FINITELY PRESENTED RESIDUALLY FREE GROUPS

MARTIN R. BRIDSON, JAMES HOWIE, CHARLES F. MILLER III,
AND HAMISH SHORT

ABSTRACT. We establish a general criterion for the finite presentability of subdirect products of groups and use this to characterize finitely presented residually free groups. We prove that, for all $n \in \mathbb{N}$, a residually free group is of type FP_n if and only if it is of type F_n.

New families of subdirect products of free groups are constructed, including the first examples of finitely presented subgroups that are neither FP_∞ nor of Stallings-Bieri type. The template for these examples leads to a more constructive characterization of finitely presented residually free groups up to commensurability.

We show that the class of finitely presented residually free groups is recursively enumerable and present a reduction of the isomorphism problem. A new algorithm is described which, given a finite presentation of a residually free group, constructs a canonical embedding into a direct product of finitely many limit groups. The (multiple) conjugacy and membership problems for finitely presented subgroups of residually free groups are solved.

1. INTRODUCTION

This article is part of a project to understand the finitely presented residually free groups. The prototypes for these groups are the finitely presented subgroups of finite direct products of free and surface groups, and in general such a group is a full subdirect product of finitely many limit groups, i.e. it can be embedded in a finite direct product of limit groups so that it intersects each factor non-trivially and projects onto each factor (cf. Theorem A). In our earlier studies [10], [7], [8], [9], we proved that these full subdirect products have finite index in the ambient product if they are of type FP_∞. We also proved that in general they virtually contain a term of the lower central series of the product. These tight restrictions set the *finitely presented* subdirect

products of limit groups apart from those that are merely *finitely generated*, since the finitely generated subgroups of the direct product of two free groups are already hopelessly complicated [28]. Nevertheless, a thorough understanding of the finitely presented subdirect products of free and limit groups has remained a distant prospect, with only a few types of examples known.

In this article we pursue such an understanding in a number of ways. We characterize finitely presented residually free groups among the full subdirect products of limit groups in terms of their projections to the direct factors. A revealing family of finitely presented full subdirect products of free groups is constructed; this gives rise to a more constructive characterization of finitely presented residually free groups. We give algorithms for finding finite presentations when they exist, for constructing certain canonical embeddings, for enumerating finitely presented residually free groups and for solving their conjugacy and membership problems.

Residually free groups provide a context for a rich and powerful interplay among group theory, topology and logic. By definition, a group G is *residually free* if, for every $1 \neq g \in G$, there is a homomorphism ϕ from G to a free group F such that $1 \neq \phi(g)$ in F. In other words G is isomorphic to a subgroup of an unrestricted direct product of free groups. In general, one requires infinitely many factors in this direct product, even if G is finitely generated. For example, the fundamental group of a closed orientable surface Σ is residually free but it cannot be embedded in a finite direct product if $\chi(\Sigma) < 0$, since $\pi_1\Sigma$ does not contain $\mathbb{Z}^2$ and is not a subgroup of a free group. However, Baumslag, Myasnikov and Remeslennikov [3, Corollary 19] proved that one can force the enveloping product to be finite at the cost of replacing free groups by $\exists$-*free groups* (see also [22, Corollary 2] and [30, Claim 7.5]). In [23] Kharlampovich and Myasnikov describe an algorithm to find such an embedding, based on the deep work of Makanin [27] and Razborov [29]. We shall describe a new algorithm that does not depend on [27] and [29]; the embedding that we construct is canonical in a strong sense (see Theorem A).

By definition, $\exists$-free groups have the same universal theory as a free group; they are now more commonly known as *limit groups*, a term coined by Sela [30]. They have been much studied in recent years in connection with Tarski's problems on the first order logic of free groups [30], [22]. They have been shown to enjoy a rich geometric structure. A useful characterisation of limit groups is that they are the finitely generated groups G that are *fully residually free*: for every finite subset $A \subset G$, there is a homomorphism from G to a free group that restricts to an injection on A.

For the most part, we treat finitely generated residually free groups S as subdirect products of limit groups. There are at least two obvious

drawbacks to this approach: the ambient product of limit groups is not canonically associated to S; and given a direct product of limit groups, one needs to determine which finitely generated subgroups are finitely presented.

The first of these drawbacks is overcome by items (1), (3) and (4) of the following theorem. Item (2) is based on Theorem 4.2 of [9]. We remind the reader that a subgroup of a direct product of groups is termed a *subdirect product* if its projection to each factor is surjective. A subdirect product is said to be *full* if it intersects each of the direct factors non-trivially.

Theorem A. *There is an algorithm that, given a finite presentation of a residually free group S, will construct an embedding $\iota : S \hookrightarrow \exists\mathrm{Env}(S)$, so that*

(1) $\exists\mathrm{Env}(S) = \Gamma_{\mathrm{ab}} \times \exists\mathrm{Env}_0(S)$ *where* $\Gamma_{\mathrm{ab}} = H_1(S, \mathbb{Z})/(\mathrm{torsion})$ *and* $\exists\mathrm{Env}_0(S) = \Gamma_1 \times \cdots \times \Gamma_n$ *is a direct product of non-abelian limit groups* Γ_i. *The intersection of S with the kernel of the projection* $\rho : \exists\mathrm{Env}(S) \to \exists\mathrm{Env}_0(S)$ *is the centre $Z(S)$ of S, and $\rho(S)$ is a full subdirect product.*

(2) $L_i := \Gamma_i \cap S$ *contains a term of the lower central series of a subgroup of finite index in Γ_i, for $i = 1, \ldots, n$, and therefore* $\mathrm{Nilp}_{\exists}(S) := \exists\mathrm{Env}(S)/(L_1 \times \cdots \times L_n)$ *is virtually nilpotent.*

(3) [Universal Property] *For every homomorphism* $\phi : S \to D = \Lambda_1 \times \cdots \times \Lambda_m$, *with $\phi(S)$ subdirect and the Λ_i non-abelian limit groups, there exists a unique homomorphism* $\hat{\phi} : \exists\mathrm{Env}_0(S) \to D$ *with* $\hat{\phi} \circ \rho|_S = \phi$;

(4) [Uniqueness] *moreover, if $\phi : S \hookrightarrow D$ embeds S as a full subdirect product, then $\hat{\phi} : \exists\mathrm{Env}_0(S) \to D$ is an isomorphism that respects the direct sum decomposition.*

The group $\exists\mathrm{Env}(S)$ in Theorem A is called the *existential envelope* of S and the associated factor $\exists\mathrm{Env}_0(S)$ is the *reduced existential envelope*. The projection ρ embeds $S/Z(S)$ in $\exists\mathrm{Env}_0(S)$, and $\rho(S) \subset \exists\mathrm{Env}_0(S)$ is always a full subdirect product. The subgroup $S \subset \exists\mathrm{Env}(S)$ is always a subdirect product but it is full if and only if S has a non-trivial centre.

The second of the drawbacks we identified in the discussion preceding Theorem A is resolved by item (4) of the following theorem. In order to state this theorem concisely we introduce the following temporary definition: an embedding $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$ of a residually free group S as a full subdirect product of limit groups is said to be *neat* if Γ_0 is abelian (possibly trivial), $S \cap \Gamma_0$ is of finite index in Γ_0, and Γ_i is non-abelian for $i = 1, \ldots, n$.

Theorem B. *Let S be a finitely generated residually free group. Then the following conditions are equivalent:*

(1) *S is finitely presentable;*

(2) S *is of type* $\mathrm{FP}_2(\mathbb{Q})$;
(3) $\dim H_2(S_0; \mathbb{Q}) < \infty$ *for all subgroups* $S_0 \subset S$ *of finite index;*
(4) *there exists a neat embedding* $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$ *into a product of limit groups such that the image of* S *under the projection to* $\Gamma_i \times \Gamma_j$ *has finite index for* $1 \le i < j \le n$;
(5) *for every neat embedding* $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$ *into a product of limit groups, the image of* S *under the projection to* $\Gamma_i \times \Gamma_j$ *has finite index for* $1 \le i < j \le n$.

Corollary C. *For all* $n \in \mathbb{N}$, *a residually free group* S *is of type* $\mathrm{F_n}$ *if and only if it is of type* $\mathrm{FP_n}(\mathbb{Q})$.

Subsequent to our work, D. Kochloukova [21] has obtained results concerning the question of which subdirect products of limit groups are FP_k for $2 < k < n$.

It follows from Theorem B that any subgroup $T \subset \exists \mathrm{Env}(S)$ containing S is again finitely presented. More generally we prove:

Theorem D. *Let* $n \ge 2$ *be an integer, let* $S \subset D := \Gamma_1 \times \cdots \times \Gamma_k$ *be a full subdirect product of limit groups, and let* $T \subset D$ *be a subgroup that contains* S. *If* S *is of type* $\mathrm{FP_n}(\mathbb{Q})$ *then so is* T.

The proof of Theorem B relies on our earlier work concerning the finiteness properties of subgroups of direct products of limit groups [9] and the following new criterion for the finite presentability of subdirect products.

Theorem E. *Let* $S \subset G_1 \times \cdots \times G_n$ *be a subgroup of a direct product of finitely presented groups. If for all* $i, j \in \{1, \ldots, n\}$, *the projection* $p_{ij}(S) \subset G_i \times G_j$ *has finite index, then* S *is finitely presented.*

An essential ingredient in the proof of this result is the following asymmetric version of the 1-2-3 Theorem of [2].

Theorem F (Asymmetric 1-2-3 Theorem). *Let* $f_1 : \Gamma_1 \to Q$ *and* $f_2 : \Gamma_2 \to Q$ *be surjective group homomorphisms. Suppose that* Γ_1 *and* Γ_2 *are finitely presented, that* Q *is of type* F_3, *and that at least one of* $\ker(f_1)$ *and* $\ker(f_2)$ *is finitely generated. Then the fibre-product of* f_1 *and* f_2,

$$P = \{(g, h) \mid f_1(g) = f_2(h)\} \subset \Gamma_1 \times \Gamma_2,$$

is finitely presented.

In Theorem 3.4 and Theorem 2.2 we shall describe *effective* versions of Theorems E and F that yield an explicit finite presentation for S. In the final section we shall use these algorithmic versions to prove:

Theorem G. *The class of finitely presented, residually free groups is recursively enumerable. More explicitly, there exists a Turing machine that generates a list of finite group presentations so that each of the groups presented is residually free and every finitely presented residually free group is isomorphic to at least one of the groups presented.*

In Section 4 we turn our attention to the construction of new families of finitely presented residually free groups. We construct the first examples of finitely presented subgroups of direct products of free groups that are neither $\mathrm{FP}_\infty(\mathbb{Q})$ nor of Stallings-Bieri type, thus answering a question in [10]. Subdirect products of free and surface groups demand particular attention because, in addition to their historical interest, the work of Delzant and Gromov [16] shows that such subgroups play an important role in the problem of determining which finitely presented groups arise as the fundamental groups of compact Kähler manifolds.

We use the standard notation $\gamma_n(G)$ to denote the n-th term of the lower central series of a group.

Theorem H. *If c and n are positive integers with $n \geq c + 2$, and $D = F_1 \times \cdots \times F_n$ is a direct product of free groups of rank 2, then there exists a finitely presented subgroup $S \subset D$ with $S \cap F_i = \gamma_{c+1}(F_i)$ for $i = 1, \ldots, n$.*

$\mathrm{Nilp}_\exists(S)$ was defined in Theorem A(2).

Corollary I. *For all positive integers c and $n \geq c + 2$, there exists a finitely presented residually free group S for which $\mathrm{Nilp}_\exists(S)$ is a direct product of n copies of the 2-generator free nilpotent group of class c.*

The proof that the group S in Theorem H is finitely presented relies on our earlier structural results. Our proof of the equality $S \cap F_i = \gamma_{c+1}(F_i)$ exploits the Magnus embedding of the free group of rank 2 into the group of units of $\mathbb{Q}[[\alpha, \beta]]$, the algebra of power series in two non-commuting variables with rational coefficients.

Theorem B describes the finitely presented residually free groups. A description of a quite different nature is given in Theorem 5.5: using a template inspired by the examples in Section 4 we prove that every finitely presented residually free group is commensurable with a particular type of subdirect product of limit groups.

In Section 7 we apply Theorem A to elucidate the algorithmic structure of finitely presented residually free groups. The restriction to finitely presented groups is essential, since decision problems for arbitrary finitely generated residually free groups are hopelessly difficult. For example, there are finitely generated subgroups of a direct product of two free groups for which the conjugacy problem and membership problem are unsolvable; and the isomorphism problem is unsolvable amongst such groups [28].

The following statement includes the statement that the conjugacy problem is solvable in every finitely presented residually free group.

Theorem J. *Let S be a finitely presented residually free group. There exists an algorithm that, given an integer n and two n-tuples of words in the generators of S, say $(u_1, \ldots, u_n)$ and $(v_1, \ldots, v_n)$, will determine*

whether or not there exists an element $s \in S$ *such that* $su_i s^{-1} = v_i$ *in* S *for* $i = 1, \ldots, n$.

Theorem K. *If* S *is a finitely presented residually free group and* $H \subset S$ *is a finitely presented subgroup, then there is an algorithm that given an arbitrary word* w *in the generators of* S *can determine whether or not* w *defines an element of* H.

Since the completion of our work, alternative approaches to the conjugacy and membership problems have been developed in [12] and [14]. In the final section of this paper we make a few remarks about the isomorphism problem for finitely presented residually free groups, taking account of the canonical embeddings $S \hookrightarrow \exists \mathrm{Env}(S)$.

This paper is organised as follows. Our first goal is to prove an effective version of the Asymmetric 1-2-3 Theorem; this is achieved in Section 2. In Section 3 we establish Theorem E. In Section 4 we construct the groups described in Theorem H. In Section 5 we establish the two characterisations of finitely presented residually free groups promised earlier: we prove Theorem B and Theorem 5.5. Section 6 is devoted to the proof of Theorem A and other aspects of the canonical embedding $S \hookrightarrow \exists \mathrm{Env}(S)$. Finally, in Section 7, we turn our attention to decidability and enumeration problems, proving Theorems G, J and K.

