

An Introduction to Limit Groups

H J R Wilton

March 3, 2005

1 Introduction

It is very natural in the study of discrete groups to ask for a description of the ‘variety’ $\text{Hom}(G, H)$, for finitely generated groups G and H . For example:

1. Recall that a group G is *Hopfian* if every epimorphism $G \rightarrow G$ is an automorphism, and *co-Hopfian* if every monomorphism $G \rightarrow G$ is an automorphism. For example, free groups are Hopfian, but not co-Hopfian. Setting $G = H$ and understanding $\text{Hom}(G, G)$ could lead to a proof that G is Hopfian or co-Hopfian.
2. An algorithmic understanding of $\text{Hom}(G, H)$ could be useful in solving the isomorphism problem.
3. Consider a system of equations Φ over a group H , given by words

$$w_i(x_1, \dots, x_n)$$

in n unknowns. Then it’s easy to see that solutions to Φ are precisely in bijection with $\text{Hom}(G(\Phi), H)$, where $G(\Phi)$ is the group with presentation

$$\langle x_1, \dots, x_n \mid w_1, w_2, \dots \rangle.$$

This so-called ‘algebraic geometry’ over H is also the starting point for the study of the first-order theory of H .

One of the oldest and hardest problems in the first-order theory of groups is the ‘Tarski problem’, which asks which groups have the same first-order theory as the free group. This problem was recently solved by Zlil Sela (see [18], [19], [12], [13], [14], [15] and [16]). The aim of this series of talks is to explain his description of $\text{Hom}(G, \mathbb{F})$, for \mathbb{F} a free group of rank at least 2. (In fact, much of the theory carries through to describe $\text{Hom}(G, H)$ for H torsion-free hyperbolic.) This description takes the form of a Makanin-Razborov diagram, which will be defined in the second talk.

Limit groups are a crucial feature of Makanin-Razborov diagrams. There are many different definitions of limit groups, and another aim of these talks is

to explain all the different definitions, why they are equivalent, and how they fit into the theory of $\text{Hom}(G, \mathbb{F})$.

Rather than plunging straight into Makanin-Razborov diagrams, therefore, we start with a class of groups that will turn out to be equivalent to limit groups, but are easier to define and to work with.

2 Fully residually free groups

2.1 Definition and easy examples

Fix \mathbb{F} a free group of rank $r > 1$.

Definition 2.1 *A finitely generated group G is residually free if, for any non-trivial $g \in G$, there exists a homomorphism $f : G \rightarrow \mathbb{F}$ with $f(g) \neq 1$.*

G is fully residually free or ω -residually free if, for any finite subset $X \subset G$, there exists a homomorphism $f : G \rightarrow \mathbb{F}$ whose restriction to X is injective.

Note that the choice of \mathbb{F} does not matter. Here are some easy examples.

Example 2.2 (Free groups) *If F is a free group then $F \hookrightarrow \mathbb{F}$. In particular, F is fully residually free.*

Example 2.3 (Free abelian groups) *If A is a free abelian group, then any finite subset can be embedded by a homomorphism in \mathbb{Z} . In particular, A is fully residually free. The proof is left as an exercise.*

2.2 Elementary properties

A group G is called *commutative transitive* if the centralizer of any element is abelian; equivalently, for $a, b, g \in G$, if a commutes with g and g commutes with b then a commutes with b .

Recall that a subgroup $H \subset G$ is *malnormal* if, whenever $g \in G - H$,

$$gHg^{-1} \cap H = 1.$$

A group G is *completely separated abelian (CSA)* if every maximal abelian subgroup is malnormal.

Exercise 2.4 *A CSA group is commutative transitive.*

It's easy to see some elementary properties of fully residually free groups.

Proposition 2.5 *Let G be a fully residually free group.*

1. *Any finitely generated subgroup of G is fully residually free.*
2. *G is torsion-free.*

3. Any pair of elements of G generates either a free group or a free abelian group.
4. G is commutative transitive.
5. G is CSA.

Proof: Property 1 is trivial.

Let $g \in G$. Then there exists a homomorphism $f : G \rightarrow G$ with $f(g)$ non-trivial. So $f(g^n) \neq 1$ for all n , so $g^n \neq 1$. This proves 2.

To prove 3, consider $g, h \in G$, and assume $[g, h] \neq 1$. Then there is an epimorphism

$$f : \langle g, h \rangle \rightarrow F_2.$$

Therefore g, h generate a free non-abelian group.

Consider $a, b, g \in G$ with $[g, a] = [g, b] = 1$. There exists a homomorphism $f : G \rightarrow \mathbb{F}$ which is injective on the set

$$\{1, g, [a, b]\}.$$

Now $f([g, a]) = f([g, b]) = 1$ so $f(a)$ and $f(b)$ and $f(g)$ must all lie in the same cyclic subgroup of \mathbb{F} ; in particular, $f([a, b]) = 1$. Therefore $[a, b] = 1$. This gives 4.

Let $H \subset G$ be a maximal abelian subgroup, consider $g \in G$, and suppose there exists non-trivial $h \in gHg^{-1} \cap H$. Let $f : G \rightarrow \mathbb{F}$ be injective on the set

$$\{1, g, h, [g, h]\}.$$

Then $f([h, ghg^{-1}]) = 1$, which implies that $f(h)$ and $f(ghg^{-1})$ lie in the same cyclic subgroup. But in a free group, this is only possible if $f(g)$ also lies in that cyclic subgroup; so $f([g, h]) = 1$, and hence $[g, h] = 1$. By 4 it follows that g commutes with every element of H , so $g \in H$. This proves 5. *QED*

These properties immediately give some examples of groups that aren't fully residually free.

- By 2, any group with torsion is not fully residually free.
- By 3, the fundamental group of the Klein bottle is not fully residually free.
- Direct products are not fully residually free. Specifically, suppose $G = A \times B$, where A is non-trivial and B is non-abelian. Then B is contained in the centralizer of any element in A , so G is not commutative transitive and hence not fully residually free.

A slightly less trivial non-example is the fundamental group of the surface Σ of Euler characteristic -1, which has presentation

$$\langle a, b, c \mid a^2 b^2 c^2 \rangle.$$

By work of Lyndon (see [9]) that, whenever three elements of \mathbb{F} satisfy $a^2 b^2 c^2 = 1$, in fact $abc = 1$. It follows that $\pi_1(\Sigma)$ is not fully residually free.

We shall see that, apart from the fundamental groups of the three simplest non-orientable surfaces, all surface groups are fully residually free.

2.3 Difficult properties

Fully residually free groups have other properties that are much less obvious. They are summarized in the following theorem.

Theorem 2.6 *Suppose G is fully residually free.*

1. G is finitely presented. Indeed, there exists a finite $K(G, 1)$.
2. All abelian subgroups of G are finitely generated.
3. If G is non-abelian then it has a non-trivial cyclic splitting.
4. G is $CAT(0)$ with isolated flats.

We shall prove 1, 2 and 3 in this series of talks, following Sela. There is a simpler independent proof that fully residually free groups are finitely presented due to Guirardel, who shows that fully residually free groups act freely on \mathbb{R}^n -trees (see [8]). Alibegovic and Bestvina proved 4 in [2]; Alibegovic (in [1]) and Dahmani (in [7]) had already independently proved that fully residually free groups were hyperbolic relative to their maximal abelian subgroups.

Note that, by property 3 of the theorem, the only 3-manifold groups that are limit groups are free products of \mathbb{Z} and \mathbb{Z}^3 .

2.4 A criterion in free groups

To prove that a group G is fully residually free, it suffices to show that for any finite $X \subset G - \{1\}$ there exists a homomorphism $f : G \rightarrow \mathbb{F}$ with $1 \notin f(X)$. So a criterion to show that an element of \mathbb{F} is not the identity will be useful.

In the short term, this will make it possible to prove that surface groups are fully residually free. Eventually, it will give a complete constructive characterization of limit groups.

Lemma 2.7 *Let $z \in \mathbb{F} - \{1\}$, and consider an element g of the form*

$$g = u_0 z^{i_1} u_1 z^{i_2} u_2 \dots u_{n-1} z^{i_n} u_n$$

where $n \geq 1$ and, whenever $0 < k < n$, $[u_k, z] \neq 1$. Then $[g, z] \neq 1$ whenever $\min_k |i_k|$ is sufficiently large. In particular, $g \neq 1$.

