

THE STRUCTURE OF LIMIT GROUPS OVER HYPERBOLIC GROUPS

DANIEL GROVES AND HENRY WILTON

ABSTRACT. Let Γ be a torsion-free hyperbolic group. We study Γ -limit groups which, unlike the fundamental case in which Γ is free, may not be finitely presentable or geometrically tractable. We define *model* Γ -limit groups, which always have good geometric properties (in particular, they are always relatively hyperbolic). Given a strict resolution of an arbitrary Γ -limit group L , we canonically construct a strict resolution of a model Γ -limit group, which encodes all homomorphisms $L \rightarrow \Gamma$ that factor through the given resolution. We propose this as the correct framework in which to study Γ -limit groups algorithmically. We enumerate all Γ -limit groups in this framework.

Limit groups over a group Γ (otherwise known as Γ -limit groups) arise naturally when studying algebraic geometry over Γ ; that is to say, sets of homomorphisms $\text{Hom}(G, \Gamma)$, where G is a finitely generated group. We will be concerned with the case in which Γ is a non-elementary, torsion-free hyperbolic group, a case which has been studied extensively by Sela [46] and others [33, 38, 39, 20]. The results of [46] extend Sela's solution to Tarski's problems in the case when Γ is free in [45] *et seq.* (see also [32] *et seq.*).

Since every finitely generated subgroup of Γ is a Γ -limit group, the hyperbolic case immediately presents new problems that do not arise in the free case. One such problem is that hyperbolic groups typically contain subgroups that are finitely generated but not finitely presented [40]. Thus, it is not immediately clear how to give a finite description of a Γ -limit group. Moreover, Γ -limit groups do not in general have the nice geometric properties of limit groups over free groups (which are all toral relatively hyperbolic [1, 8], in particular finitely presented). The geometry of hyperbolic and relatively hyperbolic groups has been integral to the solutions of many algorithmic problems [41, 43, 10, 12,

Date: May 8, 2017.

The work of the first author was supported by the National Science Foundation and by a grant from the Simons Foundation (#342049 to Daniel Groves).

The second author was supported by the EPSRC.

11]) and has been very useful in the algorithmic study of limit groups (see, for example, [19, 20]).

Our goal in this paper is to associate to a Γ -limit group L a *model* Γ -limit group M , which is toral relatively hyperbolic, contains L as a subgroup, and most importantly, encodes a large set of homomorphisms $L \rightarrow \Gamma$ in a natural way. (The reader should bear in mind that the interest of Γ -limit groups arises precisely because they encode large sets of homomorphisms to Γ ; indeed, Sela proved that Γ -limit groups are precisely the finitely generated fully residually Γ groups.) To give the reader a flavour of our results, we state some properties of a model M associated to a Γ -limit group L .

Theorem 0.1. *Let L be a freely indecomposable Γ -limit group. There exists a Γ -limit group M such that:*

- (1) *there is an injection $\eta: L \hookrightarrow M$;*
- (2) *M is toral relatively hyperbolic, in particular finitely presented;*
- (3) *there is a natural homomorphism of modular groups*

$$\Phi : \text{Mod}(L) \rightarrow \text{Mod}(M);$$

- (4) *Φ intertwines η (that is,*

$$\eta(\alpha(g)) = \Phi(\alpha)(\eta(g))$$

for $\alpha \in \text{Mod}(L)$ and $g \in L$).

Recall that the *modular automorphisms* of a Γ -limit group are those that arise naturally from the JSJ decomposition of Γ . They play a special role in the theory of Γ -limit groups, and so items (3) and (4) above demonstrate that M captures much of the information contained in L . Items (1) and (2) are proved in Theorem 4.10 below. Items (3) and (4) are proved in Theorem 5.1.

We now elaborate on the details of our construction. According to [46, Section 2], all of the homomorphisms from L to Γ can be described using a collection of *strict* resolutions (see Definition 2.21 below). These take the form of a sequence of *strict* maps

$$L_0 \rightarrow L_1 \rightarrow \cdots \rightarrow L_n$$

where L_0 is (a quotient of) L , the L_i are all Γ -limit groups and L_n admits a (fixed) strict map to Γ . In particular, Sela shows that every Γ -limit group admits a strict resolution, and we use these to build model Γ -limit groups canonically.

It is important to note that we need to relax two of Sela's conditions. Sela insists that the maps in resolutions are surjective, but we allow our resolutions to include non-surjective maps. Moreover, Sela terminates

with a Γ -limit group which is a free product of a free group and freely indecomposable subgroups of Γ . We do not know how to ensure this and keep the features of our construction. These changes seems to be necessary both in order to make model Γ -limit groups geometrically tractable and also to have their structure closely mimic that of the L_i .

Theorem 0.2. *To every strict resolution $L_0 \rightarrow L_1 \rightarrow \dots \rightarrow L_n$ there is canonically associated a strict resolution that fits into a commutative diagram*

$$\begin{array}{ccccccc} L_0 & \longrightarrow & L_1 & \longrightarrow & \cdots & \longrightarrow & L_n \\ \downarrow \eta_0 & & \downarrow \eta_1 & & & & \downarrow \eta_n \\ M_0 & \longrightarrow & M_1 & \longrightarrow & \cdots & \longrightarrow & M_n \end{array}$$

where each (M_i, η_i) satisfies the properties of Theorem 0.1.

This result follows from Theorems 4.10 and 5.1, along with our construction of models from a resolution of L_0 .

Suppose that $\mathcal{R} = L_0 \xrightarrow{\lambda_0} L_1 \xrightarrow{\lambda_1} \dots \xrightarrow{\lambda_{n-1}} L_n$ is a strict resolution of L_0 . The way in which \mathcal{R} describes a large set of homomorphisms from L_0 to Γ is via homomorphisms which *factor through* \mathcal{R} . This means that a homomorphism can be described as a composition of modular automorphisms of the L_i and the given maps λ_i , finally composed with a particular embedding of each of the freely indecomposable free factors of L_n into Γ and an arbitrary map from the free free factor of L_n to Γ . It is important to note that one needs modular automorphisms of L_n in this description. From a homomorphism factoring through \mathcal{R} , the intertwining maps $\Phi_i: \text{Mod}(L_i) \rightarrow \text{Mod}(M_i)$ canonically induce a homomorphism factoring through the model resolution.

We give a brief outline of this paper. In Section 1 we discuss some differences between the work in this paper and a construction that appears in [33] and also in [31]. In Section 2 we recall some background required for the results in this paper. In Section 3 we give the definition of the (sequence of) model(s) built from a strict resolution of a Γ -limit group L , in terms of the universal property that a model satisfies. In Section 4 we provide the construction of models inductively in terms of a graph of groups built from the Grushko and JSJ decompositions of the limit groups appearing in the resolution of L , and prove this construction satisfies the universal property from Section 3. In Section 5 we prove the remaining properties of model Γ -limit groups. In Section 6 we introduce the notion of *relatively immutable* subgroups of toral relatively hyperbolic groups, prove basic results about them and provide an effective enumeration of them. In Section 7 we prove

that quasi-convex enclosures (see Subsection 2.6) can be algorithmically computed. Finally, in Section 8 (See Theorem 8.3) we apply the construction of models to algorithmically enumerate Γ -limit groups via an enumeration of all of the pairs of resolutions as in Theorem 0.2. This enumeration is analogous to the enumeration of limit groups (over free groups) which we provided in [19].

Acknowledgements. We thank the referee for a very careful reading and for many very helpful comments and corrections which improved the manuscript in numerous ways.

1. A CAUTIONARY EXAMPLE

A somewhat similar construction to the one we perform in this paper is described in [33, §§3.5] and [31, §§6.6]. (The construction in these two papers appears to be identical; we refer to the numbering in [33].)

In order to highlight some differences between the construction in [33, 31] and the one in this paper, in this section we describe certain examples. It is worth remarking that these examples are very elementary, and the issues that they exemplify are entirely generic – they will occur for many torsion-free hyperbolic groups Γ and many Γ -limit groups.

We start with a basic fact about hyperbolic groups. For brevity, we call a torsion-free hyperbolic group *rigid* if it admits no non-trivial free or cyclic splitting.

Lemma 1.1. *Any torsion-free word-hyperbolic group G embeds in a rigid word-hyperbolic group Γ . Furthermore, if G is virtually special (in the sense of Haglund–Wise [29]) then we can choose Γ to be virtually special.*

Proof. We may assume that G is non-elementary. Choose pairwise non-conjugate elements $h_1, \dots, h_n \in G$ (with $n \geq 2$) such that, in every nontrivial free or virtually cyclic splitting of G , some h_i is hyperbolic. (Finitely many h_i suffice by standard accessibility results.) By the ping-pong lemma, replacing the elements h_i with proper powers, we may assume that the subgroup H that they generate is free and quasiconvex. Now fix any non-elementary hyperbolic group G_0 without free or cyclic splittings, and let H_0 be a free, quasiconvex, malnormal subgroup of G_0 of the same rank as H . (Examples of such pairs abound.) The amalgam

$$\Gamma := G *_{{H \cong H_0}} G_0$$

is now the fundamental group of a 2-acylindrical graph of hyperbolic groups with quasiconvex edge groups, and hence is hyperbolic by [4].

In any free or cyclic splitting of Γ , G_0 is elliptic by construction, hence G is too. Since $n \geq 2$, no edge is stabilized by H , and so the splitting is trivial.

Finally, if G is virtually special, we may also choose G_0 to be virtually special, and then Γ is virtually special by the main theorem of [50]. \square

We now describe our examples.

Example 1.2. Suppose that F is a finitely generated nonabelian free group, and that w is an element in F which is not a proper power, and so that F admits no splittings over $\{1\}$ or \mathbb{Z} in which w is elliptic. Such elements w are entirely generic and easy to find—they are called *rigid* by Cashen–Macura [7], and the procedure they describe can be used to certify that a given element is rigid.

Let D be the double of F over w , and note that $D = F *_{\langle w \rangle} F$ is the JSJ of H . The group D is an ordinary limit group, since the natural retraction $D \rightarrow F$ is strict. Applying Lemma 1.1, we obtain a torsion-free hyperbolic group Γ that admits no splittings over $\{1\}$ or \mathbb{Z} and in which D embeds.

Now let u be a primitive element of D so that D admits no nontrivial splitting over $\{1\}$ or \mathbb{Z} in which u is elliptic, and let D_1 be the double of D over $\langle u \rangle$. Let $\rho_1: D_1 \rightarrow D$ be the map that identifies the two copies of D in D_1 . This is a strict map.

Let G_1 be a finite-index subgroup of D_1 which does not contain u , and let $C = \langle u^k \rangle = G_1 \cap \langle u \rangle$ (so $k > 1$). Let $G_2 = \rho_1(G_1)$ and note that it is straightforward to check that $G_1 \twoheadrightarrow G_2$ is a strict map, so since G_2 is a subgroup of Γ ,

$$G_1 \rightarrow G_2$$

is a strict resolution of G_1 . In the description from [33], the group G_2 is given a fixed (conjugacy class of) embedding into Γ , which can be chosen to be the inclusion described above.

In this example there are many homomorphisms from G_2 to Γ that are not injections. However, by pre-composing the inclusion of G_2 into Γ with a sequence of Dehn twists in the edge group of the JSJ of G_2 , we get an infinite sequence of non-conjugate embeddings of G_2 into Γ . Needless to say, since these maps are all injections, they do not factor through a proper shortening quotient of G_2 . Thus, in order to describe all of the homomorphisms from G_2 to Γ , and therefore all of the homomorphisms from G_1 to Γ that factor through $G_1 \rightarrow G_2$, one cannot forget about the JSJ decomposition of G_2 .

Remark 1.3. It is also straightforward to come up with an example of a pair $G_1 \rightarrow G_2$ with a strict map, G_2 a subgroup of a torsion-free hyperbolic group with nontrivial JSJ such that G_2 has no proper shortening quotients. One can do this with 3-manifolds, Property (T) hyperbolic groups, the Rips construction, or other ways.

We now describe the construction from [33, §§3.5] for Example 1.2. In their notation, $n = 2$ and $G_2 = H_1$. Note that they *explicitly* insist that all of the maps between the G_i are proper quotients, and that the final H_i are given with fixed embeddings of the H_i into Γ (up to conjugacy).

The primary JSJ of G_1 contains the edge group $\langle u^k \rangle$, whose image in G_2 is u^k . (See §§2.2 for the definition of the canonical primary JSJ decomposition.) The construction from [33] extends the centralizer of u^k in G_2 . This extension of centralizers obviously has presentation

$$\langle G_2, t \mid [t, u^k] = 1 \rangle$$

This procedure is performed for each of the edges in the JSJ of G_1 (which are G_1 -conjugacy classes of intersections of $\langle u \rangle$ with G_1). Since there are no abelian or QH subgroups in the JSJ decomposition of G_1 , and since $n = 2$, we finish with the group \overline{G}_1 , which is this iterated extension of centralizers. Clearly \overline{G}_1 is a Γ -limit group and contains G_1 as a subgroup.

It is then stated in [33]:

“We now extend each subgroup H_i by its quasi-convex closure Γ_i . Denote the obtained group by N . Therefore N is NTQ and total (*sic*) relatively hyperbolic. Each Γ_i is freely indecomposable.”

There is no proof that N is NTQ or total relatively hyperbolic in [33] or [31], only the above assertion. We now explain that when applying this construction to Example 1.2 neither of these assertions is true. We do not provide the definition of an ‘NTQ group’ here, but only note that it is a particular kind of Γ -limit group. The group N constructed in [33] for Example 1.2 is not a Γ -limit group, let alone an NTQ-group.

For Example 1.2, we have $H_1 = G_2$, and only one H_i . It is not clear precisely what “extend each subgroup H_i by its quasi-convex closure” is supposed to mean, but it seems the only reasonable interpretation in this case is to form the graph of groups

$$N = \overline{G}_1 *_{G_2} \Gamma_1.$$

In this example, Γ admits no nontrivial splittings over $\{1\}$ or \mathbb{Z} so the description of how to find Γ_i in [33, §3.4] yields $\Gamma_1 = \Gamma$

However, the element u^k has a nontrivial root in Γ , so if we glue Γ onto \overline{G}_1 along G_2 then the resulting group N is not CSA. Both Γ -limit groups and toral relatively hyperbolic groups are CSA, so this means that it is not a Γ -limit group (so certainly not NTQ) and also not toral relatively hyperbolic. Thus the claim from [33] is false.

Moreover, as described above, the group N cannot be used to describe all of the homomorphisms from G_2 to Γ , nor those from G_1 to Γ which factor through $G_1 \rightarrow G_2$, since not all injective homomorphisms from G_2 to Γ can be extended to homomorphisms from Γ to Γ .

As we stated above, this is a very simple example, and it exhibits behavior that is entirely generic to the situation for Γ -limit groups for torsion-free hyperbolic Γ .

The construction in [33, §§3.5] is very important to the construction in that paper, and the above Example shows that Proposition 11 from [33] (which is [31, Theorem 2]) is not proved in either of these papers, because the construction they make does not have the properties that they claim it does. It appears to be a fundamental construction upon which many of the other algorithms rely (see for example, [33, §§4.4], where they use it to ‘construct algorithmically a finite number of Γ -NTQ systems corresponding to branches b of the canonical Hom-diagram...’). In summary, this appears to be a very serious error in [33].

The construction in this paper takes these issues (and others) into account, and embeds a Γ -limit group (equipped with a strict resolution) into a toral relatively hyperbolic Γ -limit group, as described in Theorems 0.1 and 0.2.

Remark 1.4. Since the first version of this paper was circulated, new versions of [33] and [31] have been posted to the arXiv, which claim to address the issues outlined here. In the interests of pointing out some of the subtleties inherent in the construction made in this paper, we decided to leave this section largely unchanged.

2. PRELIMINARIES

Let Q be a group. A Q -limit group is a limit of finitely generated subgroups of Q in the space of k -generated marked groups, for any fixed k . It is convenient to work with the following notions, which help to elucidate the connection between Q -limit groups and homomorphisms from a finitely generated group G to Q .

Definition 2.1. Suppose that Q and G are groups. A sequence of homomorphisms $\{\rho_n: G \rightarrow Q\}$ is *convergent*¹ if for any $g \in G$ either (i) $g \in \ker(\rho_n)$ for all but finitely many n ; or (ii) $g \notin \ker(\rho_n)$ for all but finitely many n .

The *stable kernel* of $\{\rho_n\}$ is the set $\underline{\text{Ker}}(\rho_n)$ of elements $g \in G$ which are in the kernel of ρ_n for all but finitely many n .

A Q -*limit group* is a group of the form $L = G/\underline{\text{Ker}}(\rho_n)$ where G is a finitely generated group and $\{\rho_n: G \rightarrow Q\}$ is a convergent sequence of homomorphisms.

Definition 2.2. A group G is *equationally noetherian* if for any finitely generated group H there is a finitely presented group \widehat{H} along with an epimorphism $\rho: \widehat{H} \rightarrow H$ so that the map

$$\rho^*: \text{Hom}(H, G) \rightarrow \text{Hom}(\widehat{H}, G)$$

induced by precomposition with ρ is a bijection.

In other words, every homomorphism from \widehat{H} to G factors through ρ . Yet another way of expressing this is that once a tuple of elements of G satisfies the finitely many relations for \widehat{H} it automatically also satisfies the infinitely many relations for H . Thus, the ‘Hilbert Basis Theorem’ holds for equations over G .

Theorem 2.3 (Sela; [46], Theorem 1.22). *Torsion-free hyperbolic groups are equationally noetherian.*

Definition 2.4. Suppose that G and D are groups. We say that D is *fully residually- G* if for every finite set $A \subset D$ there is a homomorphism $\phi_A: D \rightarrow G$ which is injective on A .

It is well known (see, for example, [37, Theorem 2.1]) that if G is equationally noetherian then the class of G -limit groups is exactly the class of finitely generated fully residually- G groups. From this, the following is straightforward.

Corollary 2.5. *Let Γ be a torsion-free hyperbolic group and let L be a Γ -limit group. Then there exists a convergent sequence $\{\rho_n: L \rightarrow \Gamma\}$ with $\underline{\text{Ker}}(\rho_n) = \{1\}$.*

2.1. Abelian subgroups. For the study of Γ -limit groups, the following definition and result are particularly useful.

Definition 2.6. A subgroup H of a group G is *malnormal* if for all $g \in G \setminus H$ we have $H \cap H^g = \{1\}$.