We thank G. Baumslag, W. Dison, D. Kochloukova, A. Myasnikov, Z. Sela, H. Wilton and, most particularly, M. Vaughan-Lee for helpful comments and suggestions relating to this work.

2. The Asymmetric 1-2-3 Theorem

Our proofs of Theorems E and B rely crucially on the following asymmetric version of the 1-2-3 Theorem from [2]. The adaptation from [2] is straightforward and has been written out in complete detail by W. Dison in his doctoral thesis [17]. We recall the main points of the proof here, largely because the reader will need them to hand in order to follow the proof of the *effective* version of the theorem that is proved in Subsection 2.1. The basic Asymmetric 1-2-3 Theorem states that a certain type of fibre product is finitely presented, whereas the effective version provides an algorithm that, given natural input data, constructs a finite presentation for the fibre product. This enhanced version of the theorem will play a crucial role in our proof that the class of finitely presented residually free groups is recursively enumerable.

We remind the reader that a group G is said to be of type F_3 if it there is a $K(G, 1)$ with finitely many cells in the 3-skeleton.

Theorem 2.1 (= Theorem F). *Let* $f_1 : \Gamma_1 \to Q$ *and* $f_2 : \Gamma_2 \to Q$ *be surjective group homomorphisms. Suppose that* Γ_1 *and* Γ_2 *are finitely presented, that* Q *is of type* F_3*, and that at least one of* $\ker(f_1)$ *and*

$\ker(f_2)$ *is finitely generated. Then the fibre-product of f_1 and f_2,*

$$P = \{(g, h) \mid f_1(g) = f_2(h)\} \subset \Gamma_1 \times \Gamma_2,$$

is finitely presented.

Proof. Without loss of generality, we may assume that $N := \ker(f_1)$ is finitely generated.

As in [2, §1.4], we start with a finite presentation $\langle X|R \rangle$ for Q and from this construct a finite presentation

$$\mathcal{P}_1 \equiv \langle A, X|S_1(A, X), S_2(A, X), S_3(A) \rangle$$

for Γ_1 such that

(1) A generates N;
(2) S_1 consists of relators $x^\varepsilon a x^{-\varepsilon} V_{x,a,\varepsilon}(A)$, one for each $x \in X$, $a \in A$ and $\varepsilon = \pm 1$, where $V_{x,a,\varepsilon}(A)$ is a word in the letters $A^{\pm 1}$;
(3) S_2 consists of relators $r(X)U_r(A)$, one for each $r \in R$, where $U_r(A)$ is a word in the letters $A^{\pm 1}$;
(4) S_3 is a finite set of words in the letters $A^{\pm 1}$.

We do not assume that $\ker(f_2)$ is finitely generated. Nevertheless, we can also choose a finite presentation for Γ_2 of the form $\mathcal{P}_2 \equiv \langle B, X|T_2(B, X), T_3(B, X) \rangle$, in which:

(1) $B \subset \ker(f_2)$;
(2) T_2 consists of a word $r(X)W_r(B^*)$ for each $r \in R$, where B^* denotes the set of formal conjugates $b^{w(X)}$ of the letters $B^{\pm 1}$ by words in $X^{\pm 1}$ and $W_r(B^*)$ is a word in these conjugates;
(3) T_3 is a finite set of words in the symbols B^*.

Now since $Q = \langle X|R \rangle$ is of type F_3, there is a finite set of $\mathbb{Z}Q$-module generators σ for the second homotopy group of this presentation, each of which can be expressed as an identity

$$\sigma := \prod_{j=1}^{m} w_j^{-1} r_j^{\varepsilon_j} w_j =_{F(X)} 1.$$

Following [2, §1.5] we translate σ into a relation $z_\sigma(A)$ among the generators A of N as follows: first replace each r_j by the corresponding relator $r_j(X).U_{r_j}(A)$ from S_2 above to get a word

$$\zeta_\sigma = \prod_{j=1}^{m} w_j(X)^{-1} \left(r_j(X).U_{r_j}(A) \right)^{\varepsilon_j} w_j(X);$$

then apply a sequence of relations (from S_1 above) of the form

$$x^\varepsilon a x^{-\varepsilon} = V_{x,a,\varepsilon}(A)^{-1}$$

to cancel all the x-letters from ζ_σ to leave a word z_σ involving only letters from $A^{\pm 1}$.

Let $Z = Z(A)$ be the finite set of words $z_\sigma(A)$ in the letters $A^{\pm 1}$ arising from our fixed finite set of π_2-generators σ. The crucial claim

is that the fibre product P of f_1 and f_2 is isomorphic to the quotient of the free group F on $A \cup B \cup \mathcal{X}$ modulo the following finite set of defining relators. (Here we use functional notation for words: given a set of words $\Sigma(Y)$ in symbols $Y^{\pm 1}$ and a set of symbols $\underline{y}$ in 1-1 correspondence with Y, we write $\Sigma(\underline{y})$ for the set of words obtained from $\Sigma(Y)$ by replacing each $y \in Y$ with the corresponding symbol from $\underline{y}$.)

  I) $S_1(A, \mathcal{X}) \cup S_3(A) \cup Z(A) \cup T_3(B, \mathcal{X})$;
 II) $\{r(X).U_r(A).W_r(B^*) \mid r \in R\}$;
III) $\{[a, b] \mid a \in A, b \in B\}$;
IV) $\{[a, r(\mathcal{X})U_r(A)] \mid a \in A, r \in R\}$.

We must argue that this really is a presentation of P. We map F to $P \subset \Gamma_1 \times \Gamma_2$ by a homomorphism θ defined as follows:

 (1) $\theta(a) = (a, 1)$ for $a \in A$;
 (2) $\theta(b) = (1, b)$ for $b \in B$;
 (3) $\theta(x) = (x, x)$ for $x \in \mathcal{X}$.

Let G be the quotient of F by the given relations. Since θ maps each of these relations to $(1, 1)$, it induces a homomorphism $\overline{\theta}$ from G to P. It is easy to see that θ is surjective, so it only remains to prove that $\ker(\overline{\theta}) = \{1\}$.

Suppose $W(A, B, X) \in F$ belongs to $\ker(\theta)$. Killing the generators A in G gives a presentation for Γ_2, and the associated map $G \to \Gamma_2$ factors through $\overline{\theta}$, so $W(1, B, X) = 1$ in $G/\langle\!\langle A \rangle\!\rangle$. It follows that having modified W by applying a finite sequence of the relations of G, we can assume that it is a finite product of conjugates of elements of $A^{\pm 1}$.

Now the relators $S_1(A, \mathcal{X})$ and (2) combine to show that each element of A commutes in G with each element of B^*. Thus if $a^{u(B,X)}$ is a conjugate of $a \in A$ by a word in $(B \cup X)^{\pm 1}$, we may apply the relators to replace it by $a^{u(1,X)}$, a conjugate of a by a word in $X^{\pm 1}$. But then the relators S_1 may be applied once more to replace $a^{u(1,X)}$ by a word in $A^{\pm 1}$.

At this stage we have succeeded in using the given defining relators of G to replace the initial word $W(A, B, X)$ by a word $W_0(A)$. Now $W_0(A) = 1$ in N, and [2, Theorem 1.2] tells us that the equality $W_0(A) = 1$ in N is a consequence of the defining relators $S_1(A, X)$, $S_3(A)$, $Z(A)$ and (2). Thus $W_0(A) = 1$ in G. Hence $\ker(\overline{\theta}) = 1$ and θ is an isomorphism from the finitely presented group G onto the fibre product of f_1 and f_2, as required.   □

2.1. An effective version of the Asymmetric 1-2-3 Theorem.
Given a finite presentation $\mathcal{Q} \equiv \langle X \mid R \rangle$ for a group Q, one can define the second homotopy group of $\pi_2 \mathcal{Q}$ to be π_2 of the standard 2-complex K of the presentation, regarded as a module over $\mathbb{Z}Q$ via the identification $Q = \pi_1 K$. But in the present context it is better

to regard elements of $\pi_2 \mathcal{Q}$ as equivalence classes of identity sequences $[(w_1, r_1), \ldots, (w_m, r_m)]$, where the w_i are elements of the free group $F(X)$, the $r_i \in R^{\pm 1}$, and where $\prod_{i=1}^{m} w_i^{-1} r_i w_i$ is equal to the empty word in $F(X)$; equivalence is defined by Peiffer moves, and the action of Q is induced by the obvious conjugation action of $F(X)$; see [31].

Theorem 2.2. *There exists a Turing machine that, given the following data describing group homomorphisms $f_i : \Gamma_i \to Q$ ($i = 1, 2$) will output a finite presentation of the fibre product of these maps provided that both the f_i are surjective and at least one of the kernels $\ker(f_i)$ is finitely generated. (If either of these conditions fails, the machine will not halt.)*

   Input Data:

    (1) *A finite presentation $\mathcal{Q} \equiv \langle X \mid R \rangle$ for Q.*
    (2) *A finite presentation $\langle \underline{a}^{(i)} \mid \underline{r}^{(i)} \rangle$ for Γ_i ($i = 1, 2$).*
    (3) *$\forall a \in \underline{a}_i$, a word $\overline{a} \in F(X)$ such that $\overline{a} = f_i(a)$ in Q.*
    (4) *A finite set of identity sequences that generate $\pi_2 \mathcal{Q}$ as a $\mathbb{Z}Q$-module.*

Proof. The proof of Theorem 2.1 above describes a simple, explicit process for constructing a finite presentation of P from presentations $\mathcal{P}_1, \mathcal{P}_2$ of a special form and identities σ (in the notation of the proof of Theorem 2.1). We shall now describe an effective process that, given the input data, will do the following in order:

    (i) verify that the f_i are onto, then proceed to (ii) (but fail to halt if they are not onto);
    (ii) construct $\mathcal{P}_1$ if $\ker f_1$ is finitely generated, output it, then proceed to (iii);
    (iii) construct $\mathcal{P}_2$ and output it.

If the process reaches stage (iii), it must eventually halt.

Once we have this process in hand, we apply it simultaneously to the given input data and to the data with the indices $1, 2$ permuted; one of these processes will halt if the f_i are onto and one of the kernels is finitely generated. The output data is then translated into a presentation of P by writing the relations (I) to (IV) of the preceding proof. The Turing machine implementing this parallel process and subsequent translation is the one whose existence is asserted in the theorem.

It remains to explain how steps (i) to (iii) are achieved.

(i): Suppose $X = \{x_1, \ldots, x_l\}$. To verify that f_i is onto, the process enumerates the words w in the free monoid on $\underline{a}^{(i)}$ in order of increasing length and, proceeding diagonally through this enumeration and that of all products Π in the free group $F(X)$ of conjugates of the relations R of Q, the process searches for an equality $x_1 \overline{w} = \Pi$ in $F(X)$, where $\overline{w}$ is the word obtained from w by replacing each $a \in \underline{a}^{(i)}$ by $\overline{a}$. Once such an equality is found, the process is repeated with x_2 in place of

x_1. When an equality is found for x_2, the process proceeds for x_3 and so on until an equality has been found for each of $x_1, \ldots, x_l$, at which point process (i) halts.

(ii): This stage of the programme implements a countable number of sub-programmes in a diagonal manner. The m'th involves a fixed set A_m of cardinality m. The sub-programme itself implements a countable number of sub-programmes, drawing words $V_{x,a,\epsilon}$ and U_r from a length-increasing enumeration of the free monoid on $X^{\pm 1}$ and working with presentations $\mathcal{P}$ of the form given as $\mathcal{P}_1$ in the proof of Theorem 2.1 (condition (1) concerning A_m being ignored). Let $\Gamma(\mathcal{P})$ be the group defined by $\mathcal{P}$. The surjection $F(X \cup A_m) \to Q$ defined by $\pi(x) = x$ for $x \in X$ and $\pi(a) = 1$ for $a \in A_m$ induces a surjection $\pi : \Gamma(\mathcal{P}) \to Q$.

By enumerating all Tietze transformations, the sub-programme searches for an isomorphism $q : \Gamma_1 \to \Gamma(\mathcal{P})$ such that $\pi \circ q = f_1$ (see remark 2.3). When $\mathcal{P}$ is found, the process halts and outputs $\mathcal{P}$.

(iii): This is identical to stage two except that one considers presentations with the form of $\mathcal{P}_2$ instead of those with the form of $\mathcal{P}_1$. $\square$

Remark 2.3. In the above proof we made use of an instance of the following very general observation: if one is given an arbitrary finite presentation $\mathcal{P}$ of a group Γ and one knows that Γ has a "special" presentation drawn from some recursively enumerable class $\{C_1, C_2, \ldots\}$, one can find a special presentation of Γ by proceeding as follows: enumerate the finite presentations P_n obtained from $\mathcal{P}$ by finite sequences of Tietze moves and proceed searching the finite diagonals through the rectangular array (P_n, C_m) looking for a coincidence.

3. Subdirect products

Throughout this section we consider subdirect products of arbitrary finitely presentable groups. In later sections we restrict attention to the case where the direct factors are limit groups.

Given a direct product $D := G_1 \times \cdots \times G_n$, we shall consistently write p_i and p_{ij} for the projection homomorphisms $p_i : D \to G_i$ and $p_{ij} : D \to G_i \times G_j$ $(i, j = 1, \ldots, n)$

Theorem 3.1 (= Theorem E). *Let $S \subset G_1 \times \cdots \times G_n$ be a subgroup of a direct product of finitely presented groups. If for all $i, j \in \{1, \ldots, n\}, i \neq j$, the projection $p_{ij}(S) \subset G_i \times G_j$ has finite index, then S is finitely presentable.*

We will deduce this theorem from the Asymmetric 1-2-3 Theorem by combining some well-known facts about virtually nilpotent groups with the following proposition, which generalises similar results in [10] and [9].