Choose a generating set for \mathbb{F} so the corresponding Cayley graph is a tree T , and fix $x \in T$. An element $a \in F$ specifies a geodesic $[x, ax] \subset T$. Likewise, a string of elements $a_0, a_1, \dots, a_n \in \mathbb{F}$ defines a path

$$[x, a_0 x] \cdot [a_0 x, a_0 a_1 x] \cdot \dots \cdot [a_0 \dots a_{n-1} x, a_0 \dots a_n x]$$

in T , where \cdot denotes concatenation of paths.

Here is the idea behind the lemma. In the path corresponding to the defining expression for g , a subarc corresponding to a power of z lies in a translate of




Figure 1: The proof of lemma 2.7.

$\text{Axis}(z)$. A high power of z pushes the path a long way along this translate; the next u_k term then pushes the path onto a new translate.

Proof of lemma 2.7: Note that the case where either u_0 or u_n commutes with z follows from the case where neither do; therefore assume $[u_0, z], [u_n, z] \neq 1$.

For $0 \leq k \leq n$, let

$$g_k = u_0 z^{i_1} u_1 z^{i_2} u_2 \dots u_{k-1} z^{i_k} u_k.$$

Note that $g_n = g$. Also, set $g_{-1} = 1$. Every non-trivial element of \mathbb{F} acts hyperbolically on T ; fix some $x \in \text{Axis}(z) \subset T$.

Let γ be the path in T associated with the expression for g , as above. To be precise, for $1 \leq k \leq n$ let

$$\alpha_k = [g_{k-1}x, g_{k-1}z^{i_k}x]$$

and

$$\beta_k = [g_{k-1}z^{i_k}x, g_k x].$$

Also write $\beta_0 = [x, g_0 x]$. Then

$$\gamma = \beta_0 \cdot \alpha_1 \cdot \beta_1 \cdot \dots \cdot \beta_{n-1} \cdot \alpha_n \cdot \beta_n.$$

Let $l_c(w)$ denote the cyclically reduced length of an element $w \in \mathbb{F}$ (with respect to the fixed generating set). There are some useful observations to be made about α_k and β_k .

1. The length of α_k is $l_c(z^{i_k}) = |i_k|l_c(z)$. The length of β_k is $d(g_{k-1}z^{i_k}x, g_k x) = d(x, u_k x)$.
2. Each α_k is contained in $g_{k-1}\text{Axis}(z)$.

3. No β_k is contained in any single translate of $\text{Axis}(z)$. The start point of β_k lies in $g_{k-1}\text{Axis}(z)$. The end point of β_k lies in $g_k\text{Axis}(z)$. Let x_k be the last point of β_k in $g_{k-1}\text{Axis}(z)$. Let y_k be the first point of β_k in $g_k\text{Axis}(z)$.

Henceforth, assume $|i_k|l_c(z) > d(x, u_{k-1}x) + d(x, u_kx)$ whenever $1 \leq k \leq n$. The precise statement from which the lemma follows is

$$g\text{Axis}(z) \neq \text{Axis}(z).$$

Suppose $g\text{Axis}(z) = \text{Axis}(z)$. Since T is a tree, $x_k = y_{k-1}$ for some k ; otherwise,

$$[x, x_0] \cdot [x_0, y_0] \cdot [y_0, x_1] \cdot \dots \cdot [y_{n-1}, x_n] \cdot [x_n, y_n] \cdot [y_n, x]$$

is a non-trivial loop.

By definition y_{k-1} lies in the image of β_{k-1} ; since $g_{k-1}x$ is an end-point of β_{k-1} it is immediate that

$$d(y_{k-1}, g_{k-1}x) \leq d(x, u_{k-1}x).$$

Likewise the image of β_k contains x_k . Therefore

$$d(x_k, g_{k-1}z^{i_k}x) \leq d(x, u_kx).$$

Now the triangle inequality gives

$$\begin{aligned} |i_k|l_c(z) &= d(g_{k-1}x, g_{k-1}z^{i_k}x) \\ &\leq d(y_{k-1}, g_{k-1}x) + d(x_k, g_{k-1}z^{i_k}x) \\ &\leq d(x, u_{k-1}x) + d(x, u_kx) \end{aligned}$$

contradicting the previous assumption on $|i_k|$. *QED*

2.5 Surface groups

This subsection is devoted to showing that the fundamental groups of closed orientable surfaces are fully residually free.

In fact, a larger class of groups than just surface groups are shown to be limit groups. Let F be a free group of rank at least 2, and fix a cyclic subgroup $\langle z \rangle \subset F$ that is closed under taking roots. Consider the amalgam

$$F *_{\langle z \rangle} F.$$

Such groups are called *pinched word groups*.

Example 2.8 *Let Σ be the closed surface of genus 2. Then there's an obvious simple closed curve that realizes $\pi_1(\Sigma)$ as a pinched word group.*

Since all orientable surfaces of higher genus cover the surface of genus 2, it suffices to show that $\pi_1(\Sigma)$ is fully residually free.

Proposition 2.9 *Pinched word groups are fully residually free.*

The assumption that $\langle z \rangle$ is closed under taking roots is necessary; otherwise G is not commutative transitive.

A key part of the proof will be the automorphisms naturally associated to an abelian splitting, called Dehn twists.

Definition 2.10 *Suppose $G = A *_C B$, and $z \in B$ centralizes C . Then the Dehn twist associated to z is the automorphism δ_z defined by*

$$\delta_z(a) = a$$

for $a \in A$ and

$$\delta_z(b) = zbz^{-1}$$

for $b \in B$.

Similarly, if $G = A *_C B$ and $z \in A$ centralizes C , the Dehn twist associated to z is the automorphism δ_z defined by

$$\delta_z(a) = a$$

for $a \in A$ and

$$\delta_z(t) = tz$$

for t the stable element.

Proof of proposition 2.9: Without loss, assume $F = \mathbb{F}$. Let $X \subset G - \{1\}$ be finite. It suffices to show that there exists a homomorphism $h : G \rightarrow \mathbb{F}$ such that $1 \notin h(X)$.

Let $f : G \rightarrow \mathbb{F}$ be the obvious retraction; let $\delta_z : G \rightarrow G$ be the Dehn twist in z which is the identity on the first factor.

Consider $x \in X$. By the structure theorem for amalgamated free products, x has normal form

$$x = a_0 z^{i_1} b_1 z^{j_1} a_1 \dots b_{n-1} z^{j_{n-1}} a_n z^{i_n} b_n$$

where the a_k lie in the first copy of F , the b_k lie in the second copy, $i_k, j_k \in \mathbb{Z}$, and furthermore:

1. for $k > 0$, $a_k \in F - \langle z \rangle$;
2. for $k < n$, $b_k \in F - \langle z \rangle$.

Then for $m \in \mathbb{Z}$,

$$\delta_z^m(x) = a_0 z^{i_1+m} b_1 z^{j_1-m} a_1 \dots b_{n-1} z^{j_{n-1}-m} a_n z^{i_n+m} b_n z^{-m},$$

so

$$f \circ \delta_z^m(x) = a_0 z^{i_1+m} b_1 z^{j_1-m} a_1 \dots b_{n-1} z^{j_{n-1}-m} a_n z^{i_n+m} b_n z^{-m}$$

where all terms are now thought of as elements of \mathbb{F} . Since the a_k and b_k don't commute with z (for $k > 0$ and $k < n$ respectively), this expression for $f \circ \delta_z^m(x)$ satisfies the conditions of lemma 2.7. So for sufficiently large m , the result follows. *QED*

Note that this proof has not used the full power of the free group criterion. The same proof would work for any double

$$G = L *_Z L$$

where L is fully residually free and Z is maximal abelian in L . Eventually, continuing inductively in this vein, we will obtain a complete recursive characterization of fully residually free groups.

3 Makanin-Razborov Diagrams and Algebraic Limit Groups

3.1 Main results and definitions

We return now to our study of $\text{Hom}(G, \mathbb{F})$. To state the main theorem properly, though, a definition is needed.

Definition 3.1 *A sequence of homomorphisms $(f_i : G \rightarrow \mathbb{F})$ is convergent if, for any $g \in G$, $f_i(g)$ is eventually either always trivial or always non-trivial. To a convergent sequence is associated the stable kernel $\underline{\ker} f_i$ of elements that are eventually trivial.*

An (algebraic) limit group (over \mathbb{F}) is any group of the form

$$L = G / \underline{\ker} f_i$$

for (f_i) a stable sequence of homomorphisms $G \rightarrow \mathbb{F}$.