¹Other authors call these ‘stable’ sequences. See, for example, [5, Definition 1.6].

A group G is called a *CSA group* if any maximal abelian subgroup of G is malnormal.

A group G is called *commutative transitive* if whenever $g_1, g_2, g_3 \in G$ with $g_2 \neq 1$ and $[g_1, g_2] = [g_2, g_3] = 1$ then $[g_1, g_3] = 1$.

Our interest in torsion-free CSA groups comes from the following result, which is due to Sela and implicit in [46, Section 1].

Theorem 2.7. *Suppose that Γ is a torsion-free hyperbolic group and that L is a Γ -limit group. Then L is torsion-free, CSA and all abelian subgroups of L are finitely generated.*

It is easy to see that a CSA group is commutative transitive.

Remark 2.8. Being CSA is a closed condition in the space of marked groups, as is being torsion-free. So these properties are straightforward to prove. However, the fact that abelian subgroups are finitely generated is much deeper, and is due to Sela. The reader is referred to [18, Corollary 5.12] for a proof.

2.2. Primary JSJ decompositions. JSJ decompositions are one of the fundamental tools needed to study Γ -limit groups. They were developed in this context by Sela [46]; various other groups of authors have developed related theories [42, 15, 16, 6], and the whole theory was put into a unified context by Guirardel–Levitt [23, 24, 25, 28].

The specific decomposition that we will work with is called the *canonical primary JSJ decomposition*. In this section, we shall adapt the definitions and results of Guirardel–Levitt to define the canonical primary JSJ decomposition and to explain some of its properties.

Definition 2.9. [12, Definition 3.31] Suppose that G is a torsion-free CSA group. An abelian splitting of Γ is called *essential* if whenever E is an edge group and $\gamma \in G$ satisfies $\gamma^k \in E$ for some $k \neq 0$ we already have $\gamma \in E$. An essential splitting is called *primary* if every noncyclic abelian subgroup is elliptic.

In Guirardel–Levitt’s terminology, if we let \mathcal{E} denote the set of root-closed abelian subgroups and \mathcal{A}_{nc} the class of non-cyclic abelian subgroups then the Bass–Serre tree of a primary splitting is an *\mathcal{E} -tree relative to \mathcal{A}_{nc}* . Our goal is to construct a canonical, primary JSJ tree for Γ -limit groups. A related JSJ tree for toral relatively hyperbolic groups was used in [26, §5], in which the Bass–Serre trees of primary splittings are called *RC trees*. Our JSJ decomposition will differ slightly from the RC-JSJ decomposition of [26].

We give the definition and state and prove some of the properties of the primary JSJ decomposition here.

Definition 2.10. [28, Introduction] A primary tree is *universally elliptic* if every edge stabilizer is elliptic in every primary splitting. A *primary JSJ tree* is a universally elliptic primary tree which is *maximal* among all universally elliptic primary trees, meaning that it dominates any other universally elliptic primary tree.

When we speak of a *canonical* primary JSJ splitting, we mean one that is invariant under the natural action of the outer automorphism group. The construction of the canonical primary JSJ tree goes via the canonical *abelian* JSJ tree. An abelian JSJ tree is defined in the same way as a primary JSJ tree, using the set \mathcal{A} of all abelian subgroups in the place of \mathcal{E} ; again, we only consider \mathcal{A} -trees relative to \mathcal{A}_{nc} . Guirardel–Levitt proved that a CSA group G (such as a Γ -limit group) has a canonical abelian JSJ tree, and described its structure [25].

Theorem 2.11 (Guirardel–Levitt). *Let G be a finitely generated, one-ended, torsion-free, CSA group. There exists a canonical abelian JSJ tree T_{ab} , which is bipartite with the following structure:*

- (1) *One class of vertices has maximal abelian stabilizers.*
- (2) *The stabilizers of the other class of vertices are either rigid or quadratically hanging (QH) (defined below).*
- (3) *If e is an edge of T_{ab} and x is an adjacent, non-abelian vertex, then G_e is maximal abelian in G_x .*
- (4) *Furthermore, the incident edge stabilizers at x form a malnormal family.*

In particular, T_{ab} is 2-acylindrical.

Proof. Existence of the canonical abelian JSJ tree is stated as [25, Theorem 5]; see also Proposition 6.3 of the same paper, where it is proved that T_{ab} (denoted there by T_c) is 2-acylindrical. The bipartite structure follows from its description as a tree of cylinders. \square

The *rigid* vertices have the property that they admit no non-trivial abelian splitting relative to \mathcal{A}_{nc} and their incident edge groups. Note that below we will consider vertex groups which are rigid with respect to primary splittings, rather than with respect to all abelian splittings, and this will be the meaning of the word rigid after Definition 2.13.

The *quadratically hanging* vertices are fundamental groups of surfaces, with edge groups attached isomorphically to the infinite cyclic subgroups corresponding to boundary components.

In order to describe the primary JSJ decomposition Δ , we must first introduce the notion of a *socket*.

Definition 2.12. Let Σ be compact hyperbolic surface with $k \geq 0$ boundary components and let d_1, \dots, d_k be generators for the subgroups corresponding to the boundary components. Let n_1, \dots, n_k be positive integers. The group S obtained from $\pi_1 \Sigma$ by attaching an n_i -th root to d_i for each i is called a *socket*. When considered as a subgroup, a socket will always refer to a vertex group in a graph-of-groups decomposition. If S is a vertex group in a graph of groups, we insist that the incident edge groups are infinite cyclic with generators identified with the roots attached to boundary components, and that every root has an incident edge group attached to it. A socket is called *maximal* if it is not properly contained in any larger socket which appears as a vertex in a graph-of-groups decomposition.

It is a usual feature of JSJ decompositions that the flexible vertices are QH subgroups. In the case of essential splittings, however, the flexible vertices are sockets, because of the requirement that the edges be root-closed.

There is an important difference between socket groups for abelian splittings and socket groups for primary splittings, which is the difference between our primary JSJ splitting and the RC-JSJ from [26]. Edge stabilizers in the RC-JSJ tree are required to be elliptic in every abelian splitting rel \mathcal{A}_{nc} , whether or not this splitting is RC (primary, in our terminology). For the primary JSJ tree, we require edge stabilizers to be elliptic in every primary splitting. The difference can be seen if a vertex group is a once-punctured Klein bottle, with the adjacent edge group corresponding to the puncture. Cutting along the non-separating, two-sided simple closed curve yields a primary splitting whose edge group is elliptic in all other primary splittings of the once-punctured Klein bottle group (relative to the puncture group), but not in all abelian splittings, as witnessed by cutting along a separating curve. In the RC-JSJ of [26], the once-punctured Klein bottle group would be a flexible vertex, whereas in our primary JSJ decomposition, we cut along this curve and replace this once-punctured Klein bottle with two vertex groups – one corresponding to an annular neighbourhood of the curve, and one corresponding to the rest of the surface, and we have two edge groups corresponding to the two (isotopic) curves that join these two vertex groups. The new vertex group corresponding to the surface will be a rigid socket group.

As noted in [28, §9.5], this distinction happens on orbifolds exactly when the surface is either a Klein bottle, a once-punctured Klein bottle, or a Klein bottle with one orbifold point. Since we are in the torsion-free CSA setting, the once-punctured Klein bottle is the only case that can

arise. Moreover, if Γ is a one-ended torsion-free CSA group and there is an abelian edge group E in some primary splitting of Γ which is elliptic in all primary splittings of Γ but hyperbolic in some abelian splitting Λ , then E must correspond to the non-separating, two-sided curve on some once-punctured Klein bottle QH subgroup. Indeed, we must get a hyperbolic-hyperbolic pair from E and some edge group in Λ , which leads to a QH subgroup. If we take a maximal QH subgroup, then there can be no non-separating, two-sided curve intersecting the curve for E , which leads us to conclude that the QH subgroup corresponds to a once-punctured Klein bottle.

We refer to these rigid vertex groups obtained from a Klein bottle in the above way as an *exceptional socket*. Note that an exceptional socket is not a maximal socket.

Definition 2.13. Suppose that K is a group and \mathcal{H} is a collection of subgroups. We say that (K, \mathcal{H}) is *rigid* if K does not admit any nontrivial primary splitting relative to \mathcal{H} . If K is a nonabelian vertex group in a graph-of-groups decomposition, we say that K is *rigid* if $(K, \mathcal{K}_e \cup \mathcal{A}_{nc})$ is rigid where \mathcal{K}_e is the set of edge groups adjacent to K .

A vertex group of a primary decomposition that is neither abelian nor rigid is called *flexible*.

Proposition 2.14. *Suppose that G is freely indecomposable, finitely generated, torsion-free and CSA. Any flexible vertex of a primary JSJ decomposition of G is a maximal socket.*

Proof. This can be proved by following the proof from [28, §6], which builds a QH subgroup, then noting that in order to find a primary splitting we need to adjoin the roots to the QH subgroup in order to build a socket. This is also similar to [44, 13, 12], and can be proved by similar techniques to any of those papers. \square

Note that the converse to the above proposition does not hold: there can also be rigid vertices that are sockets, necessarily built by attaching roots to the boundary components of a thrice-punctured sphere. When discussing the JSJ below, it will still be important to insist that any such vertices are either exceptional sockets or maximal sockets.

Remark 2.15. Note that our terminology differs from that of [26], in which the non-abelian vertices are called *rigid with sockets* and *QH with sockets*.

We apply these ideas in our setting. Let Γ be a torsion-free hyperbolic group and let L be a Γ -limit group. In particular, L is a torsion-free CSA group and all abelian subgroups of L are finitely generated.

Proposition 2.16. *If L is a freely indecomposable Γ -limit group then L has a canonical primary JSJ decomposition Δ . Furthermore, Δ has the following properties.*

- (1) *Vertices of Δ fall into three classes (denoted by u_i, v_j and w_k respectively) with the following properties:*
 - (a) Δ_{u_i} *is rigid for each i (and, if it is a socket, is either an exceptional socket or a maximal socket);*
 - (b) Δ_{v_j} *is a maximal socket, for each j ;*
 - (c) Δ_{w_k} *is abelian.*
- (2) Δ *is bipartite, with only abelian vertices adjacent to rigid or socket vertices.*
- (3) *If v is a vertex associated to a nonabelian vertex group Δ_v then every edge group adjacent to Δ_v is maximal abelian in Δ_v .*
- (4) Δ *is 2-acylindrical.*

Conversely, any primary splitting Δ that satisfies the above properties is the canonical primary JSJ.

Proof. We begin by showing that the primary JSJ decomposition exists. We start with the canonical abelian JSJ tree T_{ab} , guaranteed by Theorem 2.11, and explain how to construct T , the Bass–Serre tree of the canonical primary JSJ decomposition Δ , following [26, §5].

For each edge e incident at some abelian vertex w , let \overline{G}_e be the subgroup consisting of the roots of elements of the edge stabilizer G_e . The union of the translates $\overline{G}_e.e$ is a subtree with every edge adjacent to w . For each edge e , fold the corresponding subtree to a single edge \bar{e} incident at w , with stabilizer \overline{G}_e . The tree T' is then the minimal G -invariant subtree of the result of performing these folds. (More precisely: T' is trivial if G acts with a global fixed point on the folded tree, and otherwise it is *the* minimal G -invariant subtree.)

The other modification to T_{ab} which is required is as follows. Suppose that Δ_{v_i} is a QH-subgroup of T_{ab} with underlying surface a once-punctured Klein bottle. (See the above discussion for an explanation of why the once-punctured Klein bottle is the only relevant case.) Then Δ_{v_i} admits a single primary splitting relative to the adjacent edge group, corresponding to cutting the Klein bottle along the two-sided simple closed curve. Such a splitting is universally elliptic for primary splittings, but not for abelian splittings. Thus, cut along this curve and replace the once-punctured Klein bottle with an exceptional socket group, along with an extra abelian vertex group corresponding to the curve.

Having performed these modifications we have a tree T , which we claim satisfies properties (1)–(4).

We first prove that the resulting tree satisfies properties (1)–(4) from the statement of the result, and then that any splitting which satisfies these properties is a canonical primary JSJ decomposition.

The vertex groups of T come in different kinds, depending on how they arose from the tree T_{ab} . It is clear from the construction that the vertex groups of T can be partitioned into the required types from the statement of (1), and that the resulting structure is bipartite according to the scheme described in (2). Statement (3) follows immediately from the construction and the corresponding property for T_{ab} , and statement (4) follows immediately from statement (3). It follows from conditions (2)–(4) that T is equal to its own (collapsed) tree of cylinders, and then it follows from [28, Corollary 7.4] that T is canonical.

We next prove the converse. Let Δ be a primary splitting satisfying items (1) to (4). We will first show that Δ is a primary JSJ splitting; to do this, we check first that it is universally elliptic among all primary decompositions, and then that it is maximal among all universally elliptic primary decompositions.

To check that Δ is universally elliptic, let Υ be a (without loss of generality one-edge) primary splitting and let e be an edge of Δ . If Δ_e is non-cyclic then it is elliptic in Υ by hypothesis, so we need only consider the case in which Δ_e is cyclic. Suppose that Δ_e is not elliptic in Υ . Let Λ be a primary JSJ decomposition of L . By [28, Lemma 2.8] some refinement Λ_0 of Λ dominates Δ . Since edge stabilizers in Δ and Λ_0 are both abelian and Λ_0 is primary, there is an edge group in Λ_0 which is conjugate to Δ_e . However, Δ_e is not universally elliptic, so this edge group in Λ_0 must correspond to a simple closed curve c in a surface associated to a socket subgroup of Λ .

Let Δ_v be the nonabelian vertex group adjacent to Δ_e . The above argument shows that every cyclic edge group of Δ is elliptic in Λ_0 , and thus every edge group is elliptic, because Λ_0 is primary. If Δ_v is rigid, this implies that Δ_v is elliptic in Λ , and hence must be contained in part of the socket containing the group associated to c . Thus, in this case, Δ_v is a socket which is neither exceptional nor maximal, which contradicts property (1a). On the other hand, if Δ_v is a socket, then it is clearly not maximal, which contradicts property (1b). This proves that Δ_e is elliptic in Υ , and therefore Δ is a universally elliptic splitting.

To prove that Δ is maximal universally elliptic, we must prove that all vertex groups of Δ are elliptic in any universally elliptic primary splitting Θ (see [28, Definition, p.4]). Since the edge groups are elliptic in Θ , it is clear that the rigid vertex groups of Δ must be elliptic in Θ . That abelian vertex groups of Δ are elliptic in Θ follows quickly

from the fact that G is one-ended and CSA, and the fact that Θ is a primary splitting.

Finally, suppose that Δ_{v_j} is a socket vertex group of Δ . Then there is a natural subgroup $H \leq \Delta_{v_j}$ (the fundamental group of the surface Σ from the definition, before the roots are added). It is clear that H is a *QH-subgroup*, in the sense of [28, Definition 5.13]. Therefore, by [28, Theorem 5.27], H is elliptic in the JSJ deformation space corresponding to *all* abelian edge groups relative to noncyclic abelian edge groups. However, since we are only considering primary splittings, this implies that the socket of H is elliptic in Θ , as required. Thus Δ is maximal universally elliptic, and so Δ is a primary JSJ decomposition.

Finally, we note that conditions (2) to (4) imply that the Bass–Serre tree of Δ is its own (collapsed) tree of cylinders, and hence is a canonical primary JSJ tree by the argument applied to T above. This completes the proof. \square

For brevity, we will call the decomposition Δ produced by the above theorem the *JSJ decomposition* of L . When a freely indecomposable Γ -limit group L is not a subgroup of Γ (and often when it is) it admits a nontrivial primary splitting. Therefore, unless L is a (socket of a) surface group, the primary JSJ decomposition of L is nontrivial.

Let \mathcal{H} be a family of subgroups of L , closed under conjugacy. We will also sometimes need to work with the Grushko and JSJ decompositions of L *relative to* \mathcal{H} . These decompositions exist for the same reasons as above – for example, instead of taking the JSJ decomposition for $(\mathcal{E}, \mathcal{A}_{\text{nc}})$ -trees, take the JSJ decomposition for $(\mathcal{E}, \mathcal{A}_{\text{nc}} \cup \mathcal{H})$ -trees. The work in [28] works in the relative setting also.

2.3. Modular automorphisms and envelopes. One important role of the JSJ is to encode an associated collection of automorphisms of a limit group L , namely the *modular group* $\text{Mod}(L)$ (cf. [45, Definition 5.2]). We briefly recall the definition of a Dehn twist. Suppose a group $G = A *_C B$ and z is in the centralizer $Z(C)$. Then the assignments $\delta_z(a) = a$ for all $a \in A$ and $\delta_z(b) = zbz^{-1}$ for all $b \in B$ define an automorphism δ_z of G , called a *Dehn twist* in the splitting $G = A *_C B$. Similarly, if $G = A *_C$ with stable letter t , and z is in the centralizer $Z(C)$, then the assignments $\delta_z(a) = a$ for all $a \in A$ and $\delta_z(t) = tz$ define an automorphism δ_z of G , called a *Dehn twist* in the splitting $G = A *_C$.

Definition 2.17. Let L be a Γ -limit group, and let Δ be its canonical primary JSJ decomposition. An automorphism α of L is called a *generalized Dehn twist* if it is of one of the following forms.

- (1) A Dehn twist in a one-edge splitting of L obtained by contracting all but one edges of Δ .
- (2) A Dehn twist in a one-edge splitting of L obtained by cutting along a two-sided simple closed curve in a socket vertex of Δ .
- (3) Let Δ_w be an abelian vertex group of Δ . Then α restricted to Δ_w fixes every edge group incident at Δ_w (equivalently, α fixes the peripheral subgroup $\overline{P}(\Delta_w)$ defined in Definition 2.18), and α acts as the identity on every other vertex group and stable letter of Δ .

The *modular group* of L , denoted by $\text{Mod}(L)$, is the subgroup of $\text{Out}(L)$ generated by the equivalence classes of generalized Dehn twists. The elements of the modular group are called *modular automorphisms*.

Note that modular automorphisms act as inner automorphisms on rigid vertex groups. Furthermore, certain slightly larger subgroups associated to a rigid vertex group are also preserved (up to conjugacy) by modular automorphisms. These larger subgroups were named the *envelopes* of rigid vertices by Bestvina–Feighn [5].

Working with the canonical JSJ, we are able to simplify slightly the definition of an envelope.