Proposition 3.2. *Let $G_1, \ldots, G_n$ be groups and let $S \subset G_1 \times \cdots \times G_n$ be a subgroup. If $p_{ij}(S) \subset G_i \times G_j$ is of finite index for all $i, j \in \{1, \ldots, n\}, i \neq j$, then*

(1) there exist finite-index subgroups $G_i^0 \subset G_i$ such that $\gamma_{n-1}(G_i^0) \subset S$.

If, in addition, the groups G_i are all finitely generated, then

(2) $L_i := S \cap G_i$ is finitely generated as a normal subgroup of S,
(3) $N_i := S \cap \ker(p_i)$ is finitely generated, and
(4) S is itself finitely generated.

Proof. The conditions imply that $p_i(S)$ is a finite index subgroup of G_i, and by passing to subgroups of finite index we may assume without loss that S is subdirect.

Let

$$G_1^0 = \{g \in G_1 \mid \forall j \neq 1 \, \exists (g, *, \ldots, *, 1, * \ldots) \in N_j\} = \bigcap_{j=2}^{n} (p_{1j}(S) \cap G_1)$$

and define G_i^0 similarly. As $p_{ij}(S) \subset G_i \times G_j$ is of finite index, G_i^0 has finite index in G_i for $i = 1, \ldots, n$.

For notational convenience we focus on $i = 1$ and explain why $\gamma_{n-1}(G_1^0) \subset S$. The key point to observe is that for all $x_1, \ldots, x_{n-1} \in G_1^0$ the commutator $([x_1, x_2, \ldots, x_{n-1}], 1, \ldots, 1)$ can be expressed as the commutator of elements from the subgroups $N_j \subset S$; explicitly it is

$$[\, (x_1, 1, *, \ldots, *), \; (x_2, *, 1, *, \ldots, *), \ldots, (x_{n-1}, *, \ldots, *, 1) \,].$$

This proves the first assertion.

For (2), note that since S is subdirect, $S \cap G_i$ is normal in G_i and the normal closure in G_i of any set $T \subset S \cap G_i$ is the same as its normal closure in S. Since G_i is finitely generated, $G_i/(S \cap G_i)$ is a finitely generated virtually nilpotent group; hence it is finitely presented and $S \cap G_i$ is the normal closure in G_i (hence S) of a finite subset.

Towards proving (3), note that the image of $N_1 = S \cap \ker(p_1)$ in G_i under the projection p_i has finite index for $2 \leq i \leq n$, since $p_{1i}(S)$ has finite index in $G_1 \times G_i$ and N_1 is the kernel of the restriction to S of $p_1 = p_1 \circ p_{1i}$. In particular $p_i(N_1)$ is finitely generated.

Note also that $L_i = S \cap G_i$ is the normal closure of a finite subset of $p_i(N_1)$ by (2).

Now let $L := L_2 \times \cdots \times L_n$. Then N_1/L is a subgroup of the finitely generated virtually nilpotent group

$$\frac{G_2 \times \cdots \times G_n}{L} \cong \frac{G_2}{L_2} \times \cdots \times \frac{G_n}{L_n},$$

and hence is also finitely generated (and virtually nilpotent).

Putting all these facts together, we see that we can choose a finite subset X of N_1 such that:

- $p_i(X)$ generates $p_i(N_1)$ for each $i = 2, \ldots, n$;
- $X \cap L_i$ generates L_i as a normal subgroup of $p_i(N_1)$, for each $i = 2, \ldots, n$;
- $\{xL : \ x \in X\}$ generates N_1/L.

These three properties ensure that X generates N_1, and the proof of (3) is complete.

We can express S as an extension of N_1 by G_1 which are both finitely generated (using (3)), and (4) follows immediately. $\qquad\square$

Remark 3.3. We shall use this lemma in tandem with the fact that finitely generated virtually nilpotent groups are F_∞, i.e. have classifying spaces with finitely many cells in each dimension. Indeed this is true of virtually polycyclic groups P, because such a group has a torsion-free subgroup of finite index that is poly-$\mathbb{Z}$, and hence is the fundamental group of a closed aspherical manifold. If B has type F_∞ (e.g. a finite group) and A has type F_∞ (e.g. the fundamental group of an aspherical manifold), then any extension of A by B is also of type F_∞ (see [18] Theorem 7.1.10).

3.1. **Proof of Theorem E.** The hypothesis on $p_{ij}(S)$ implies that the image of S in each factor G_i is of finite index. Replacing the G_i and S with subgroups of finite index does not alter their finiteness properties. Thus we may assume that S is a subdirect product. Let $L_i = G_i \cap S$ and note that L_i is normal in both S and G_i. Proposition 3.2 tells us that $Q_i := G_i/L_i$ is virtually nilpotent; in particular it is of type F_3 (see remark 3.3).

Assuming that S is a subdirect product, we proceed by induction on n. The base case, $n = 2$, is trivial.

Let $q : G_1 \times \cdots \times G_n \to G_1 \times \cdots \times G_{n-1}$ be the projection with kernel G_n and let $T = q(S)$. By the inductive hypothesis, T is finitely presented. We may regard S as a subdirect product of $T \times G_n$. Equivalently, writing $K = T \cap S$ and noting that

$$\frac{T}{K} \cong \frac{S}{K \times L_n} \cong \frac{G_n}{L_n} = Q_n,$$

we see that S is the fibre-product associated to the short exact sequences $1 \to K \to T \to Q_n \to 1$ and $1 \to L_n \to G_n \to Q_n \to 1$. Thus, by the Asymmetric 1-2-3 Theorem, our induction is complete because according to Proposition 3.2, K is finitely generated. $\qquad\square$

3.2. **The effective version.** We next prove an effective version of Theorem E, which will play a key part in proving that the class of finitely presentable residually free groups is recursively enumerable.

Theorem 3.4. *There exists a Turing machine that, given a finite collection $G_1, \ldots, G_n$ of finitely presentable groups (each given by an explicit finite presentation) and a finite subset $Y \subset G_1 \times \cdots \times G_n$ (given*

as a set of n-tuples of words in the generators of the G_i) such that each projection $p_{ij}(Y)$ generates a finite-index subgroup of $G_i \times G_j$ ($1 \le i < j \le n$), will output a finite presentation for $S := \langle Y \rangle$.

Proof. With the effective Asymmetric 1-2-3 Theorem (Theorem 2.2) in hand, we follow the proof of Theorem E. As in Theorem E we first replace each G_i by the finite-index subgroup $p_i(S)$ to get to a situation where S is subdirect. Here we use the Todd-Coxeter and Reidemeister-Schreier processes to replace the given presentations of the G_i by presentations of the appropriate finite-index subgroups. By using Tietze transformations we may take $p_i(Y)$ to be the generators of this presentation. Thus we express the revised G_i as quotients of the free group on Y.

We argue by induction on n. The initial cases $n = 1, 2$ are easily handled by the Todd-Coxeter and Reidemeister-Schreier processes, since then S has finite index in the direct product. So we may assume that $n \ge 3$.

By Theorem 2.2 it is sufficient to find finite presentations for

(1) $T = q(S)$, where q is the natural projection from $G_1 \times \cdots \times G_n$ to $G_1 \times \cdots \times G_{n-1}$,
(2) G_n, and
(3) $Q = G_n/(G_n \cap S)$,

together with

(4) explicit epimorphisms $T \to Q$ and $G_n \to Q$, and
(5) a finite set of generators for π_2 of the presentation for Q, as a $\mathbb{Z}Q$-module.

A finite presentation for G_n is part of the input.

We may assume inductively that we have found a finite presentation for T, with generators $q(Y)$. We write this presentation as $\langle Y \mid r_1(Y), \ldots, r_m(Y) \rangle$.

To obtain a finite presentation for Q, we proceed as follows. The words $r_j(p_n(Y))$ normally generate $G_n \cap S$. Adding these words as relations to the existing presentation of G_n gives a finite presentation of Q, together with the natural quotient map $G_n \to Q$.

The epimorphism $T \to Q$ is induced by the identity map on Y.

We would now be done if we could compute a finite set of π_2-generators for our chosen finite presentation $\mathcal{P}$ of the virtually nilpotent group Q. But it is more convenient to proceed in a slightly different manner, modifying $\mathcal{P}$.

First, we search among finite-index normal subgroups Q' of Q for an isomorphism $Q' \to P$, for some group P given by a poly-$\mathbb{Z}$ presentation $\mathcal{P}'$. The latter presentation defines an explicit construction for a

finite $K(P, 1)$-complex X, and in particular a finite set of generators of $\pi_2(X^{(2)})$ (the attaching maps of the 3-cells).

We next replace our initial presentation $\mathcal{P}$ for Q by a new presentation $\mathcal{Q}$ that contains $\mathcal{P}'$ as a sub-presentation. Indeed, we know that such presentations exist, so we can find one, together with an explicit isomorphism that extends the given isomorphism $P \to Q'$, by a naïve search procedure (see Remark 2.3).

Let Y denote the 2-dimensional complex model of the presentation $\mathcal{Q}$, $\widehat{Y}$ the regular cover of Y corresponding to the normal subgroup $P = Q'$, and Z the preimage of $X^{(2)} \subset Y$ in $\widehat{Y}$. Then Z consists of one copy of $X^{(2)}$ at each vertex of $\widehat{Y}$; these are indexed by the elements of the finite quotient group $H = Q/Q'$.

We then have an exact homotopy sequence

$$\cdots \to \mathbb{Z}Q \otimes_{\mathbb{Z}Q'} \pi_2(X^{(2)}) \to \pi_2(\widehat{Y}) \to \pi_2(\widehat{Y}, Z) \to 0$$

(since the map $P \to Q$ is injective by hypothesis), together with a finite set B of generators for $\pi_2(X^{(2)})$ as a $\mathbb{Z}Q'$-module. But $\pi_2(\widehat{Y}, Z) \cong H_2(\widehat{Y}/Z)$, since the quotient complex $\widehat{Y}/Z$ is simply connected. Hence $\pi_2(Y) = \pi_2(\widehat{Y})$ is generated as a $\mathbb{Z}Q$-module by B together with any finite set C that maps onto a generating set for the finitely generated abelian group $H_2(\widehat{Y}/Z)$. Such a set C can be found by a naïve search over finite sets of identity sequences over $\mathcal{Q}$. $\square$

4. Novel Examples

¿From [9] (or [10] in the case of surface groups) we know that a finitely presented full subdirect product S of n limit groups Γ_i must virtually contain the term γ_{n-1} of the lower central series of the product. So the quotient groups $\Gamma_i/(S \cap \Gamma_i)$ are virtually nilpotent of class at most $n - 2$. In particular for $n = 3$ the quotients $\Gamma_i/(S \cap \Gamma_i)$ are virtually abelian.

A question left unresolved in [10] is whether a finitely presented subdirect product S of n free groups Φ_i can have $\Phi_i/(S \cap \Phi_i)$ nilpotent strictly of class 2 or more (necessarily $n \geq 4$ for this to happen). Theorem 4.2 below settles this question and shows that the bounds on the nilpotency class given in [9] and [10] are optimal.

4.1. **The groups** $S(E, c)$. Let $F = \langle a, b \rangle$ be the free group of rank 2, and let $\Phi = F^{\mathbb{Z}}$ denote the unrestricted direct product of a countably infinite collection of copies of F, thought of as the set of functions $f : \mathbb{Z} \to F$ endowed with pointwise multiplication.

Let $\Gamma = \langle w, x, y, z \rangle$ be a free group of rank 4, and define a homomorphism $\phi : \Gamma \to \Phi$ by $\phi(w)(n) = a$, $\phi(x)(n) = b$, $\phi(y)(n) = a^n$, $\phi(z)(n) = b^n$ for all $n \in \mathbb{Z}$.

Given a finite subset $E \subset \mathbb{Z}$, we may regard the direct product of $|E|$ copies of F as the set F^E of functions $E \to F$. We then obtain a projection $p_E : \Phi \to F^E$ by restriction: $p_E(f) = f|_E : E \to F$.

Notice that when $E = \{n\}$ is a singleton $p_E \circ \phi$ is surjective. It will be convenient to write Φ_n for $F^{\{n\}}$, p_n for the projection $p_{\{n\}} : \Phi \to \Phi_n$, and a_n, b_n for the copy of a, b respectively in Φ_n. The surjectivity of $p_n \circ \phi$ means that, for any finite subset $E \subset \mathbb{Z}$, the image of $p_E \circ \phi$ is a finitely generated subdirect product of the free groups Φ_n ($n \in E$).

This subdirect product is not in general finitely presented.

Now let c be a positive integer. We may choose a finite set $R = R(a, b)$ of normal generators for the c'th term $\gamma_c(F)$ of the lower central series of F. We then define $S(E, c)$ to be the subgroup of F^E that is generated by $(p_E \circ \phi)(\Gamma)$ together with the sets $R(a_n, b_n) \subset \Phi_n$ for each $n \in E$.

As a concrete example we note that $S(\{1, 2, 3, 4\}, 3)$ is the subgroup of $\Phi_1 \times \Phi_2 \times \Phi_3 \times \Phi_4$ generated by the following 12 elements: the four images of the generators of Γ

$$(a_1, a_2, a_3, a_4) \ , \ (b_1, b_2, b_3, b_4)$$
$$(a_1, a_2^2, a_3^3, a_4^4) \ , \ (b_1, b_2^2, b_3^3, b_4^4)$$

together with the eight elements

$$([[a_1, b_1], a_1], 1, 1, 1) \ , \ ([[a_1, b_1], b_1], 1, 1, 1) \ , \ (1, [[a_2, b_2], a_2], 1, 1) \ , \ \dots$$
$$\dots \ , \ (1, 1, 1, [[a_4, b_4], a_4]) \ , \ (1, 1, 1, [[a_4, b_4], b_4])$$

which are normal generators for the subgroups $\gamma_3(\Phi_i)$ for $1 \le i \le 4$.