Example 3.2 *Any fully residually free group is a limit group. Suppose G is fully residually free, and let*

$$S_1 \subset S_2 \subset \dots \subset G$$

be an exhaustion of G by finite sets. Then for each i there exists a homomorphism $f_i : G \rightarrow \mathbb{F}$ injective on S_i . Then (f_i) is a convergent sequence with trivial stable kernel.

Here are the two main theorems of this series of talks.

Theorem 3.3 (Finite width) *For any finitely generated group G that isn't free there exists a finite collection of proper epimorphisms*

$$\{q_i : G \rightarrow L_i\}$$

with each L_i a limit group, such that any homomorphism $G \rightarrow \mathbb{F}$ factors through one of the q_i , after precomposing with some automorphism of G . If G is not a limit group, no precomposing is necessary.

Theorem 3.4 (Finite length) *Any sequence of epimorphisms of limit groups*

$$L_1 \rightarrow L_2 \rightarrow \dots$$

eventually stabilizes.

Iterating the construction of theorem 3.3 gives a tree of epimorphisms through which any homomorphism $G \rightarrow \mathbb{F}$ factors, ‘twisting’ with automorphisms at each stage. By theorem 3.4, the branches of the tree end in free groups after finitely many epimorphisms.

3.2 Finite length

We’ll start with the proof of theorem 3.4, which, over free groups, admits a neat simplification.

Lemma 3.5 *Consider a sequence of epimorphisms*

$$G_1 \rightarrow G_2 \rightarrow \dots$$

The corresponding sequence of monomorphisms

$$\mathrm{Hom}(G_1, \mathbb{F}) \leftarrow \mathrm{Hom}(G_2, \mathbb{F}) \dots$$

eventually stabilizes.

Proof: Identifying \mathbb{F} with the fundamental group of a hyperbolic punctured sphere exhibits an embedding

$$\mathbb{F} \hookrightarrow \mathrm{PSL}_2(\mathbb{R}),$$

which lifts to an embedding $\mathbb{F} \hookrightarrow \mathrm{SL}_2(\mathbb{R})$; this in turn induces an embedding

$$\mathrm{Hom}(G_i, \mathbb{F}) \hookrightarrow \mathrm{Hom}(G_i, \mathrm{SL}_2(\mathbb{R}))$$

for each i .

Pick a presentation

$$\langle s_1, \dots, s_n \mid r_1, r_2, \dots \rangle$$

for G . Any $f \in \mathrm{Hom}(G, \mathrm{SL}_2(\mathbb{R}))$ can be thought of as a set of choices for $f(s_1), \dots, f(s_n)$ that satisfy the relation r_1, r_2, \dots . These relations are polynomial conditions in $\mathrm{SL}_2(\mathbb{R})$, so $\mathrm{Hom}(G_i, \mathrm{SL}_2(\mathbb{R}))$ is identified with a subvariety of $\mathrm{SL}_2(\mathbb{R})^n$. The decreasing sequence of subvarieties

$$\mathrm{Hom}(G_1, \mathrm{SL}_2(\mathbb{R})) \supset \mathrm{Hom}(G_2, \mathrm{SL}_2(\mathbb{R})) \supset \dots$$

terminates by Hilbert’s Basis Theorem. *QED*

Proposition 3.6 *Limit groups are fully residually free.*

Proof: Let L be a limit group, and let G and f_i be as in the definition. Consider a (generally infinite) sequence of epimorphisms

$$G = G_0 \rightarrow G_1 \rightarrow \dots \rightarrow L$$

obtained by adding one relation at a time. Let G' be such that

$$\text{Hom}(G', \mathbb{F}) = \text{Hom}(L, \mathbb{F}).$$

Then all but finitely many f_i factor through G' , since each added relation is killed by almost all f_i . Therefore all but finitely many f_i factor through L , and each non-trivial element of L is killed by only finitely many f_i . Therefore L is fully residually free. *QED*

This shows that the classes of algebraic limit groups and fully residually free groups coincide, and allows us to prove theorem 3.4.

Proof of theorem 3.4: Consider a proper epimorphism

$$q : L \rightarrow L'$$

for L residually free. Then

$$\text{Hom}(L, \mathbb{F}) \supsetneq \text{Hom}(L', \mathbb{F});$$

for given $k \in L$ with $q(k) = 1$, there exists $f \in \text{Hom}(L, \mathbb{F})$ with $f(k) \neq 1$, so $f \notin \text{Hom}(L', \mathbb{F})$. The theorem now follows from lemma 3.5 and the fact that limit groups are residually free. *QED*

3.3 An attempt to prove finite width

Having proved the finite length of the Makanin-Razborov diagram, we turn to its finite width. The idea is to compactify $\text{Hom}(G, \mathbb{F})$, following ideas contained in [6].

To this end, let $\mathcal{Q}(G)$ be the set of epimorphisms

$$q : G \rightarrow H,$$

where $q_1 : G \rightarrow H_1$ and $q_2 : G \rightarrow H_2$ are regarded as equivalent if there is an isomorphism $i : H_1 \rightarrow H_2$ with $q_2 = i \circ q_1$. Alternatively, $\mathcal{Q}(G)$ can be regarded as the set of normal subgroups of G , and we will usually work from this point of view. There's an obvious map

$$\text{Hom}(G, \mathbb{F}) \rightarrow \mathcal{Q}(G)$$

that sends each homomorphism to its kernel. To compactify $\text{Hom}(G, \mathbb{F})$, therefore, we start by topologising $\mathcal{Q}(G)$.

The interpretation of $\mathcal{Q}(G)$ as the normal subgroups of G identifies it with a subset of 2^G , the power set of G . Equip 2^G with the product topology –

the weakest topology with respect to which the projections onto the factors are continuous. To work with this topology, we need to understand it better. Fix a finite generating set for G , and let $B_n \subset G$ be the ball of radius n in the word metric. For a subset $X \subset G$, and a positive integer n , consider the set of subsets

$$U(X, n) = \{Y \subset G \mid B_n \cap X = B_n \cap Y\}.$$

Lemma 3.7 *The topology induced by the neighbourhoods $U(X, n)$ is the product topology.*

Proof: The product topology is generated by the sub-basis consisting of the open sets

$$V(g) = \{Y \subset G \mid g \in Y\}$$

together with

$$V'(g) = \{Y \subset G \mid g \notin Y\}$$

for all $g \in G$. Now, for any $X \subset G$ and n ,

$$U(X, n) = \bigcap_{g \in B_n \cap X} V(g) \cap \bigcap_{g \in B_n - X} V'(g).$$

Conversely, for $g \in G$ of word-length n ,

$$V(g) = \bigcup_{g \in X \subset B_n} U(X, n)$$

and

$$V'(g) = \bigcup_{g \notin X \subset B_n} U(X, n).$$

QED

These equivalent perspectives make it easy to read off the properties of the topology.

Theorem 3.8 *Consider 2^G endowed with the product topology.*

1. 2^G is compact.
2. 2^G is metrizable.
3. A sequence $X_i \subset G$ converges if and only if every $g \in G$ is eventually always in X_i or eventually not in X_i . The limit is the set of elements that are eventually in X_i .
4. $\mathcal{Q}(G) \subset 2^G$ is closed, hence compact.

Proof: Part 1 is Tychonoff's theorem.

Given $X, Y \subset G$, define

$$d(X, Y) = e^{-n}$$

where n is the greatest integer such that $B_n \cap X = B_n \cap Y$. This gives 2.

Part 3 is just exactly what it means to converge in the usual definition of the product topology.

Suppose now that X_i is a sequence of normal subgroups, converging to $X \subset G$. Then for $g, h \in G$, let n be greater than the length of g, h and gh . Then for all sufficiently large i ,

$$g, h \in X \cap B_n \subset X_i$$

and

$$gh \in X_i \cap B_n \subset X,$$

so X is a subgroup. Similarly, for $g \in G$ and $h \in X$, $ghg^{-1} \in X$, so X is normal. This proves 4. *QED*

Note that it follows from 3 that

$$G/\ker f_i \rightarrow G/\underline{\ker} f_i$$

in $\mathcal{Q}(G)$. In particular, the set of limit groups for G could equivalently be defined as the closure of the free groups in $\mathcal{Q}(G)$.

Remark 3.9 *In a sense, the philosophically correct interpretation of parts 2,3 and 5 of proposition 2.5 is that being torsion-free, commutative transitive and CSA are closed properties in $\mathcal{Q}(G)$.*

We are now ready to attack theorem 3.3, in the case where G is not a limit group.