Definition 2.18. Let Δ be the JSJ of L , and A an abelian vertex stabilizer. The subgroup $P(A) \subseteq A$ is generated by the incident edge stabilizers in A , and the *peripheral subgroup* $\overline{P}(A)$ is the minimal direct factor of A containing $P(A)$.

The *envelope* of V , $E(V)$, is then generated by V together with the peripheral subgroups of the adjacent (abelian) vertex stabilizers.

2.4. Strict homomorphisms and resolutions. Together with the JSJ decomposition, the second pillar of the structure theory of Γ -limit groups is provided by the notion of a strict homomorphism. The following definition is rather lengthy, but the idea is that there exists a homomorphism $L \rightarrow R$, where R is a ‘simpler’ limit group, and the homomorphism has prescribed behaviour on each part of the JSJ of L .

Definition 2.19 (Strict Homomorphism; Sela [45], Definition 5.9). Let G be a finitely generated group, and let $G = G_1 * G_2 * \cdots * G_n * \mathbb{F}$ be the Grushko decomposition of G , where \mathbb{F} is a free group and the G_i are freely indecomposable. Suppose that Λ_i is a primary splitting of G_i .

Let H be a finitely generated group. A homomorphism $\eta: G \rightarrow H$ is *strict with respect to the splittings* Λ_i if there is a free product decomposition $H = H_1 * \cdots * H_n$ of H (no assumptions are made

about the decomposability of the H_i) so that $\eta(G_i) = H_i$ and for each i we have that:

- (1) η is injective when restricted to envelopes of rigid vertex groups of Λ_i ;
- (2) The image under η of each socket vertex group of Λ_i is non-abelian in H_i ;
- (3) If E is an edge group of Λ_i corresponding to a boundary component of a socket-type vertex then $\eta|_E$ is injective; and
- (4) If A is an abelian vertex group of Λ_i , and $\overline{P}(A)$ is the peripheral subgroup of A then $\eta|_{\overline{P}(A)}$ is injective.

If G is a freely indecomposable Γ -limit group and Λ is the primary JSJ decomposition of G then a homomorphism $\eta: G \rightarrow H$ is *strict* if it is strict with respect to Λ . Note that in the case the Λ_i are the canonical primary JSJ splittings, item 4 implies item 3.

Definition 2.20. Let G be a finitely generated group, L a Γ -limit group and $\{f_n: G \rightarrow \Gamma\}$ a convergent sequence of homomorphisms. We say that $\{f_n\}$ *converges to L* if

$$L \cong G / \underline{\text{Ker}} \{f_n\}.$$

Definition 2.21. Suppose that G is a finitely generated group, equipped with its Grushko decomposition and primary splittings Λ_i of the free factors as in Definition 2.19. A *strict resolution of G* is a sequence of epimorphisms

$$G \twoheadrightarrow L_0 \twoheadrightarrow L_1 \twoheadrightarrow \dots \twoheadrightarrow L_n,$$

where L_n is a free product of subgroups of Γ and each map $G \twoheadrightarrow L_0$ and $L_i \twoheadrightarrow L_{i+1}$ is a strict map (with respect to some free and primary splittings of L_i).

The next theorem now characterizes Γ -limit groups in terms of strict resolutions. Although this characterization is significantly more complicated than the original definition, it comes equipped with much more information. In particular, it will allow us to describe a Γ -limit group G inductively, in terms of ‘simpler’ Γ -limit groups.

Theorem 2.22. [46, Theorem 1.31], *cf.* [45, Theorem 5.12] *A finitely generated group G is a Γ -limit group if and only if it admits a strict resolution.*

Furthermore, each of the groups G and L_i in the resolution are Γ -limit groups, and we can take the abelian splittings of the free factors of the L_i to be the canonical primary JSJ decomposition of the free factors.

Remark 2.23. By Theorem 2.22, we lose nothing by assuming up front that the groups L_i in a strict resolution of a finitely generated group G are Γ -limit groups.

Once we assume that the groups involved are Γ -limit groups, there is no reason to assume that each of the homomorphisms in a resolution are surjective (since finitely generated subgroups of Γ -limit groups are themselves Γ -limit groups).

Moreover, the terminal group L_n may be assumed merely to admit a strict map to Γ , rather than be a free product of subgroups of Γ , since Theorem 2.22 still implies that L_n is a Γ -limit group. In this case, we may have two resolutions like $\lambda: L_n \rightarrow \Gamma$ and $L_n \rightarrow \lambda_n(L_n) \rightarrow \Gamma$. In case the freely indecomposable free factors of $\lambda(L_n)$ admit nontrivial JSJ decompositions, this second resolution encodes more homomorphisms from L_n to Γ than the first one.

These two changes give us some extra flexibility which is central to the constructions in this paper.

In this paper, we use a strict resolution \mathcal{R} of a freely indecomposable Γ -limit group L in order to embed L in a Γ -limit group $\mathcal{M}_{\mathcal{R}}$ so that $\mathcal{M}_{\mathcal{R}}$ is finitely presented, and furthermore toral relatively hyperbolic. This allows us to give an algorithmic enumeration of Γ -limit groups, even though not all Γ -limit groups need be finitely presented.

We make the following definition in order to capture one of the features of strict homomorphisms in a more general setting.

Definition 2.24. Suppose that G is a group and \mathcal{P} is a finite collection of subgroups of G . We say that a homomorphism $\eta: G \rightarrow H$ (for some group H) is *non-degenerate with respect to \mathcal{P}* if $P \cap \ker(\eta) = \{1\}$ for all $P \in \mathcal{P}$. Otherwise, we say that η is *degenerate with respect to \mathcal{P}* .

Observation 2.25. *Suppose that L is a freely indecomposable Γ -limit group and that \mathcal{D} is the set of edge groups of the primary JSJ decomposition of L . If $\eta: L \rightarrow R$ is a strict map then η is non-degenerate with respect to \mathcal{D} .*

2.5. Relative hyperbolicity. One of the key properties of the model Γ -limit groups which we construct in this paper is that they are toral relatively hyperbolic, which means that there are many available algorithmic tools which can be applied to them (see, for example [9, 11, 12, 13]). In this section, we recall the basic definitions, and some results which we need in order to prove that model Γ -limit groups are toral relatively hyperbolic.

Definition 2.26. A group is called *toral relatively hyperbolic* if it is torsion-free and hyperbolic relative to finitely generated abelian subgroups. There are many different equivalent definitions of relative hyperbolicity; we refer to [30] for a summary of many of them, and proofs that they are equivalent.

A toral relatively hyperbolic group admits a canonical set of peripheral subgroups, namely a set of representatives for the conjugacy classes of maximal noncyclic abelian subgroups. We always take this to be the peripheral structure, and so we often leave the peripheral structure unmentioned.

We will construct our relatively hyperbolic models by gluing. Combination theorems for relatively hyperbolic groups were proved by Alibegović [2] and Dahmani [8], and subsequently by Mj–Reeves [36], and others. We state Dahmani’s combination theorem here.

Theorem 2.27. [8, Theorem 0.1]

- (1) *Let Γ be the fundamental group of a finite acylindrical graph of relatively hyperbolic groups, whose edge groups are fully quasi-convex subgroups of the adjacent vertex groups. Let \mathcal{G} be the collection of images of the maximal parabolic subgroups of the vertex groups in Γ , and their conjugates. Then (Γ, \mathcal{G}) is a relatively hyperbolic group.*
- (2) *Let G be a group which is hyperbolic relative to \mathcal{G} , and let $P \in \mathcal{G}$. Let A be a finitely generated group which has P as a subgroup. Then $\Gamma = A *_P G$ is hyperbolic relative to $(\mathcal{H} \cup \mathcal{A})$, where \mathcal{H} is the set of Γ -conjugates of elements of $\mathcal{G} \setminus P^G$ and \mathcal{A} is the set of Γ -conjugates of A .*
- (3) *Let G_1 and G_2 be relatively hyperbolic groups, let P be a maximal parabolic subgroup of G_1 and suppose that P can be embedded in a parabolic subgroup of G_2 . Then $\Gamma = G_1 *_P G_2$ is relatively hyperbolic.*
- (4) *Let G be a relatively hyperbolic group and let P, P' be non-conjugate isomorphic parabolic subgroups of G . Then $\Gamma = G *_P$ is relatively hyperbolic.*

The following converse to the combination theorem will also be useful.

Proposition 2.28. [27, Proposition 3.4] *Suppose that G is a toral relatively hyperbolic group and that Λ is a graph of groups decomposition in which the edge groups are abelian and all noncyclic abelian subgroups are elliptic.*

Then the vertex groups of Λ are relatively quasiconvex, and hence toral relatively hyperbolic.

2.6. Strong accessibility. Having decomposed a (relatively hyperbolic) group G into the freely indecomposable factors G_i of its Grushko decomposition, and having computed the JSJ decompositions Δ_i of those factors, one has in a sense reduced the study of G to the study of the rigid vertex groups $G_{i,j}$ of the Δ_i . It is now natural to apply this procedure again to each $G_{i,j}$, and to continue recursively. The assertion that this process terminates is called *strong accessibility*.

We now make this precise. To a toral relatively hyperbolic group G and a family of subgroups \mathcal{P} , we associate a tree $\mathcal{T}_{G,\mathcal{P}}$ of subgroups as follows.

- (1) We start by labelling a node by (G, \mathcal{P}) .
- (2) For a node v labelled by a relatively freely decomposable pair (H, \mathcal{Q}) , we compute the Grushko decomposition $H_1 * \dots * H_k * \mathbb{F}$ of H (with the H_i well defined subgroups up to conjugacy) relative to \mathcal{Q} ; if $k = 1$ and $\mathbb{F} = \{1\}$ we proceed to the next step; otherwise, we attach k edges to v , and label the other ends of the edges by the (H_i, \mathcal{Q}_i) , where \mathcal{Q}_i consists of all the elements of \mathcal{Q} contained in H_i .
- (3) For a node v labelled by a relatively freely indecomposable subgroup (H, \mathcal{Q}) , we compute the primary JSJ decomposition Δ of H relative to \mathcal{Q} .
 - (a) If Δ is trivial, then add no edges to v .
 - (b) If Δ is non-trivial, and H_1, \dots, H_k are the rigid vertices of Δ (again, well defined up to conjugacy), then attach k edges to v , and label the other ends of the edges by the (H_i, \mathcal{Q}_i) , where \mathcal{Q}_i consists of all the elements of \mathcal{Q} contained in H_i .
- (4) Apply steps 2 and 3 to any remaining nodes.

The Strong Accessibility theorem [35, Theorem 2.5] now asserts that this construction terminates after finitely many steps.

Theorem 2.29 (Strong Accessibility, Louder–Touikan). *If G is a toral relatively hyperbolic group and \mathcal{P} is any family of subgroups then the tree $\mathcal{T}_{G,\mathcal{P}}$ is finite.*

We apply this to deduce that any subgroup of a toral relatively hyperbolic group has a *quasi-convex enclosure*. Let H be a non-abelian subgroup of a relatively hyperbolic G and \mathcal{D} a finite collection of subgroups of H . We suppose that H does not admit a nontrivial primary splitting relative to \mathcal{D} . We may now use the hierarchy $\mathcal{T}_{G,\mathcal{D}}$ to define a

sequence of subgroups

$$H \subseteq G_n \subseteq G_{n-1} \subseteq \dots \subseteq G_0 = G$$

and, for each i , an G_i -tree T_i . Indeed, we set v_n to be the unique leaf of $\mathcal{T}_{G,\mathcal{D}}$ that contains (a conjugate of) H and let the embedded path from v_n to the root of $\mathcal{T}_{G,\mathcal{D}}$ be v_n, \dots, v_0 . For each i , let G_i be such that the label of v_i is (G_i, \mathcal{D}) (where if necessary we replace G_i by a suitable conjugate so that $H \subseteq G_i$) and let T_i be the Bass–Serre tree of the splitting of G_i given in the definition of $\mathcal{T}_{G,\mathcal{D}}$ (taking T_n to be a point). Note that each T_i has a unique vertex u_i stabilized by H , and G_{i+1} is the stabilizer of u_i .

Definition 2.30. With the above notation, the *quasiconvex enclosure* of H (relative to \mathcal{D}) is (with a slight abuse of notation) denoted by \overline{H} and defined to be the subgroup G_n .

While this paper was in preparation, the authors discovered that a similar definition is made in [33, §§3.4], where a “quasi-convex closure” is defined similarly, although that definition is not relative to a set of subgroups \mathcal{D} . The application of the definition in [33] is similar in spirit to the one in this paper, although, as we explain in §1 above, considerable care is needed to ensure that the groups constructed really do have the good geometric features one would like, and also that simultaneously one does not lose control over the sets of homomorphisms one considers.

Corollary 2.31. *Suppose that G is a toral relatively hyperbolic group, that H is a finitely generated non-abelian subgroup, and \mathcal{D} a finite collection of subgroups of H . Suppose that H admits no abelian splitting relative to \mathcal{D} . Then:*

- (1) *the quasiconvex enclosure \overline{H} is the maximal subgroup of G that contains H and does not admit any essential abelian splitting relative to \mathcal{D} ; and*
- (2) *\overline{H} is relatively quasiconvex (and hence itself toral relatively hyperbolic).*

Proof. The first assertion is immediate from the definition. The second follows from Proposition 2.28 and induction. \square

2.7. Basics on algorithms with toral relatively hyperbolic groups.

Proposition 2.32. [9, Theorem 0.2] *There is a one-sided algorithm which takes as input finite presentations \mathcal{P} and terminates if and only if the group defined by \mathcal{P} is toral relatively hyperbolic. In case it is toral*

relatively hyperbolic, the algorithm provides a basis for each peripheral subgroup.

Suppose that we are given a finite presentation of a toral relatively hyperbolic group Υ , which is witnessed to be such by the above algorithm. Then we have an explicit solution to the word problem for Υ . One way to see this is [10, Theorem 0.1], which proves that the existential theory of Υ is decidable.

The next theorem is not immediately implied by the statement of [19, Theorem 3.5]. However the proof there applies verbatim to prove the following stronger result.

Theorem 2.33 (cf. Theorem 3.5, [19]). *There exists an algorithm that, given as input a presentation for a toral relatively hyperbolic group Γ , and an element $\gamma \in \Gamma$, outputs a minimal set of generators for $Z_\Gamma(\gamma)$.*

We remark that in case Γ is toral relatively hyperbolic, $Z_\Gamma(\gamma)$ is a finitely generated free abelian group. Therefore the cardinality of a minimal generating set determines the group up to isomorphism. In particular, it tells us the rank and a presentation can be easily (and algorithmically) written down for $Z_\Gamma(\gamma)$.

3. MODEL Γ -LIMIT GROUPS

In this section we make the required definitions of model Γ -limit groups, which are the central object of concern in this paper.

We start by introducing the main category we work with. Recall the definition of a *non-degenerate* map from Definition 2.24.

Definition 3.1. The category \mathcal{C}_{csa} is defined as follows:

- The objects of \mathcal{C}_{csa} are pairs (G, \mathcal{P}) , where G is a torsion-free CSA group and \mathcal{P} is a finite collection of non-trivial abelian subgroups of G . Denote the set of objects by $\text{Obj}(\mathcal{C}_{\text{csa}})$.
- The morphisms from (G, \mathcal{P}) to (H, \mathcal{Q}) are the homomorphisms $\phi: G \rightarrow H$ that are non-degenerate with respect to \mathcal{P} and such that, for each $P \in \mathcal{P}$, $\phi(P)$ is conjugate into an element of \mathcal{Q} . If $X_1, X_2 \in \text{Obj}(\mathcal{C}_{\text{csa}})$ then $\text{Mor}_{\mathcal{C}_{\text{csa}}}(X_1, X_2)$ denotes the set of morphisms between X_1 and X_2 .

Lemma 3.2. \mathcal{C}_{csa} is a category.

Proof. The only thing that needs to be checked is that morphisms compose. They do, since the distinguished subgroups of one group map into the distinguished subgroups of another, so the composite homomorphism remains injective on distinguished subgroups. \square

The construction of model Γ -limit groups we give in the next section relies on colimits in \mathcal{C}_{csa} .

In this section we give the definition of models in terms of a universal property, leaving open the question of existence. In Section 4 we provide the construction of models.

Let L_0 be a Γ -limit group and let

$$L_0 \xrightarrow{\lambda_0} L_1 \xrightarrow{\lambda_1} \cdots \xrightarrow{\lambda_{n-1}} L_n$$

be a strict resolution of L_0 , where L_n admits a (given) strict map $\lambda_n: L_n \rightarrow \Gamma$. Call this resolution (along with the map λ_n) \mathcal{R}_0 . For $i \geq 0$, let \mathcal{R}_i denote the resolution obtained from \mathcal{R}_0 by starting at L_i , and finally appending the map $\lambda_n: L_n \rightarrow \Gamma$.

Below we define the *model Γ -limit groups of L_i with respect to \mathcal{R}_i* , denoted $\mathcal{M}_{\mathcal{R}_i}$. The definition is made by (reverse) induction on i .

Before we describe the construction, we note a few important properties of models, which will be proved throughout the subsequent sections.

Theorem 3.3. *Let \mathcal{R}_0 be as above, and let \mathcal{C}_i be the edge groups of the JSJ decomposition of the freely indecomposable factors of the Grushko decomposition of L_i . For $i \geq 0$, there exists a canonical group $\mathcal{M}_{\mathcal{R}_i}$, and these groups satisfy the following properties:*

- (1) $\mathcal{M}_{\mathcal{R}_i}$ is a toral relatively hyperbolic Γ -limit group;
- (2) There is a strict map $\mu_i: \mathcal{M}_{\mathcal{R}_i} \rightarrow \mathcal{M}_{\mathcal{R}_{i+1}}$;
- (3) There is a distinguished family $\mathcal{D}_{\mathcal{R}_i}$ of abelian subgroups of $\mathcal{M}_{\mathcal{R}_i}$ so that $(\mathcal{M}_{\mathcal{R}_i}, \mathcal{D}_{\mathcal{R}_i}) \in \text{Obj}(\mathcal{C}_{\text{csa}})$;
- (4) There are injective maps $\eta_i: L_i \rightarrow \mathcal{M}_{\mathcal{R}_i}$ so that $\mu_i \circ \eta_i = \eta_{i+1} \circ \lambda_i$;
- (5) The map η_i takes \mathcal{C}_i to $\mathcal{D}_{\mathcal{R}_i}$ and is non-degenerate with respect to \mathcal{C}_i . In other words, $\eta_i \in \text{Mor}(\mathcal{C}_{\text{csa}})$.