Proposition 4.1. *The groups $S(E, c)$ have the following properties.*

(1) $S(E, c)$ *contains* $\gamma_c(F^E)$.
(2) $S(E, c)$ *is finitely presentable.*
(3) *If* $E' = E + t = \{e + t; \ e \in E\}$ *is a translate of E in $\mathbb{Z}$, then*

$$S(-E, c) \cong S(E, c) \cong S(E', c).$$

(4) *If* $E \subset E'$, *then the projection $F^{E'} \to F^E$ induces an epimorphism $S(E', c) \to S(E, c)$.*

Proof.

(1) Since $R(a_n, b_n) \subset S(E, c) \cap \Phi_n$ by construction, and since $p_n \circ \phi$ is surjective for all $n \in E$, it follows that $S(E, c) \supset \gamma_c(\Phi_n)$ for each $n \in E$, and hence $S(E, c) \supset \gamma_c(F^E)$.

(2) Clearly $S(E, c)$ is finitely generated. For any 2-element subset $T = \{m, n\}$ of E, the image of the projection of $S(E, c)$ to $F^T = \Phi_m \times \Phi_n$ is precisely $S(T, c)$. Since $S(T, c)$ contains the elements $p_T(\phi(w)) = (a_m, a_n)$, $p_T(\phi(x)) = (b_m, b_n)$, $p_T(\phi(yw^{-m})) = (1, a_n^{n-m})$ and $p_T(\phi(zx^{-m})) = (1, b_n^{n-m})$, together with $\gamma_c(\Phi_m \times \Phi_n)$, we see that the quotient of each of the direct factors $\Phi_m \cong F \cong \Phi_n$ by its intersection with $S(T, c)$ is a nilpotent group of class at most c, generated by

two elements of finite order, and hence is finite. Thus $S(T, c)$ has finite index in F^T. In other words, the projection of the subdirect product $S(E, c) < F^E$ to each product of two factors F^T has finite index. Hence by Theorem E, $S(E, c)$ is finitely presentable.

(3) It is clear that $S(-E, c) \cong S(E, c)$ via the isomorphism $F^E \to F^{-E}$ defined by $a_n \mapsto a_{-n}$, $b_n \mapsto b_{-n}$.

To show that $S(E, c) \cong S(E', c)$, it is clearly enough to consider the case $t = 1$. The isomorphism $\theta : F^E \to F^{E'}$ defined by $a_n \mapsto a_{n+1}$, $b_n \mapsto b_{n+1}$ is induced by the shift automorphism $\overline{\theta} : \Phi \to \Phi$ defined by $\overline{\theta}(f)(k) := f(k - 1)$, in the sense that $p_{E'} \circ \overline{\theta} = \theta \circ p_E$.

Similarly, $\overline{\theta}$ commutes with the automorphism $\widehat{\theta}$ of Γ defined by $w \mapsto w$, $x \mapsto x$, $y \mapsto yw^{-1}$, $z \mapsto zx^{-1}$, in the sense that $\overline{\theta} \circ \phi = \phi \circ \widehat{\theta}$. It follows immediately from the definitions that θ maps $S(E, c)$ onto $S(E', c)$

(4) This is immediate from the definitions. $\square$

We can now state and prove the main result of this section. We thank Mike Vaughan-Lee for several helpful suggestions concerning this proof.

Theorem 4.2 (= Theorem H). *For any positive integer c, and any finite subset $E \subset \mathbb{Z}$ of cardinality at least $c + 1$, the group $S(E, c)$ is a finitely presentable subdirect product of the non-abelian free groups Φ_n $(n \in E)$ and $S(E, c) \cap \Phi_n = \gamma_c(\Phi_n)$ for each $n \in E$.*

Proof. By construction, $S(E, c)$ is a subdirect product of the Φ_n for $n \in E$, and by Proposition 4.1 (2) it is finitely presentable. By Proposition 4.1 (1) we have

$$S(E, c) \cap \Phi_n \supset \gamma_c(\Phi_n)$$

for each $n \in E$, so it only remains to prove the reverse inclusion.

Let $A = \mathbb{Q}[[\alpha, \beta]]$ be the algebra of power series in two non-commuting variables α, β with rational coefficients, and for each n let $\eta_n : \Phi_n \to U(A)$ be the Magnus embedding of Φ_n into the group of units $U(A)$ of A, defined by $\eta_n(a_n) = 1 + \alpha$, $\eta_n(b_n) = 1 + \beta$. By Magnus' Theorem [25] (or [26, Chapter 5]), $\eta_n^{-1}(1 + J^c) = \gamma_c(\Phi_n)$. Here J is the ideal generated by the elements with 0 constant term and J^c is its c-th power.

Now define $\eta : \Gamma \to U(\mathbb{Q}[t] \otimes_{\mathbb{Q}} A)$ by $\eta(w) = 1 + \alpha$, $\eta(x) = 1 + \beta$, $\eta(y) = (1 + \alpha)^t$, $\eta(z) = (1 + \beta)^t$, where for example $(1 + \alpha)^t$ means the power series

$$(1 + \alpha)^t = \sum_{k=0}^{\infty} \binom{t}{k} \alpha^k = \sum_{k=0}^{\infty} \frac{t(t-1) \cdots (t - k + 1)}{k!} \alpha^k.$$

Note that $\eta_n \circ \phi_n = \psi_n \circ \eta$, where $\psi_n : \mathbb{Q}[t] \otimes_{\mathbb{Q}} A \to A$ is defined by $f(t) \otimes a \mapsto f(n)a$ and where $\phi_n = p_n \circ \phi$.

Note also that, for any $g \in \Gamma$, $\eta(g)$ has the form

$$\eta(g) = \sum_{W \in \Omega} p_W(t) \cdot W(\alpha, \beta),$$

where Ω is the free monoid on $\{\alpha, \beta\}$ and $p_W(t) \in \mathbb{Q}[t]$ has degree at most equal to the length of W. Hence, for each $n \in \mathbb{Z}$, we have

$$\eta_n(\phi_n(g)) = \psi_n(\eta(g)) = \sum_{W \in \Omega} p_W(n) \cdot W(\alpha, \beta).$$

Suppose now that $E \subset \mathbb{Z}$ is a finite set of integers of cardinality at least $c+1$, and that $g \in \Gamma$ such that $p_E(\phi(g)) \in S(E,c) \cap \Phi_n$ for some $n \in E$. Then, for each $m \in E \smallsetminus \{n\}$, we have

$$\psi_m(\eta(g)) = \eta_m(\phi_m(g)) = \eta_m(1) = 1.$$

It follows that, in the expression $\eta(g) = \sum_{W \in \Omega} p_W(t) \cdot W(\alpha, \beta)$ for $\eta(g)$, the c elements of $E \smallsetminus \{n\}$ are roots of all the polynomials $p_W(t)$. In particular, for words W of length less than c, the polynomials p_W are identically zero. Hence $\psi_m(\eta(g)) \in 1 + J^c$ for all $m \in \mathbb{Z}$, in particular for $m = n$. Hence $\phi_n(g) \in \eta_n^{-1}(1 + J^c) = \gamma_c(\Phi_n)$.

Thus

$$S(E,c) \cap \Phi_n \subset \gamma_c(\Phi_n),$$

completing the proof that

$$S(E,c) \cap \Phi_n = \gamma_c(\Phi_n).$$

$\square$

4.2. **Sample calculations.** We use the explicit form of the map η from the proof of Theorem 4.2 to make some calculations that illuminate the preceding proof.

Remark 4.3. Suppose that $U, V \in \Gamma$ and $\alpha \in J^k$, $\beta \in J^\ell$ are such that $\eta(U) = 1 + \alpha \bmod J^{k+1}$, $\eta(V) = 1 + \beta \bmod J^{\ell+1}$. Then $\eta(UV) - \eta(VU) = \alpha\beta - \beta\alpha \bmod J^{k+\ell+1}$, while $\eta(U^{-1}V^{-1}) = 1 \bmod J^2$, so

$$\eta([U,V]) - 1 = \eta(U^{-1}V^{-1})(\eta(UV) - \eta(VU)) = \alpha\beta - \beta\alpha \bmod J^{k+\ell+1}.$$

Example 4.4. For each integer k, we calculate that

$$\eta(zx^{-k}) = 1 + (t-k)\beta \bmod J^2.$$

Also

$$\eta(Y) = 1 + t\alpha \bmod J^2.$$

Repeatedly applying Remark 4.3, we see that

$$\eta([y, zx^{-1}, zx^{-2}, \ldots, zx^{-m}]) = 1 + t(t-1)\cdots(t-m)V_m(\alpha, \beta) \bmod J^{m+2},$$

where

$$V_m := \sum_{k=1}^{m} \binom{m}{k} \beta^k \alpha \beta^{m-k}$$

is a non-trivial $\mathbb{Z}$-linear combination of homogeneous monomials of degree $m + 1$.

Notice that the coefficient of $V_m(\alpha, \beta)$ is a polynomial of degree $m + 1$ in t with roots $0, 1, \dots, m$. In particular this gives an example of an element in $S(\{0, \dots, m + 1\}, m + 2) \cap \gamma_{m+1}(\Phi_{m+1})$ which is not in $\gamma_{m+2}(\Phi_{m+1})$.

Example 4.5. As another application of Remark 4.3, we see inductively that, for any basic commutator C of weight c in the generators of Γ,

$$\eta(C) \in \mathbb{Z}[t][[\alpha, \beta]] + J^{c+1},$$

and hence

$$\eta(\gamma_c(\Gamma)) \subset \mathbb{Z}[t][[\alpha, \beta]] + J^{c+1}.$$

On the other hand, if we put $U = [w, z][x, y] \in \gamma_2(\Gamma)$, then

$$\eta(U) = 1 + \binom{t}{2}(\alpha\beta^2 + \beta^2\alpha + \beta\alpha^2 + \alpha^2\beta - 2\alpha\beta\alpha - 2\beta\alpha\beta) \bmod J^4.$$

Thus $\phi(U)$ is an element of $\gamma_3(S(E, c))$ for any E, c. On the other hand, since $\binom{t}{2} \notin \mathbb{Z}[t]$, $\eta(U) \notin \eta(\gamma_3(\Gamma))$, so for sufficiently large E, c the element $\phi(U) \in \gamma_3(S(E, c))$ does not belong to $\phi(\gamma_3(\Gamma))$.

5. Characterizations

In this section we discuss the structure of finitely presentable residually free groups, and prove some results concerning their classification.

5.1. **Subdirect products and homological finiteness properties.**
We remind the reader of the shorthand we introduced in order to state Theorem B concisely: an embedding $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$ of a residually free group S as a full subdirect product of limit groups is said to be *neat* if Γ_0 is abelian, $S \cap \Gamma_0$ is of finite index in Γ_0, and Γ_i is non-abelian for $i = 1, \dots, n$.

Theorem 5.1 (=Theorem B). *Let S be a finitely generated residually free group. Then the following conditions are equivalent:*

    (1) *S is finitely presentable;*
    (2) *S is of type $\mathrm{FP}_2(\mathbb{Q})$;*
    (3) *$\dim H_2(S_0; \mathbb{Q}) < \infty$ for all subgroups $S_0 \subset S$ of finite index;*
    (4) *there exists a neat embedding $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$ such that the image of S under the projection to $\Gamma_i \times \Gamma_j$ has finite index for $1 \le i < j \le n$;*
    (5) *for every neat embedding $S \hookrightarrow \Gamma_0 \times \cdots \times \Gamma_n$, the image of S under the projection to $\Gamma_i \times \Gamma_j$ has finite index for $1 \le i < j \le n$.*

Proof. The implications (1) implies (2) implies (3) are clear. Theorem E shows that (4) implies (1).

In order to establish the remaining implications, we first argue that every finitely generated residually free group has a neat embedding.

The embedding theorem from [3] tells us that S embeds into the direct product of a finite collection of limit groups. Since finitely generated subgroups of limit groups are limit groups, we may assume that S is a subdirect product of finitely many limit groups. Moreover, by projecting away from any factor with which S has trivial intersection, we may assume that S is a full subdirect product of limit groups, say $S < \Gamma_0 \times \cdots \times \Gamma_n$. Moreover, if two or more of the factors Γ_i are abelian, we may regard their direct product as a single direct factor, so we may assume that Γ_0 is abelian (possibly trivial), and that Γ_i is non-abelian for $i > 0$. Finally, the intersection $S \cap \Gamma_0$ has finite index in some direct summand of Γ_0: by projecting away from a complement of such a direct summand, we may assume that $S \cap \Gamma_0$ has finite index in Γ_0. Thus we obtain a neat embedding of S. With this existence result in hand, it is clear that (5) implies (4). To complete the proof we shall argue that (3) implies (5).

Since the given embedding is neat, the image of the projection of S to $\Gamma_0 \times \Gamma_i$ has finite index for any $i > 0$, and the quotient $\overline{S}$ of S by $Z(S) = S \cap \Gamma_0$ is a full subdirect product of the non-abelian limit groups $\Gamma_1, \ldots, \Gamma_n$. Moreover, since $S \cap \Gamma_0$ is finitely generated, (3) implies that $H_2(\overline{S}_0; \mathbb{Q})$ is finite dimensional for all subgroups $\overline{S}_0 < \overline{S}$ of finite index in $\overline{S}$. It then follows from Theorem 4.2 of [9] that the image of the projection of S to $\Gamma_i \times \Gamma_j$ has finite index for any i, j with $0 < i < j \leq n$. Thus (3) implies (5). $\qquad \square$

It follows easily from Theorem 5.1 that any subdirect product of limit groups that contains a finitely presentable full subdirect product is again finitely presentable. More generally we prove:

Theorem 5.2 (= Theorem D). *Let $n \geq 2$ be an integer, let $S \subset D := \Gamma_1 \times \cdots \times \Gamma_k$ be a full subdirect product of limit groups, and let $T \subset D$ be a subgroup that contains S. If S is of type $\mathrm{FP}_n(\mathbb{Q})$ then so is T.*

Proof. We have $S < T < D = \Gamma_1 \times \cdots \times \Gamma_k$ where the Γ_i are limit groups and S is a full subdirect product of type $\mathrm{FP}_n(\mathbb{Q})$ with $n \geq 2$.

