action of L on $X(\rho - D)$ is proper and cocompact. We then conclude that the set $H' = \{f_i([s,t]) : s, t \in H\}$ has to be finite, hence $f_i(H)$ is virtually abelian subgroup of L [15, Lemma 1.A]. Since L is torsion free limit group, $f_i(H)$ is in fact abelian.
- $d(y_i, X(\rho)) \to \infty$. In this case $f_i([H,H])$ is contained in an abelian subgroup of L. Tits alternative stands true for limit groups, so either $f_i(H)$ contains a free subgroup or it is virtually abelian. If it did contain a free subgroup, then so would $[f_i(H), f_i(H)]$, which is not true. We conclude, as above, that $f_i(H)$ is abelian.

In either case $f_i([s,t]) = 1$ for all $s,t \in H$, for almost all i, and hence $[s,t] \in K_\infty$.

We prove (2) by contradiction. Suppose the action of L_∞ is not stable, hence there is a nondegenerate subtree $T_1 \subset T_\infty$ which does not contain a stable subtree. Therefore we can find a decreasing sequence of subtrees $T_1 \supset T_2 \supset T_3 \supset \cdots$ whose stabilizers form a strictly increasing sequence of subgroups of L_∞. Let $H_j < G$ be a subgroup of G which fixes the tree T_j. For almost all i we have $f_i(H_j)$ is an abelian subgroup of L and $\{f_i(H_j) : j \in \mathbb{N}\}$ is an increasing sequence of abelian subgroups of the limit group L. According to Lemma 2.1(L3) every such sequence stabilizes, ie there is $n \in \mathbb{N}$ so that for all $m \geq n$ $f_i(H_n) = f_i(H_m)$. Therefore for all $m \geq n$ we have $\mathrm{Stab}_{L_\infty}(T_n) = \mathrm{Stab}_{L_\infty}(T_m)$. Contradiction. □

Instead of considering actions on $\mathbb{R}$–trees, we will consider laminations on simplicial complexes. We will state the definition and relevant results as they appear in [6], but we refer an interested reader to [5] for more details.

Definition 2.6 A *measured lamination* Λ on a finite simplicial complex K consists of a closed subset $|\Lambda|$ and a transverse measure μ. $|\Lambda|$ is disjoint from the vertex set of K, intersects each edge either in Cantor or empty set, and intersects each triangle in 0, 1, 2 or 3 families of straight line segments connecting distinct sides. The measure μ assigns a nonnegative number $\mu(I)$ to each interval contained in an edge whose endpoints do not belong to $|\Lambda|$ so that the following is satisfied:

- If two intervals I and J contained in the sides of the same triangle $\triangle$ intersect the same components of $|\Lambda| \cap \triangle$ then $\mu(I) = \mu(J)$.
- The measure μ restricted to an edge is equivalent under the Cantor function to the Lebesgue measure on an interval in $\mathbb{R}$.

Two measured laminations on K are considered equivalent if they assign the same value to each edge. Further, the space $\mathcal{ML}(K)$ of nonempty measured laminations on K can be identified with the closed cone without 0 in $\mathbb{R}_+^E$, where E is the set of edges of K, given by the triangle inequalities for each triangle of K. The projectivized space $\mathcal{PML}(K)$ is compact.

Proposition 2.7 (Morgan–Shalen [14]) *Let Λ be a measured lamination on K. Then*

$$\Lambda = \Lambda_1 \sqcup \Lambda_2 \sqcup \ldots \sqcup \Lambda_k$$

so that each Λ_i is either minimal (each leaf is dense in $|\Lambda_i|$) or simplicial (each leaf is compact and a regular neighborhood of $|\Lambda_i|$ is an I–bundle over a leaf and $|\Lambda_i|$ is a Cantor set subbundle).

It turns out that there are only three types of minimal laminations: surface, toral and thin, and whenever we have a minimal component we can "replace" it by one of these types:

Theorem 2.8 [6, Theorem 3.6] *Let Λ be a measured lamination on a finite 2–complex K, and let Λ_i be a minimal component of Λ. There are a neighborhood N of $|\Lambda_i|$, a finite 2–complex N' with measured lamination Λ' as in one of the above types, and a measure preserving π_1–isomorphism $f\colon N \to N'$ such that $f^*(\Lambda') = \Lambda_i$.*

From now on we will call N the *standard neighborhood* of Λ_i. For precise description of the three different types of minimal laminations the reader should see Bestvina [4, Examples 3.3, 3.4, 3.5].

For every homomorphism $f\colon G \to L$ we get a measured lamination Λ_f on the complex K whose fundamental group is G. We define a resolution $\Phi\colon \widetilde{K} \to X$ by defining it f–equivariantly on the vertices of the triangulation of $\widetilde{K}$. We extend, again equivariantly, it to edges so that each edge of the triangulation is mapped into a geodesic between the images of its endpoints. Finally we extend the map equivariantly to 2–cells.

The image of each triangle has a unique measured lamination on it. Let ABC be a triangle in $\widetilde{K}$ whose image under Φ is a triangle $A'B'C'$ in X. There are points $c_{A'} \in [B'C']$, $c_{B'} \in [A'C']$ and $c_{C'} \in [A'B']$ such that $d(A', c_{B'}) = d(A', c_{C'})$, $d(B', c_{A'}) = d(B', c_{C'})$ and $d(C', c_{A'}) = d(C', c_{B'})$. There is a unique isometry $f_{A'B'C'}$ from the sides of the triangle $A'B'C'$ onto a tripod with valence one vertices A'', B'' and C'', and a valence three vertex W for which the following holds: $d(A'', W) = d(A', c_{B'}) = d(A', c_{C'})$, and similarly for other vertices; see Figure 3. Since X is δ–hyperbolic, we have if $f_{A'B'C'}(u) = f_{A'B'C'}(v)$, then $d_{X'}(u, v) \leq 6\delta$ (see Alonso et al [2, page 11]). The triangle $A'B'C'$ acquires the following lamination: two points u and v belong to a leaf if they are mapped to the same point in the interior of an edge of the tripod $A''B''C''$ under $f_{A'B'C'}$. We will pull back this lamination to get $\widetilde{\Lambda_f}$; see Figure 3.