The claims of Theorem 3.3 are contained in Theorems 4.10 and 5.1 below.

We now give the definition of models. First suppose that L_0 is freely indecomposable.

Suppose inductively that we have defined $\mathcal{M}_{\mathcal{R}_1}$, along with a homomorphism $\eta_1: L_1 \rightarrow \mathcal{M}_{\mathcal{R}_1}$. Consider the primary JSJ decomposition Δ of L_0 , and let $\{\Delta_v\}$ be the collection of rigid vertex groups. Let Δ_v be such a rigid vertex group. Let $\mathcal{C}_v = \{\Delta_{e_1}, \dots, \Delta_{e_s}\}$ be the set of adjacent edge groups. We remark that these edge groups are maximal abelian subgroups of Δ_v .

We define the *relative model* of Δ_v , denoted $\mathcal{M}(\Delta_v)$, to be the maximal subgroup of $\mathcal{M}_{\mathcal{R}_1}$ which does not admit any essential abelian splitting relative to $(\eta_1 \circ \lambda_0)(\mathcal{C}_v)$. If we suppose by induction that $\mathcal{M}_{\mathcal{R}_1}$

is a toral relatively hyperbolic group, such a maximal subgroup exists by Corollary 2.31, and indeed is the quasiconvex enclosure $\overline{\eta_1 \circ \lambda_0(\Delta_v)}$ (relative to $(\eta_1 \circ \lambda_0)(\mathcal{C}_v)$). In particular, $\mathcal{M}(\Delta_v)$ is toral relatively hyperbolic, and hence finitely presented, torsion-free and CSA by that same result.

Consider the elements $\{(\Delta_v, \mathcal{C}_v)\}, \{(\mathcal{M}(\Delta_v), \overline{\mathcal{C}}_v)\}$ of $\text{Obj}(\mathcal{C}_{\text{csa}})$, where $\overline{\mathcal{C}}_v$ is the set $\{(\eta_1 \circ \lambda_0)(C) \mid C \in \mathcal{C}_v\}$. Also consider the set \mathcal{C} of all peripheral subgroups of edge groups of the JSJ of L_0 , and the element $(L_0, \mathcal{C}) \in \text{Obj}(\mathcal{C}_{\text{csa}})$.

The model $\mathcal{M}_{\mathcal{R}_0}$ is then the group that appears in the colimit of the following diagram (labelling only groups and not the attendant subgroups), in the category \mathcal{C}_{csa} (supposing for the moment that such a colimit exists):

$$\begin{array}{ccc}
 \Delta_{v_1} & \xrightarrow{\iota} & L_0 \\
 \downarrow \eta_1 \circ \lambda_0 & \searrow \text{dotted} & \nearrow \iota \\
 & & \Delta_{v_n} \\
 & & \downarrow \eta_1 \circ \lambda_0 \\
 \mathcal{M}(\Delta_{v_1}) & \searrow \text{dotted} & \mathcal{M}(\Delta_{v_n})
 \end{array}$$

The map $\eta_0: L_0 \rightarrow \mathcal{M}_{\mathcal{R}_0}$ is the one furnished by the definition of colimit. Further, the groups $\mathcal{M}(\Delta_{v_i})$ are subgroups of $\mathcal{M}_{\mathcal{R}_1}$, and we have a map $\eta_1 \circ \lambda_0: L_0 \rightarrow \mathcal{M}_{\mathcal{R}_1}$. Temporarily define a set of maximal abelian subgroups \mathcal{D} of $\mathcal{M}_{\mathcal{R}_1}$ by taking images in $\mathcal{M}_{\mathcal{R}_1}$ of \mathcal{C} and of the $\overline{\mathcal{C}}_v$. Note that this will typically not be the set of maximal abelian subgroups of $\mathcal{M}_{\mathcal{R}_1}$ that it was furnished with when *it* was defined as a pushout in \mathcal{C}_{csa} .

Since the maps $\eta_1 \circ \lambda_0$ and inclusion of $\mathcal{M}(\Delta_{v_i})$ into $\mathcal{M}_{\mathcal{R}_1}$ agree on each Δ_{v_i} , and these maps are all in $\text{Mor}(\mathcal{C}_{\text{csa}})$, the universal property of the colimit yields a map $\mu_0: \mathcal{M}_{\mathcal{R}_0} \rightarrow \mathcal{M}_{\mathcal{R}_1}$ so that $\mu_0 \circ \eta_0 = \eta_1 \circ \lambda_0$. Note also that (so long as it exists) the universal property of $\mathcal{M}_{\mathcal{R}_0}$, together with the distinguished subgroups of the groups Δ_v , $\mathcal{M}(\Delta_v)$ and L_0 , furnish $\mathcal{M}_{\mathcal{R}_0}$ with a collection of distinguished subgroups.

Now suppose that L_0 is freely decomposable

$$L_0 = H_1 * \dots * H_r * F_k$$

where each H_i is freely indecomposable, and F_k is a free group. Let $\mathcal{R}_{0,i}$ denote the resolution obtained by restricting \mathcal{R}_0 to the subgroup H_i of L_0 .

The definition above gives the models $\mathcal{M}_{\mathcal{R}_{0,i}}$ along with maps to the model of L_1 . The model of L_0 is

$$\mathcal{M}_{\mathcal{R}_0} = \mathcal{M}_{\mathcal{R}_{0,1}} * \dots * \mathcal{M}_{\mathcal{R}_{0,r}} * F_k$$

and the map η_0 is defined in the obvious way.

Remark 3.4. At the moment, it is not obvious that the colimit above exists. If it does, it is not clear that the model should be a Γ -limit group, or what its properties should be (except that if it exists it is clearly torsion-free and CSA). Much of the rest of the paper is spent proving that models exist and understanding the properties of models, in order to see what they are and why they are useful.

If we apply the definition of model to the strict resolution $\Gamma \rightarrow \Gamma$ given by the identity map, we can see that with respect to any collection \mathcal{P} of abelian subgroups, the model of (Γ, \mathcal{P}) is (Γ, \mathcal{P}) .

Proposition 3.5. *For any collection of nontrivial abelian subgroups \mathcal{P} of Γ , the model of (Γ, \mathcal{P}) (with respect to the trivial resolution) is (Γ, \mathcal{P}) .*

Proof. In case Γ is freely indecomposable, this is immediate from the observation that $\mathcal{M}(\Delta_v) = \Delta_v$. The freely decomposable case is also immediate. \square

4. ABELIAN GRAPHS OF GROUPS IN \mathcal{C}_{csa} AND THE CONSTRUCTION OF MODELS

Let L_0 be a freely indecomposable Γ -limit group, equipped with a strict resolution \mathcal{R}_0 . In this section, we give a construction of $\mathcal{M}_{\mathcal{R}_0}$ which is amenable to understanding the structure of the model of L_0 .

Let \mathcal{C}_{csa} be the category defined in Definition 3.1 above. In order to prove our structural results about the model of L_0 , we investigate certain kinds of colimits in \mathcal{C}_{csa} . Along the way, we exhibit various subtleties about building torsion-free CSA groups by taking graphs of groups. We will then show (in case L_0 is freely indecomposable) that we can construct $\mathcal{M}_{\mathcal{R}_0}$ as a graph of groups with abelian edge groups.

4.1. Some colimits. The following simple example illustrates that the class of torsion-free CSA groups is not closed under taking pushouts in the category of groups.

Example 4.1. Let $A_1 = A_2 = C = \mathbb{Z}$ and let $\iota_i: C \rightarrow A_i$ be index 2 embeddings for $i = 1, 2$. Then the pushout in the category of groups is $A_1 *_C A_2 = \mathbb{Z} *_2 \mathbb{Z}$ which is not commutative transitive.

In the category of abelian groups, the pushout is $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$, which is not torsion-free.

In the category of free abelian groups, the pushout is \mathbb{Z} , with the maps from the A_i being the same isomorphism and the map from C being the index 2 embedding.

The next example shows that we do not want to merely take pushouts in the category of torsion-free CSA groups, where all homomorphisms are allowed, but required our maps to be non-degenerate.

Example 4.2. Let $A = \langle a_1, a_2 \mid [a_1, a_2] \rangle \cong \mathbb{Z}^2$, $B = \langle b_1, b_2 \mid [b_1, b_2] \rangle \cong \mathbb{Z}^2$ and $C = \langle c \mid \emptyset \rangle \cong \mathbb{Z}$. Let $\iota_1: C \rightarrow A$ be defined by $\iota_1(c) = a_1$ and $\iota_2: C \rightarrow B$ by $\iota_2(c) = b_1$. Then the pushout in the category of groups is $G = \langle x, y, z \mid [x, y], [y, z] \rangle \cong \mathbb{F}_2 \times \mathbb{Z}$, which is not CSA.

Certainly, $G_{ab} = \mathbb{Z}^3$ is a torsion-free CSA quotient of G . If we added peripheral structures $\mathcal{P}_C = \{C\}$ and the images of C in A and B to A and B , then this would be the pushout in \mathcal{C}_{csa} .

On the other hand, killing y gives $\mathbb{F}_2 = \langle x, z \mid \emptyset \rangle$ as a torsion-free CSA quotient of G . These are incomparable quotients, in the sense that there is no torsion-free CSA quotient of G which ‘lies above’ them both.²

To remedy these difficulties, we consider the pushout in the category \mathcal{C}_{csa} , and in particular we restrict to non-degenerate maps.

First, we explicitly describe the pushout in the category of finitely generated abelian groups. It is easy to see that pushouts exist in the category of abelian groups (the pushout is the abelianization of the pushout in the category of groups). However, we are interested in free abelian groups rather than merely abelian groups.

Lemma 4.3. *Let A_1, A_2, C be finitely generated nontrivial free abelian groups, and let $\iota_i: C \rightarrow A_i$ be embeddings, for $i = 1, 2$. There exists a pushout of the diagram*

²Of course, this phenomenon is well known from the construction of Makanin–Razborov diagrams since G_{ab} and \mathbb{F}_2 , equipped with the canonical quotient maps from G form the first level of the MR diagram for G (over a nonabelian free group \mathbb{F}). It is also one of the key reasons for the distinguished abelian subgroups and the requirement that homomorphisms be non-degenerate.

$$\begin{array}{ccc} C & \longrightarrow & A_1 \\ & & \downarrow \\ & & A_2 \end{array}$$

in the category of free abelian groups.

Proof. This is straightforward. One way to proceed is to pass to direct factors of the A_i which contain the image of C as finite-index subgroups, embed as lattices in \mathbb{R}^k , and take the span of these two lattices. \square

The following lemma is straightforward.

Lemma 4.4. *If M is the pushout of a collection of finitely generated free abelian groups $\{A_i\}$ (with appropriate maps), and H is a finitely generated free abelian group with injective maps $A_i \hookrightarrow H$ for each i , making the appropriate diagram commute, then the induced map from M to H is injective.*

This is useful because of the following.

Corollary 4.5. *If all groups are free abelian and all maps are injective then the pushout of the diagram*

$$\begin{array}{ccc} C & \longrightarrow & A \\ & & \downarrow \\ & & B \end{array}$$

in the category of finitely generated free abelian groups is also the pushout of that diagram in \mathcal{C}_{CSA} so long as the set of distinguished subgroups of C is nonempty, and each map is a morphism in \mathcal{C}_{CSA} .

In Subsection 4.2 below, we will need a slightly more general construction than a colimit, involving a graph of groups.

4.2. Construction of models.

Definition 4.6. Suppose that G is a finitely generated, torsion-free CSA group. A graph of groups decomposition Δ of G is *JSJ-like* if it satisfies the following properties:

- (1) Δ is universally elliptic, i.e. the edge stabilizers of Δ are elliptic in every primary splitting of G .
- (2) Every non-cyclic abelian subgroup of L_0 is elliptic in Δ .
- (3) Δ is bipartite, with one class of vertex stabilizers abelian and the other either (nonabelian) rigid (in the sense of Definition 2.13) or a maximal socket (with every boundary component used). We denote the set of abelian vertices by $\{w\}$, the rigid vertices by $\{u\}$ and the socket vertices by $\{v\}$.

- (4) Every edge group is maximal abelian in the adjacent vertex group which is not abelian. Furthermore, Δ is 2-cylindrical.

Note that in case G is freely indecomposable, it follows immediately from Proposition 2.16 that the canonical primary JSJ decomposition of G is JSJ-like, and that every JSJ-like decomposition is in fact a canonical primary JSJ. We introduce this definition for the slight extra flexibility of applying it in case that we do not know that G is freely indecomposable.

Suppose that L is a freely indecomposable Γ -limit group, that Δ is a JSJ-like decomposition of L and that $\rho: L \rightarrow M$ is a strict map with respect to Δ , where M is a toral relatively hyperbolic group. In our applications, M will be the model of L_1 , where $L \rightarrow L_1$ is the beginning of a strict resolution of L .

Given Δ and ρ , we build a new graph of groups $\widehat{\Delta}$ as follows:

For a rigid vertex u of Δ , we consider the map $\rho|_{\Delta_u}: \Delta_u \rightarrow M$. Since λ_0 is strict, $\rho|_{\Delta_u}$ is injective. We define $\widehat{\Delta}_u$ to be the quasiconvex enclosure of $\rho(\Delta_u)$ in M (relative to the images under ρ of the adjacent edges groups). Let $\eta_u = \rho|_{\Delta_u}$ be the natural inclusion from Δ_u to $\widehat{\Delta}_u$.

Define an equivalence relation on the edges of Δ as follows: Two edges e_1 and e_2 adjacent to a rigid vertex u are equivalent if $Z_{\widehat{\Delta}_u}(\eta_u(\Delta_{e_1}))$ and $Z_{\widehat{\Delta}_u}(\eta_u(\Delta_{e_2}))$ are conjugate in $\widehat{\Delta}_u$. Edges not adjacent to rigid vertices are equivalent only to themselves.

Suppose that e_1 and e_2 are adjacent to a nonabelian vertex y and are equivalent. Let $Z_1 = Z_{\widehat{\Delta}_y}(\eta_y(\Delta_{e_1}))$ and $Z_2 = Z_{\widehat{\Delta}_y}(\eta_y(\Delta_{e_2}))$. By CSA, any $\gamma \in \widehat{\Delta}_y$ which conjugates Z_1 to Z_2 induces the same isomorphism, so there is a canonical isomorphism between Z_1 and Z_2 .

The underlying graph of $\widehat{\Delta}$ is equal to the quotient graph of the underlying graph of Δ obtained by identifying edges which are equivalent according to the above relation. Note that there is a natural way of labeling each vertex of this graph as exactly one of abelian, rigid or socket. If y is a vertex of Δ we let $[y]$ denote the corresponding vertex in $\widehat{\Delta}$ under the quotient map; likewise, the edges of $\widehat{\Delta}$ are the equivalence classes $[e]$ where e is an edge of Δ . Note, however, that the equivalence classes $[u]$ of rigid vertices and $[v]$ of socket vertices are singletons, and so we will usually write u for $[u]$ and v for $[v]$.

We now define the edge and vertex groups of $\widehat{\Delta}$ as follows:

- (1) A rigid vertex u is labelled by the group $\widehat{\Delta}_u$ defined above.
- (2) A socket vertex v is labelled by the socket group $\widehat{\Delta}_v$, obtained from Δ_v by keeping the same surface, but adding any extra roots that exist in M to the boundary components.

- (3) Let $[e]$ be an edge joining an abelian vertex $[w]$ to a vertex y . Note that y is not an abelian vertex. Then $\widehat{\Delta}_{[e]} = Z_{\widehat{\Delta}_y}(\eta_y(\Delta_e))$.
- (4) Finally, let $[w]$ be an abelian vertex, and let $[e_1], \dots, [e_l]$ be the adjacent edges. If $[w] = \{w_k\}$ and $[e_i] = \{e_{i,j}\}$ then we have two sets of inclusions of free abelian groups:
- (a) $\iota_{i,j}: \Delta_{e_{i,j}} \rightarrow \widehat{\Delta}_{[e_i]}$ for each i, j ;
 - (b) $\kappa_{i,j}: \Delta_{e_{i,j}} \rightarrow \Delta_{w_k}$ whenever $e_{i,j}$ is incident at w_k .

We let these data define a graph of groups $\Lambda_{[w]}$ in the natural way: the underlying graph is bipartite, with vertex set the disjoint union of the $\{w_k\}$ and $\{[e_i]\}$, and edge set $\{e_{i,j}\}$ equipped with the natural incidence relations; the vertex group of w_k is Δ_{w_k} , the vertex group of $[e_i]$ is $\widehat{\Delta}_{[e_i]}$, the edge group of $e_{i,j}$ is $\Delta_{e_{i,j}}$, and the edge maps are given by $\iota_{i,j}$ and $\kappa_{i,j}$. The vertex group $\widehat{\Delta}_{[w]}$ is now defined to be the abelianization of $\pi_1\Lambda_{[w]}$.

We remark that the definition of $\widehat{\Delta}_{[e]}$ used the representative e of $[e]$; however we have already noted that the appropriate edge groups are *canonically* isomorphic, so there is no ambiguity in this definition.

Finally, the edge maps of $\widehat{\Delta}$ are defined in the natural way using the canonical isomorphisms of conjugacy representatives of edge groups and either inclusion (in the case of rigid and socket vertices) or the natural inclusion of a vertex group in the fundamental group of a graph of groups (in the abelian case).

We call the graph of groups $\widehat{\Delta}$ the *expansion* of Δ with respect to $\rho: L \rightarrow M$.

Lemma 4.7. *Suppose that L is a freely indecomposable Γ -limit group, and that Δ is a JSJ-like decomposition of L . Suppose further that $\rho: L \rightarrow M$ is a strict map, that M is toral relatively hyperbolic, and that $\widehat{\Delta}$ is the expansion of Δ with respect to ρ . Then $\pi_1(\widehat{\Delta})$ is toral relatively hyperbolic, and in particular torsion-free and CSA.*

Proof. By construction, the edge groups incident at each rigid vertex form a malnormal family. In particular, we may apply Dahmani's Combination Theorem 2.27 together with the fact that the vertex groups $\widehat{\Delta}_v$ are toral relatively hyperbolic (by Corollary 2.31) to deduce that $\pi_1\widehat{\Delta}$ is toral relatively hyperbolic, in particular torsion-free and CSA. \square

There is a natural homomorphism $\eta: L \rightarrow \pi_1(\widehat{\Delta})$ defined as follows. For rigid vertices Δ_u , the map η is given by $\eta_u: \Delta_u \rightarrow \widehat{\Delta}_u$. For socket vertex groups Δ_v , η is defined to be the inclusion of the socket Δ_v into the socket $\widehat{\Delta}_v$. For abelian vertex group Δ_w , note that Δ_w is a vertex group of $\Lambda_{[w]}$, and so η is defined to be the composition of the

inclusion map $\Delta_w \rightarrow \pi_1\Lambda_{[w]}$ with abelianization. Finally, after picking a maximal tree in the underlying graph of Δ , each stable letter either maps to a stable letter of $\widehat{\Delta}$ or to an element of $\pi_1\Lambda_{[w]}$ for some abelian vertex group w .