Proof of theorem 3.3 when G is not limit group: Consider $\mathcal{F}(G)$, the closure of the set of free groups in $\mathcal{Q}(G)$. Since G is not a limit group, every epimorphism in $\mathcal{F}(G)$ is proper. So $\{V(g)|g \in G\}$ is an open cover for $\mathcal{F}(G)$. Let

$$\{V(g_1), \dots, V(g_n)\}$$

be a finite subcover. Then every homomorphism

$$f : G \rightarrow \mathbb{F}$$

factors through one of the quotients

$$q_i : G \rightarrow G/\langle\langle g_i \rangle\rangle.$$

The only complication is that the quotients $q_i(G)$ may not be limit groups.

Replacing G with its residually free quotient

$$G/\bigcap_{f:G \rightarrow \mathbb{F}} \ker f$$

it can be assumed that G is residually free. Now $\text{Hom}(q_i(G), \mathbb{F})$ is strictly contained in $\text{Hom}(G, \mathbb{F})$, so one may apply lemma 3.5 and induction to conclude that each $q_i(G)$ has a collection of finitely many proper quotients through which any homomorphism to \mathbb{F} factors. *QED*

This argument fails when G is a limit group, since then there exist convergent sequences with trivial stable kernel, so $\{V(g)|g \in G\}$ is no longer a cover of $\mathcal{F}(G)$. We need to find a subspace $\mathcal{F}'(G) \subset \mathcal{F}(G)$, covered by $\{V(G)|g \in G\}$, such that any $q \in \mathcal{F}(G)$ is related to some $q' \in \mathcal{F}'(G)$ by an automorphism. To do this, we use a geometric analogue of the techniques of this section.

4 Geometric Limit Groups

4.1 The space of trees

Recall the basic definitions and results of the theory of trees.

Definition 4.1 *A (real) tree is a geometric metric space (T, d) in which every geodesic triangle is a tripod. Equivalently, T is geodesic and satisfies Gromov's four-point condition:*

$$d(w, x) + d(y, z) \geq \min(d(w, y) + d(x, z), d(w, z) + d(x, y))$$

whenever $w, x, y, z \in T$. For details of the equivalence of these definitions see, for example, [5].

A (real) G -tree is a real tree with an action of G by isometries.

A tree T is non-degenerate if it is not a point. A G -tree T is *trivial* if G fixes a point of T . T is *minimal* if it contains no proper G -invariant subtrees.

Lemma 4.2 *If T is non-trivial then T contains a separable unique minimal subtree.*

For the proof of this lemma see, for example, [3].

Consider the set $\mathcal{A}(G)$ of non-trivial minimal real G -trees. $\mathcal{A}(G)$ is endowed with the *equivariant Gromov-Hausdorff topology*, defined as follows. A sequence (T_n, d_n) of G -trees converges to a G -tree (T, d) if and only if, for any $\epsilon > 0$ and any finite subsets $K \subset T$, $P \subset G$, there exist $K_n \subset T_n$ and bijections $b_n : K_n \rightarrow K$ such that

$$|d_n(gb_n(x_n), b_n(y_n)) - d(gx_n, y_n)| < \epsilon$$

whenever $x_n, y_n \in K_n$ and $g \in P$. This can be thought of as saying that larger and larger subtrees of T_i coincide with subtrees of T .

Let $\mathbb{P}\mathcal{A}(G)$ be the projectivization of $\mathcal{A}(G)$, so (T, d) is identified with $(T, \lambda d)$ for all $\lambda > 0$.

Fix a set of generators for \mathbb{F} with respect to which the Cayley graph is a tree $T_{\mathbb{F}}$ with the word metric $d_{\mathbb{F}}$. A homomorphism $f : G \rightarrow \mathbb{F}$ defines a G -tree, with the action given by left multiplication. Since $G/\ker f$ acts freely, if $f \neq 1$ then T_f is non-trivial. Let $T_f \in \mathbb{P}\mathcal{A}(G)$ be the equivalence class of the minimal G -invariant subtree. The space of interest is $\mathcal{J}(G)$, the closure of

$$\{T_f | 1 \neq f \in \text{Hom}(G, \mathbb{F})\}$$

in the equivariant Gromov-Hausdorff topology.

Compactness of the space $\mathcal{T}(G)$ was proved first in [10], in which the limits are constructed using convex hulls. A quicker and more general method, though, is to use non-standard analysis.

4.2 Ultraproducts

An *ultrafilter* is a finitely additive probability measure

$$\omega : 2^{\mathbb{N}} \rightarrow \{0, 1\}$$

An ultrafilter ω is *non-principal* if, whenever $S \subset \mathbb{N}$ is finite, $\omega(S) = 0$.

Lemma 4.3 *Non-principal ultrafilters exist.*

For this and all subsequent results in this section, see chapter I.5 of [5].

Let X be any topological space. For points $x_n, x \in X$, write $\lim_{\omega} x_n = x$ if, for any open set $U \subset X$,

$$\omega\{n \in \mathbb{N} | x_n \in U\} = 1.$$

The point x is called the *ultralimit* of the sequence x_n (with respect to ω).

Lemma 4.4 *Fix ω a non-principal ultrafilter. If X is a compact metric space then every sequence has an ultralimit (with respect to ω).*

Proof: Let x_n be a sequence in X , and suppose x_n has no ultralimit. Then every $x \in X$ has an open neighbourhood U_x with $\omega\{n \in \mathbb{N} | x_n \in U_x\} = 0$. Now $\{U_x | x \in X\}$ is an open cover of X ; let $\{U_{x_1}, \dots, U_{x_n}\}$ be a finite subcover. But then

$$1 = \omega(\mathbb{N}) \leq \sum_i \omega\{n | x_n \in U_{x_i}\} = 0$$

a contradiction. *QED*

Let (X_n, d_n, x_n) be a sequence of pointed metric G -spaces (ρ_n denotes the G -action). Let

$$Y \subset \prod_n X_n$$

be the subspace of sequences (y_n) such that $d_n(x_n, y_n)$ is uniformly bounded. The space Y admits a pseudometric defined by

$$D((y_n), (z_n)) = \lim_{\omega} d_n(y_n, z_n).$$

The associated metric space is denoted (X_{ω}, d_{ω}) and is called the *ultraproduct* of the sequence (X_n, d_n, x_n) .

Lemma 4.5 *Let ω be an ultrafilter and $(X_n, d_n, x_n), (X_{\omega}, d_{\omega})$ as above.*

1. *If the (X_n, d_n) are geodesic then so is (X_{ω}, d_{ω}) .*

2. If each (X_n, d_n) is an \mathbb{R} -tree then (X_ω, d_ω) .

Proof: Consider elements $y = [(y_n)], z = [(z_n)] \in X_\omega$. Assuming each X_n is geodesic, let $\gamma_n : [0, d_n(y_n, z_n)] \rightarrow X_n$ be a geodesic from y_n to z_n . Define $\gamma : [0, d_\omega(y, z)] \rightarrow X_\omega$ by

$$t \mapsto [\gamma_n(\frac{d_n(y_n, z_n)}{d_\omega(y, z)}t)].$$

Note that this is well-defined since, for any point w_n on γ_n ,

$$d_n(x_n, w_n) \leq 2d_n(x_n, y_n) + d_n(x_n, z_n).$$

Moreover, γ is a geodesic, since for $s, t \in [0, d_\omega(y, z)]$,

$$d_\omega(\gamma(s), \gamma(t)) = \lim_\omega d_n(\gamma_n(\frac{d_n(y_n, z_n)}{d_\omega(y, z)}s), \gamma_n(\frac{d_n(y_n, z_n)}{d_\omega(y, z)}t)) = |s - t|$$

as required. This proves 1.

Assertion 2 follows immediately from Gromov's four-point condition. *QED*

Suppose each (X_n, d_n) admits a G -action. Then the action $G \times Y \rightarrow Y$ given by

$$g(y_n) = (g \cdot y_n)$$

descends to an action on (X_ω, d_ω) by isometries.

Ultraproducts are useful in this context because they provide limits in the Gromov topology.