Suppose b' and c' are the points on $A'B'$ and $A'C'$, respectively, the same distance from A' for which $f_{A'B'C'}(b') = f_{A'B'C'}(c')$. Let b and c be the points on AB and AC, respectively, for which $\Phi(b) = b'$ and $\Phi(c) = c'$, and let $[bc]$ be the segment for which $\Phi([bc]) = [b'c']$. The segment $[bc]$ is contained in a leaf of our lamination $\widetilde{\Lambda_f}$. Notice that the length of the image of $[bc]$ is smaller than or equal to 6δ. Λ_f is then a projection of $\widetilde{\Lambda_f}$ to K.

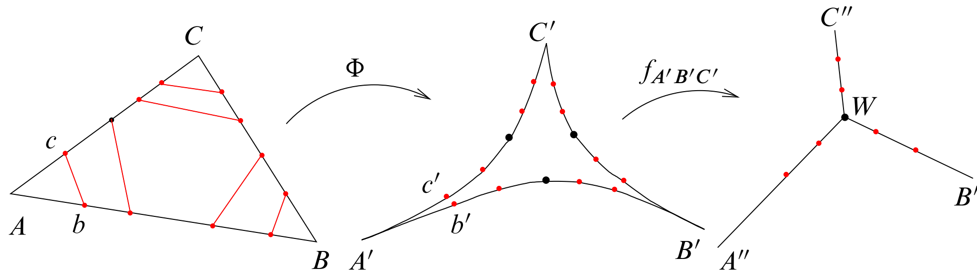

Figure 3: Defining the lamination $\widetilde{\Lambda_f}$ on $\widetilde{K}$

Definition 2.9 For a measured lamination (Λ,μ) on K we define the *length* of Λ to be

$$L(\Lambda) = \sum_{e \in K^{(1)}} \mu(e).$$

The length of a resolution $\Phi\colon \widetilde{K} \to X$, $L(\Phi)$, is the length of the induced lamination.

For each homomorphism $f\colon G \to L$ we choose an (f)–resolution Φ_f so that

$$L(\Phi_f) = \inf\{L(\Phi) : \Phi \text{ is an } (f)-\text{resolution}\}.$$

We need to verify that such a resolution Φ_f exists.

Lemma 2.10 *Let $\{\Phi_i\}$ be a sequence of resolutions with bounded lengths for a homomorphism $f\colon G \to L$ with $d_{[f]} > 0$ and such that the* $\lim L(\Phi_i)$ *exists. Then there exists a resolution $\Phi\colon \widetilde{K} \to X$ such that* $L(\Phi) = \lim L(\Phi_i)$.

Proof Let K' be the fundamental domain for the action of G on $\widetilde{K}$. We first note that $\lim L(\Phi_i)$ has to be positive. If it were not, we would have that $d(X(\rho), \Phi_i(K')) \to \infty$ as $i \to \infty$. On the other hand, since the image of f is not abelian, there has to exist $g \in G$ so that $\Phi_i(gK')$ belongs to a horoball disjoint from the one containing $\Phi_i(K')$. Hence, there should also exist a translate of $\Phi_i(K')$ that intersects $X(\rho)$ nontrivially, since $\Phi_i(\widetilde{K})$ is connected, but that is impossible under these hypotheses. Therefore, $L = \lim L(\Phi_i) > 0$. By definition,

$$L(\Phi_i) = \sum_{j=1}^{E} \mu_i(e_j),$$

where E is the number of edges in $K^{(1)}$, and (Λ_i, μ_i) is a measured lamination induced by Φ_i. After passing to a subsequence, we may assume that for every $j = 1, \dots, E$, the sequence $\{\mu_i(e_j)\}$ converges, and we denote its limit by ℓ_j. Hence,

$$L = \sum_{j=1}^{E} \ell_j.$$

Since $\{L_i\}$ is a bounded sequence and its limit is positive, we know that all of the vertices of K' are mapped into a bounded neighborhood of $X(\rho)$ by each Φ_i. Pick one of these vertices, say v. We can assume that v is contained in $X(\rho)$ (if it is not, we can always enlarge $X(\rho)$ so as to cover $\Phi_i(v)$). If we denote by X' the fundamental domain of the action of L on $X(\rho)$, then there exists $l_i \in L$ such that $l_i \Phi_i(v) \in X'$. Furthermore, $l_i \Phi_i(K')$ is contained in a bounded neighborhood of X'. Hence, for every vertex v of K' we can find a convergent subsequence of $l_i \Phi_i(v)$. We define

$$\Phi(v) = \lim_{i \to \infty} l_i \Phi_i(v), \quad v \text{ a vertex of } K'.$$

We extend Φ f–equivariantly, and get a resolution. Clearly, Φ has the desired length, since l_i's are isometries and preserve the lengths of Φ_i's. □

We are interested in the closure, $\mathrm{LIM}(K)$, of all Λ_f's corresponding to Φ_f, for all $f\colon G \to L$, inside the space of projectivized measured laminations, $\mathcal{PML}(K)$. The space $\mathrm{LIM}(K)$ is compact.

Remark 2.11 Let us back up to the sequence $\{f_i\}$ which converges to an action ρ of G on a tree T_∞. For each of f_i's we take a short resolution Φ_{f_i} and a corresponding lamination Λ_{f_i}. This sequence will converge in $\mathcal{PML}(K)$ to some Λ. There exists a resolution $\Phi\colon \widetilde{K} \to T_\infty$ so that the leaves of Λ are mapped to points.

Definition 2.12 Let $f\colon G \to L$ be a homomorphism between two limit groups. We define a collection of moves that are allowed to be performed on f:

(M1) Precompose f by an element of $\mathrm{Mod}(G)$.

(M2) Conjugate f by an element of L.

(M3) If there is an abelian vertex group A in Δ_G such that $f(A)$ is contained in a parabolic subgroup of L, then redefine f on A so that the new homomorphism coincides with f on the adjacent edge groups.

(M4) Bending: suppose there is an edge group E in Δ_G whose image is nontrivial and contained in a parabolic subgroup of L. If the edge X_E corresponding to E separates Δ_G, we conjugate the image of one of the connected components of $\Delta_G \backslash X_E$ by an element of L that commutes with $f(E)$. If edge X_E is nonseparating, ie corresponds to an HNN extension, we multiply the image of the Bass–Serre generator by an element of L that commutes with $f(E)$.