Proposition 4.8. *In the above situation $\eta: L \rightarrow \pi_1\widehat{\Delta}$ is an injection.*

Proof. In order to prove injectivity of the map η it is convenient to modify Δ to produce a new graph of groups Δ' . The underlying graph of Δ' is the same as the underlying graph of Δ . If v is a (flexible) socket vertex of Δ then we take $\Delta'_v = \Delta_v$, and likewise if w is an abelian vertex we take $\Delta'_w = \Delta_w$. For u a rigid vertex of Δ , we set $\Delta'_u = E(\Delta_u) \cap \rho^{-1}(\widehat{\Delta}_u)$ (recalling from Definition 2.18 that $E(\Delta_u)$ denotes the envelope of Δ_u). Finally, for e an edge adjoining an abelian vertex w and a socket or rigid vertex y , we take $\Delta'_e = \Delta'_w \cap \Delta'_y$. We make three remarks about the construction of Δ' .

First, Δ' was obtained from Δ by pulling certain elements across edges, and so $\pi_1\Delta' = L$.

Second, if y is a rigid or socket vertex in Δ , then it remains the case that the incident edge groups in Δ' form a malnormal family in Δ'_y ; in particular, Δ' is also 2–acylindrical.

Third, for each vertex or edge y of Δ , we have that $\eta(\Delta'_y) \subseteq \widehat{\Delta}_{[y]}$, by the definition of Δ'_y in the different cases. It follows that there is a morphism of trees $\alpha: T' \rightarrow \widehat{T}$ (the Bass–Serre trees of Δ' and $\widehat{\Delta}$ respectively) that intertwines η – that is, $\alpha(g.x) = \eta(g)\alpha(x)$ for any $g \in L$ and $x \in T'$.

To show that η is injective, it now suffices to prove that η is injective on vertex groups of Δ' (which follows from strictness of ρ) and that α is a local injection (i.e. does not factor through a fold). In order to prove that α does not factor through a fold, we must prove that for each edge e and each vertex y we have $\eta^{-1}(\widehat{\Delta}_{[e]}) \cap \Delta'_y = \Delta'_e$. Of course, it is clear that $\Delta'_e \subseteq \eta^{-1}(\widehat{\Delta}_{[e]}) \cap \Delta'_y$, and it is the reverse inclusion that must be proved.

So, suppose that $g \in \eta^{-1}(\widehat{\Delta}_{[e]}) \cap \Delta'_y$. As usual, there are three cases, depending on the type of y . In case y is a socket vertex, the only difference between Δ'_y and $\widehat{\Delta}_y$ is that there may be extra roots of boundary components of the surface which exist in M but not in L . Thus in this case we clearly have $g \in \Delta'_e$, as required. Suppose then that y is a rigid vertex group. By definition $\Delta'_y = E(\Delta_y) \cap \rho^{-1}(\widehat{\Delta}_y)$ and $\Delta'_e = \Delta'_w \cap \Delta'_y$, where w is the vertex at the other end of e . Also, by definition, $\widehat{\Delta}_{[e]} = Z_{\widehat{\Delta}_y}(\eta_y(\Delta_e))$ (for a suitable choice of representative

edge e). But $g \in \Delta'_y$, so $g \in Z_{\Delta'_y}(\Delta_e) = \Delta'_e$, by construction and malnormality.

Finally, suppose that y is an abelian vertex (so $\Delta'_y = \Delta_y$), and let x be the vertex at the other end of e . If x is a socket vertex then there is nothing to prove, so we assume that x is rigid.

The defining graph of groups $\pi_1\Lambda_{[y]}$, of which $\widehat{\Delta}_{[y]}$ is the abelianization, has a quotient obtained from killing each peripheral subgroup of each Δ_{y_i} , and also each vertex group of the form $\widehat{\Delta}_{[e_i]}$. By abelianizing, we obtain a quotient of $\widehat{\Delta}_{[y]}$. In this quotient, the image of each $\widehat{\Delta}_{[e_i]}$ is trivial, and the image of Δ_y is $\Delta_y/\overline{P}(\Delta_y)$. Since $g \in \eta^{-1}(\widehat{\Delta}_{[e]})$, its image in this quotient is trivial. Therefore, $g \in \overline{P}(\Delta_y) \subseteq E(\Delta_x)$ and so, since $\eta(g)$ is contained in $\widehat{\Delta}_{[e]} \subseteq \widehat{\Delta}_x$, we have

$$g \in E(\Delta_x) \cap \eta^{-1}(\widehat{\Delta}_x) = E(\Delta_x) \cap \rho^{-1}(\overline{\Delta}_x) = \Delta'_x$$

whence, since g is also contained in Δ'_y , we have that $g \in \Delta'_x \cap \Delta'_y = \Delta'_e$ as required. \square

Proposition 4.9. *Suppose that L_0 is a freely indecomposable Γ -limit group and that \mathcal{R}_0 is a strict resolution of L_0 , starting with the strict map $\lambda_0: L_0 \rightarrow L_1$ and continuing with the strict resolution \mathcal{R}_1 of L_1 . Suppose further that $\mathcal{M}_{\mathcal{R}_1}$ exists, and satisfies:*

- (1) *The natural map $\eta_1: L_1 \rightarrow \mathcal{M}_{\mathcal{R}_1}$ is an injection; and*
- (2) *$\mathcal{M}_{\mathcal{R}_1}$ is toral relatively hyperbolic.*

Let $\rho: L_0 \rightarrow \mathcal{M}_{\mathcal{R}_1}$ be the map $\rho = \eta_1 \circ \lambda_0$, let Δ be a JSJ-like decomposition of L_0 and let $\widehat{\Delta}$ be the expansion of Δ with respect to ρ .

Equip $\pi_1\widehat{\Delta}$ with the distinguished set $\mathcal{D}_{\widehat{\Delta}}$ of peripheral subgroups of abelian vertex groups of $\widehat{\Delta}$. Then $(\pi_1\widehat{\Delta}, \mathcal{D}_{\widehat{\Delta}})$ is the model of L_0 with respect to \mathcal{R}_0 .

Proof. We briefly recall what it means for $(\pi_1\widehat{\Delta}, \mathcal{D}_{\widehat{\Delta}})$ to be the required model and what we have to prove. Recall that we have, for each rigid vertex v of Δ , two pairs:

- (1) $(\Delta_v, \mathcal{C}_v)$, where Δ_v is the vertex group, and \mathcal{C}_v is the family of adjacent edge groups; and
- (2) $(\mathcal{M}_{\mathcal{R}_1}(\Delta_v), \overline{\mathcal{C}}_v)$, where $\mathcal{M}_{\mathcal{R}_1}(\Delta_v)$ is the relative model (quasiconvex enclosure) of Δ_v and $\overline{\mathcal{C}}_v$ is the set of centralizers in $\mathcal{M}_{\mathcal{R}_1}(\Delta_v)$ of the images of \mathcal{C}_v .

We also have the pair $(L_0, \mathcal{D}_{\Delta})$, where \mathcal{D}_{Δ} is the set of peripheral subgroups of abelian vertex groups of Δ .

Equip $\pi_1\widehat{\Delta}$ with the collection $\mathcal{D}_{\widehat{\Delta}}$ of peripheral subgroups of abelian vertex groups of $\widehat{\Delta}$. We claim that $(\pi_1\widehat{\Delta}, \mathcal{D}_{\widehat{\Delta}})$ is the required colimit

in \mathcal{C}_{CSA} . Indeed, suppose that (H, \mathcal{Q}_0) is an object in \mathcal{C}_{CSA} and $\alpha \in \text{Mor}_{\mathcal{C}_{\text{CSA}}}((L_0, \mathcal{D}_\Delta), (H, \mathcal{Q}_0))$, $\beta_v \in \text{Mor}_{\mathcal{C}_{\text{CSA}}}((\mathcal{M}_{\mathcal{R}_1}(\Delta_v), \bar{\mathcal{C}}_v), (H, \mathcal{Q}_0))$ are morphisms that make the diagram below commute.

$$\begin{array}{ccc}
 \Delta_{u_1} & \xrightarrow{\iota} & L_0 \\
 \downarrow \eta_1 \circ \lambda_0 & \searrow \text{dotted} & \nearrow \iota \\
 & & \Delta_{u_n} \\
 & & \downarrow \eta_1 \circ \lambda_0 \\
 \mathcal{M}_{\mathcal{R}_1}(\Delta_{u_1}) & \xrightarrow{\beta_{v_1}} & \mathcal{M}_{\mathcal{R}_1}(\Delta_{u_n}) \xrightarrow{\beta_{v_n}} H \\
 & \searrow \text{dotted} & \nearrow \\
 & & H
 \end{array}$$

α (vertical arrow from L_0 to H)

The map $\pi: \pi_1 \widehat{\Delta} \rightarrow H$ is defined as follows.

- (1) For each rigid vertex u , $\widehat{\Delta}_u = \mathcal{M}_{\mathcal{R}_0}(\Delta_u)$ and $\pi|_{\widehat{\Delta}_u} = \beta_u$.
- (2) For each socket vertex v , $\widehat{\Delta}_v$ is obtained from attaching extra roots to boundary elements of the socket Δ_v . This means that $\widehat{\Delta}_v$ is the colimit of a diagram obtained by restricting the maps ι and $\eta_1 \circ \lambda_0$ to subgroups. The universal property of this colimit defines the map $\pi|_{\widehat{\Delta}_{u_i}}$.
- (3) Consider an abelian vertex $[w]$. The commutative diagram above shows that the maps α and β_u naturally induce a unique, non-degenerate map $\pi_1 \Lambda_{[w]} \rightarrow H$. Non-degeneracy, together with the hypothesis that H is CSA, implies that this map factors through the abelianization.
- (4) We may choose a maximal tree for Δ that maps to a maximal tree for $\widehat{\Delta}$. Therefore, each stable letter \widehat{t} of $\widehat{\Delta}$ corresponds to a stable letter t in Δ and we define $\pi(\widehat{t}) = \alpha(t)$.

It is easy to check that π is a homomorphism. It remains to check that π respects the associated families of distinguished subgroups. But this is clearly the case, because the homomorphisms α and β_v do, and the distinguished subgroups of $\pi_1 \widehat{\Delta}$ are images of distinguished subgroups of L_0 and $\mathcal{M}_{\mathcal{R}_1}(\Delta_v)$. \square

We can now prove the following theorem by induction on the length of a strict resolution.

Theorem 4.10. *As above, let L_0 be a Γ -limit group with strict resolution \mathcal{R}_0 .*

- (1) $\mathcal{M}_{\mathcal{R}_0}$ exists.

- (2) The natural map $\eta_0: L_0 \rightarrow \mathcal{M}_{\mathcal{R}_0}$ is an injection.
- (3) $\mathcal{M}_{\mathcal{R}_0}$ is a toral relatively hyperbolic group, in particular is finitely presentable.

Proof. We have already observed in Proposition 3.5 that, when \mathcal{R}_0 has length 0 and $L_0 = \Gamma$, the model for Γ , with respect to any collection of distinguished subgroups, is itself Γ . Therefore, we may assume that \mathcal{R}_0 is of length at least one. We further assume by induction that $\mathcal{M}_{\mathcal{R}_1}$ exists and satisfies the three properties in the statement of the theorem.

Also, it is clear that the theorem is stable under passing to free products, so we may assume that L_0 is freely indecomposable.

Let Δ be the canonical primary JSJ of L_0 , and recall that Δ is JSJ-like. Let $\lambda_0: L_0 \rightarrow L_1$ be the first map from \mathcal{R}_0 and $\eta_1: L_1 \rightarrow \mathcal{M}_{\mathcal{R}_1}$ be the natural map from L_1 to its model. By induction, we know that η_1 is injective, so $\rho = \eta_1 \circ \lambda_0$ is a strict map. Let $\widehat{\Delta}$ be the expansion of Δ with respect to ρ .

It follows from Proposition 4.9 that $\pi_1 \widehat{\Delta}$ is the model of L_0 with respect to \mathcal{R}_0 . The other two properties in the statement of the theorem follow immediately from Lemma 4.7 and Proposition 4.8. \square

5. FURTHER PROPERTIES OF MODEL Γ -LIMIT GROUPS

The purpose of this section is to prove further properties of models, which are summarized in the next result. In order to state the result, we adopt the following notation. Let L be a Γ -limit group and \mathcal{R} a strict resolution of L . Let Δ be the canonical primary JSJ decomposition of L . Let $\lambda: L \twoheadrightarrow L_1$ be the first map in \mathcal{R} . Let $\mathcal{M}_{\mathcal{R}}$ be the model of L with respect to \mathcal{R} , let $\mu: \mathcal{M}_{\mathcal{R}} \rightarrow \mathcal{M}_{\mathcal{R}_1}$ be the associated map between models and let $\eta: L \hookrightarrow \mathcal{M}_{\mathcal{R}}$ be the canonical inclusion. Let $\eta_1: L_1 \hookrightarrow \mathcal{M}_{\mathcal{R}_1}$ be the canonical inclusion and let $\widehat{\Delta}$ be the expansion of Δ with respect to $\eta_1 \circ \lambda$, so that $\mathcal{M}_{\mathcal{R}} = \pi_1 \widehat{\Delta}$.

Theorem 5.1. *Let L be a freely indecomposable limit group with a strict resolution \mathcal{R} and let $\mathcal{M}_{\mathcal{R}}$ be the associated model for M .*

- (1) $\mathcal{M}_{\mathcal{R}}$ is freely indecomposable.
- (2) $\widehat{\Delta}$ is the primary JSJ decomposition for $\mathcal{M}_{\mathcal{R}}$.
- (3) The map $\mu: \mathcal{M}_{\mathcal{R}} \rightarrow \mathcal{M}_{\mathcal{R}_1}$ is strict.
- (4) The model $\mathcal{M}_{\mathcal{R}}$ is a Γ -limit group.
- (5) There is a canonical homomorphism $\Phi: \text{Mod}(L) \rightarrow \text{Mod}(\mathcal{M}_{\mathcal{R}})$ that intertwines the injection η ; that is, for $\alpha \in \text{Mod}(L)$ and $g \in L$, we have that $\eta(\alpha(g)) = \Phi(\alpha)(\eta(g))$.

The proof of this theorem is spread out over a variety of results in this section. The reader should note the key implication: the universal

property of $\mathcal{M}_{\mathcal{R}}$ allows one to take a factorization of a homomorphism $f: L \rightarrow \Gamma$ through \mathcal{R} and induce a factorization of a homomorphism $\tilde{f}: \mathcal{M}_{\mathcal{R}} \rightarrow \Gamma$ through the induced resolution $\widehat{\mathcal{R}}$ of $\mathcal{M}_{\mathcal{R}}$. *This is the key point of the construction of models.*

We start with item (1), which is a consequence of a standard result that characterizes when graphs of groups admit free splittings.

Lemma 5.2. *With the above notation, $\mathcal{M}_{\mathcal{R}}$ is freely indecomposable.*

Proof. For a contradiction, we suppose that $\mathcal{M}_{\mathcal{R}}$ acts on a tree T with trivial edge stabilizers and without a global fixed point. Let $\widehat{\Delta}_1, \dots, \widehat{\Delta}_k$ be the components of $\widehat{\Delta}$ created when we delete all non-cyclic abelian edge groups. For some i , $\pi_1 \widehat{\Delta}_i$ acts on T without a global fixed point. By [49, Theorem 18], there is a vertex v of $\widehat{\Delta}_i$ (and hence of $\widehat{\Delta}$) whose vertex group admits a non-trivial free splitting relative to the incident cyclic edge groups. However, non-cyclic abelian groups do not split freely, and so this splitting is also relative to the incident non-cyclic abelian edges. But this contradicts the definition of $\widehat{\Delta}$. \square

The next result implies item (2).

Proposition 5.3. *If Δ_u is a rigid vertex group of Δ (in the sense of Definition 5.3) then the corresponding vertex group $\widehat{\Delta}_u$ of $\widehat{\Delta}$ is rigid.*

Proof. Consider an $\mathcal{M}_{\mathcal{R}}$ -tree T which has abelian edge stabilizers and is so that every noncyclic abelian subgroup is elliptic. Consider the induced L -action coming from the canonical embedding of L into $\mathcal{M}_{\mathcal{R}}$. This is an action of the same type, and Δ_u is a rigid vertex group Δ , so Δ_u fixes a point in T . However, the group $\widehat{\Delta}_u$ was chosen so that it did not admit any essential abelian splittings relative to (the images of) the edge groups of Δ adjacent to Δ_u . Thus $\widehat{\Delta}_u$ fixes a point of T , as required. \square

Corollary 5.4. *The graph of groups $\widehat{\Delta}$ is the canonical primary JSJ decomposition of $\mathcal{M}_{\mathcal{R}}$.*

Proof. We have proved that $\mathcal{M}_{\mathcal{R}}$ is freely indecomposable. Therefore, Proposition 2.16 implies that in order to show that $\widehat{\Delta}_{\mathcal{R}}$ is a canonical primary JSJ decomposition we have to verify Conditions (1)–(4) from that result. Proposition 5.3 shows that the vertices labelled ‘rigid’ are indeed rigid. Suppose that Δ_v is a (maximal) socket group of Δ , with corresponding vertex group $\widehat{\Delta}_v$ of $\widehat{\Delta}$. Clearly $\widehat{\Delta}_v$ is a socket, and we claim that it is a maximal socket. Suppose not. Then there is an abelian vertex group adjacent to $\widehat{\Delta}_v$ of valence 2 which is isomorphic to

each of its edge groups, and the other nonabelian vertex group adjacent to this abelian vertex group is also a socket. By construction, it cannot be a flexible socket, so it must be a rigid socket. However, in this case, it arose from a vertex group of Δ which contained the boundary curve from $\widehat{\Delta}_v$, and also powers of the other two boundary curves. This would mean that the socket group Δ_v was not maximal, contrary to assumption.