Lemma 4.6 *If (X_n) is a sequence of G -spaces let X_ω be the ultraproduct for some choice of base-points. Suppose $X \subset X_\omega$ is a separable G -equivariant subspace. Then some subsequence of (X_n) converges to X in the Gromov topology.*

Proof: Let $S \subset X$ be a countable dense subset, and let

$$S_1 \subset S_2 \subset \dots \subset S$$

be an exhaustion of S by finite subsets. Let

$$P_1 \subset P_2 \subset \dots \subset G$$

be an exhaustion of G by finite subsets. Define $I_n \subset \mathbb{N}$ to consist of those $i \in \mathbb{N}$ for which

$$|d_i(gx_i, y_i) - d(gx, y)| < \frac{1}{n}$$

for all $g \in P_n$ and $x, y \in S_n$. By definition, $\omega(I_n) = 1$. Let n_1 be the least element of I_1 , and inductively define n_i to be the least element of I_i not to be contained in $\{n_1, \dots, n_{i-1}\}$. The subsequence (X_{n_i}) now converges to X in the Gromov topology. *QED*

4.3 Compactness of $\mathcal{T}(G)$

The only remaining tricky detail is to ensure that the ultralimit is non-trivial. This is achieved by carefully controlling the base-points and scaling the metric, and is the key trick of the ‘Bestvina-Paulin’ method.

Let $f_n : G \rightarrow \mathbb{F}$ be a sequence of homomorphisms, and T_n the corresponding sequence of minimal G -trees with the usual word metric $d_{\mathbb{F}}$. Fix a generating set S for G . Consider the function $\sigma_n : T_n \rightarrow \mathbb{R}$ given by

$$\sigma_n(x) = \max_{g \in S} d_{\mathbb{F}}(x, f_n(g)x);$$

let

$$\delta_n = \inf_{x \in T_n} \sigma_n(x).$$

Each tree T_n is simplicial, and the function σ_n is integer-valued on vertices and mid-points of edges, and linear in between. Therefore σ_n attains its infimum on T_n , say at x_n . Since S is finite, for some $g_0 \in S$,

$$\omega\{n \in \mathbb{N} \mid d_n(x_n, g_0 x_n) = 1\} = 1;$$

that is, g_0 ω -almost always realizes the maximum in the definition of $\sigma_n(x_n)$.

Equip T_n with the modified metric $d_n = d_{\mathbb{F}}/\delta_n$. By the results of the previous section, the ultralimit (T_ω, d_ω) of the sequence

$$(T_n, d_n, x_n)$$

is a metric G -tree. Furthermore, consider a point $y = [(y_n)] \in T_\omega$. Then

$$d_\omega(y, g_0 y) = \lim_{\omega} d_n(y_n, g_0 y_n) \geq \lim_{\omega} d_n(x_n, g_0 x_n) = 1;$$

in particular, $g_0 y \neq y$, so T_ω is non-trivial. Let $T \subset T_\omega$ be the minimal G -invariant subtree, which is separable. Then a subsequence of T_n converges to T in the equivariant Gromov-Hausdorff topology. This is a limit for the sequence T_n in $\mathcal{T}(G)$. Henceforth, denote the metric on T by d .

4.4 Geometric limit groups

Definition 4.7 *Let T be a real G -tree in $\mathcal{T}(G)$ as constructed in the previous section; let $\ker T$ be the kernel of the action of G on T . A geometric limit group is any group of the form*

$$G/\ker T$$

for such a T .

To extract information from this geometric picture, a careful analysis of equivariant Gromov-Hausdorff is needed. The results of this analysis are summarized in the following technical theorem.

Theorem 4.8 *Let T_n be a sequence of non-trivial minimal G -trees arising from homomorphisms $f_n : G \rightarrow \mathbb{F}$ as above, converging to a non-trivial minimal tree T . Then the following hold.*

1. *If T is not a line then $\varprojlim (f_n) = \ker T$.*
2. *The stabilizer in $G/\ker T$ of a tripod is trivial.*
3. *Stabilizers in $G/\ker T$ of non-degenerate arc in T are free abelian.*
4. *If $J \subset I$ are non-degenerate arcs in T then $\text{Stab}_{G/\ker T}(I) = \text{Stab}_{G/\ker T}(J)$.*
5. *T is a line if and only if, for all sufficiently large n , f_n has non-trivial abelian image.*

The proof of this theorem is omitted.

The first consequence of the theorem is that the two different notions of limit group coincide.

Corollary 4.9 *The set of algebraic limit groups and the set of geometric limit groups coincide.*

Proof: If L is a finitely generated free abelian group then L is both an algebraic and a geometric limit group. It follows from parts 1 and 5 of the theorem that, for L non-abelian, the two notions coincide. *QED*

Furthermore, the theorem ensures that T has the properties we shall require to apply Rips theory to study the action in detail.

Definition 4.10 *A real G -tree T is stable if, for every descending sequence of non-degenerate subtrees*

$$T \supset T_1 \supset T_2 \supset \dots$$

the corresponding sequence of pointwise stabilizers

$$\text{Stab}_G(T) \subset \text{Stab}_G(T_1) \subset \text{Stab}_G(T_2) \subset \dots$$

eventually stabilizes.

A G -tree is very small if it is non-trivial, minimal, stable, has abelian (non-degenerate) arc stabilizers, and trivial tripod stabilizers.

Corollary 4.11 *For a geometric limit group $L = G/\ker T$, the tree T is a very small L -tree.*

5 The Shortening Argument

5.1 Definitions and the statement

The shortening argument is a difficult trick that closely analyses the action of G on T , to force the kernel to be non-trivial. In this section, we'll just try to outline some of the argument, and complete the proof of the finite width theorem.

Definition 5.1 A generalized abelian decomposition for group G is a finite graph of groups Δ with abelian edge groups and three classes of vertices.

1. Surface vertices are the fundamental groups of compact surfaces with boundary. It is required that the surface carry a pseudo-Anosov automorphism; that is, it is either a torus with a single boundary component, or has Euler characteristic at most -2 . Edges adjoining surface vertices are infinite cyclic, and identified with the fundamental groups of boundary components.
2. Abelian vertices are finitely generated abelian groups. For A an abelian vertex, let $P(A)$ be the subgroup generated by incident edge groups. Then define the peripheral subgroup to be

$$\overline{P(A)} = \bigcap \{\ker(f) \mid f \in \text{Hom}(A, \mathbb{Z}), f(P(A)) = 0\}.$$

3. All other vertices are designated rigid.

Definition 5.2 Let Δ be a generalized abelian decomposition for G . The associated modular group $\text{Mod}(\Delta)$ is the subgroup of $\text{Aut}(G)$ generated by:

1. inner automorphisms of G ;
2. Dehn twists of edges of Δ ;
3. unimodular (that is, determinant 1) automorphisms of abelian vertices A , which are the identity on the peripheral subgroup;
4. automorphisms of surface vertices arising from automorphisms of the underlying surface that fix boundary components (note that these induce well-defined group automorphisms, since the base-point can be taken in a boundary component).

The modular group of G , $\text{Mod}(G)$, is the group of automorphisms of G generated by the modular automorphisms of all generalized abelian decompositions of G .

Definition 5.3 A homomorphism $f : G \rightarrow \mathbb{F}$ is equivalent to all homomorphisms of the form $i \circ f \circ \alpha$ where $i \in \text{Inn}(\mathbb{F})$ and $\alpha \in \text{Mod}(G)$.

Fix a generating set S for G . The length of f is defined to be

$$|f| = \max_{g \in S} l(f(g)),$$

where l is word-length in \mathbb{F} .

A homomorphism $f : G \rightarrow \mathbb{F}$ is short if its length is minimal in its equivalence class.

This is precisely the notion we need.

Theorem 5.4 (See [11]) Suppose G is freely indecomposable. Let $f_i : G \rightarrow \mathbb{F}$ be a convergent sequence of short homomorphisms. Then

$$\varprojlim f_i \neq 1.$$

Suppose $\varprojlim f_i = 1$. Then G has a very small action on T . The plan is to use this action to find an automorphism to shorten the f_i .

5.2 A flavour of the proof

The proof relies heavily on the Rips theory of finitely-generated-group actions on trees. Bestvina and Feighn, in [4], give a more delicate version of the shortening argument that only requires Rips theory for finitely presented groups. An alternative proof would involve quoting Guirardel's result that limit groups are finitely presented, then applying the usual shortening argument, using only the finitely presented Rips theory.

We start with the simplest examples of non-simplicial free group actions on trees.