We say that homomorphisms $f, g\colon G \to L$ are equivalent, $f \sim g$, if there is a sequence $f = f_0, f_1, \dots, f_n = g\colon G \to L$ such that f_{i+1} is obtained from f_i by performing one of the moves (M1)–(M4).

Remark 2.13 If we take L to be a free group then the equivalent homomorphisms are the ones that differ by an element of $\mathrm{Mod}(G)$ and are conjugates of each other. Why did we need to add moves (M3) and (M4)? Let us look at the following example.

Example 2.14 Let $G = \langle a, b, s, t \mid w(a, b) = s, [s, t] = 1\rangle = F_2 *_{\mathbb{Z}} \mathbb{Z}^2$, where the word $w(a, b)$ is chosen so that G is a limit group. Consider the sequence of homomorphisms $f_i\colon G \to G$ defined by $f_i(a) = a$, $f_i(b) = b$, $f_i(s) = s$, and $f_i(t) = t^i$. This sequence has the property that $d_i \to \infty$, but all of its members are embeddings. This is where we will need move (M4).

Remark 2.15 Note that move (M3) is not a special case of (M1). In the previous example all f_i's are equivalent to the homomorphism that has t in its kernel and fixes all other generators of G.

Definition 2.16 A homomorphism f is *short* if Φ_f is shortest among all Φ_g when $g \sim f$. $\mathrm{LIM}'(K)$ will denote the closure of the set of Λ_f's, for short f's, in $\mathcal{PML}(K)$.

That the shortest homomorphisms exist (not that they are not necessarily unique) can be proved in the exactly same way as Proposition 2.4.

3 Proof of the Main Theorem

It was shown in Section 2 of [6] that, when L is a free group, the Main Theorem can be reduced to the case where G is a limit group. The idea is that if G was not a limit group, then quotients of it could be found that are either proper or whose group of homomorphisms into L was strictly contained in that of G, and hence induction could be applied. Identical argument can be used here, and we from now on consider G to be a limit group. We fix a complex K that reflects the decomposition of G as in Definition 2.2. We will prove our theorem by considering the space $\mathrm{LIM}'(K)$. Our first concern is to find suitable proper quotients of G. The series of lemmas that follow will show how to obtain these quotients depending on the type of the individual laminations in $\mathrm{LIM}'(K)$. We first concentrate on a lamination Λ which is a limit of laminations $\Lambda_{f_i} \in \mathrm{LIM}'(K)$, where $f_i\colon G \to L$ are homomorphisms with nonabelian images and the property that $d_i \to \infty$.

Lemma 3.1 *If Λ is entirely simplicial we form the following quotients:*

- *Abelianize the subgroup carried by a generic leaf in each Cantor set bundle, if these subgroups are nonabelian.*
- *If all of these subgroups are abelian, for each of them we make the following quotients:*
  - *Make a quotient of G by their normal closures.*
  - *Abelianize each vertex group in the splitting of G inherited from Λ.*

At least one of these quotients is proper. Further, for all but finitely many f_i's, a homomorphism equivalent to f_i factors through one of the proper quotients.

Proof We will assume that there is only one family of parallel leaves and let N be its regular neighborhood. N is homeomorphic to $\ell \times [0,1]$, where ℓ is a leaf of the lamination Λ. Denote by ℓ_1 and ℓ_2 the boundary components of N; see Figure 4. Let H be a subgroup of G carried by ℓ_1. H cannot be trivial since that would imply that G is freely decomposable. Suppose $H = \langle h_1, \dots, h_k \rangle$.

(i) If H is nonabelian, so are its conjugates, hence abelianization of a subgroup carried by a generic leaf yields a proper quotient.

Further, we showed in Proposition 2.5 that if a subgroup of G fixes an arc in the limiting tree T_∞, then the image of that subgroup under almost all f_i is abelian. Hence $f_i(H)$ is abelian, and f_i will factor through the abelianization of H, for almost all i.

(ii) If H is abelian and $f_i(H) = 1$, then our homomorphism factors through a proper quotient $G/\langle\langle H \rangle\rangle$. Hence we may assume that there is at least one $j \in \{1, \dots, k\}$ for which $f_i(h_j) \neq 1$. Since H is an abelian group, $f_i(H)$ is contained either in:

(ii1) a parabolic subgroup of L of rank n, or

(ii2) a cyclic subgroup of L generated by a hyperbolic element.

We would like to know what the general position of $\Phi_i(\widetilde{N})$ is within X. Due to the equivariance, the first thing we conclude is that a connected component of $\Phi_i(\widetilde{N})$ is contained in a bounded neighborhood of a horoball B that corresponds to the parabolic subgroup of L which contains the image of H. Further, we can also conclude that, since $\mu_i(\ell_1)/d_i \to 0$, all the points of $\Phi_i(\widetilde{\ell_1})$ are approximately at the same distance from P (same for $\Phi_i(\widetilde{\ell_2})$), where P is the boundary hyperplane of the horoball B.

(ii1) $\Phi_i(\widetilde{\ell_1})$ and $\Phi_i(\widetilde{\ell_2})$ are completely contained in B.

Suppose that $\Phi_i(\widetilde{\ell_2})$ is at a distance proportional to d_i from P. Assume that the fundamental group of the vertex space, call it K_2, which is adjacent to ℓ_2, is not