It is clear from the construction that $\widehat{\Delta}$ is bipartite with the required structure. Moreover, by construction, any edge group adjacent to a nonabelian vertex group Δ_v is maximal abelian in Δ_v . It follows that $\widehat{\Delta}$ is 2-acylindrical and so, by Proposition 2.16, $\widehat{\Delta}$ is the canonical primary JSJ decomposition of $\mathcal{M}_{\mathcal{R}}$, as required. \square

We turn our attention to item (3).

Proposition 5.5. *The map μ is strict.*

Proof. There are three conditions to check, for the different flavours of vertex groups of $\widehat{\Delta}$.

For socket vertices, we must show that the image under μ is non-abelian, and this is immediate from the construction, because the socket vertices of L embed in the socket groups of $\widehat{\Delta}$ and $\lambda: L \rightarrow L_1$ is strict.

Suppose now that $\widehat{\Delta}_{[w]}$ is an abelian vertex group of $\widehat{\Delta}$, and let $\overline{P}(\widehat{\Delta}_{[w]})$ be the peripheral subgroup. Then $\overline{P}(\widehat{\Delta}_{[w]})$ is a colimit of the peripheral groups $\overline{P}(\Delta_{w_k})$ and the incident edge groups $\widehat{\Delta}_{[e_i]}$ with maps from the edge groups of Δ to each of these. That μ is injective on $\overline{P}(\widehat{\Delta}_{[w]})$ follows from the injectivity of the maps from $\overline{P}(\Delta_{w_k})$ and $\widehat{\Delta}_{[e_i]}$.

So it remains to check that μ is injective on the envelopes of rigid vertices. We therefore consider a rigid vertex v and its envelope $E = E(\widehat{\Delta}_v)$. Note that the Bass–Serre tree \widehat{T} of $\widehat{\Delta}$ induces a splitting of E : the subgroup E acts cocompactly on its minimal invariant subtree S , and the quotient graph is a star, with a central vertex labelled $\widehat{\Delta}_v$ and the leaves labelled by the peripheral subgroups of the adjacent abelian vertices.

For convenience we write $M = \mathcal{M}_{\mathcal{R}_1}$ and consider the subgroup $N = \mu(E)$. The hierarchy $\mathcal{T}_{M, \mu(\overline{c}_v)}$ induces a hierarchy $\mathcal{T}_{N, \mu(\overline{c}_v)}$ which, as in the paragraph after Theorem 2.29, defines a chain of subgroups

$$N_n \subseteq N_{n-1} \subseteq \dots \subseteq N_0 = N$$

and abelian N_i -trees T_i so that N_{i+1} is the stabilizer of a unique vertex u_i of T_i . By definition, the restriction of μ to $\widehat{\Delta}_v$ is an isomorphism

$\widehat{\Delta}_v \rightarrow N_n$. Let $E_i = \mu^{-1}(N_i)$ and let S_i be the minimal E_i -invariant subtree of S .

It is not hard to see that the N_i are all freely indecomposable, and hence that the N_i -trees T_i induced from Grushko trees are trivial. We therefore remove these redundancies from this list in the obvious way, and henceforth assume that each T_i is induced from a primary JSJ tree.

We define canonical equivariant maps $q_i: S_i \rightarrow T_i$ by sending the vertex of S_i stabilized by $\widehat{\Delta}_v$ to u_i , and by sending a vertex w of S_i with abelian stabilizer to the unique vertex of T_i whose stabilizer is the maximal abelian subgroup of N_i containing $\mu(\text{Stab}_{E_i}(w))$.

We now prove that $\mu: E \rightarrow N$ is injective by induction. It suffices to show that if $\mu|_{E_i}: E_i \rightarrow N_i$ is injective then $q_{i-1}: S_{i-1} \rightarrow T_{i-1}$ does not factor through a fold. Indeed, the malnormality of the family of edge groups incident at the central vertex of S_{i-1} shows that a folded edge has non-abelian stabilizer, which contradicts the fact that edge stabilizers in T_{i-1} are abelian. This completes the proof. \square

Remark 5.6. Consider a rigid vertex v of $\widehat{\Delta}$. The proof of Proposition 5.5 uses the hierarchy induced on the image N of the envelope $E = E(\widehat{\Delta}_v)$ under the map μ , which the proof shows is injective on E . The hierarchy $\mathcal{T}_{N, \mu(\overline{c}_v)}$ for N therefore yields a hierarchy $\mathcal{T}_{E, \overline{c}_v}$ for E , which we will make use of later. As in the proof of the proposition, we may assume that every level of $\mathcal{T}_{E, \overline{c}_v}$ is induced by a JSJ tree.

Item (4) is an immediate consequence of Theorem 2.22 and Proposition 5.5.

Corollary 5.7. *The group $\mathcal{M}_{\mathcal{R}}$ is a Γ -limit group.*

Before proving item (5), we first construct a map $\Phi: \text{Mod}(L) \rightarrow \text{Out}(\mathcal{M}_{\mathcal{R}})$. Let $\alpha \in \text{Mod}(L)$. After suitably conjugating the natural maps $\mathcal{M}_{\mathcal{R}}(\Delta_v) \rightarrow \mathcal{M}_{\mathcal{R}}$, we obtain a commutative diagram, and the universal property of $\mathcal{M}_{\mathcal{R}}$ guarantees a unique map $\widehat{\alpha}: \mathcal{M}_{\mathcal{R}} \rightarrow \mathcal{M}_{\mathcal{R}}$. We set $\Phi(\alpha) = \widehat{\alpha}$.

Proposition 5.8. *The map $\Phi: \text{Mod}(L) \rightarrow \text{Out}(\mathcal{M}_{\mathcal{R}})$ is a homomorphism and, for $\alpha \in \text{Mod}(L)$ and $g \in L$, we have that $\eta(\alpha(g)) = \Phi(\alpha)(\eta(g))$. Furthermore, its image lies in $\text{Mod}(\mathcal{M}_{\mathcal{R}})$.*

Proof. Uniqueness in the universal property of $\mathcal{M}_{\mathcal{R}}$ implies in the usual way that the assignment $\Phi(\alpha) = \widehat{\alpha}$ is a homomorphism. From the commutativity of the diagram,

$$\eta(\alpha(g)) = \widehat{\alpha}(\eta(g))$$

for all $g \in L$, as required.

It therefore remains to prove that $\hat{\alpha} \in \text{Mod}(\mathcal{M}_{\mathcal{R}})$. By Definition 2.17, it suffices to show that if $\alpha \in \text{Mod}(L)$ is a generalized Dehn twist then so is $\hat{\alpha} \in \text{Out}(\mathcal{M}_{\mathcal{R}})$. We need to check the three cases of Definition 2.17. The second is immediate: if α is a Dehn twist in an essential two-sided simple closed curve in a socket vertex of Δ , the same curve arises on a socket vertex of $\hat{\Delta}$, and we can also Dehn twist in that.

We deal with Cases (2) and (3) of Definition 2.17 simultaneously, but first we need some notation. For a group G with a collection of subgroups $\mathcal{P} = \{P_1, \dots, P_n\}$, we write $\text{Aut}(G, \mathcal{P})$ for the subgroup of $\text{Aut}(G)$ consisting of those automorphisms α so that, for each i , there is an inner automorphism $\iota_i \in \text{Aut}(G)$ so that $\alpha|_{P_i} = \iota_i$. Let $\text{Out}(G, \mathcal{P})$ be the image of $\text{Aut}(G, \mathcal{P})$ in $\text{Out}(G)$.

Let $[w]$ be an abelian vertex of $\hat{\Delta}$ and let $[e]$ be an incident edge. Then we may assume that α is a generalized Dehn twist either coming from Case (2) of Definition 2.17 and associated to an edge e of Δ , or coming from Case (3) of Definition 2.17 and associated to an abelian vertex w of Δ . Let $\Lambda_{[w]}$ be the graph of abelian groups that arises in the construction of $\hat{\Delta}_{[w]}$, as in §§4.2.

We next delete the equivalence class $[w]$, thought of as a set of abelian vertices of Δ , and let $\Delta_1, \dots, \Delta_n$ be the graphs of groups corresponding to the resulting connected components. Let $P_i = \pi_1 \Delta_i$ for each i , and let $\mathcal{P} = \{P_1, \dots, P_n\}$, and note that α is contained in $\text{Out}(L, \mathcal{P})$.

We next equip $\mathcal{M}_{\mathcal{R}} = \pi_1 \hat{\Delta}$ with a similar family of subgroups. Let $\hat{\Delta}_1, \dots, \hat{\Delta}_m$ be the connected components obtained from $\hat{\Delta}$ after deleting the vertex $[w]$, let $Q_j = \pi_1 \hat{\Delta}_j$ for each j , and let $\mathcal{Q} = \{Q_1, \dots, Q_m\}$.

Since $\hat{\Delta}_{[w]}$ is defined to be the abelianization of $\pi_1 \Lambda_{[w]}$, there is a natural map $\text{Out}(L, \mathcal{P}) \rightarrow \text{Out}(\mathcal{M}_{\mathcal{R}}, \mathcal{Q})$. But the latter is generated by Dehn twists associated to edges incident at $[w]$ and by generalized Dehn twists associated to $[w]$. In particular, α is indeed contained in $\text{Mod}(\mathcal{M}_{\mathcal{R}})$, as claimed. □

6. RELATIVELY IMMUTABLE SUBGROUPS

6.1. Limiting actions on \mathbb{R} -trees and the Rips machine. We recall from [17, 18] that if Υ is a toral relatively hyperbolic group, G is finitely generated group and $\{\rho_i: G \rightarrow \Upsilon\}$ a sequence of non-conjugate homomorphisms with stable kernel $K = \underline{\text{Ker}}(\rho_i)$ then there is a limit action of $L = G/K$ on an \mathbb{R} -tree T_∞ with no global fixed point, abelian edge stabilizers and trivial tripod stabilizers (see [17, Theorem 6.5]). Since the L -action on T_∞ is superstable, the Rips machine (see [22,

Main Theorem]) applies and L admits a nontrivial splitting over an abelian subgroup.

Lemma 6.1. *The group $L = G/\underline{\text{Ker}}(\rho_i)$ described above admits a nontrivial primary splitting.*

Proof. This follows from the way that the L -action on T_∞ is turned into a simplicial action of L on a tree. There are essentially three cases to consider: axial, surface and simplicial. The axial and surface cases are straightforward and the simplicial case follows quickly from the fact that tripod stabilizers are trivial. One needs to be slightly careful using the tree of actions decomposition that [22, Main Theorem] produces, since the splitting may come from a point which is the intersection of two vertex trees, so the edge group needn't fix an arc in the tree. This splitting may not be primary, but the classification of the vertex trees in this decomposition makes it straightforward to find a different nontrivial primary splitting of L . \square

Below, when we consider relatively immutable subgroups, it will be useful to consider homomorphisms that act as inner automorphisms on particular subgroups. To this end, we make the following definition.

Definition 6.2. Suppose that G is a group, and $H \leq G$ a subgroup. Suppose further that \mathcal{P} is a collection of subgroups of H . Denote by $\text{Hom}_{\mathcal{P}}(H, G)$ the set of maps $\lambda: H \rightarrow G$ so that for each $P \in \mathcal{P}$ the restriction of λ to P is the restriction of an inner automorphism of G to P .

Note that the inner automorphism associated to each P might be different for different P .

Lemma 6.3. *Let Υ be a toral relatively hyperbolic group, H a finitely generated subgroup and \mathcal{P} a finite collection of finitely generated subgroups of H . Suppose that $\{\rho_i\}$ is a convergent sequence of non-conjugate homomorphisms from $\text{Hom}_{\mathcal{P}}(H, \Upsilon)$ with trivial stable kernel. Each element of \mathcal{P} fixes a point in the limiting \mathbb{R} -tree T_∞ .*

Moreover, H admits a nontrivial primary splitting relative to \mathcal{P} .

Proof. Consider the homomorphisms $\{\rho_i: H \rightarrow \Upsilon\}$ as in the statement of the lemma, let $P \in \mathcal{P}$ and suppose $p \in P$. Then $\rho_i(p)$ is always conjugate to p in Υ , which implies that there is some $x_i \in \Upsilon$ (the conjugating element) so that $d_\Upsilon(\rho_i(p).x_i, x_i) \leq K$, where K is the word length of p , and in particular doesn't depend on i . By [12, Lemma 5.4] this means that p fixes a point in T_∞ . If we take another $q \in P$ then it will fix the same point of T_∞ . Applying this argument to a finite generating set of each $P \in \mathcal{P}$ shows that there are all elliptic in T_∞ .

It now follows from the relative version of the Rips machine in [21] that H admits a nontrivial abelian splitting relative to \mathcal{P} , and a nontrivial primary relative splitting can be found in the same way as for Lemma 6.1 above. \square

6.2. Relatively immutable subgroups.

Definition 6.4. Let G be a group, $H \leq G$ a subgroup and \mathcal{P} a collection of subgroups of H . The pair (H, \mathcal{P}) is *relatively immutable in G* if the set $\text{Hom}_{\mathcal{P}}(H, G)$ contains only finitely many conjugacy classes of injective homomorphisms.

Lemma 6.5. *Suppose that Υ is toral relatively hyperbolic, that $H \leq \Upsilon$ is finitely generated and that \mathcal{P} is a finite nonempty collection of nontrivial finitely generated subgroups of H . Then (H, \mathcal{P}) is relatively immutable if and only if H admits no nontrivial primary splitting relative to \mathcal{P} .*

Proof. If $\text{Hom}_{\mathcal{P}}(H, \Upsilon)$ contains infinitely many conjugacy classes of injective homomorphisms then Lemma 6.3 implies that H admits a nontrivial primary splitting relative to \mathcal{P} .

On the other hand, if H does admit a nontrivial primary splitting relative to \mathcal{P} then there are infinitely many conjugacy classes of automorphisms of H which are contained in $\text{Hom}_{\mathcal{P}}(H, H)$. This can be seen, for example, by a straightforward adaptation to the relative case of [12, Lemma 3.34]. \square

The following result is similar to results in [12, § 4,5] and also to [20, Lemma 7.5]. In [12] it is assumed that the domain is finitely presented, and in [20] it is assumed that the target is hyperbolic and it is considering immutable rather than relatively immutable subgroups. However, the technical details are similar.

Proposition 6.6. *Suppose that Υ is a toral relatively hyperbolic group, that H is a finitely generated subgroup of Υ and that \mathcal{P} is a finite, nonempty collection of nontrivial finitely generated subgroups of H . Then (H, \mathcal{P}) is relatively immutable if and only if there exist finitely many maps $\rho_1, \dots, \rho_k \in \text{Hom}_{\mathcal{P}}(H, \Upsilon)$ and a positive integer D so that for any map $\eta \in \text{Hom}_{\mathcal{P}}(H, \Upsilon)$ either*

- (1) η is conjugate to ρ_i for some i ; or
- (2) η is not injective on the ball of radius D about 1 in the Cayley graph of H (with respect to the chosen finite generating set).

Proof. If (H, \mathcal{P}) is not relatively immutable then there are infinitely many conjugacy classes of injective maps in $\text{Hom}_{\mathcal{P}}(H, \Upsilon)$, so in this

case there does not exist a D and a finite list as in the statement of the result.

Conversely, suppose that the statement of the result does not hold. Then there is a sequence $\{\phi_i\}$ of non-conjugate maps in $\text{Hom}_{\mathcal{P}}(H, \Upsilon)$ so that ϕ_i is injective on the ball of radius i about 1 in H . We can pass to a convergent subsequence of $\{\phi_i\}$ and obtain a faithful action of H on a limiting \mathbb{R} -tree T_∞ . By Lemma 6.3 H admits a nontrivial primary splitting relative to \mathcal{P} and so (H, \mathcal{P}) is not relatively immutable, by Lemma 6.5. \square

6.3. Equations with rational constraints. In this subsection we recall some results from [10] and [12], following the point of view of [20]. What we do in this subsection is similar to [20, §6], working with toral relatively hyperbolic groups instead of torsion-free hyperbolic groups. In particular, see [20, §2] for an introduction to systems of equations and inequations over groups, and their relationship with homomorphisms. In this section, we will always consider a solution to a system of equations as a homomorphism.

Throughout this section, we fix a toral relatively hyperbolic group Υ . The main theorem in this section is Theorem 6.13 which asserts that we can algorithmically recognize when a finite system of equations and inequations over Υ has finitely many conjugacy classes of solutions (and list conjugacy-representatives of solutions in this case). In fact, the algorithm is *uniform* in the sense that it takes a finite presentation for Υ as input, along with the equations and inequations.

Proving Theorem 6.13 involves understanding certain conjugacy classes of homomorphisms from a group H_Σ (where $\Sigma = 1$ are the equations) to Υ . It is straightforward to check (using [10, Theorem 0.1]) whether or not every homomorphism from H_Σ to Υ has abelian image, and in this case Theorem 6.13 is straightforward. Therefore, we henceforth deal with situations where H_Σ has nonabelian images in Υ , and thereby can (effectively) choose a pair of elements $a, b \in H_\Sigma$ which do not commute (this will be witnessed by non-commuting images in Υ).

Definition 6.7. Let H be a group with a fixed generating set X . A homomorphism $\psi: H \rightarrow \Upsilon$ is *compatible* if it is injective on the ball of radius 8 in H (with respect to X).

Definition 6.8. Let k be a natural number and let $B_{F(X)}(k)$ be the ball of radius k in the free group $F(X)$. A system of equations and inequations $\Sigma = 1, \Lambda \neq 1$ *forces the ball of radius k* if there is a subset $S \subseteq B_{F(X)}(k)$ such that $S \subseteq \Sigma$ and $B_{F(X)}(k) \setminus S \subseteq \Lambda$.

Suppose that $\Sigma = 1, \Lambda \neq 1$ is a system of equations and inequations. If X is finite, there are finitely many subsets of $B_{F(X)}(8)$. For any $S \subseteq B_{F(X)}(8)$, let $\Sigma_S = \Sigma \cup S$ and $\Lambda_S = \Lambda \cap (B_{F(X)}(8) \setminus S)$. For any solution ψ to $\Sigma = 1, \Lambda \neq 1$ there is a unique S so that ψ is a solution to $\Sigma_S = 1, \Lambda_S \neq 1$. Conversely, for any S , a solution to $\Sigma_S = 1, \Lambda_S \neq 1$ is obviously a solution to $\Sigma = 1, \Lambda \neq 1$. Moreover, for any S , any solution ϕ to $\Sigma_S = 1, \Lambda_S \neq 1$ corresponds to a compatible homomorphism $H_{\Sigma_S} \rightarrow \Upsilon$.