Example 5.5 (Abelian type) *Suppose $T \cong \mathbb{R}$ and $G \cong \mathbb{Z}^n$. There exists a faithful homomorphism $G \rightarrow \text{Isom}(\mathbb{R})$ mapping each generator to something algebraically independent of the other generators.*

Example 5.6 (Surface type) *Let Σ be a surface (or, more generally, a 2-orbifold), and \mathcal{F} a minimal foliation on Σ with transverse measure μ . This induces a pseudometric on $\tilde{\Sigma}$ given by*

$$d(x, y) = \inf_{\alpha} \mu(\alpha)$$

where α lifts to a path from x to y . The associated metric space is a tree, on which $\pi_1(\Sigma)$ acts freely.

In fact, these are the only examples we need to worry about.

Theorem 5.7 *Let G be a finitely generated, freely indecomposable group with an very small action on a minimal tree T . Then T is covered by orbits of a finite collection of subtrees T_1, \dots, T_n such that:*

1. $gT_i \cap T_j$ is at most one point if $i \neq j$;
2. $gT_i \cap T_i$ is either T_i or at most one point;
3. the action of $G_i = \text{Stab}(T_i)$ on T_i is either of abelian type, or of surface type, or simplicial.

G has a graph of groups decomposition with:

1. vertices corresponding to orbits of points of T with non-trivial stabilizer;
2. vertices corresponding to orbits of T_i ;
3. edges corresponding to orbits of simplicial edges of T_i ;
4. edges corresponding to points of intersection of orbits of T_i .

The idea is to find automorphisms that shorten the length of a geodesic $[x, gx]$, for any generator g . We'll just do a couple of cases.

Suppose T is a tree of abelian type. Fix $\epsilon > 0$. Suppose g_1 is the generator with the longest translation length, and g_2 has the second longest translation length. Then there exists k such that

$$|g_1 + kg_2| < |g_1|.$$

Now replace g_1 by $g_1 + kg_2$. Proceeding in this manner, every generator can be given translation length less than ϵ . Since the translation length of a generator in the approximating simplicial trees converges to the translation length in the limit tree, this contradicts the assumption that the f_i are short.

It is also true that, in the surface case, the translation lengths of generators in the limit can be made arbitrarily small.

Now suppose T is the Bass-Serre tree of a splitting of the form

$$G = A *_C B$$

for abelian C . In this case there's a uniform lower bound on the translation lengths of elements in T , so we have to shorten in the approximating spaces.

Let e be the edge fixed by C . Assume x is fixed by A , and y is the vertex fixed by B .

Any generator is of the form

$$g = a_0 b_1 a_1 \cdots b_n a_n$$

where $|g| = 2n$. Fix $z \in C$. The segments approximating e become arbitrarily close to the axes of z . Then there exist $m(n)$ so that

$$f_n(z^{m(n)})y_n \rightarrow x.$$

For large n , $d_n(x_n, f_n(g)x_n)$ is approximately $2n$, while $d_n(x_n, f_n \circ \delta_z^{m(n)}(g)x_n)$ is approximately 0.

5.3 The proof of finite width

Armed with Sela's shortening argument, we are now in a position to prove the finite width theorem.

Proof of theorem 3.3: If G is abelian, then any homomorphism to \mathbb{F} factors through projection onto a factor. So it can be assumed that G is non-abelian.

Suppose the theorem is proved for freely decomposable G . If $G = G_1 * G_2$ is a non-trivial free product decomposition of G then $\text{Hom}(G, \mathbb{F}) \supseteq \text{Hom}(G_1, \mathbb{F})$. By induction, every homomorphism $G_1 \rightarrow \mathbb{F}$ factors through some finite set

$$\{q_i : G_1 \rightarrow L_i\}$$

of proper factors, so every homomorphism $G \rightarrow \mathbb{F}$ factors through

$$\{q_i * \text{id}_{G_2} : G \rightarrow L_i * G_2\}.$$

So G can be assumed freely indecomposable.

Let $\mathcal{L} \subset \mathcal{T}(G)$ be the subspace of linear subtrees. Theorem 4.8 implies that $\mathcal{L}(G)$ is open in \mathcal{T} , and

$$V(g) = \{T \in \mathcal{T}(G) - \mathcal{L} \mid g \in \ker T\}$$

is open. Consider the subspace $\mathcal{T}'(G) \subset \mathcal{T}(G)$, the closure of the space of trees arising from short homomorphisms $G \rightarrow \mathbb{F}$. By the shortening argument, $\{U(G) \mid g \in G\} \cup \{\mathcal{L}\}$ is an open cover of $\mathcal{T}'(G)$. By compactness, therefore, there exists a finite subcover

$$\{U(g_1), \dots, U(g_n), \mathcal{L}\}.$$

Now any short homomorphism $G \rightarrow \mathbb{F}$ factors through one of

$$\{G \rightarrow G/\langle\langle g_i \rangle\rangle\} \cup \{G \rightarrow G/[G, G]\}.$$

The argument is now concluded as in our first attempt. *QED*

6 JSJ Decompositions

6.1 Splittings of limit groups

A JSJ decomposition is, loosely, a universal splitting for a group G . To construct JSJ splittings for limit groups, therefore, we will need to understand their splittings, or equivalently, their actions on simplicial trees. Because limit groups are CSA, the picture is greatly simplified. Here is the key lemma.

Lemma 6.1 *Consider a one-edge splitting of a limit group G over an abelian subgroup, and $M \subset G$ a non-cyclic maximal abelian subgroup.*

1. *If $G = A *_C B$ then M is conjugate into A or B .*
2. *If $G = A *_C$ and M is not conjugate into a A then for some conjugate M^g of M ,*

$$G = A *_C M^g.$$

Note: we don't yet know that abelian subgroups of G are finitely generated.

Proof: Suppose $G = A *_C B$, and let T be the Bass-Serre tree. Assume M is not conjugate into either A or B . Because M is abelian, it follows from a coarse classification of group actions on trees that H either fixes a line in T , or a point on the boundary. If it fixed a point on the boundary, then there would be an increasing sequence of edge stabilizers

$$C_1 \subset C_2 \subset \dots \subset M.$$

But each C_i is conjugate to C and M is malnormal, a contradiction; so H fixes a line in T , called the axis of M . Conjugating if necessary, C is the stabilizer of an edge in the axis. But M acts as

$$M = M' \oplus \mathbb{Z}$$

where M' fixes the axis, so $M' \subset C$ and, by commutative transitivity, $C \subset M$. So C fixes the whole axis. But there is only one orbit of edges, so there exists $a \in A - C$ with $aCa^{-1} = C$. Since M is malnormal, $a \in M$, a contradiction.

Now suppose $G = A*_C$, and M is not conjugate into A . As before, M preserves an axis in the Bass-Serre tree T , C can be assumed to lie in M and fix the axis, and $A \cap M = C$. Now if t is the stable letter of the splitting, then without loss of generality, $tCt^{-1} = C$. Since M is malnormal, it $t \in M$ and the result follows. *QED*

We will see that this lemma tells us that non-cyclic abelian splittings of limit groups only interact in a very simple way.

Definition 6.2 *A simplicial G -tree is k -acylindrical if the fixed point set of any $g \in G$ has diameter at most k .*

Lemma 6.3 *If G is a limit group and T is a simplicial G -tree with abelian edge stabilizers then, without loss of generality, T is 2-acylindrical.*

Proof: Let $C \subset G$ be the stabilizer of an edge e of T . Since G is commutative transitive, C lies in a unique maximal abelian subgroup M . By the previous lemma, M can be assumed to fix a vertex v of T . Now after a 'slide', e adjoins v . So every element of C can only fix edges adjacent to v . *QED*

6.2 The definition

For a one-edge splitting Γ , $g \in G$ is called Γ -elliptic if g acts elliptically on the associated Bass-Serre tree, or equivalently if g is conjugate into a vertex group of Γ . If \mathcal{S} is a set of splittings of G then $g \in G$ is \mathcal{S} -elliptic if it is Γ -elliptic for all $\Gamma \in \mathcal{S}$.

Let Δ be a generalized abelian decomposition for G . Then $g \in G$ is Δ -elliptic if:

1. g is conjugate into a vertex group of Δ ;
2. if the vertex is surface then g is conjugate into a boundary component;
3. if the vertex is abelian then g is conjugate into the peripheral subgroup.

If g is not Δ -elliptic then there is an obvious one-edge splitting Γ of G such that g is no Γ -elliptic.

Consider the set \mathcal{A} of one-edge splittings of G satisfying:

1. the edge group is abelian;
2. the edge group is closed under taking roots.