In particular, S is of type $\mathrm{FP}_2(\mathbb{Q})$, so by [9, Theorem 4.2] the quotient group S/L is virtually nilpotent, where $L = (S \cap \Gamma_1) \times \cdots \times (S \cap \Gamma_k)$.

By [9, Corollary 8.2] there is a finite index subgroup $S_0 < S$, and a subnormal chain $S_0 \lhd S_1 \lhd \cdots \lhd S_\ell = T$ such that each quotient S_{i+1}/S_i is either finite or infinite cyclic.

Since S is of type $\mathrm{FP}_n(\mathbb{Q})$, so is S_0, and by the obvious induction so are $S_1, \ldots, S_\ell = T$. $\qquad \square$

Note that the condition $n \geq 2$ in Theorem 5.2 is essential. For example, if $G = \langle x, y \mid r_1, r_2, \ldots \rangle$ is a 2-generator group that is not finitely presentable, then the subgroup T of $\langle x, y \rangle \times \langle x, y \rangle$ generated by $\{(x, x), (y, y), (1, r_1), (1, r_2), \ldots\}$ is a full subdirect product that is not

finitely generated, while the finitely generated subgroup S of T generated by $\{(x,x),(y,y),(1,r_1)\}$ is also a full subdirect product (provided $r_1 \neq 1$ in $\langle x,y \rangle$). This is another example of the notable divergence in behaviour between finitely presentable residually free groups and more general finitely generated residually free groups.

5.2. **The three factor case.** Thereom B tells us which full subdirect products of non-abelian limit groups are finitely presentable. In the case of two factors, the criterion is particularly simple: the subgroup must have finite index in the direct product. Our next result, which extends Theorem E of [10], shows that the criterion also takes a particularly simple form in the case of a full subdirect product of three non-abelian limit groups. Our results in Section 4 show that the situation is noticeably more subtle for subdirect products of four or more factors.

Theorem 5.3. *Let* $\Gamma_1, \Gamma_2, \Gamma_3$ *be non-abelian limit groups, and let* $S < \Gamma_1 \times \Gamma_2 \times \Gamma_3$ *be a full subdirect product. Then* S *is finitely presentable if and only if there are subgroups* $\Lambda_i < \Gamma_i$ *of finite index, an abelian group* Q, *and epimorphisms* $\phi_i : \Lambda_i \to Q$, *such that*

$$S \cap (\Lambda_1 \times \Lambda_2 \times \Lambda_3) = \ker(\phi),$$

where

$$\phi : \Lambda_1 \times \Lambda_2 \times \Lambda_3 \to Q, \ \phi(\lambda_1, \lambda_2, \lambda_3) := \phi_1(\lambda_1) + \phi_2(\lambda_2) + \phi_3(\lambda_3).$$

Proof. The criterion in the statement is clearly sufficient, by Theorem E, since each ϕ_i is an epimorphism. For example, given $\lambda_1 \in \Lambda_1$ and $\lambda_2 \in \Lambda_2$, there exists $\lambda_3 \in \Lambda_3$ such that $\phi_3(\lambda_3) = -\phi_1(\lambda_1) - \phi_2(\lambda_2)$. Thus $(\lambda_1, \lambda_2, \lambda_3) \in \ker(\phi)$ so the projection $p_{12} : \Gamma_1 \times \Gamma_2 \times \Gamma_3 \to \Gamma_1 \times \Gamma_2$ maps $\ker(\phi)$ onto the finite-index subgroup $\Lambda_1 \times \Lambda_2$ of $\Gamma_1 \times \Gamma_2$. Similar arguments apply to the projections p_{13} and p_{23}, so the finite-index subgroup $\ker(\phi)$ of S is finitely presentable, by Theorem E, and hence S is also finitely presentable.

Conversely, suppose that S is finitely presentable. By [9, Theorem 4.2] the image of each of the projections $p_{ij} : S \to \Gamma_i \times \Gamma_j$ ($1 \leq i < j \leq 3$) has finite index. The images of p_{12} and p_{13} intersect in a finite index subgroup $K_1 < \Gamma_1$. For each $a \in K_1$ there are elements $(a, 1, x_a), (a, y_a, 1) \in S$. So given $a, b \in K_1$, we have $([a,b], 1, 1) = [(a, 1, x_a), (b, y_b, 1)] \in [S, S]$. Thus $[K_1, K_1] < ([S, S] \cap \Gamma_1)$. Similarly there are finite-index subgroups $K_2 < \Gamma_2$ and $K_3 < \Gamma_3$ such that $[K_i, K_i] < ([S, S] \cap \Gamma_i)$ for $i = 2, 3$. Let A denote the abelian group

$$A = \frac{K_1 \times K_2 \times K_3}{S \cap (K_1 \times K_2 \times K_3)},$$

let $\phi : K_1 \times K_2 \times K_3 \to A$ be the canonical epimorphism, and let ϕ_i be the restriction of ϕ to K_i for $i = 1, 2, 3$. Since $p_{23}(S)$ has finite index in

$\Gamma_2 \times \Gamma_3$, the same is true of $p_{23}(S \cap (K_1 \times K_2 \times K_3))$ in $K_2 \times K_3$. Now let $\alpha = (x, y, z) \cdot (S \cap (K_1 \times K_2 \times K_3)) \in A$. For some positive integer N we have $(y^N, z^N) \in p_{23}(S \cap (K_1 \times K_2 \times K_3))$, so $(w, y^N, z^N) \in S$ for some $w \in K_1$. But then $\alpha^N = \phi_1(x^N w^{-1})$, so $\phi_1(K_1)$ has finite index in A. Similarly, $\phi_2(K_2)$ and $\phi_3(K_3)$ have finite index in A. Let Q be the finite-index subgroup $\phi_1(K_1) \cap \phi_2(K_2) \cap \phi_3(K_3)$ of A, and define $\Lambda_i = \phi_i^{-1}(Q)$ for $i = 1, 2, 3$. Then Λ_i has finite index in Γ_i, $S \cap (\Lambda_1 \times \Lambda_2 \times \Lambda_3)$ is the kernel of the restriction $\phi : \Lambda_1 \times \Lambda_2 \times \Lambda_3 \to Q$, and each $\phi_i : \Lambda_i \to Q$ is an epimorphism. $\qquad\square$

5.3. **Classification up to commensurability.** We construct a collection of examples of finitely presentable residually free groups which is complete up to commensurability.

Definition 5.4. Let $\mathcal{G} = \{\Gamma_1, \ldots, \Gamma_n\}$ be a finite collection of 2 or more limit groups, let $c \geq 2$ be an integer, and let $\underline{g} = \{(g_{k,1}, \ldots, g_{k,n}),\ 1 \leq k \leq m\}$ be a finite subset of $\Gamma := \Gamma_1 \times \cdots \times \Gamma_n$.

Define $T = T(\mathcal{G}, \underline{g}, c)$ to be the subgroup of Γ generated by $\underline{g}$ together with the c'th term $\gamma_c(\Gamma)$ of the lower central series of Γ.

Theorem 5.5. *Let $T(\mathcal{G}, \underline{g}, c)$ be defined as above.*

(1) *If, for all $1 \leq i < j \leq n$, the images in $H_1\Gamma_i \times H_1\Gamma_j$ of the ordered pairs $(g_{k,i}, g_{k,j})$ generate a subgroup of finite index, then the residually free group $T(\mathcal{G}, \underline{g}, c)$ is finitely presentable.*

(2) *Every finitely presentable residually free group is either a limit group or else is commensurable with one of the groups $T(\mathcal{G}, \underline{g}, c)$.*

Proof. To see that $T = T(\mathcal{G}, \underline{g}, c)$ is finitely presentable, it is sufficient in the light of Theorem E to note that the projection of T to $\Gamma_i \times \Gamma_j$ is virtually surjective for each $i < j$. This in turn follows from the observation that a subgroup of a finitely generated nilpotent group N has finite index whenever its image in H_1N has finite index.

Conversely, suppose that S is a finitely presentable residually free group. If S is not itself a limit group, then Theorem B tells us that S may be expressed as a full subdirect product of limit groups $\Delta_1, \ldots, \Delta_n$ such that the projection of S to $\Delta_i \times \Delta_j$ is virtually surjective for each $i < j$. By Theorem 4.2 of [9], each Δ_i contains a finite-index subgroup Γ_i such that $\gamma_{n-1}(\Gamma_i) \subset S$. Set $\mathcal{G} = \{\Gamma_1, \ldots, \Gamma_n\}$, and $c = n - 1$. We choose any finite set $\{(g_{1,1}, \ldots, g_{1,n}), \ldots, (g_{m,1}, \ldots, g_{m,n})\}$ in the direct product $D := \Gamma_1 \times \cdots \times \Gamma_n$ whose image in $D/\gamma_{n-1}(D)$ generates $(S \cap D) \cdot \gamma_{n-1}(D)/\gamma_{n-1}(D)$. Finally, take $\underline{g}$ to be the collection of coordinates $g_{k,i}$, and note that $T = T(\mathcal{G}, \underline{g}, n - 1) = S \cap D$ is a finite-index subgroup of S. $\qquad\square$

6. The Canonical Embedding Theorem

The purpose of this section is to prove Theorem A: we shall describe an effective construction for $\exists\mathrm{Env}(S)$, hence $\exists\mathrm{Env}_0(S)$, then establish the universal property of the latter. We shall see that the direct factors of $\exists\mathrm{Env}(S)$ are the maximal limit group quotients of S: the maximal free abelian quotient $H_1(S, \mathbb{Z})/(\text{torsion})$ is one of these, and the remaining (non-abelian) quotients form $\exists\mathrm{Env}_0(S)$. At the end of the section we shall discuss how $\exists\mathrm{Env}(S)$ is related to the Makanin-Razborov diagram for S.

Our first goal is to prove Theorem A(1).

Theorem 6.1. *There is an algorithm that, given a finite presentation of a residually free group S will construct an embedding*

$$S \hookrightarrow \exists\mathrm{Env}(S) = \Gamma_{\mathrm{ab}} \times \exists\mathrm{Env}_0(S)$$

where $\Gamma_{\mathrm{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$ and $\exists\mathrm{Env}_0(S) = \Gamma_1 \times \cdots \times \Gamma_n$ with each Γ_i $(i \geq 1)$ a non-abelian limit group. The intersection of S with the kernel of the projection $\rho : \exists\mathrm{Env}(S) \to \exists\mathrm{Env}_0(S)$ is the centre $Z(S)$ of S.

In outline, our proof of this theorem proceeds as follows. First we define a finite set of data — a *maximal centralizer system* — that encodes a canonical system of subgroups in S. Then, in Lemma 6.7, we prove that every finitely presented residually free group possesses such a system; the proof, which is not effective, relies on Proposition 3.2 and results from [9]. In Lemma 6.9 we establish the existence of a simple algorithm that, given a maximal centralizer system, will construct $S \hookrightarrow \exists\mathrm{Env}(S)$. Finally, in Subsection 6.3, we describe an algorithm that, given a finite presentation of a residually free group, will construct a maximal centralizer system for that group (termination of the algorithm is guaranteed by Lemma 6.7).

The description of $Z(S)$ given in Theorem 6.1 is covered by the following lemma.

Lemma 6.2. *Let S be a residually free group and let $Z(S)$ be its centre.*

(1) *The restriction of $S \to H_1(S, \mathbb{Z})/(\text{torsion})$ to $Z(S)$ is injective.*
(2) *If Γ is a non-abelian limit group and $\psi : S \to \Gamma$ has non-abelian image, then $\psi(Z(S)) = \{1\}$.*

Proof. Let $\gamma \in Z(S)$. Since S is residually free, there is an epimorphism ψ from S to a free group such that $\psi(\gamma) \neq 1$. But the only free group with a non-trivial centre is $\mathbb{Z}$, so $\psi([S, S]) = 1$ and hence $\gamma \notin [S, S]$. This observation, together with the fact that residually free groups are torsion-free, proves (1).

Item (2) follows easily from the fact that limit groups are commutative-transitive. □

6.1. **Centralizer systems.** Before pursuing the strategy of proof outlined above, we present an auxiliary result that motivates the definition of a maximal centralizer system. Recall that a set of subgroups of a group H is said to be *characteristic* if any automorphism of H permutes the subgroups in the set.

Proposition 6.3. *Let $D = \Gamma_1 \times \cdots \times \Gamma_n$ be a direct product of non-abelian limit groups, let $S \subset D$ be a full subdirect product, let $L_i = S \cap \Gamma_i$ and let*

$$M_i = S \cap (\Gamma_1 \times \cdots \times \Gamma_{i-1} \times 1 \times \Gamma_{i+1} \times \cdots \times \Gamma_n).$$

The sets of subgroups $\{L_1, \ldots, L_n\}$ and $\{M_1, \ldots, M_n\}$ are characteristic in S.

Proof. If Γ is a non-abelian limit group, and if γ_1 and γ_2 are two non-commuting elements of Γ, then the centralizer $C_\Gamma(\gamma_1, \gamma_2)$ of the pair is trivial, by commutative-transitivity.

The collection of centralizers of non-commuting pairs of elements of S has a finite set of maximal elements, namely the centralizers of pairs x_i and y_i which are non-commuting pairs in L_i. These maximal elements are exactly the M_i, which therefore form a characteristic set. Moreover the L_i are the centralizers of the M_i and hence the set of these is also characteristic (cf. [11]). $\square$

Remark 6.4. Applying the proposition with $S = D$ one sees that if $D = \Gamma_1 \times \cdots \times \Gamma_n$ is the direct product of non-abelian limit groups, then the set of subgroups Γ_i is characteristic. In particular, the decomposition of D as a direct product of limit groups is unique.

The example $D = \mathbb{Z} \times F_2$ shows that this uniqueness fails if abelian factors are allowed.

Definition 6.5. Let S be a finitely presented, non-abelian residually free group. A finite list $(Y_i; Z_i) = (Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ of finite subsets of S will be called a *maximal centralizer structure (MCS)* for S if it has the following properties.