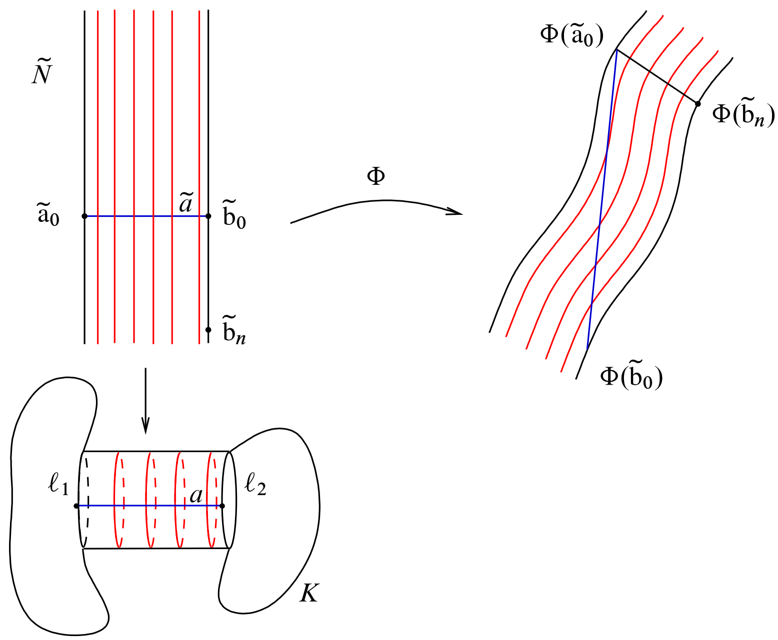

Figure 4: Shortening the simplicial component

abelian. If $f_i(\pi_1(K_2))$ is abelian then f_i factors through the proper quotient of G obtained by abelianizing $\pi_1(K_2)$ and our claim holds. Otherwise, we can find a closed path $p \subset K_2$ whose measure is close to 0 and for which $[f_i(p), f_i(\ell_2)] \neq 1$ (we use the same notation for the elements of the fundamental group and loops they represent). Consider the image under Φ_i of a lift $\widetilde{p}$ of p. One of the endpoints of $\Phi_i(\widetilde{p})$ lies on $\Phi_i(\widetilde{\ell}_2)$, but $\Phi_i(\widetilde{p}) \nsubseteq B$, for if it were we could conclude that $f_i(\ell_2)$ and $f_i(p)$ commute. Therefore the length of $\Phi_i(\widetilde{p})$ is proportional to $2d_i$. On the other hand, $\mu_i(p)$ is virtually zero, contradiction. Hence, if $f_i(\pi_1(K_2))$ is not abelian, the distance from $\Phi_i(\widetilde{\ell}_2)$ to P is much smaller than d_i. The same holds for $\Phi_i(\widetilde{\ell}_1)$.

We now have two cases to consider:

(1) $f_i(\pi_1(K_j))$, $j = 1, 2$, is not abelian.

(2) At least one of $f_i(\pi_1(K_j))$, $j = 1, 2$ is abelian.

Our strategy is as follows. We consider an edge a of the triangulation of K that intersects leaves of the lamination Λ.Further, f_i factors through at least one of the quotients we had made, for otherwise we will be able to decrease the measure of the edge a, hence obtaining a lamination shorter than Λ_{f_i}.

(1) Let a be an edge in $K^{(1)}$ that intersects N nontrivially, and let $a \cap \ell_1 = \{\mathrm{a}_0\}$ and $a \cap \ell_2 = \{\mathrm{b}_0\}$. We know now that $\Phi_i(\widetilde{\ell}_1)$ and $\Phi_i(\widetilde{\ell}_2)$ are contained in the n_i–neighborhood of P, and $n_i/d_i \to 0$ as $i \to \infty$. If we assume that the Euclidean height of P is 1, which we may without loss of generality, then for $x, y \in N_{n_i}(P)$ we will have $|\ln x_{n+1} - \ln y_{n+1}| \le n_i$.

Let τ_i denote the maximal displacement of points on $\Phi_i(\widetilde{\ell}_1)$ under the action of $f_i(\widehat{H})$, where $\widehat{H}$ denotes the centralizer of H. We know that $\tau_i \le \mu_i(\ell_1)$ and that $d(\Phi_i(\widetilde{\mathrm{a}}_0), \Phi_i(\widetilde{\mathrm{b}}_0)) \sim d_i$. If there were a point in the $f_i(h_j)$–orbit of $\Phi_i(\widetilde{\mathrm{b}}_0)$ whose distance to $\Phi_i(\widetilde{\mathrm{a}}_0)$ is smaller than that of $\Phi_i(\widetilde{\mathrm{b}}_0)$, then we could precompose f_i by an appropriate Dehn twist determined by h_j and obtain a resolution shorter than Φ_i, a contradiction (see Figure 4). Hence,

$$d(\Phi_i(\widetilde{\mathrm{a}}_0), \Phi_i(\widetilde{\mathrm{b}}_0)) \le d(\Phi_i(\widetilde{\mathrm{a}}_0), f_i(h_j)^m \Phi_i(\widetilde{\mathrm{b}}_0)) \quad \text{for all } m \in \mathbb{Z}.$$

The orthogonal projections of $\Phi_i(\widetilde{\ell}_1)$ and $\Phi_i(\widetilde{\ell}_2)$ to P will be contained in small Hausdorff neighborhoods of parallel real lines. There exists an element g of the parabolic subgroup of L corresponding to B which will translate these lines towards each other. For i large enough, the translation length of g, call it $\tau_P(g)$, while acting on P is going to be *much* smaller than d_i. We can find $m, k \in \mathbb{Z}$ and a point y on $\Phi_i(\widetilde{\ell}_1)$ so that

$$d(g^m \Phi_i(\widetilde{\mathrm{b}}_0), y') \le \tau_P(g) \quad \text{and} \quad d(f_i(h_j)^k y, \Phi_i(\widetilde{\mathrm{a}}_0)) \le \tau_i.$$

where y' is the orthogonal projection of y onto the horizontal hyperplane in B containing $\Phi_i(\widetilde{\mathrm{b}}_0)$; see Figure 5. We then have

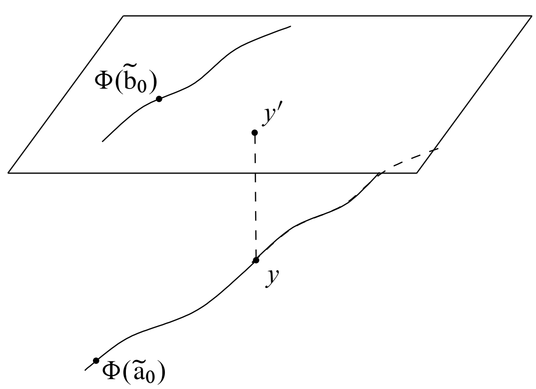

Figure 5: Shortening via bending

$$\begin{aligned} d(f_i(h_j)^k g^m \Phi_i(\widetilde{\mathrm{b}}_0), \Phi_i(\widetilde{\mathrm{a}}_0)) &\le d(g^m \Phi_i(\widetilde{\mathrm{b}}_0), y') + d(f_i(h_j)^k y', \Phi_i(\widetilde{\mathrm{a}}_0)) \\ &\le \tau_P(g) + d(f_i(h_j)^k y', f_i(h_j)^k y) \\ &\qquad + d(f_i(h_j)^k y, \Phi_i(\widetilde{\mathrm{a}}_0)) \\ &\le \tau_P(g) + \left| \ln \frac{(\Phi_i(\widetilde{\mathrm{a}}_0))_{n+1}}{(\Phi_i(\widetilde{\mathrm{b}}_0))_{n+1}} \right| + \tau_i \ll d_i. \end{aligned}$$