Definition 6.4 *An abelian JSJ decomposition for G is a generalized abelian decomposition Δ for G such that the set of Δ -elliptics is the set of \mathcal{A} -elliptics.*

The remainder of this section is devoted to explaining why the existence of abelian JSJ decompositions for limit groups is much simpler than it is for general groups. In fact, the cyclic splittings are as hard as in the general case, but the non-cyclic abelian subgroups are covered by lemma 6.1.

By lemma 6.1, every splitting of \mathcal{A} is, at least, closely related to a splitting in the subset $\mathcal{A}' \subset \mathcal{A}$ of splittings in which every non-cyclic abelian subgroup is elliptic.

6.3 Non-intersecting splittings

Let Γ_1, Γ_2 be a pair of non-trivial one-edge splittings of a group G , with free abelian edge groups C_1, C_2 respectively. Γ_1 is Γ_2 -elliptic if C_1 is conjugate into a vertex group of Γ_2 ; otherwise, Γ_1 is Γ_2 -hyperbolic.

Lemma 6.5 *If Γ_1 is Γ_2 -elliptic then there exists a splitting Γ such that the set of Γ -elliptics is equal to the intersection of the sets of Γ_1 - and Γ_2 -elliptics.*

Proof: Assume that Γ_1 is an amalgamated free product; the HNN-extension case is similar. So

$$G = A_1 *_{C_1} B_1 = A_2 *_{C_2} B_2.$$

Suppose C_1 is conjugate into A_2 ; conjugating A_2 if necessary, it can be assumed that $C_1 \subset A_2$. Γ_2 induces splittings of A_1 and B_2 , denoted Γ_2^A and Γ_2^B respectively. Without loss of generality, $A_1 \cap A_2$ can be assumed to be a vertex group of Γ_2^A ; and likewise $B_1 \cap A_2$ can be assumed to be a vertex group of Γ_2^B . Since $C_1 \subset A_1 \cap A_2, B_1 \cap A_2$, Γ_2^A and Γ_2^B can be joined by an edge stabilized by C_1 to produce Γ , as required. *QED*

The graph of groups Γ is called the *refinement* of Γ_1 by Γ_2 . Note that a similar construction can be carried out when Γ_1 and Γ_2 have more than one edge.

The pair (Γ_1, Γ_2) is *elliptic-elliptic* if Γ_1 is Γ_2 -elliptic and vice versa. Likewise, such a pair can be *elliptic-hyperbolic*, *hyperbolic-elliptic*, and *hyperbolic-hyperbolic*. Elliptic-elliptic pairs should be thought of as disjoint; hyperbolic-hyperbolic pairs should be thought of as intersecting. The other two possibilities can be ruled out, under certain assumptions.

Lemma 6.6 *Let Γ_1 and Γ_2 be as above, and suppose G does not split over any infinite-index subgroups of C_2 . Then (Γ_1, Γ_2) is either elliptic-elliptic or hyperbolic-hyperbolic.*

Proof: Suppose Γ_1 is Γ_2 -elliptic, and Γ_2 is Γ_1 -hyperbolic. As before, assume Γ_1 is an amalgamated free product; the HNN-extension case is similar. Let Γ be the refinement of Γ_1 by Γ_2 .

Suppose the graphs of groups Γ_2^A, Γ_2^B are both trivial; so A_1 and B_1 are both conjugate into vertex groups of Γ_2 . These vertex groups must be distinct, since Γ_2 is a non-trivial splitting; therefore C_1 is conjugate into C_2 , and must be conjugate to a finite index subgroup; but then C_2 must be conjugate into a vertex of Γ_1 .

Assume, therefore, that Γ_2^A is non-trivial; an edge group is a conjugate of a subgroup of C_2 , and also a subgroup of A_1 . This subgroup of C_2 must be of infinite index; since it is also an edge group of Γ , this contradicts the assumption that G doesn't split over infinite-index subgroups. *QED*

The would like to iteratively apply lemma 6.5 to construct the JSJ decomposition; but we don't yet know how to deal with hyperbolic-hyperbolic pairs of splittings.

6.4 Intersecting splittings

A set of splittings is called *intersecting* if, for any splittings Γ, Γ' , there exists a finite chain of splittings

$$\Gamma = \Gamma_1, \dots, \Gamma_n = \Gamma'$$

such that (Γ_i, Γ_{i+1}) is a hyperbolic-hyperbolic pair of splittings.

Henceforth, assume G is freely indecomposable. Suppose $\mathcal{S} \subset \mathcal{A}'$ is a set of intersecting abelian splittings. If Γ has cyclic edge group and Γ' has non-cyclic abelian edge group, then Γ' is assumed to be Γ -hyperbolic. Now by lemma 6.6, Γ is Γ' -hyperbolic.

The model example of an intersecting pair of splittings is a pair of intersecting simple closed curves on a surface. In fact, this is the only example!

Theorem 6.7 *Let \mathcal{S} be a maximal set of intersecting infinite-cyclic splittings of finitely generated G . Then there exists a graph of groups $\Gamma_{\mathcal{S}}$ for G with cyclic edge groups and an orbifold vertex S such that every splitting in \mathcal{S} arises by splitting S along a simple closed curve, and every edge group is identified with a boundary component or cone point of S .*

The intersection of the sets of Γ -elliptics for all $\Gamma \in \mathcal{S}$ is the set of elements that are conjugate into a vertex of $\Gamma_{\mathcal{S}}$ other than S .

S is called an *enclosing vertex* for \mathcal{S} . By assumption, every pair of non-cyclic splittings in \mathcal{A}' is elliptic-elliptic.

6.5 Acylindrical accessibility

To construct the decomposition, now, we simply proceed by induction, splitting repeatedly and replacing intersecting sets of splittings with surface vertices. We need to know that this process terminates. If we knew G was finitely presented, then we could apply a relatively simple accessibility result of Bestvina and Feighn. Since we don't know that G is finitely presented, we have to use a different approach.

Theorem 6.8 (Sela, [17]) *Let G be a non-cyclic freely indecomposable finitely generated group and Γ a k -acylindrical graph of groups for G . Then there is a uniform bound on the number of vertices in the core of Γ .*

By lemma 6.3, every splitting of \mathcal{A}' , their refinements the graphs of their enclosing vertices can be assumed to be 2-acylindrical.

6.6 The construction

In this subsection, we prove the existence of abelian JSJ decompositions for limit groups.

Theorem 6.9 *Every freely indecomposable limit group G has an abelian JSJ decomposition.*

Proof: Associated to any one-edge splitting $\Gamma \in \mathcal{A}'$, there is the maximal intersecting set of splittings \mathcal{S} . If $|\mathcal{S}| > 1$ then Γ is a splitting over \mathbb{Z} , and there is the associated graph $\Gamma_{\mathcal{S}}$. Otherwise, set $\Gamma_{\mathcal{S}}$ to be Γ .

We construct the JSJ decomposition Δ inductively, as follows. Let Γ_1 be a one-edge splitting in \mathcal{A}' . Let \mathcal{S}_1 be the maximal intersecting set of one-edge splittings containing Γ_1 . Set $\Delta = \Gamma_{\mathcal{S}_1}$.

Now let Γ_2 be a one-edge splitting in $\mathcal{A}' - \mathcal{S}_1$. Let \mathcal{S}_2 be the maximal intersecting set of splittings containing Γ_2 . Then by definition, Δ and $\Gamma'_{\mathcal{S}_2}$ are an elliptic-elliptic pair of splittings. Define the new Δ to be the core of the refinement of this pair. Now continue in this way iteratively.

Because every splitting of Γ is without loss 2-acylindrical, it follows that this process eventually terminates. The result is a JSJ-decomposition for G . *QED*

7 Constructive Limit Groups

7.1 The main theorem

Definition 7.1 *Constructive limit groups (CLGs) are finitely generated and defined inductively. A CLG of level 0 is a finitely generated free group. A group G is a CLG of level at most n if one of the following holds.*

1. $G = G_1 * G_2$, for G_1, G_2 CLGs of level at most $n - 1$.
2. G has a generalized abelian decomposition Δ . Δ is assumed to have finitely generated vertex groups and edge groups; furthermore, each edge group is assumed to be maximal abelian on one side of the associated one-edge splitting. There exists a homomorphism $\rho : G \rightarrow G'$, for G' a CLG of level at most $n - 1$, satisfying the following properties:
 - (a) ρ is injective on edge groups;
 - (b) ρ has non-abelian image on surface vertices;
 - (c) ρ is injective on the peripheral subgroups of abelian vertices;
 - (d) ρ is injective on the envelopes of rigid vertices.