Therefore, by replacing $\Sigma = 1, \Lambda \neq 1$ by finitely many systems of equations and inequations, we may consider only solutions corresponding to compatible homomorphisms.

Fix once and for all a pair $a, b \in H$ of noncommuting elements, which exist if H is nonelementary. Given such a and b , [12, Remark 4.8] defines a condition Ω , and [12, Definition 4.5] (first defined in [10]) defines the notion of an *acceptable lift*, which is a choice of words representing the images of the ball of radius 2 under a homomorphism. See [12], [10] for more details.

Definition 6.9. A homomorphism $\psi: H_\Sigma \rightarrow \Upsilon$ is *fairly short* if it is compatible and has an acceptable lift $\tilde{\psi}: B_{H_\Sigma}(2) \rightarrow F$ that satisfies Ω . A homomorphism $\psi: H_\Sigma \rightarrow \Upsilon$ is *very short* if it is compatible and every acceptable lift $\tilde{\psi}: B_{H_\Sigma}(2) \rightarrow F$ satisfies Ω .

The following result is implied by [12, Lemma 4.6, Proposition 4.7].

Lemma 6.10. *Any conjugacy class of compatible homomorphism $H_\Sigma \rightarrow \Upsilon$ contains at least one very short homomorphism and at most finitely many fairly short homomorphisms.*

As explained in [20, §6], the following is now a consequence of [10, Proposition 1.5].

These notions are useful because of the following proposition, which is an immediate consequence of a theorem of Dahmani [10]. We refer the reader to [12, Theorem 3.22] for the statement.

Proposition 6.11. *There exists an algorithm that takes as input:*

- (1) *a finite presentation for a toral relatively hyperbolic group Υ and a pair of non-commuting elements $a, b \in \Upsilon$,*
- (2) *two finite sets of variables $\underline{x} = \{x_i\}, \underline{y} = \{y_j\}$,*
- (3) *a finite system (*) consisting of equations $\Sigma(\underline{x}) = 1$ and inequations $\Lambda(\underline{x}) \neq 1$, and equations $\Theta(\underline{x}, \underline{y}) = 1$ with coefficients in Υ*

and always terminates with answer ‘Yes’ or ‘No’. In case the answer is ‘Yes’, there exists a solution to () in Υ with an acceptable lift of \underline{x}*

that satisfies Ω . In case the answer is ‘No’, there is no solution to (*) in Υ for which every acceptable lift of \underline{x} satisfies Ω .

A special case of this result (where the variables \underline{y} and the system Θ are not required) is the following.

Proposition 6.12. *There exists an algorithm that, given a finite presentation for a toral relatively hyperbolic group, a finite system of equations $\Sigma = 1$ and a finite system of inequations $\Lambda \neq 1$, each with coefficients in Υ , always terminates with answer ‘Yes’ or ‘No’. In case the answer is ‘Yes’, there exists a fairly short solution to $\Sigma = 1$, $\Lambda \neq 1$ in Υ . In case the answer is ‘No’, there does not exist any very short solution.*

The following now has exactly the same proof as [20, Theorem 6.9]. We remark that the analogous result there is stated as if the algorithm depends on which group Υ is being considered. However, as noted in [20, Remark 1.5] all algorithms in that paper are uniform, in the sense that the finite presentation of Υ can be taken as input (with a single algorithm). We state the uniform version below.

Theorem 6.13. *There is an algorithm that takes as input a finite presentation for a toral relatively hyperbolic group Υ , a finite system of equations and inequations $\Sigma = 1$, $\Lambda \neq 1$ and terminates if and only if there are finitely many conjugacy classes of solutions to the system. In case the algorithm terminates, it outputs a list consisting of exactly one representative of each conjugacy class of solutions.*

Note that the proof from [20] uses the solution to the simultaneous conjugacy problem in Υ . One way to see that the simultaneous conjugacy problem is solvable in toral relatively hyperbolic groups is that this is a special case of an equation, which can be solved by [10, Theorem 0.1].

Question 6.14. *In [20], we proved that there exists a ‘complementary’ algorithm to the one above, which terminates if and only if the system has infinitely many conjugacy classes of solutions. Is there such a complementary result in the setting of toral relatively hyperbolic groups?*

6.4. Enumerating relatively immutable subgroups. In this section we adapt the proof of [20, Proposition 7.8] to enumerate *relatively* immutable subgroups (in the toral relatively hyperbolic setting). Immutable subgroups H of a hyperbolic group Γ were characterized as those subgroups for which there is a positive integer D such that the set $\text{Hom}(H, \Gamma)$ contained only finitely many Γ -conjugacy classes of homomorphisms injective on a ball of radius D . We then recognized this

condition by translating it into a set of equations and inequations over Γ , and using the rational constraint Ω to restrict conjugacy classes to finitely many representatives.

Our strategy here is similar. Proposition 6.6 characterizes relatively immutable subgroups (H, \mathcal{P}) of a toral relatively hyperbolic group Υ as those for which the set $\text{Hom}_{\mathcal{P}}(H, \Upsilon)$ contains only finitely many Υ -conjugacy classes injective on a ball of radius D (for some D). We will translate this into a system of equations and inequations over Υ and again use the rational constraint Ω to restrict conjugacy classes to finitely many representatives.

Proposition 6.15 (cf. [20], Proposition 7.8). *There exists a Turing machine that takes as input a finite presentation $\langle X \mid R \rangle$ for a toral relatively hyperbolic group Υ , a finite subset \mathcal{A} of Υ , and a tuple of finite subsets $\{Q_1, \dots, Q_k\}$ so that $Q_i \subseteq \langle \mathcal{A} \rangle$ for each i , and terminates if and only if $(\langle \mathcal{A} \rangle, \mathcal{Q})$ is relatively immutable in Υ , where $\mathcal{Q} = \{\langle Q_i \rangle\}_i$.*

Proof. Let $H = \langle \mathcal{A} \rangle \leq \Upsilon$. We run the following arguments in parallel for increasing values of a positive integer D . For such a positive integer D , let \widehat{H}_D be the group generated by \mathcal{A} with relations equal to all the loops in the ball of radius D in H . These relations can be enumerated using a solution to the word problem in Υ . Let $\{Q_1^D, \dots, Q_k^D\}$ be the tuples of elements of \widehat{H}_D corresponding to $\{Q_1, \dots, Q_k\}$.

A homomorphism $\widehat{H}_D \rightarrow \Upsilon$ which is injective on the ball of radius D in \widehat{H}_D can be characterized by a finite system of equations and inequations over Υ . The equations correspond to the relations in \widehat{H}_D , and the inequations separate the elements of the ball of radius D in \widehat{H}_D . To be precise, we take variables \underline{x} corresponding to the elements of \mathcal{A} , a set of elements $\Sigma_D(\underline{x})$ in the free group $F(\underline{x})$ on \underline{x} corresponding to the loops in the ball of radius D in \widehat{H}_D , and a set of elements $\Lambda_D(\underline{x}) \subseteq F(\underline{x})$ corresponding to the differences between distinct elements of the ball of radius D in \widehat{H}_D . Homomorphisms in $\text{Hom}(\widehat{H}_D, \Upsilon)$ which are injective on the ball of radius D are then in natural bijection with solutions to the system

$$(1) \quad \Sigma_D(\underline{x}) = 1, \quad \Lambda_D(\underline{x}) \neq 1$$

over Υ .

We now need to restrict attention to elements of $\text{Hom}_{\mathcal{Q}_D}(\widehat{H}_D, \Upsilon)$, where $\mathcal{Q}_D = \{\langle Q_1^D \rangle, \dots, \langle Q_k^D \rangle\}$. To do this, we add one variable y_i for each i , together with extra equations which say that, for each i , the conjugate of each element of Q_i^D by y_i is equal to the corresponding element in Q_i . To be precise, for each i , let $Q_i = \{q_{ij}\} \subseteq \Upsilon$ and for

each j let $\eta_{ij}(\underline{x})$ be the corresponding element of $Q_i^D \subseteq \widehat{H}_D$ (written as a word in the generators \underline{x}). For each i and j we add the equation with coefficients

$$(2) \quad \Theta_D(\underline{x}, \underline{y}) := y_i \eta_{ij}(\underline{x}) y_i^{-1} q_{ij}^{-1} = 1$$

over Υ to the system (1), to obtain a system that is denoted by (3).

We consider the projection map from the set of solutions to the system (3) to the set of solutions to the system (1) that we obtain by forgetting the values of the variables y_i . When next argue that we can recognize when the image of this projection consists of finitely many Υ -conjugacy classes.

As long as $D \geq 8$ the system of equations (1) forces the ball of radius 8, and therefore any solution necessarily corresponds to a compatible homomorphism. We now apply the algorithm from Proposition 6.11 to the system (3). If it terminates with ‘No’, there is no very short solution \underline{x} to the system (1) in the image of the projection map from the solutions to (3). In particular, there are no conjugacy classes in the image of the projection, by Lemma 6.10, and it follows that $\text{Hom}_{\mathcal{Q}_D}(\widehat{H}_D, \Upsilon)$ contains no conjugacy classes that are injective on the ball of radius D .

If the algorithm terminates with ‘Yes’, then there is a fairly short solution to (1) in the image of the projection, and a simpleminded search will find a solution $(\underline{x}_0, \underline{y}_0)$ to (3) which projects to a fairly short solution to (1). We now add further inequations (with coefficients) to (1), stipulating that $\underline{x} \neq \underline{x}_0$, and repeat the procedure.

Since the natural quotient map $\widehat{H}_D \rightarrow H$ induces an Υ -equivariant injection $\text{Hom}_{\mathcal{Q}}(H, \Upsilon) \hookrightarrow \text{Hom}_{\mathcal{Q}_D}(\widehat{H}_D, \Upsilon)$, if this algorithm terminates then there is a D such that $\text{Hom}_{\mathcal{Q}}(H, \Upsilon)$ contains only finitely many conjugacy classes that are injective on the ball of radius D , and hence (H, \mathcal{Q}) is relatively immutable.

Conversely, suppose that (H, \mathcal{Q}) is relatively immutable. Then for large enough D the conclusion of Proposition 6.6 is satisfied. Since Υ is equationally Noetherian, for sufficiently large D every homomorphism from \widehat{H}_D to Υ factors through the natural quotient map from \widehat{H}_D to H . Therefore, for sufficiently large D we know that the projection of the solutions of system (3) has only finitely many conjugacy classes. It follows that the algorithm terminates exactly when (H, \mathcal{Q}) is relatively immutable, as required. \square

In our enumeration of Γ -limit groups using models in Section 8, we will require the following elementary observations.

Lemma 6.16. *Suppose that Υ is a toral relatively hyperbolic group, that $H \leq \Upsilon$ and that \mathcal{P} is a collection of nontrivial abelian subgroups of H . If (H, \mathcal{P}) is relatively immutable then so is (H, \mathcal{P}') , where $\mathcal{P}' = \{Z_H(P) \mid P \in \mathcal{P}\}$.*

Also, if (H, \mathcal{P}) is relatively immutable, then so is (H, \mathcal{P}_0) , where \mathcal{P}_0 consists of a collection of H -conjugacy representatives of the elements of \mathcal{P} .

7. CALCULATION OF QUASI-CONVEX ENCLOSURES

Suppose that Υ is a toral relatively hyperbolic group. In this section, we explain how to take a finite subset \mathcal{A} of Υ , along with finitely many finite subsets Q_1, \dots, Q_k (all given as words in the generators of Υ) and algorithmically calculate the quasi-convex enclosure of $\langle \mathcal{A} \rangle$ relative to $\{\langle Q_1 \rangle, \dots, \langle Q_k \rangle\}$. Let $\mathcal{Q} = \{\langle Q_1 \rangle, \dots, \langle Q_k \rangle\}$.

The main difficulty is resolved by Theorem 7.2, which asserts that we can compute *relative* Grushko and JSJ decompositions for toral relatively hyperbolic groups. The absolute version of this theorem follows immediately from [11, Theorem 1.4] and [12, Theorem D]. However, the strategy employed in [11] to compute Grushko decompositions (using connectivity of the Bowditch boundary) does not extend easily to the setting of relatively one-ended groups. Instead, we use the following theorem, which enables us to handle free and abelian splittings simultaneously.

Theorem 7.1. [49, Theorem 18] *Let Υ be finitely generated, and the fundamental group of a graph of groups with infinite cyclic edge groups. Then Υ is one-ended if and only if every vertex group is freely indecomposable relative to the incident edge groups.*

Theorem 7.1 generalizes a result of Shenitzer, and various similar statements and proofs have appeared in the literature [47, 14, 16, 34, 48]. The first author learned of Theorem 7.1 from Sela, who intended to use it to prove Theorem 7.2 in the setting of torsion-free hyperbolic groups, as part of his unpublished proof of the isomorphism problem for those groups, and Fujiwara explained a proof to the second author.

Theorem 7.2. *There is an algorithm which takes as input a finite presentation $\langle \mathcal{A} \mid \mathcal{R} \rangle$ of a toral relatively hyperbolic group Υ and a finite collection of finite tuples $\{Q_1, \dots, Q_n\}$ of elements of Υ and outputs a Grushko decomposition for Υ relative to the subgroups $\mathcal{Q}_i = \langle Q_i \rangle$ and, for each free factor Υ_0 which is freely indecomposable relative to those \mathcal{Q}_i conjugate into Υ_0 , it also outputs the relative primary JSJ decomposition for Υ_0 .*

The proof of Theorem 7.2 follows the strategy of the proof of [12, Theorem D], with two important modifications. We outline the strategy and explain the modifications.

The algorithm in parallel attempts to find non-trivial primary relative splittings of Υ (using the obvious adaptation of [12, Theorem 5.15] to the relative setting) and to prove that Υ admits no such splittings using the algorithm from Proposition 6.15. If a non-trivial splitting is found then we pass to the vertex groups, add the incident edge groups to the relative structure, and repeat. By [3], this procedure eventually terminates in a primary decomposition Λ for Υ in which every vertex group is immutable relative to those \mathcal{Q}_i that are conjugate into it and also to the incident edge groups.

Collapsing the non-trivial edge groups of Λ , one obtains a free splitting Λ_G of Υ in which every vertex group is the fundamental group of a graph of groups with non-trivial abelian edge groups and relatively immutable vertex groups. Note that Λ_G is the relative Grushko decomposition of Υ . Indeed, if some vertex group Υ_v were to admit a relative free splitting then, since non-cyclic abelian subgroups would necessarily be elliptic in such a splitting, it would follow from Theorem 7.1 that *every* edge group incident at Υ_v would be elliptic. Hence, Υ_v would admit a non-trivial relative free splitting, contradicting the relative immutability of Υ_v .

Since we have now found the Grushko decomposition, we restrict attention to a freely indecomposable vertex group, which we shall for notational brevity also call Υ . The splitting Λ of Υ that we have found has the property that every vertex group is relatively immutable; no vertex group admits a non-trivial relative primary splitting. Such a decomposition is obtained from a relative primary JSJ of Υ by cutting each socket vertex along a maximal non-peripheral essential multicurve in the associated surface. The procedure described after [12, Proposition 6.3] explains how to reassemble the socket pieces into a primary JSJ decomposition as required.

Finally, we need to compute the *canonical* primary JSJ decomposition Δ associated with Λ by computing the associated tree of cylinders. We explain how to do this in the following result, which completes the proof of Theorem 7.2.

Proposition 7.3 (cf. [13], Lemma 2.34). *There is an algorithm that takes as input a presentation for a toral relatively hyperbolic group Υ and a primary graph-of-groups decomposition Λ for Υ and outputs the decomposition Δ corresponding to the associated (collapsed) tree of cylinders.*

Proof. Using the word problem in Υ , we can determine which vertex groups of Λ are abelian. We list the non-abelian vertices x_1, \dots, x_k together with presentations of their associated stabilizers $\Delta_{x_i} = \Lambda_{x_i}$. For each i , we also record the images of incident edge maps in Λ_{x_i} .

We next compute conjugacy representatives of the cylinders of Λ as follows. For each pair of edges e_1 and e_2 , we determine whether or not there exists $\gamma \in \Upsilon$ such that $\gamma\Lambda_{e_1}\gamma^{-1}$ commutes with Λ_{e_2} , by encoding this as a system of equations over Υ . If so, we write $e_1 \sim e_2$. We now introduce vertices w_1, \dots, w_l , one for each equivalence class of edges $[e_j]$. For each w_j , we choose a representative edge e_j and label w_j with the centralizer $\Delta_{w_j} = C(\Lambda_{e_j})$.

Finally, we need to describe the edges of Δ . For each non-abelian vertex x_i , we partition the centralizers of incident edge groups into Δ_{x_i} -conjugacy classes. Each such class is conjugate into a unique cylinder corresponding to some w_j . These define the edge groups and the attaching maps. \square

It is now straightforward to compute quasiconvex enclosures: one iteratively computes relative Grushko and JSJ decompositions, and Theorem 2.29 ensures that the procedure terminates. Therefore we have the following result.

Proposition 7.4. *There is an algorithm which takes as input a finite presentation $\langle \mathcal{A} \mid \mathcal{R} \rangle$ of a toral relatively hyperbolic group Υ and a finite subset S of Υ (given as a set of words in \mathcal{A}^\pm) and outputs a finite presentation for the quasi-convex enclosure of $\langle S \rangle$ in Υ .*

8. ENUMERATING Γ -LIMIT GROUPS

We now turn to the enumeration of Γ -limit groups. First, we need to describe the data that we will use to specify a Γ -limit group $L = L_0$. Our enumeration will be non-unique, and indeed our limit group L_0 will come equipped with a strict resolution \mathcal{R}_0 . The philosophy is to specify L_0 by specifying its model $M_0 := \mathcal{M}_{\mathcal{R}_0}$, together with certain nicely specified subsets of M_0 that generate L_0 . The data will also include the models of all the limit groups in the resolution \mathcal{R}_0 , forming a strict resolution of M_0 .