Therefore, performing the bending (a power of it, to be more precise) determined by g and precomposing the obtained homomorphism with a power of the Dehn twist determined by h_j will, contrary to our assumption, give us a homomorphism in the same class with f_i which has a shorter resolution.

(2) In the case that one of $f_i(\pi_1(K_j))$, $j = 1, 2$, is abelian we may in fact assume that $\pi_1(K_j)$ is abelian itself. If not, f_i will factor through the quotient of G obtained by abelianizing $\pi_1(K_j)$. Further, only one K_j can have abelian fundamental group, and without loss of generality we may assume it is K_1. As noted before $\Phi_i(\widetilde{\ell}_2)$ lies in the n_i–neighborhood of P. If we further have that $\Phi_i(\widetilde{\ell}_1)$ also lies in the n_i neighborhood of P, we apply (1). We therefore may assume that the distance from $\Phi_i(\widetilde{\mathrm{a}}_0)$ is proportional to d_i. $\Phi_i(\widetilde{\ell}_1)$ being mapped so deep into the horoball is a consequence of $\mu_i(p)$, for every loop p in K_1, being very short compared to both $\mu_i(\ell_2)$ and d_i. We consider a new f_i–resolution $\Phi_i'\colon \widetilde{K} \to X$ which coincides with Φ_i on $\widetilde{\ell}_2$, but lowers $\widetilde{\ell}_1$. We explain the term "lowers" formally: there are constants $c_1, \dots, c_n \in \mathbb{R}$ so that for every point $x \in \widetilde{\ell}_1$ there is a point $x' \in \widetilde{\ell}_2$ such that $(\Phi_i(x))_j = (\Phi_i(x'))_j + c_j$, $j = 1, \dots, n$. We define $\Phi_i'(x)$ to be the point with coordinates

$$((\Phi_i(x))_1, \dots, (\Phi_i(x))_n, (\Phi_i(x'))_{n+1}).$$

Notice that we did not change the homomorphism f_i in any way, we only changed the resolution. By doing that we shortened the edge a substantially. Unfortunately we increased the measures of closed paths $p \subset K_1$ and ℓ_1. However, if we perform move (M3) on f_i we obtain a homomorphisms with a shorter resolution. Namely, the measure of the path a will be the same as under the new resolution for f_i, but the paths p will now have measures 0. This contradicts our assumption.

This concludes our discussion when $\Phi_i(\widetilde{\ell}_1)$ and $\Phi_i(\widetilde{\ell}_2)$ are contained in B. Further, these arguments remain valid when the intersection of $\Phi_i(\widetilde{N})$ with the boundary horosphere P is a disjoint union of segments.

(ii2) $\Phi_i(\widetilde{\ell}_j)$, $j = 1, 2$, are contained in $X(\rho)$.

If $\Phi_i(\widetilde{\ell}_j)$, $j=1,2$, belong to different connected components of $X(\rho)$, or if $\Phi_i(\widetilde{N})\cap B$ is nonempty we can shorten the resolution by shortening, as above, the distance between the connected components of $\Phi_i(\widetilde{N})\cap P$. Suppose then that $\Phi_i(\widetilde{N})$ is contained in a single connected component of $X(\rho)$. $\Phi_i(\widetilde{N})/L$ is an annulus F in a compact part of X/L with boundary components of length $\ll d_i$ and of length approximately d_i. Without loss of generality we may assume that N had a drum triangulation, and $\Phi_i(\widetilde{N})$ has the same. Further we may assume that the triangulation on N is formed out of k triangles. Each of these triangles has two sides that are extremely long, meaning their lengths are greater than or equal to $d_i-\mu_i(\ell_1)$, or $d_i-\mu_i(\ell_2)$, depending on which boundary component of N the short side of the triangle lies. Let $s_0,\dots,s_k$ denote the lifts of these long sides of the triangles belonging to the same lift of F in $\Phi_i(\widetilde{N}')$, where s_k is a translate of s_0 under an element $g\in L$. Let x be a point on s_0 at distance at least $\mu_i(\ell_1)+k\delta$ and $\mu_i(\ell_2)+k\delta$, along s_0, from $\Phi_i(\widetilde{\ell}_1)$ and $\Phi_i(\widetilde{\ell}_2)$, respectively. There is a point x_1 on s_1 at a distance less than or equal to δ from x (δ–hyperbolicity of X). We continue moving through all the triangles in $\Phi_i(\widetilde{N}')$, and conclude that there is a point $x_k\in s_k$ at a distance less than or equal to $k\delta$ from x. If the distance along s_0 between x and $g^{-1}x_k$ is greater than $k\delta$ then performing a Dehn twist will shorten the side s_0, and all the others, and hence we obtain a shorter resolution.