MCS(1) Each Y_i contains at least two elements x_i and y_i which do not commute.

MCS(2) Each Z_i contains all of the Y_j with $j \neq i$.

MCS(3) For each i, the elements of Z_i commute with the elements of Y_i. (Hence the elements in Y_i commute with those in Y_j for all $i \neq j$.)

MCS(4) Each Z_i generates a normal subgroup of S.

MCS(5) For each i, the quotient group $S/\langle Z_i \rangle$ admits a splitting (as an amalgamated free product or HNN extension) either over the trivial subgroup or over a non-normal, infinite cyclic subgroup.

MCS(6) There is a subgroup S_0 of finite index in S such that each $Y_i \subset S_0$ and $S_0/\langle\langle Y_1, \ldots, Y_n \rangle\rangle$ is nilpotent of class at most $n - 2$.

For the case $n = 1$ we require that $\langle Z_1 \rangle = Z(S)$ and that Y_1 be the given generating set for S.

Remark 6.6. One of the basic properties of non-abelian limit groups is that that they split as in MCS(5). Conversely, we shall see in Lemma 6.9 that, in the presence of the other conditions, MCS(5) implies the following condition:

MCS(5′) For each i, the quotient $S/\langle Z_i \rangle$ is a non-abelian limit group.

Lemma 6.7. *Every finitely presented non-abelian residually free group possesses a maximal centralizer structure.*

Proof. Let S be a finitely presented non-abelian residually free group, and define $H = S/Z(S)$. We shall first construct an MCS for H.

As in the proof of Theorem B, H can be embedded as a full subdirect product in some $D = \Gamma_1 \times \cdots \times \Gamma_n$ where the Γ_i are non-abelian limit groups. Let $p_i : D \to \Gamma_i$ denote the projection.

If $n = 1$, then H itself is a non-abelian limit group. In this case, we follow the directions in the definition of MCS: Y_1 is the given set of generators for H, $Z_1 = \{1\}$, and $H_0 = H$. Then MCS(1-4) and MCS(6) are trivially satisfied, as is MCS(5)', hence MCS(5).

¿From now on we assume that $n > 1$. Then $\Gamma_i/(H \cap \Gamma_i)$ is virtually nilpotent by [9], so $(H \cap \Gamma_i)$ is finitely generated as a normal subgroup of Γ_i. Choose a finite set Y_i of normal generators for $H \cap \Gamma_i$ containing at least two elements that do not commute.

Let M_i denote the centralizer of Y_i in H (this is consistent with the notation in Proposition 6.3). Note that $M_i = H \cap \ker(p_i)$, which by Proposition 3.2(3) is a finitely generated subgroup of H. Note that $\Gamma_i \cong H/M_i$. Choose Z_i to be a finite generating set for M_i containing Y_j for all $j \neq i$.

This provides an MCS $(Y_i; Z_i)$ for H: each of the properties MCS(1-4) is explicit in the construction, as are MCS(5)' and MCS(6).

It remains to construct an MCS for S from the one just constructed for $H = S/Z(S)$. We know from Lemma 6.2 that $Z(S)$ is a finitely generated free abelian group. To obtain an MCS $(\hat{Y}_i; \hat{Z}_1)$ for S, we lift each $Y_i \subset H$ to a finite subset $\hat{Y}_i$ of S, and take a finite subset $\hat{Z}_i$ in the preimage of each Z_i containing (i) $\hat{Y}_j$ for all $j \neq i$, and (ii) a finite generating set for $Z(S)$.

To see that $(\hat{Y}_i; \hat{Z}_1)$ satisfies MCS(1), note that $Z(S) \cap [S, S] = 1$. Modulo this observation, it is clear that $(\hat{Y}_i; \hat{Z}_1)$ inherits the properties MCS(1-6) from $(Y_i; Z_i)$. □

6.2. **Two useful lemmata.** The following are the two principal lemmata used in the proof of Theorem A. We first prove a technical lemma about splittings which allows us to detect when a given quotient of S is a non-abelian limit group rather than a direct product.

Lemma 6.8. *Let Γ be a torsion–free group, H a group, and $G \hookrightarrow \Gamma \times H$ a subdirect product such that $G \cap \Gamma$ contains a free group of rank 2. Let N be a normal subgroup of G with $N < K = G \cap H$. If G/N admits a cyclic splitting, and $N \neq K$, then K/N is cyclic and the splitting is over K/N.*

Proof. The quotient $G/N \hookrightarrow \Gamma \times H/N$ is a subdirect product.

The cyclic splitting gives a G/N action on a tree T which is edge-transitive and has cyclic edge-stabilisers. A free subgroup $F = \langle x, y \rangle$ of $G \cap \Gamma$ either fixes a vertex v or contains an element w acting hyperbolically (with axis A, say). In the first case v is unique (since F cannot fix an edge), so v is K/N-invariant since K/N commutes with F. But K/N is normal so K/N also fixes $g(v)$ for all $g \in G$. Pick g with $g(v) \neq v$, then K/N fixes more than one vertex, and hence fixes an edge.

In the second case, the axis A is K/N-invariant since K/N commutes with w. If the action of K/N on A is non-trivial, then A is the (unique) minimal K/N–invariant subtree of T. But then T is F-invariant since F commutes with K/N. Thus F acts non–trivially on A with cyclic edge-stabilisers, which is impossible. Hence K/N fixes an edge.

In both cases, K/N fixes an edge, hence fixes all edges since K/N is normal and the action is edge-transitive. Thus K/N is a cyclic group acting trivially on T. The induced action of $\Gamma = G/K$ has finite cyclic edge stabilisers of the form $Stab_G(e)/K$. But Γ is torsion-free so these are all trivial. $\square$

As above, we write $G_{\mathrm{ab}} = H_1(G, \mathbb{Z})/(\text{torsion})$.

Lemma 6.9. *Suppose S is a finitely presented residually free group and that $(Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ is an MCS for S. Then:*
   (0) *each of the groups $S_i/\langle Z_i \rangle$ is a non-abelian limit group;*
   (1) *the natural homomorphism $S \to S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$ has kernel $Z(S)$ and so embeds $S/Z(S)$ as a full subdirect product of n non-abelian limit groups;*
   (2) *the natural homomorphism $S \to \Gamma_{\mathrm{ab}} \times S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$ is an embedding, where $\Gamma_{\mathrm{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$.*

Definition 6.10. To obtain the *reduced existential envelope* of S we fix an MCS $(Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ and define $\exists\mathrm{Env}_0(S) := S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$. The existential envelopeof S is then defined to be $\exists\mathrm{Env}(S) = \Gamma_{\mathrm{ab}} \times \exists\mathrm{Env}_0(S)$, where $\Gamma_{\mathrm{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$.

Remark 6.11. The above definition makes sense in the light of Lemmas 6.9 and Lemma 6.7. In the proof of Lemma 6.7, we chose the Z_i so that $M_i = \langle Z_i \rangle$, in the notation of Proposition 6.3, and we shall see in a moment that this equality is forced by the definition of an MCS alone. The canonical nature of the M_i makes envelopes more canonical than

they appear in the definition — Theorem A (4-5) makes this assertion precise.

Proof. Suppose that $(Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ is an MCS for the finitely presented residually free group S. Then by MCS(3) we know $\langle Z_i \rangle \subseteq C_S(Y_i)$. Now there are $x_i, y_i \in Y_i$ such that $[x_i, y_i] \neq_S 1$. Moreover $[x_i, y_i] \notin C_S(Y_i)$ because S is residually free. Hence the images of x_i and y_i in $S/\langle Z_i \rangle$ form a non-commuting pair. Writing S as a subdirect product of some collection $\Gamma_1, \ldots, \Gamma_n$ of limit groups, the projections of x_i and y_i into one of the factors Γ_j, say, do not commute. Now we see that S is a subdirect product of $\Gamma \times H$, where $\Gamma = \Gamma_j$ is a non-abelian limit group, H is a subdirect product of the Γ_i ($i \neq j$), and $Z_i \subset H$ (by commutative transitivity in Γ).

Now put $N = \langle Z_i \rangle \triangleleft S$ (by MCS(4)), and note that $N \subset K := S \cap H$. It follows from MCS(5) that S/N admits a splitting either over the trivial subgroup or a non-normal, infinite cyclic subgroup. Then by Lemma 6.8, if $K \neq N$, then the splitting is over K/N - a contradiction since K/N is normal in S/N.

Hence $\langle Z_i \rangle = N = K = S \cap H$, so $S/\langle Z_i \rangle \cong \Gamma$ is a non-abelian limit group, which proves (0).

Since limit groups are fully residually free, the centralizer of any non-commuting pair of elements in $S/\langle Z_i \rangle$ is trivial. Thus $\langle Z_i \rangle$ is maximal among the centralizers of non-commuting pairs of elements of S (cf. Proposition 6.3). In particular $\langle Z_i \rangle = C_S(Y_i)$ and $\langle\langle Y_i \rangle\rangle \subseteq C_S(\langle Z_i \rangle)$. Clearly each $\langle Z_i \rangle \supseteq Z(S)$.

Suppose now that $1 \neq u \in \langle Z_1 \rangle \cap \cdots \cap \langle Z_n \rangle$ but $u \notin Z(S)$. Then there is some other element v with $[u, v] \neq 1$. Since S is residually free, u and v freely generate a free subgroup of rank 2. Thus u and $v^{-1}uv$ freely generate a free subgroup of $\langle Z_1 \rangle \cap \cdots \cap \langle Z_n \rangle$ which centralizes each $\langle\langle Y_i \rangle\rangle$. So their images in $S/\langle\langle Y_1, \ldots, Y_n \rangle\rangle$ freely generate a free subgroup which contradicts MCS(6). Thus $\langle Z_1 \rangle \cap \cdots \cap \langle Z_n \rangle = Z(S)$. This proves (1).

The existence of the embedding in (2) follows immediately from (1), in the light of Lemma 6.2. $\square$

6.3. **Proofs of Theorem A(1) and (2).** We are given a finite presentation $\langle A \mid R \rangle$ for a residually free group S. In order to prove Theorem 6.1, we must describe an algorithm that will construct an MCS for S from this presentation: we know by Lemma 6.7 that S has an MCS and we know from Lemma 6.9 (and Definition 6.10) how to embed S in its envelopes once an MCS is constructed.

We shall repeatedly use the fact that one can use the given presentation of S to solve the word problem explicitly: one enumerates homomorphisms from S to the free group of rank 2 by choosing putative images for the generators $a \in A$, checking that each of the relations $r \in R$ is mapped to a word that freely reduces to the empty word; if

a word w in the letters $A^{\pm 1}$ is non-trivial is S, one will be able to see this in one of the free quotients enumerated, since S is residually free. (Implementing a naive search that verifies if w does equal the identity is a triviality in any recursively presented group.)

Using this solution to the word problem, we can recursively enumerate all finite collections $\Delta = (Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ of finite subsets of S satisfying conditions MCS(1), MCS(2) and MCS(3). Next we enumerate all equations in S and look for those of the form $a^{-1}za =_S w(Z_i)$ where $z \in Z_i$ and $a^{\pm 1}$ is a generator of S (and w any word on Z_i). If a given Δ satisfies MCS(4), we will eventually discover this by checking the list of equations. (As ever with such processes, one runs through the finite diagonals of an array, checking all equations against all choices of Δ.) Thus we obtain an enumeration of those Δ satisfying MCS(1-4).

Next, we must describe a process that, given

$$\Delta = (Y_1, \ldots, Y_n; Z_1, \ldots, Z_n),$$

can determine if it satisfies MCS(5), i.e. if each of the groups $S/\langle Z_i \rangle$ has a splitting of the required form. Again we only need a process that will terminate if Δ does indeed satisfy MCS(5) — we are content for it not to terminate if MCS(5) is not satisfied.

We have a finite presentation $\langle A \mid R, Z_i \rangle$ for $S/\langle Z_i \rangle$. By applying Tietze moves (or searching naively for inverse pairs of isomorphisms) we can enumerate finite presentations of $S/\langle Z_i \rangle$ that have one of the following two forms

$$\langle A_1, A_2 \mid R_1, R_2, u_1 u_2 \rangle, \quad \langle A_1, t \mid R_1, t u_1 t^{-1} v \rangle,$$

where A_1, A_2 and $\{t\}$ are disjoint sets, $R_i \cup \{u_i\}$ is a set of words in the letters $A_i^{\pm 1}$, and v is a word in the letters $A_1^{\pm 1}$. These are the standard forms of presentation for groups that split over (possibly trivial or finite) cyclic groups. When we find such a presentation, we can use the solution to the word problem in S to determine if at least one of the generators from A_1 and (for the first form) one from A_2 are non-trivial in S. We proceed to the next stage of the argument only if non-trivial elements are found. In the next stage, we use the solution to the word problem to check if $u_1 = u_2 = 1$ in S (or $u_1 = v = 1$). If these equalities hold, we have found the desired splitting over the trivial group. If not, then we have a splitting over a non-trivial cyclic group, and since S is torsion-free, this cyclic group $C = \langle u_1 \rangle$ must be infinite. In a residually free group, each 2-generator subgroup is free of rank 1 or 2 (consider the image of $[x, y]$ in a free group). Thus C is normal if and only if it is central, and this can be determined by applying the solution of the word problem to all commutators $[u, a]$ with $a \in A_1 \cup A_2$ (resp. $a \in A_1$). In the case of amalgamated free products, we require that there is a generator in each of A_1 and A_2 that does not commute with C, in order that the splitting be non-degenerate. This concludes the description of the process that will correctly determine if a given

$\Delta = (Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$ satisfies MCS(5), halting if it does (but not necessarily halting if it does not).

Finally, we use coset enumeration to get presentations $\langle A' \mid R' \rangle$ of subgroups of finite index $S_0 \subset S$ with $Y_i \subset S_0$, and we enumerate equations in the quotients $\langle A' \mid R', Y_1, \ldots, Y_n \rangle$ to see if the generators satisfy the defining relations of the free nilpotent group of class $n - 2$ on $|A'|$ generators (and we need only look for a positive answer). As an MCS for S exists (Lemma 6.7) this process will eventually terminate, yielding an explicit Δ satisfying MCS(1-6).