Definition 8.1. A *model pair of (strict) resolutions* is a diagram of maps of strict resolutions of the form

$$\begin{array}{ccccccc} L_0 & \xrightarrow{\lambda_0} & L_1 & \xrightarrow{\lambda_1} & \cdots & \xrightarrow{\lambda_{n-1}} & L_n & \xrightarrow{\lambda_n} & \Gamma \\ \downarrow \eta_0 & & \downarrow \eta_1 & & & & \downarrow \eta_n & & \downarrow \text{id} \\ M_0 & \xrightarrow{\mu_0} & M_1 & \xrightarrow{\mu_1} & \cdots & \xrightarrow{\mu_{n-1}} & M_n & \xrightarrow{\mu_n} & \Gamma \end{array}$$

where the first line is a strict resolution \mathcal{R}_0 of L_0 followed by a strict map $\lambda_n: L_n \rightarrow \Gamma$, and the second line consists of the canonically associated resolution of models (along with a strict map $\mu_n: M_n \rightarrow \Gamma$). The *length* of this model pair is n .

In such a model pair of strict resolutions, the pair of Γ -limit groups $\eta_0: L_0 \rightarrow M_0$ is said to be a *model pair of Γ -limit groups built over \mathcal{R}_1* , where \mathcal{R}_1 is the resolution of L_1 induced by truncating \mathcal{R}_0 .

By abuse of indices, we refer to the identity map $\Gamma \rightarrow \Gamma$ as a model pair of (strict) resolutions of length -1 . Note that by Proposition 3.5, when Γ is considered with respect to the trivial resolution, the model of Γ is Γ .

Definition 8.2. An *effective pair of resolutions of length n* (specifying a resolution \mathcal{R}_0) consists of a commutative diagram of the form

$$\begin{array}{ccccccc} L_0 & \xrightarrow{\lambda_0} & L_1 & \xrightarrow{\lambda_1} & \cdots & \xrightarrow{\lambda_{n-1}} & L_n & \xrightarrow{\lambda_n} & \Gamma \\ \downarrow \eta_0 & & \downarrow \eta_1 & & & & \downarrow \eta_n & & \downarrow \text{id} \\ M_0 & \xrightarrow{\mu_0} & M_1 & \xrightarrow{\mu_1} & \cdots & \xrightarrow{\mu_{n-1}} & M_n & \xrightarrow{\mu_n} & \Gamma \end{array}$$

where each M_i is specified by a finite presentation $\langle \mathcal{A}_i \mid \mathcal{R}_i \rangle$, each L_i is specified by a finite set \mathcal{S}_i , each map η_i is specified by a map $\mathcal{S}_i \rightarrow F(\mathcal{A}_i)$, each map $\mu_i: M_i \rightarrow M_{i+1}$ is specified by a map $\mathcal{A}_i \rightarrow F(\mathcal{A}_{i+1})$, and each map λ_i is specified by a map $\mathcal{S}_i \rightarrow F(\mathcal{S}_{i+1})$.

An *effective model pair of resolutions* is an effective pair of resolutions that defines a model pair.

We remark that, on the face of it, there is no effective procedure for checking whether or not the maps λ_i and μ_i that we specify are strict, or that M_i is the model of L_i (with respect to \mathcal{R}_i). However, we will describe a Turing machine that enumerates effective model pairs for which these properties do hold.

Theorem 8.3. *There is a Turing machine that takes as input a presentation for a torsion-free hyperbolic group Γ and outputs a list of effective pairs of resolutions over Γ such that:*

- every effective resolution on the list defines a model pair of strict resolutions;
- for every strict resolution \mathcal{R} over Γ , every effective pair of resolutions that specifies the model pair canonically associated to \mathcal{R} appears on the list.

Proof. We begin our enumeration with the unique effective model pair of resolutions of length -1 , which is specified by the given presentation

for Γ (considered as the model), the given generating set for Γ (considered as the Γ -limit group) and the identity map on this generating set (specifying the map from the Γ -limit group to its model).

We now provide a parallel enumeration of all effective model pairs of resolutions. Suppose that we have an effective pair of resolutions of length $n - 1$ (for some $n \geq 0$) given by

$$\begin{array}{ccccccc} L_1 & \xrightarrow{\lambda_1} & \cdots & \xrightarrow{\lambda_{n-1}} & L_n & \xrightarrow{\lambda_n} & \Gamma \\ \downarrow \eta_1 & & & & \downarrow \eta_n & & \downarrow \text{id} \\ M_1 & \xrightarrow{\mu_1} & \cdots & \xrightarrow{\mu_{n-1}} & M_n & \xrightarrow{\mu_n} & \Gamma \end{array}$$

along with data \mathcal{A}_i , \mathcal{R}_i and maps as in Definition 8.2. (We start the indices with 1 so that we can append a model pair $\eta_0: L_0 \rightarrow M_0$ built over the model pair of resolutions.)

We now begin the enumeration of the model pairs $\eta_0: L_0 \rightarrow M_0$ built over the above effective pair. We proceed to enumerate in parallel all of the possible collections of data as follows:

- (1) A finite bipartite graph Δ_0 with *red* and *blue* vertices, and a choice of maximal tree in Δ_0 .
- (2) A decomposition of the blue vertices of Δ_0 into two disjoint families: the ‘socket’ vertices, and the ‘rigid’ vertices. The red vertices of Δ_0 are all ‘abelian’ vertices.
- (3) For each edge e of Δ_0 , a finite set of words Q_e in S_1^\pm which define nontrivial commuting elements of M_1 . If an edge e is adjacent to a socket vertex of Δ_0 , the set Q_e consists of a single element q_e .
- (4) For each rigid vertex u , with adjacent edges e_1, \dots, e_k , a finite set of words W_u in S_1^\pm so that
 - (a) Q_{e_1}, \dots, Q_{e_k} are all in the subgroup $\langle W_u \rangle$; and
 - (b) $(\langle W_u \rangle, \{\langle Q_{e_i} \rangle\})$ is relatively immutable in M_1 .
- (5) For a socket vertex v , a presentation of the socket of the fundamental group Σ_v of a surface with boundary, with boundary components in bijection with certain powers of generators q_e of edge groups for edges adjacent to v . There is also a map from Σ_v to M_1 with image lying in $\langle S_1 \rangle$ so that the generator of a boundary component maps to a power of the corresponding q_e (under the bijection between boundary components and edges adjacent to v) and the image of Σ_v is nonabelian.
- (6) For each edge e of Δ_0 we choose elements $\tau_e \in M_1$, and insist that $\tau_e = 1$ if e is in the maximal tree of Δ_0 .

- (7) For each abelian vertex group w , a finite set of words W_w in S_1^\pm which define nontrivial commuting elements of M_1 so that $\langle W_w \rangle$ contains each of the $\tau_e Q_e \tau_e^{-1}$ for edges e adjacent to w .

Note that each such piece of data can be enumerated in a simple-minded manner, with the verification that it satisfies the appropriate property also being verified by the appropriate simple-minded enumeration, or an application of Proposition 6.15 in the case of relatively immutable subgroups.

We will define a model Γ -limit group M as the fundamental group of a graph of groups $\widehat{\Delta}$ which will be constructed from the above data as in Section 4. The subgroup L of which M is the model will also be defined as the fundamental group of a graph of groups Δ . We will construct $\widehat{\Delta}$ algorithmically. However, Δ need not be equal to Δ_0 , and we do not know how to effectively construct Δ . On the other hand, we will obtain an explicit finite list of generators for L , as required.

We now define the graph of groups $\widehat{\Delta}$ as follows. Suppose that u is a rigid vertex of Δ_0 , and consider the subgroup $\langle W_u \rangle$ of M_1 . Let $\widehat{\Delta}_u$ be the quasiconvex enclosure of $\langle W_u \rangle$ in M_1 . A finite presentation for $\widehat{\Delta}_u$ can be calculated by Proposition 7.4. If e is an edge in Δ adjacent to u , we define $\widehat{\Delta}_e$ to be the centralizer in $\widehat{\Delta}_u$ of $\langle Q_e \rangle$, which can be calculated by Theorem 2.33 (note that since $\widehat{\Delta}_u$ is toral relatively hyperbolic, it is CSA, so the centralizer of $\langle Q_e \rangle$ is the centralizer of any nontrivial element of Q_e).

We define an equivalence relation on the edges of Δ_0 by saying that two edges e_1 and e_2 , adjacent to a rigid vertex u are equivalent if $\widehat{\Delta}_{e_1}$ and $\widehat{\Delta}_{e_2}$ are conjugate in $\widehat{\Delta}_u$. Since these subgroups are nontrivial abelian subgroups, and $\widehat{\Delta}_u$ is toral relatively hyperbolic, we can determine this equivalence using CSA and the algorithm to solve equations over $\widehat{\Delta}_u$ (note that both $\widehat{\Delta}_{e_1}$ and $\widehat{\Delta}_{e_2}$ are maximal abelian subgroups of $\widehat{\Delta}_u$, by construction).

As in Section 4, the graph of $\widehat{\Delta}$ is the quotient graph of Δ_0 obtained by identifying equivalent edges. Each vertex in $\widehat{\Delta}$ is naturally exactly one of rigid, socket or abelian. The definition of the graph of groups $\widehat{\Delta}$ is now exactly as in Section 4. We have already defined the rigid vertex groups. The socket vertex groups in $\widehat{\Delta}$ are the sockets obtained from the socket vertices of Δ_0 by attaching maximal roots in M to generators of boundary components. The edge groups are centralizers in the same way as in Section 4, and the abelian vertex groups are abelianizations of the fundamental group of a graph of groups exactly as in Section 4. Then we let $M = \pi_1 \widehat{\Delta}$.

We next specify the finite set S and the map $S \rightarrow M$ that defines L . The data listed above naturally defines a graph of groups structure on Δ_0 . Note that we do *not* have access to presentations for the vertex and edge groups of Δ_0 (indeed, the vertex groups may not be finitely presentable). However, we may take S to be a natural generating set for $\pi_1\Delta_0$ corresponding to the disjoint union of:

- the subsets Q_e (as e ranges over the edges of Δ_0);
- the W_u (as u ranges over all rigid vertices);
- the generating sets specified by the chosen presentation for Σ_v (as v ranges over all socket vertices);
- the W_w (as w ranges over all abelian vertices);
- the choices of stable letters τ_e for edges e of Δ_0 not in the chosen maximal tree.

The map $S \rightarrow M$ is now given by the construction of M .

It follows immediately from Proposition 4.9 that M is the model of L (with respect to the resolution $L \rightarrow L_1 \rightarrow \dots \rightarrow L_n \rightarrow \Gamma$).

Enumerating these data enumerates all *freely indecomposable* effective model pairs built over the above effective pair of resolutions. To enumerate *all* effective model pairs $L_0 \rightarrow M_0$, we enumerate finite sets of such data, together with finitely generated free groups, and take their free products.

Proceeding in this way, recursively and in parallel, over all effective pairs of resolutions that we build, we obtain a list of effective pairs of resolutions over Γ . It is clear that every model pair arises in this list, and we have shown that every pair of resolutions that we build is a model pair. \square

We end by remarking that we do not know how to algorithmically decide if a given (strict) map $\lambda_n: L_n \rightarrow \Gamma$ is an embedding of the freely indecomposable free factors of L_n , nor do we know how to algorithmically compute the JSJ decomposition of a given finitely generated subgroup of Γ . Therefore, terminating with a strict map $\lambda_n: L_n \rightarrow \Gamma$ is the best that we can do. This introduces some extra redundancy in the list of resolutions, as in Remark 2.23, but we do not believe it affects the utility of the construction.

REFERENCES

- [1] E. Alibegović. A combination theorem for relatively hyperbolic groups. *Bull. London Math. Soc.*, 37(3):459–466, 2005.
- [2] E. Alibegović. A combination theorem for relatively hyperbolic groups. *Bull. London Math. Soc.*, 37(3):459–466, 2005.

- [3] M. Bestvina and M. Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.
- [4] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.
- [5] M. Bestvina and M. Feighn. Notes on Sela’s work: limit groups and Makanin-Razborov diagrams. In *Geometric and cohomological methods in group theory*, volume 358 of *London Math. Soc. Lecture Note Ser.*, pages 1–29. Cambridge Univ. Press, Cambridge, 2009.
- [6] B. H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [7] C. H. Cashen and N. Macura. Line patterns in free groups. preprint, [arXiv:1006.2123](#), 2010.
- [8] F. Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963 (electronic), 2003.
- [9] F. Dahmani. Finding relative hyperbolic structures. *Bull. Lond. Math. Soc.*, 40(3):395–404, 2008.
- [10] F. Dahmani. Existential questions in (relatively) hyperbolic groups. *Israel J. Math.*, 173:91–124, 2009.
- [11] F. Dahmani and D. Groves. Detecting free splittings in relatively hyperbolic groups. *Trans. Amer. Math. Soc.*, 360(12):6303–6318, 2008.
- [12] F. Dahmani and D. Groves. The isomorphism problem for toral relatively hyperbolic groups. *Publ. Math. Inst. Hautes Études Sci.*, 107:211–290, 2008.
- [13] F. Dahmani and V. Guirardel. The isomorphism problem for all hyperbolic groups. *Geom. Funct. Anal.*, 21(2):223–300, 2011.
- [14] G. Diao and M. Feighn. The Grushko decomposition of a finite graph of finite rank free groups: an algorithm. *Geometry and Topology*, 9:1835–1880, 2005.
- [15] M. J. Dunwoody and M. E. Sageev. JSJ-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.
- [16] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.
- [17] D. Groves. Limit groups for relatively hyperbolic groups, i: The basic tools. Preprint, available at [front.math.ucdavis.edu/math.GR/0412492](#), 2004.
- [18] D. Groves. Limit groups for relatively hyperbolic groups. II. Makanin-Razborov diagrams. *Geom. Topol.*, 9:2319–2358, 2005.
- [19] D. Groves and H. Wilton. Enumerating limit groups. *Groups Geom. Dyn.*, 3(3):389–399, 2009.
- [20] D. Groves and H. Wilton. Conjugacy classes of solutions to equations and inequations over hyperbolic groups. *J. Topol.*, 3(2):311–332, 2010.
- [21] V. Guirardel. Limit groups and groups acting freely on \mathbb{R}^n -trees. *Geom. Topol.*, 8:1427–1470 (electronic), 2004.
- [22] V. Guirardel. Actions of finitely generated groups on \mathbb{R} -trees. *Ann. Inst. Fourier (Grenoble)*, 58(1):159–211, 2008.
- [23] V. Guirardel and G. Levitt. JSJ decompositions: definitions, existence, uniqueness. I: The JSJ deformation space. preprint, [arXiv:0911.3173](#), 2009.
- [24] V. Guirardel and G. Levitt. JSJ decompositions: definitions, existence, uniqueness. II: Compatibility and acylindricity. preprint, [arXiv:1002.4564](#), 2010.
- [25] V. Guirardel and G. Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

- [26] V. Guirardel and G. Levitt. McCool groups of toral relatively hyperbolic groups. *Algebr. Geom. Topol.*, 15(6):3485–3534, 2015.
- [27] V. Guirardel and G. Levitt. Splittings and automorphisms of relatively hyperbolic groups. *Groups Geom. Dyn.*, 9(2):599–663, 2015.
- [28] V. Guirardel and G. Levitt. JSJ decompositions of groups. arxiv:1602.05139, 2016.
- [29] F. Haglund. Finite index subgroups of graph products. *Geom. Dedicata*, 135:167–209, 2008.
- [30] G. C. Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.
- [31] O. Kharlampovich, A. Miasnikov, and A. Taam. Effective construction of canonical hom-diagrams for equations over torsion-free hyperbolic groups. arxiv:1501.03097v2, 2015.
- [32] O. Kharlampovich and A. Myasnikov. Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz. *Journal of Algebra*, 200(2):472–516, 1998.
- [33] O. Kharlampovich and A. Myasnikov. Decidability of the elementary theory of a torsion-free hyperbolic group. arXiv:1303.0760v7, 2015.
- [34] L. Louder. Scott complexity and adjoining roots to finitely generated groups. *Groups Geom. Dyn.*, 7(2):451–474, 2013.
- [35] L. Louder and N. W. M. Touikan. Strong accessibility for finitely presented groups. preprint, 2015.
- [36] M. Mj and L. Reeves. A combination theorem for strong relative hyperbolicity. *Geom. Topol.*, 12(3):1777–1798, 2008.
- [37] A. Ould Houcine. Limit groups of equationally Noetherian groups. In *Geometric group theory*, Trends Math., pages 103–119. Birkhäuser, Basel, 2007.
- [38] C. Perin. Elementary embeddings in torsion-free hyperbolic groups. *Ann. Sci. Éc. Norm. Supér. (4)*, 44(4):631–681, 2011.
- [39] C. Perin, A. Pillay, R. Sklinos, and K. Tent. On groups and fields interpretable in torsion-free hyperbolic groups. *Münster J. Math.*, 7(2):609–621, 2014.
- [40] E. Rips. Subgroups of small cancellation groups. *The Bulletin of the London Mathematical Society*, 14(1):45–47, 1982.
- [41] E. Rips and Z. Sela. Canonical representatives and equations in hyperbolic groups. *Invent. Math.*, 120(3):489–512, 1995.
- [42] E. Rips and Z. Sela. Cyclic splittings of finitely presented groups and the canonical JSJ decomposition. *Ann. of Math. (2)*, 146(1):53–109, 1997.
- [43] Z. Sela. The isomorphism problem for hyperbolic groups I. *Annals of Mathematics*, 141:217–283, 1995.
- [44] Z. Sela. Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II. *Geom. Funct. Anal.*, 7(3):561–593, 1997.
- [45] Z. Sela. Diophantine geometry over groups. I. Makanin-Razborov diagrams. *Publ. Math. Inst. Hautes Études Sci.*, 93:31–105, 2001.
- [46] Z. Sela. Diophantine geometry over groups. VII. The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc. (3)*, 99(1):217–273, 2009.
- [47] A. Shenitzer. Decomposition of a group with a single defining relation into a free product. *Proceedings of the American Mathematical Society*, 6:273–279, 1955.

- [48] N. Touikan. On the one-endedness of graphs of groups. *Pacific J. Math.*, 278(2):463–478, 2015.
- [49] H. Wilton. One-ended subgroups of graphs of free groups with cyclic edge groups. *Geom. Topol.*, 16(2):665–683, 2012.
- [50] D. T. Wise. The structure of groups with a quasiconvex hierarchy. unpublished manuscript.

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

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

DPMMS, CENTRE FOR MATHEMATICAL SCIENCES, WILBERFORCE ROAD, CAMBRIDGE, CB3 0WB, UK

E-mail address: h.wilton@maths.cam.ac.uk