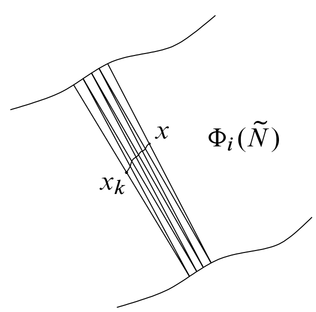

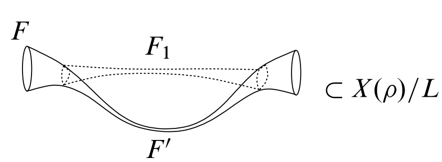

Figure 6: Shortening of the annulus F

Therefore, for the points x,x_k as above $d_{s_0}(x,g^{-1}x_k)<k\delta$. For every $j=0,\dots,k$, let s'_j denote the projection to F of a segment obtained from s_j by removing the ending

subsegments of length $\mu_i(\ell_1)+k\delta$ and $\mu_i(\ell_2)+k\delta$ that are adjacent to $\Phi_i(\widetilde{\ell}_1)$ and $\Phi_i(\widetilde{\ell}_2)$, respectively. For each $x \in s_0'$ we then have a loop based at x of length $\leq 2k\delta$. Hence, boundary loops of F are homotopic to loops of length $\leq 2k\delta$, and there are only finitely many such. Moreover, there is a constant C so that for each two homotopic loops of length $\leq 2k\delta$ there is a homotopy between them of length $\leq C$. Let γ_1 and γ_2 be two of the short loops we have found; the ones based at the initial and the end point of s_0', respectively. Also, let F' denote the part of F between them. As we said earlier there is an annulus F_1 between γ_1 and γ_2 of length $\leq C$, given by the aforementioned short homotopy. We now replace F' by F_1; see Figure 6. We obtain a shorter resolution that corresponds to a homomorphism obtained from f_i by bending (the torus $F_1 \cup F'$ gives us the element which determines the bending).

If Λ has more than one family of parallel leaves, we shorten the resolution as above working on each family at the same time. Hence, we conclude that if Λ is completely simplicial, the subgroups carried by generic leaves are all abelian, and f_i's do not factor through the quotients we made in this case, then the resolutions Φ_i were not shortest as assumed. □

Remark 3.2 We have assumed that Λ induces a splitting of G with two vertex groups. We realize that assumption is not critical, as the case of an HNN extension would be proved in the exactly same way. It is worth noting that bending in the case of HNN extensions, which we perform in order to shorten the resolutions, will yield multiplication by elements of parabolic subgroups. This is exactly what we needed in Example 2.14.

Lemma 3.3 *If Λ has a minimal component Λ_0 of either surface or thin type, we choose a leaf ℓ in it. We form the following quotients:*

- *Trivialize all the loops in ℓ.*
- *Abelianize the fundamental group of the standard neighborhood of Λ_0.*

These quotients are proper quotients of G.

Proof If Λ has a component of thin type, then the quotient of the first kind above has to be proper for otherwise G was freely decomposable.

Suppose now Λ has a surface component. If the loops that we quotiented out were nontrivial in G, we obviously get a proper quotient. Suppose that was not the case. We have a certain number of annuli attached to the surface in our complex K, and to those annuli are maybe attached different components of K, which we will call black boxes.

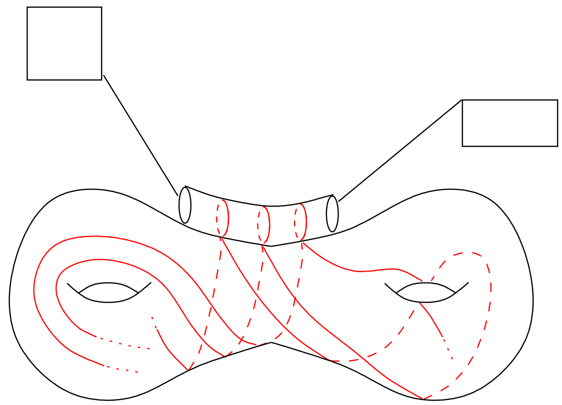

Figure 7: Lamination on the surface with the attached annuli

Let us remind ourselves that the annulus $S^1 \times I$ with a lamination $S^1 \times \{\text{Cantor set}\}$ is attached along an arc transverse to the lamination; see Figure 7.

Since all the loops in the leaves of the lamination that are contained in the annulus, ie $S^1 \times \{\text{pt}\}$, are trivial we may collapse all of them, including the circles that are not contained in the leaves of the lamination. After this collapsing the black boxes are attached either directly to the surface, that is to the components complementary to the lamination, or to a boundary component of the surface. Since the lamination on the surface is the filling one, the components complementary to the lamination are either simply connected or homotopy equivalent to a boundary component. Hence if the fundamental group of one of the black boxes attached to the simply connected complementary component is nontrivial we get a free decomposition of G, which is a contradiction. We conclude that all black boxes are either attached to the boundary components of the surface or have trivial fundamental groups (and hence can be collapsed). In both cases we get a pure surface component and a surface vertex group in a cyclic decomposition of G. This surface vertex group must be conjugate to a QH vertex group in the JSJ of G [6]. The fact that Λ was obtained as a limit of Λ_{f_i}'s and that it is filling on the surface implies that the loops that correspond to each generator of our surface group have positive measures. Proposition 5.8 in [5] guarantees the existence of the set of short generators. That is, by applying Process II to the component of K that carries Λ, we can obtain from K a new band complex K' which resolves the same tree, but so that the fundamental group of the component carrying the lamination is generated by loops short with respect to an interval s of arbitrarily small measure. That the loop p is short with respect to s means that $p = p_1 * \lambda * p_2$, where p_1 and p_2 are contained in s and λ is contained in a leaf of Λ. This now tells us that these generators are shorter with respect to the measure μ_i than the original generators were.

Since the combinatorial type of band complexes that correspond to these two sets of generators are the same we can find a modular automorphism that takes one of these sets into the other. Precomposing the f_i's by this modular automorphism will give us homomorphisms shorter than f_i's, which is a contradiction with our choice of f_i's.

Finally, if Λ has either surface or thin type component then the fundamental group of the regular neighborhood of that component is not abelian, hence its abelianization leads to a proper quotient of G. □

Lemma 3.4 *Suppose Λ has a minimal component Λ_0 which resolves a line. Let N be the standard neighborhood of Λ_0. We form the following quotients:*

- *Abelianize $\pi_1(N)$, if it is nonabelian.*
- *If $\pi_1(N)$ is abelian, make the following quotients:*
  - *Abelianize the subgroup of G generated by $\pi_1(N)$ and all vertex groups adjacent to it.*
  - *Divide by the direct summand of $\pi_1(N)$ that intersects trivially the peripheral subgroup, if such exists.*

At least one of these quotients is proper. Further, for almost all f_i, there is an equivalent homomorphism factoring through one of these quotients.