Part (2) of Theorem A follows immediately from part 1 in the light of Proposition 3.2. $\qquad\square$

6.4. Proof of Theorem A(3) [the universal property of $\exists\mathrm{Env}_0(S)$].
We first record the following general result which is also used implicitly in our discussion of how $\exists\mathrm{Env}(S)$ is related to the Makanin-Razborov diagram of S.

Proposition 6.12. *Let G be a subdirect product of a finite collection of groups: $G < G_1 \times \cdots \times G_n$. Then any homomorphism from S onto a non-abelian limit group Γ factors through one of the projection maps $p_i : G \to G_i$ $(i = 1, \ldots, n)$.*

Proof. An easy induction reduces us to the case where $n = 2$.

Define $L_i := G \cap G_i$ for $i = 1, 2$. Then L_i is normal in G for each i. Suppose that Γ is a non-abelian limit group and $\phi : G \to \Gamma$ is an epimorphism. Then $\phi(L_1)$ and $\phi(L_2)$ are mutually commuting normal subgroups of $\phi(G) = \Gamma$. If (say) $\phi(L_1)$ is non-trivial in Γ, then commutative transitivity in Γ implies that $\phi(L_2)$ is abelian. But Γ has no non-trivial abelian normal subgroups, so $\phi(L_2)$ is trivial.

Hence one or both of $\phi(L_i)$ $(i = 1, 2)$ is trivial. But if $\phi(L_1)$ is trivial, then ϕ factors through p_2, while if $\phi(L_2)$ is trivial, then ϕ factors through p_1. $\qquad\square$

To prove Theorem A (3), let S be a finitely presented, non-abelian, residually free group with MCS $(Y_1, \ldots, Y_n; Z_1, \ldots, Z_n)$. We have $S \hookrightarrow \exists\mathrm{Env}_0(S) = S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$, and we are given a homomorphism $\phi : S \to D = \Lambda_1 \times \cdots \times \Lambda_m$ with the Λ_i non-abelian limit groups and $\phi(S)$ subdirect. We must prove that ϕ extends uniquely to a homomorphism $\hat\phi : \exists\mathrm{Env}_0(S) \to D$.

For $k = 1, \ldots, m$ let ϕ_k denote the composition of ϕ with the projection $D \to \Lambda_k$. Since Λ_k is a non-abelian limit group, Proposition 6.12 says that the surjective map $\phi_k : S \to \Lambda_k$ factors through the projection $S \to S/\langle Z_i \rangle$ for some i. In particular, $\phi_k(Y_j) = 1$ for each $j \neq i$, since $Y_j \subset Z_i$. However, we must have $\phi_k(Y_i) \neq \{1\}$ by MCS(6) (else Λ is virtually nilpotent). Thus $i = i(k)$ is uniquely determined by k.

Applying the above in turn to each ϕ_k yields a unique $i(k)$ such that ϕ_k factors through a map $\zeta_k : S/\langle Z_{i(k)} \rangle \to \Lambda_k$. Putting all these maps together produces the required $\hat{\phi} : \exists\mathrm{Env}_0(S) \to \Lambda_1 \times \cdots \times \Lambda_m$. $\square$

6.5. Proof of Theorem A(4) [the uniqueness of $\exists\mathrm{Env}_0(S)$].

We are assuming that $\phi : S \hookrightarrow D = \Lambda_1 \times \cdots \times \Lambda_m$ is a full subdirect product of non-abelian limit groups, and we must prove that $\hat{\phi} : \exists\mathrm{Env}_0(S) \to D$ is an isomorphism.

As in the proof of Lemma 6.7, we can construct an MCS for S from the embedding $\phi : S \hookrightarrow D$, say $(Y_1', \ldots, Y_m'; Z_1', \ldots, Z_m')$. Here, $Y_i' \subset S$ generates $\phi(S) \cap \Lambda_i$ as a normal subgroup, Z_i' generates the centralizer of Y_i' in S, and ϕ induces an isomorphism $\overline{\phi}_i : S/\langle Z_i' \rangle \to \Lambda_i$ for $i = 1, \ldots, m$.

By using $(Y_i'; Z_i')$ in place of $(Y_i; Z_i)$ in Definition 6.10 we obtain an alternative model $\exists\mathrm{Env}_0(S)' = S/\langle Z_1' \rangle \times \cdots \times S/\langle Z_m' \rangle$ for $\exists\mathrm{Env}_0(S)$, and we have an isomorphism $\Phi = (\overline{\phi}_1, \ldots, \overline{\phi}_m) : \exists\mathrm{Env}_0(S)' \to D$ that restricts to ϕ on the canonical image of S in $\exists\mathrm{Env}_0(S)'$.

In proving Theorem A(3) we established the universal property for $\exists\mathrm{Env}_0(S)'$. We apply this to obtain a unique homomorphism $\alpha : \exists\mathrm{Env}_0(S)' \to \exists\mathrm{Env}_0(S)$ extending the inclusion $S \hookrightarrow \exists\mathrm{Env}_0(S)$. Thus we obtain a homomorphism $\alpha \circ \Phi^{-1} : D \to \exists\mathrm{Env}_0(S)$ such that $\alpha \circ \Phi^{-1} \circ \phi$ is the identity on S. But this means that $\alpha \circ \Phi^{-1} \circ \hat{\phi} : \exists\mathrm{Env}_0(S) \to \exists\mathrm{Env}_0(S)$ extends id $: S \to S$. The identity map of $\exists\mathrm{Env}_0(S)$ is also such an extension, so by the uniqueness assertion in A(3) we have that $\alpha \circ \Phi^{-1}$ is a left-inverse to $\tilde{\phi}$. By reversing the roles of $\exists\mathrm{Env}_0(S)$ and $\exists\mathrm{Env}_0(S)'$ we see that it is also a right-inverse. $\square$

6.6. Makanin-Razborov Diagrams.

We explain how existential envelopes are related to Makanin-Razborov diagrams.

The *Makanin-Razborov diagram* (or MR diagram) of a finitely generated group G is a method of encoding the collection of all epimorphisms from G to free groups. The name arises from the fact that these diagrams originate from the fundamental work of Makanin [27] and later Razborov [29] on the solution sets of systems of equations in free groups.

The MR diagram of G consists of a finite rooted tree, where the root is labelled by G and the other vertices are labelled by limit groups, with the leaves being labelled by free groups. The edges are labelled by proper epimorphisms – the epimorphism labeling $e = (u, v)$ mapping the group labeling u onto the group labeling v.

The basic property of this diagram is that each epimorphism from G onto a free group can be described using a directed path in this graph from the root to some leaf, the epimorphism in question being a composite of all the labeling epimorphisms of edges on this path, interspersed with suitable choices of 'modular' automorphisms of the

intermediate limit groups that label the vertices. Details can be found in [30, Section 7] and, in different language, [23, Section 8].

An immediate observation is that any epimorphism from G onto a free group factors through the canonical quotient $G/\mathrm{FR}(G)$, where $\mathrm{FR}(G)$ is the *free residual* of G, namely the intersection of the kernels of all epimorphisms from G to free groups. Thus the MR diagrams of G and of $G/\mathrm{FR}(G)$ are identical.

Observe that $\mathrm{FR}(G/\mathrm{FR}(G)) = 1$; in other words $G/\mathrm{FR}(G)$ is *residually free*. Thus, when studying MR diagrams for finitely generated groups, it is sufficient to restrict attention to the case of residually free groups.

For finitely generated residually free G, the top layer of the Makanin-Razborov diagram consists of the set of maximal limit-group quotients of G. These are the factors of our existential envelope $\exists\mathrm{Env}(G)$, namely the maximal free abelian quotient $\Gamma_{\mathrm{ab}}(G)$ and the non-abelian quotients $\Gamma_1, \ldots, \Gamma_n$. The fact that one can construct this effectively is contained in [23, Corollary 3.3], but our construction of the embedding $G \hookrightarrow \exists\mathrm{Env}(G)$ is of a quite different nature, and we feel that there is considerable benefit in its explicit description. It is also worth pointing out that neither the construction of our algorithm nor the proof that it terminates relies on the original results of Makanin and Razborov.

7. Decision problems

Theorem A provides considerable effective control over the finitely presented residually free groups. In this section we use this effectiveness to solve the multiple conjugacy problem for these groups and the membership problem for their finitely presented subgroups. Both of these problems are unsolvable in the finitely generated case, indeed there exist finitely generated subgroups of a direct product of two free groups for which the conjugacy and membership problems are unsolvable [28].

7.1. **The conjugacy problem.** Instead of considering the conjugacy problem for individual elements, we consider the multiple conjugacy problem, since the proof that this is solvable is no harder. The multiple conjugacy problem for a finitely generated group G asks if there is an algorithm that, given an integer l and two l-tuples of elements of G (as words in the generators), say $x = (x_1, \ldots, x_l)$ and $y = (y_1, \ldots, y_l)$, can determine if there exists $g \in G$ such that $gx_ig^{-1} = y_i$ in G, for $i = 1, \ldots, l$. There exist groups in which the conjugacy problem is solvable but the multiple conjugacy problem is not [6].

The scheme of our solution to the conjugacy problem uses an argument from [10] that is based on Theorem 3.1 of [4]. This is phrased in terms of bicombable groups. Recall that a group G with finite generating set A is said to be *bicombable* if there is a constant K and choice

of words $\{\sigma(g) \mid g \in G\}$ in the letters $A^{\pm 1}$ such that

$$d(a.\sigma(a^{-1}ga')_t, \ \sigma(g)_t) \leq K$$

for all $a, a' \in A$ and $g \in G$, where w_t denotes the image in G of the prefix of length t in w, and d is the word metric associated to A.

We shall only use three facts about bicombable groups. First, the fundamental groups of compact non-positively curved spaces are the prototypical bicombable groups, and limit groups are such fundamental groups [1]. Secondly, there is an algorithm that given any finite set $X \subset \Gamma$ as words in the generators of G will calculate a finite generating set for the centralizer of X. (This is proved in [4] using an argument from [19].) Finally, we need the fact that the multiple conjugacy problem is solvable in bicombable groups. The proof of this is a mild variation on the standard proof that bicombable groups have a solvable conjugacy problem. The key point to observe is that, given words u and v in the generators, if $g \in G$ is such that $g^{-1}ug = v$, then as t varies, the distance from 1 to $\sigma(g)_t^{-1}u\sigma(g)_t$ never exceeds $K \max\{|u|, |v|\}$. It follows that in order to check if two $(u_1, \ldots, u_k)$ and $(v_1, \ldots, v_k)$ are conjugate in G, one need only check if they are conjugated by an element g with $d(1, g) \leq |2A|^{K \max\{|u_i|, |v_i|\}}$ (cf. Algorithm 1.11 on p. 466 of [5]).

Proposition 7.1. *Let Γ be a bicombable group, let $H \subset \Gamma$ be a subgroup, and suppose that there exists a subgroup $L \subset H$ normal in Γ such that Γ/L is nilpotent. Then H has a solvable multiple conjugacy problem.*

Proof. Given a positive integer l and two l-tuples $\underline{x}, \underline{y}$ from H (as lists of words in the generators of Γ) we use the positive solution to the multiple conjugacy problem in Γ to determine if there exists $\gamma \in \Gamma$ such that $\gamma x_i \gamma^{-1} = y_i$ for $i = 1, \ldots, l$. If no such γ exists, we stop and declare that $\underline{x}$ and $\underline{y}$ are not conjugate in H. If γ does exist then we find it and consider

$$\gamma C = \{g \in \Gamma \mid g x_i g^{-1} = y_i \text{ for } i = 1, \ldots, l\},$$

where C is the centralizer of $\underline{x}$ in Γ. Note that $\underline{x}$ is conjugate to $\underline{y}$ in H if and only if $\gamma C \cap H$ is non-empty.

We noted above that there is an algorithm that computes a finite generating set for C. This enables us to employ Lo's algorithm (Lemma 7.3) in the nilpotent group Γ/L to determine if the image of γC intersects the image of H. Since $L \subset H$, this intersection is non-trivial (and hence x is conjugate to y) if and only if $\gamma C \cap H$ is non-empty. $\square$

A group G is said to have *unique roots* if for all $x, y \in G$ and $n \neq 0$ one has $x = y \iff x^n = y^n$. It is easy to see that residually free groups have this property. As in Lemma 5.3 of [10] we have:

Lemma 7.2. *Suppose G is a group in which roots are unique and $H \subset G$ is a subgroup of finite index. If the multiple conjugacy problem for H is solvable, then the multiple conjugacy problem for G is solvable.*

The final lemma that we need can be proved by a straightforward induction on the nilpotency class, but there is a more elegant argument due to Lo (Algorithm 6.1 of [24]) that provides an algorithm which is practical for computer implementation.

Lemma 7.3. *If Q is a finitely generated nilpotent group, then there is an algorithm that, given finite sets $S, T \subset Q$ and $q \in Q$, will decide if $q\langle S \rangle$ intersects $\langle T \rangle$ non-trivially.* $\qquad\square$

Theorem 7.4 (=Theorem J). *The multiple conjugacy problem is solvable in every finitely presented residually free group.*

Proof. Let Γ be a finitely presented residually free group. Theorem A allows us to embed Γ as a subdirect product in $D = \Lambda_1 \times \cdots \times \Lambda_n$, where Λ_i are limit groups, each $L_i = \Lambda_i \cap \Gamma$ is non-trivial, $L = L_1 \times \cdots \times L_n$ is normal in D, and D/L is virtually nilpotent. Let N be a nilpotent subgroup of finite index in D/L, let D_0 be its inverse image in D and let $\Gamma_0 = D_0 \cap \Gamma$.

We are now in the situation of Proposition 7.1 with $\Gamma = D_0$ and $H = \Gamma_0$. Thus Γ_0 has a solvable multiple conjugacy problem. Lemma 7.2 applies to residually free groups, so the multiple conjugacy problem for Γ is also solvable. $\qquad\square$

7.2. **The membership problem.** In the course of proving our next theorem we will need the following technical observation.