Some properties, such as the existence of a finite presentation, are much easier to prove for CLGs than for limit groups. So the next theorem has some profound consequences.

Theorem 7.2 *The sets of limit groups and constructive limit groups coincide.*

The proof that all CLGs are limit groups is tricky but not very enlightening. The idea is to twist ρ with modular automorphisms of Δ , and use lemma 2.7 to prove that G is fully residually free. The other direction, however, makes heavy use of the shortening argument and JSJ decompositions. The rest of this section is devoted to this direction of the proof.

7.2 Modular automorphisms

The most important observation is the relation of the abelian JSJ decomposition to modular automorphisms.

Definition 7.3 *Consider a one-edge splitting Δ , of the form*

$$G = A *_C B$$

or

$$G = A *_C .$$

A generalized Dehn twist associated to a splitting is a Dehn twist, or if A is abelian a unimodular automorphism of A restricting to the identity on the edge group (and B in the first case).

Lemma 7.4 *For any G , $\text{Mod}(G)$ is generated by generalized Dehn twists in one-edge splittings in A .*

Proof: Surface automorphisms are generated by Dehn twists of the surface. Unimodular automorphisms of abelian vertex groups A are generalized Dehn twists in the obvious one-edge splittings of the form

$$G = A *_{\overline{P(A)}} B.$$

Consider Γ a one-edge splitting of G . If the edge group isn't closed under taking roots on one side then there is an immediate contradiction of property 4 of lemma 2.5. Let A be a vertex group of Γ , and suppose $Z = Z_A(C) \subsetneq C$. Let Γ' be the one-edge splitting obtained by expanding Z and contracting C . For example, if Γ is the splitting $G = A *_C B$ then Γ' is

$$G = A *_Z (Z *_C B).$$

Then the edge group of Γ' is closed under taking roots, and any Dehn twist in Γ arises as a Dehn twist in Γ' . *QED*

7.3 Abelian subgroups

It's fairly easy to see that abelian subgroups of CLGs are finitely generated. To show that limit groups are constructive, we'll need to see that the same is true of limit groups.

Proposition 7.5 *Abelian subgroups of limit groups are finitely generated.*

Lemma 7.6 *If G is a limit group with factor set*

$$\{q_i : G \rightarrow G_i\}$$

and $H \subset G$ is any subgroup such that every homomorphism $H \rightarrow \mathbb{F}$ factors through some $q_i|_H$ (without pre-composing with automorphisms) then, for some i , $q_i|_H$ is injective.

Proof: Suppose not. Then there exists non-trivial $h_i \in \ker(q_i|_H)$. But G is a limit group, so there exists $f : G \rightarrow H$ injective on $\{1, h_1, \dots, h_n\}$. But $f|_H$ factors through some $q_i|_H$, so $f(h_i) = 1$, a contradiction. *QED*

Proof of proposition 7.5: All limit groups are torsion free, so it remains to show that any abelian subgroup $A \subset G$ is finitely generated. For $\alpha \in \text{Mod}(G)$ there exists a finitely generated subgroup $A_\alpha \subset A$ and a retraction $r_\alpha : G \rightarrow A_\alpha$ such that $\alpha|_A$ agrees with an inner automorphism on $A \cap \ker(r_\alpha)$. For, suppose A is non-cyclic; then by lemma 6.1, A can be assumed to be conjugate to a vertex in Δ . Now, if it's true for $\alpha, \beta \in \text{Mod}(G)$ then setting $A_{\alpha\beta} = (A_\beta)_\alpha$ and $r_{\alpha\beta} = r_\alpha \circ r_\beta$, it's true for $\alpha\beta$ too.

The homomorphism

$$\prod_{\alpha \in \text{Mod}(G)} r_\alpha : G \rightarrow \prod_{\alpha} A_\alpha$$

has finitely generated image, since G is finitely generated. Therefore $A = A_0 \oplus A_1$, where A_1 is finitely generated and each r_α is trivial on A_0 . Now by lemma 7.6, for some homomorphism q in the factor set of G , $q|_{A_0}$ is injective. Now by induction $q|_{A_0}$ is finitely generated, hence so are A_0 and A . *QED*

7.4 Limit groups are constructive

We're now ready to prove that limit groups are constructive.

Let G be a generic limit group, and fix a generating set S . Let f_i be a sequence in $\text{Hom}(G, \mathbb{F})$ such that f_i is injective on elements of length at most i in the length metric with respect to S . Then $\underline{\ker} f_i = 1$. Choose \hat{f}_i to be short maps equivalent to f_i . By theorem 5.4, $\rho : G \rightarrow G' = G/\underline{\ker} \hat{f}_i$ is a proper epimorphism, and so by induction assume that G' is a CLG. Let Δ be an abelian JSJ decomposition for G . The claim is that Δ and ρ satisfy the conditions of definition 7.1.

Let E be an edge group of the JSJ decomposition. Then it is elliptic in every one-splitting of \mathcal{A} , so all generalized Dehn twists coincide with some inner automorphism on E . Therefore for $g \in E - \{1\}$, $\hat{f}_i(g)$ is conjugate to $f_i(g)$ which is non-trivial for all sufficiently large i ; so $\rho|_E$ is injective. Moreover, E must be maximal abelian in a vertex of the associated one-edge splitting, by commutative transitivity.

Abelian vertices of Δ are finitely generated and free by proposition 7.5. Let $\overline{P(A)}$ be the peripheral subgroup. As in the edge group case, $\overline{P(A)}$ is elliptic

in all relevant one-edge splittings, so $\text{Mod}(G)$ acts as inner automorphisms and $\rho|_{\overline{P(A)}}$ is injective.

Let S be a surface vertex, and suppose $\rho(S)$ is abelian. Then for all sufficiently large i , $\hat{f}_i(S)$ is abelian. But note that every element of $\text{Mod}(G)$ maps S to (a conjugate of) itself, so eventually $f_i(S)$ is abelian, contradicting the triviality of $\ker f_i$.

The envelope of a rigid vertex B is elliptic in every splitting of G , so is preserved up to conjugacy by $\text{Mod}(G)$. Therefore $\rho|_B$ is injective. This completes the proof.

References

- [1] E. Alibegović. A combination theorem for relatively hyperbolic groups, 2003. Preprint.
- [2] E. Alibegović and M. Bestvina. Limit groups are CAT(0), 2004. Preprint.
- [3] M. Bestvina. \mathbb{R} -trees in geometry, topology and group theory. In *Handbook of geometric topology*. North-Holland, Amsterdam, 2002.
- [4] M. Bestvina and M. Feighn. Notes on Sela's work: Limit groups and Makanin-Razborov diagrams, 2003. Preprint.
- [5] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.
- [6] C. Champetier and V. Guirardel. Limit groups as limits of free groups: compactifying the set of free groups, 2004. Preprint.
- [7] F. Dahmani. Combination of convergence groups. *Geom. Topol.*, 86, 2003.
- [8] V. Guirardel. Limit groups and groups acting freely on \mathbb{R}^n -trees, 2003. Preprint.
- [9] R. S. Lyndon. The equation $a^2b^2 = c^2$ in free groups. *Michigan Math. J.*, 6, 1959.
- [10] F. Paulin. Topologie de Gromov équivariante, structures hyperboliques et arbres réels. *Invent. Math.*, 94(1), 1988.
- [11] E. Rips and Z. Sela. Structure and rigidity in hyperbolic groups I. *GAF*, 4(3), 1994.
- [12] Z. Sela. Diophantine geometry over groups III: Rigid and solid solutions. Preprint.
- [13] Z. Sela. Diophantine geometry over groups IV: An iterative procedure for validation of a sentence. Preprint.

- [14] Z. Sela. Diophantine geometry over groups V: Quantifier elimination. Preprint.
- [15] Z. Sela. Diophantine geometry over groups VI: The elementary theory of a free group. Preprint.
- [16] Z. Sela. Diophantine geometry over groups VIII: The elementary theory of a hyperbolic group. Preprint.
- [17] Z. Sela. Acylindrical accessibility for groups. *Invent. Math.*, 129(3), 1997.
- [18] Z. Sela. Diophantine geometry over groups I: Makanin-Razborov diagrams. *Publ. Inst. Hautes Études Sci.*, 93, 2001.
- [19] Z. Sela. Diophantine geometry over groups II: Completions, closures and formal solutions. *Israel J. Math.*, 134, 2003.