Proof If $H = \pi_1(N)$ is nonabelian, then its abelianization leads to a proper quotient. Every commutator of elements in H fixes an arc in T_∞, ie $[H, H]$ fixes this arc. Hence, for sufficiently large i, $f_i([H, H])$ is an abelian group. By the Tits alternative the subgroup $f_i(H)$ is either virtually abelian or contains a free group on two generators. If the latter is the case, then $[f_i(H), f_i(H)]$ could not be an abelian group, hence $f_i(H)$ is virtually abelian. Now we have $f_i(H)$ as a virtually abelian subgroup of a torsion-free limit group, which means it is abelian. We conclude that $f_i([a,b]) = 1$ for all $a, b \in H$, and f_i factors through the abelianization of H.

Assume then that H is an abelian group and N is a genuine torus. If Λ also has a simplicial, surface or thin component, we apply Lemma 3.1 and Lemma 3.3. Hence our only concern is when Λ has only toral components, each of which has a torus as a standard neighborhood. Since G is a limit group, these tori can not be glued directly to each other, ie an edge space cannot embed into both of them.

Let H' denote the peripheral subgroup of H. Suppose first H' is contained in a proper direct summand of H. We will call the smallest such H' so that $H = H' \oplus H''$. For each i we define a homomorphism $f_i'\colon G \to L$ that coincides with f_i on all vertex groups in the decomposition of G except for H where we define it to be:

$$f_i'(h) = \begin{cases} f_i(h), & h \in H' \\ 1, & h \in H'' \end{cases}$$

f_i and f_i' are equivalent under our moves, and f_i''s factor through the quotient $G/\langle\langle H''\rangle\rangle$. If, on the other hand, H' is a finite index subgroup of H there exists an edge a of the triangulation of K intersecting an edge space E adjacent to N nontrivially so that $\mu_i(a) \sim d_i$. Let V be the vertex space at the other end of E. Since $f_i(H) < A$ is abelian, $\Phi_i(\widetilde{N})$ is contained in a bounded neighborhood of the horosphere P bounding the horoball B whose stabilizer is A. We first assume that the minimum distance between points in $\Phi_i(\widetilde{N})$ and points in P is proportional to d_i. We can find a closed path p contained entirely in V and E whose measure is "virtually" 0 and whose image does not belong to the centralizer of $f_i(H)$. If such a path did not exist we would get a proper quotient of G by imposing relations that make $\pi_1(V)$ and H commute, and f_i would factor through this quotient. We consider the lift $\widetilde{p}$ of p having a nonempty intersection with $\widetilde{N}$. One of the endpoints of $\Phi_i(\widetilde{p})$ will belong to $\Phi_i(\widetilde{N})$, but $\Phi_i(\widetilde{p})$ is not completely contained in B, since $f_i(p)$ does not commute with $f_i(H)$. Hence, $\ell(\Phi_i(\widetilde{p})) \geq 2d_i$, contradicting the assumption that $\mu_i(p) \sim 0$. We know now that

$$\Phi_i(\widetilde{N}) \subset N_{n_i}(P) \quad \text{and} \quad \lim_{i\to\infty} \frac{n_i}{d_i} = 0.$$

All the hypotheses of Lemma 3.1(1) are satisfied by our resolutions. Therefore, we find the homomorphisms shorter than f_i belonging to the same equivalence classes. □

Lemma 3.5 *Suppose $\Lambda \in \mathrm{LIM}'(K)$ is a limit of Λ_{f_i}, where $f_i\colon G \to L$ is a sequence of short homomorphisms for which $d_{f_i} \to \infty$. There is a neighborhood $U \subset \mathrm{LIM}'(K)$ of Λ such that if $\Lambda_f \in U$ then f factors through one of the quotients defined in Lemma 3.1, Lemma 3.3 and Lemma 3.4.*

Proof Recall from Proposition 2.5 that if a subgroup of G fixes an arc in the limiting tree T_∞, then the image of that subgroup under almost all f_i is abelian.

If Λ is either entirely simplicial, or if all of its minimal components are toral we proceed as in Lemma 3.1 and Lemma 3.4, respectively.

If Λ has a surface or thin minimal component, we follow the argument due to M Bestvina. We consider a standard neighborhood N of the surface or thin component of Λ. Recall that in Lemma 3.3 we fixed a leaf of this component. A loop in that leaf will fix an arc I in T_∞. Choose the segment I so that every subsegment has the same stabilizer, and let $g \in G$ be the element corresponding to the loop that fixes I. We will argue that if $f_i(g) \neq 1$ for almost all i (and hence do not factor through the first type of quotient made in Lemma 3.3), then almost all f_i map $\pi_1(N)$ into an abelian subgroup of L and hence they all factor through an abelianization of $\pi_1(N)$. Suppose there is an element $h \in G$ such that $J = h(I) \cap I$ is a nonempty interval. Then both g and $g' = hgh^{-1}$ fix J, and hence I since the action on T_∞ is stable. Our remark at the beginning says that $[f_n(g), f_n(g')] = 1$ for almost all n. Since $f_n(h)$ conjugates $f_n(g)$ into $f_n(g')$ and maximal abelian subgroups of L are malnormal, we conclude that all three images must belong to the same abelian subgroup of L. Therefore, $[f_n(h), f_n(g)] = 1$, for almost all n. We get that $f_n(g)$ commutes with $f_n(h)$ whenever $h(I) \cap I \neq \varnothing$. Such elements h generate $\pi_1(N)$ (existence of "small" finite generating set of $\pi_1(N)$, which only uses minimality [5]), and so $f_i(\pi_1(N))$ is abelian for almost all i. □

Proof of the Main Theorem If $\Lambda \in \mathrm{LIM}'(K)$ is such that $\Lambda_{f_i} \to \Lambda$ and $d_{[f_i]} \to \infty$, then our previous four lemmas prove the claim: we have formed finitely many quotients of G and we have found a neighborhood U_Λ of Λ so that whenever $\Lambda_f \in U_\Lambda$ and $d_{[f]} \gg 0$ then a homomorphism equivalent to f factors through one of these quotients. U_Λ's together with neighborhoods of Λ_{k_i}'s cover $\mathrm{LIM}'(K)$. Since this space is compact, it is covered by finitely many of these neighborhoods. Hence we have finitely many quotients through which an element of the equivalence class $\sim$ of any $f\colon G \to L$ with $d_{[f]} > 0$ factors. If $d_{[f]} = 0$, then f factors through abelianization of G.

We now claim that there are only finitely many quotients that do not factor through quotients formed above. Suppose not, and $\{g_i\}$ is a sequence of homomorphisms that do not factor through any of those quotients. This implies that the sequence $\{d_{[g_i]}\}$ is bounded. Proposition 2.4 tells us that in such a sequence we can have only finitely many nonconjugate homomorphisms. We pick a representative of each conjugacy class, say $k_i\colon G \to L$, $i = 1, \dots, l$ and form quotients $K_i = G/\ker(k_i)$. If this quotient is not proper, then k_i was an embedding. This concludes the proof. □

It would be useful to know what happens if we iterate this construction, ie if we apply the Main Theorem to all the quotients we obtained. The following lemma is well known.

Lemma 3.6 *A sequence of epimorphisms between ω–residually free groups eventually stabilizes.*

Proof Take a sequence of epimorphisms between ω–residually free groups:

$$G_1 \to G_2 \to \cdots \to G_n \to \cdots$$

For a fixed free group F the sequence

$$\mathrm{Hom}(G_1, F) \leftarrow \mathrm{Hom}(G_2, F) \leftarrow \cdots \mathrm{Hom}(G_n, F) \leftarrow \cdots$$

eventually stabilizes, ie consists of bijections [6]. Suppose

$$\mathrm{Hom}(G_k, F) \cong \mathrm{Hom}(G_{k+1}, F),$$

where the isomorphism is given by the obvious inclusion of $\mathrm{Hom}(G_{k+1}, F)$ into $\mathrm{Hom}(G_k, F)$. Let us also suppose that $e\colon G_k \to G_{k+1}$ is a proper epimorphism. Let $g \in \ker(e)$. Since G_k is ω–residually free there is a homomorphisms $f\colon G_k \to F$ such that $f(g) \neq 1$. On the other hand, there exists $f'\colon G_{k+1} \to F$ such that $f = f' \circ e$. We now have

$$1 \neq f(g) = f'(e(g)) = 1.$$

Hence, $G_k \cong G_{k+1}$. □

We now form the Makanin–Razborov diagram, Figure 1, except we add an edge issuing from each group in this diagram and ending in L representing the embeddings. Furthermore, each branch of this diagram is finite, as we have just shown.

Remark 3.7 At a first glance it may appear as if the proof of the Main Theorem should hold for a broader class of groups, in particular for groups that are hyperbolic relative to a collection of their maximal noncyclic abelian subgroups. There are two problems that appear. The first one is that we do not know that a finitely generated subgroup of a relatively hyperbolic group is finitely presented, and so the proof would need to be modified in order to deal with not necessarily finite complexes. The second problem is that the proof of Lemma 3.6 does not apply in this context. Since this paper was written, Groves has solved this problem [11].

References

[1] **E Alibegović**, *A combination theorem for relatively hyperbolic groups*, Bull. London Math. Soc. 37 (2005) 459–466 MR2131400

[2] **J M Alonso**, **et al**, *Notes on word hyperbolic groups*, from: "Group theory from a geometrical viewpoint (Trieste, 1990)", World Sci. Publ., River Edge, NJ (1991) 3–63 MR1170363 Edited by H. Short

[3] **M Bestvina**, *Degenerations of the hyperbolic space*, Duke Math. J. 56 (1988) 143–161 MR932860

[4] **M Bestvina**, *$\mathbb{R}$–trees in topology, geometry, and group theory*, from: "Handbook of geometric topology", North-Holland, Amsterdam (2002) 55–91 MR1886668

[5] **M Bestvina**, **M Feighn**, *Stable actions of groups on real trees*, Invent. Math. 121 (1995) 287–321 MR1346208

[6] **M Bestvina**, **M Feighn**, *Notes on Sela's work: Limit groups and Makanin–Razborov diagrams* (2003) Available at `http://www.math.utah.edu/~bestvina/research.html`

[7] **M R Bridson**, **G A Swarup**, *On Hausdorff-Gromov convergence and a theorem of Paulin*, Enseign. Math. (2) 40 (1994) 267–289 MR1309129

[8] **I Chiswell**, *Introduction to Λ-trees*, World Scientific Publishing Co., River Edge, NJ (2001) MR1851337

[9] **F Dahmani**, *Combination of convergence groups*, Geom. Topol. 7 (2003) 933–963 MR2026551

[10] **M Gromov**, *Hyperbolic groups*, from: "Essays in group theory", Math. Sci. Res. Inst. Publ. 8, Springer, New York (1987) 75–263 MR919829

[11] **D Groves**, *Limit groups for relatively hyperbolic groups. II. Makanin-Razborov diagrams*, Geom. Topol. 9 (2005) 2319–2358 MR2209374

[12] **O Kharlampovich**, **A Myasnikov**, *Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Algebra 200 (1998) 472–516 MR1610660

[13] **O Kharlampovich**, **A Myasnikov**, *Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Algebra 200 (1998) 517–570 MR1610664

[14] **J W Morgan**, **P B Shalen**, *Valuations, trees, and degenerations of hyperbolic structures. I*, Ann. of Math. (2) 120 (1984) 401–476 MR769158

[15] **F Paulin**, *Outer automorphisms of hyperbolic groups and small actions on $\mathbb{R}$–trees*, from: "Arboreal group theory (Berkeley, CA, 1988)", Math. Sci. Res. Inst. Publ. 19, Springer, New York (1991) 331–343 MR1105339

[16] **Z Sela**, *Diophantine geometry over groups. I. Makanin-Razborov diagrams*, Publ. Math. Inst. Hautes Études Sci. (2001) 31–105 MR1863735

[17] **H Zieschang**, *Alternierende Produkte in freien Gruppen*, Abh. Math. Sem. Univ. Hamburg 27 (1964) 13–31 MR0161901

University of Michigan
530 Church St, Ann Arbor, MI 48109-1043, USA

`eminaa@umich.edu`

`http://www.math.lsa.umich.edu/~eminaa`

Proposed: Martin Bridson Received: 31 August 2004
Seconded: Benson Farb, Walter Neumann Revised: 20 November 2006