Lemma 7.5. *If Λ is a limit group, then there is an algorithm that, given a finite set $X \subset \Lambda$, will output a finite presentation for the subgroup generated by X.*

Proof. Let H be the subgroup generated by X. The lemma is a simple consequence of Wilton's theorem [32] that Λ has a subgroup of finite index that retracts onto H (using the argument of Lemma 5.5 in [10]). $\qquad\square$

Theorem 7.6 (=Theorem K). *If G is a finitely presented residually free group (given by a finite presentation) and $H \subset G$ is a finitely presentable subgroup (given by a finite generating set of words in the generators of G), then the membership problem for H is decidable, i.e. there is an algorithm which, given $g \in G$ (as a word in the generators) will determine whether or not $g \in H$.*

Note that, although we assume that H is finitely presentable, we do not assume knowledge of a finite presentation for H. Moreover, our algorithm is not *uniform* in H. That is, the algorithm depends on H (but not on $g \in G$). Indeed, the proof below describes more than one

algorithm: for any given H one of these algorithms works, but we do not claim to be able to tell which. See the remark following this proof for further discussion of this problem.

Proof. Theorem A provides a direct product D of limit groups that contains G, and a solution to the membership problem for $H \subset D$ provides a solution for $H \subset G$. Thus there is no loss of generality in assuming that G is a direct product of limit groups, say $G = \Lambda_1 \times \cdots \times \Lambda_n$. To complete the proof, we argue by induction on n. The case $n = 1$ is covered by the fact that limit groups are subgroup separable [32].

Let us assume, then, that there is a solution to the membership problem for each finitely presented subgroup of a direct product of $n - 1$ or fewer limit groups. We have $H \subset G = \Lambda_1 \times \cdots \times \Lambda_n$. Define $L_i = H \cap \Lambda_i$.

There is no loss of generality in assuming that elements $g \in G$ are given as words in the generators of the factors, and thus we write $g = (g_1, \ldots, g_n)$. We assume that the generators of H are given likewise.

We first deal with the case where some L_i is trivial, say L_1. The projection of H to $\Lambda_2 \times \cdots \times \Lambda_n$ is then isomorphic to H, so in particular it is finitely presented and our induction provides an algorithm that determines if $(g_2, \ldots, g_n)$ lies in this projection. If it does not, then $g \notin H$. If it does, then naively enumerating equalities $g^{-1}w = 1$ we eventually find a word w in the generators of H so that $g^{-1}w$ projects to $1 \in \Lambda_2 \times \cdots \times \Lambda_n$. Since $L_1 = H \cap \Lambda_1 = \{1\}$, we deduce that in this case $g \in H$ if and only if $g^{-1}w = 1$, and the validity of this equality can be checked because the word problem is solvable in G.

It remains to consider the case where H intersects each factor non-trivially. Again we are given $g = (g_1, \ldots, g_n)$. The projection H_i of H to Λ_i is finitely generated and Wilton's theorem [32] tells us that Λ_i is subgroup separable, so we can determine algorithmically if $g_i \in H_i$. If $g_i \notin H_i$ for some i then $g \notin H$ and we stop. Otherwise, we replace G by the direct product D of the H_i. Lemma 7.5 allows us to compute a finite presentation for H_i and hence D.

We are now reduced to the case where H is a full subdirect product of $G(= D)$. Theorem A(2) now tells us that $Q = G/L$ is virtually nilpotent, where $L = L_1 \times \cdots \times L_n$. Let $\phi : G \to Q$ be the quotient map.

Virtually nilpotent groups are subgroup separable, so if $\phi(g) \notin \phi(H)$ then there is a finite quotient of Q (and hence G) that separates g from H. But $\phi(g) \notin \phi(H)$ if $g \notin H$ because $L = \ker \phi$ is contained in H. Thus an enumeration of the finite quotients of G provides an effective procedure for proving that $g \notin H$ if this is the case. (Note that we need a finite presentation of G in order to make this enumeration procedure effective; hence our appeal to Lemma 7.5.)

We now have a procedure that will terminate in a proof if $g \notin H$. Once again, we run this procedure in parallel with a simple-minded enumeration of $g^{-1}w$ that will terminate with a proof that $g \in H$ if this is true. $\qquad\square$

Remark 7.7. Since we discovered the above proof, Bridson and Wilton [12] have proved that in the profinite topology of any finitely generated residually free group, all finitely presentable subgroups are closed. This gives a *uniform* solution to the membership problem for such subgroups. Using the results of [12] and [9], Chagas and Zalesski [14] proved that all finitely presented residually free groups are conjugacy separable.

7.3. **Recursive enumerablility.** In view of the insights we have gained into the structure of finitely presentable residually free groups, it seems reasonable to conjecture that the isomorphism problem for this class of groups is solvable. We have not yet succeeded in constructing an algorithm to determine isomorphism, but we are nevertheless able to prove the following partial result in this direction.

Theorem 7.8 (= Theorem G)**.** *The class of finitely presentable residually free groups is recursively enumerable. More precisely, there is a Turing machine which will output a list of finite group presentations* $\mathcal{P}_1, \mathcal{P}_2, \ldots$ *such that:*

  (1) *the group G_i presented by each $\mathcal{P}_i$ is residually free; and*
  (2) *every finitely presented residually free group is isomorphic to at least one of the groups G_i.*

Proof. First we enumerate the limit groups, using the algorithm in [20]. This leads in a standard way to an enumeration of finite subsets Y of finite direct products thereof: $Y \subset D := \Gamma_1 \times \cdots \times \Gamma_n$.

For each such Y and each pair i, j, the Todd-Coxeter procedure will tell us if $p_{ij}(Y)$ generates a finite-index subgroup of $\Gamma_i \times \Gamma_j$ (but will not terminate if it does not).

Whenever we encounter a finite collection of limit groups $\Gamma_1, \ldots, \Gamma_n$ and a finite subset $Y \subset D$ such that $p_{ij}(Y)$ generates a finite-index subgroup of $\Gamma_i \times \Gamma_j$ for all i, j, we set about constructing a finite presentation for the subgroup generated by Y, using Theorem 3.4.

Thus a list can be constructed of all finitely-presented full subdirect products of limit groups, together with a finite presentation for each one. By Theorem B this list contains (at least one isomorphic copy of) every finitely presentable residually free group. $\qquad\square$

The facts we have proved or mentioned in this paper provide recursive enumerations of various other classes of groups:

  (1) There is a recursive enumeration of the finitely generated residually free groups $S = sgp(X)$: each is given by a finite set X

that generates a full subdirect product in a finite direct product of limit groups $\Gamma_1 \times \cdots \times \Gamma_n$.

(2) One can extract from (1) a recursive enumeration of the finitely generated residually free groups with trivial centre (those for which each Γ_i is non-abelian), and a complementary enumeration of those with non-trivial centre.

(3) The subsequence of (1) consisting of those S that are finitely presentable is recursively enumerable (cf. Theorem 7.8).

(4) The subsequences of (3) consisting of those finitely presented residually free groups with trivial (resp. non-trivial) centre are recursively enumerable, as are the corresponding subsequences of the enumeration in Theorem 7.8.

7.4. **Partial results on the isomorphism problem.** Suppose we are given two finite presentations of residually free groups G and H. Can we decide algorithmically whether or not $G \cong H$?

There is a partial algorithm that will search for a mutually inverse pair of isomorphisms, expressed in terms of the given finite generating sets for G and H. This will terminate if and only if $G \cong H$, giving us the desired isomorphism in the process.

The difficult part of the problem is therefore to recognise, via invariants or otherwise, when $G \not\cong H$.

Our earlier results have provided computations of an important invariant, namely the set of maximal limit group quotients of G. Using the solution to the isomorphism problem for limit groups ([13, 15]), we can distinguish G from H unless these agree for G and H. The problem is thus effectively reduced to the case where G and H are specifically given to us as full subdirect products of limit groups $\Gamma_1, \ldots, \Gamma_n$.

Moreover, $Z(G) \cong Z(H) = Z$, say, and the Γ_i are all non-abelian if Z is trivial. In the case where Z is non-trivial, then precisely one of the Γ_i is abelian. We make the convention that in this case Γ_1 is abelian. Then $\Gamma_1 \cong H_1(G, \mathbb{Z})/(\text{torsion}) \cong H_1(H, \mathbb{Z})/(\text{torsion})$, and $Z(G) = G \cap \Gamma_1$, $Z(H) = H \cap \Gamma_1$. Under these circumstances, as a special case of Theorem A(4) we have:

Proposition 7.9. *Any isomorphism $\theta : G \to H$ is the restriction of an ambient automorphism of the direct product $\Gamma_1 \times \cdots \times \Gamma_n$. This in turn restricts to a set of isomorphisms $\Gamma_i \to \Gamma_{\sigma(i)}$ ($i = 1, \ldots, n$) for some permutation σ of $\{1, \ldots, n\}$.*

Since there are only finitely many candidate permutations σ, this proposition effectively reduces the isomorphism problem to the case where σ is the identity, in other words to the following:

Question: Given finitely presented full subdirect products G, H of a collection of limit groups $\Gamma_1, \ldots, \Gamma_n$ (at most one of which is abelian),

can we find automorphisms θ_i of Γ_i for each i, such that

$$(\theta_1, \ldots, \theta_n)(G) = H?$$

Recall that the automorphism groups of limit groups can be effectively described [13]. In particular, we can find finite generating sets X_i for each $Aut(\Gamma_i)$.

Proposition 7.10. *There is a solution to the isomorphism problem in the case when at most 2 of the Γ_i are non-abelian.*

Proof. Suppose first that no Γ_i is abelian (so that $n \leq 2$). If $n = 1$ then $G = \Gamma_1 = H$ and there is nothing to prove, so we may suppose that $n = 2$. By Theorem E, since G, H are finitely presented they have finite index in $\Gamma = \Gamma_1 \times \Gamma_2$. The index can be computed in each case using the Todd-Coxeter algorithm, and we may assume that the two indices are equal (to k, say). Now by [13] we can find a finite set $X = X_1 \times X_2$ of generators for $\Theta = Aut(\Gamma_1) \times Aut(\Gamma_2)$.

It is straightforward to construct the permutation graph for the action of Θ on the finite set of index k subgroups, and then to check whether or not G, H lie in the same component of this graph. This happens if and only if G is isomorphic to H via an automorphism of $\Gamma_1 \times \Gamma_2$ that preserves the direct factors. By Proposition 7.9, this suffices to solve the problem.

If Γ_1 is abelian, then G and H need not have finite index, so we have to amend the argument slightly. We may assume that Γ_1 is the only abelian direct summand. Moreover, Γ_1 is a torsion free abelian quotient of G and of H, while $G \cap \Gamma_1 = Z(G)$ and $H \cap \Gamma_1 = Z(H)$. By Section 6.4, we can effectively determine $Z(G)$ and $Z(H)$ as subgroups of Γ_1. By the classification of finitely generated abelian groups, we can decide whether or not there is an automorphism of Γ_1 that maps $Z(G)$ to $Z(H)$. If not, then $G \not\cong H$ and we are finished. Otherwise, we are reduced to the case where $G \cap \Gamma_1 = H \cap \Gamma_1$.

Now there is a unique direct summand A of Γ_1 such that $\Gamma_1 \cap G$ has finite index in A. Choosing an arbitrary direct complement B for A in Γ_1 gives us embeddings of G and H as finite index subgroups of $(\Gamma_1/B) \times \Gamma_2 \cong A \times \Gamma_2$ or of $A \times \Gamma_2 \times \Gamma_3$, and we may complete the argument as before. $\qquad\square$

One possible approach to the more general case is to proceed by induction on the number of direct factors. Projecting a finitely presentable subdirect product to the product of fewer factors again gives a finitely presentable group, so by induction we can assume that the corresponding projections of our two subgroups are isomorphic. But for the moment we do not see how this information might be used to complete a proof that the isomorphism problem is solvable.

References

[1] E. Alibegović and M. Bestvina, *Limit groups are CAT*(0), J. London Math. Soc. (2) **74** (2006), 259–272.

[2] G. Baumslag, M. R. Bridson, C. F. Miller III and H. Short, *Fibre Products, non-positive curvature, and decision problems*, Comment. Math. Helv. **75** (2000), 457-477.

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

[4] M. R. Bridson, *On the subgroups of semihyperbolic groups*, in "Essays on geometry and related topics", pp. 85–111, Monogr. Enseign. Math. **38**, Geneva, 2001.

[5] M. R. Bridson and A. Haefliger, "Metric Spaces of Non-Positive Curvature", Grund. Math. Wiss. **319**, Springer-Verlag, Heidelberg-Berlin, 1999.

[6] M. R. Bridson and J. Howie, *Conjugacy of finite subsets in hyperbolic groups.* Internat. J. Algebra Comput. **15** (2005), no. 4, 725–756.

[7] M. R. Bridson and J. Howie, *Subgroups of direct products of elementarily free groups*, Geom. Funct. Anal. (GAFA) **17** (2007), 385–403.

[8] M. R. Bridson, J. Howie, C. F. Miller III and H. Short, *The subgroups of direct products of surface groups*, Geometriae Dedicata **92** (2002), 95–103.

[9] M. R. Bridson, J. Howie, C. F. Miller III and H. Short, *Subgroups of direct products of limit groups.* Ann. of Math., in press.

[10] M. R. Bridson and C. F. Miller III, *Structure and finiteness properties of subdirect products of groups.* Proc. London Math. Soc., in press.

[11] M. R. Bridson and C. F. Miller III, *Recognition of subgroups of direct products of hyperbolic groups*, Proc. of Amer. Math. Soc., **132** (2003), 59–65.

[12] M. R. Bridson and H. Wilton, *Subgroup separability in residually free groups*, Math. Z., DOI 10.1007/s00209-007-0256-7.

[13] I. Bumagin, O. Kharlampovich and A. Miasnikov, *The isomorphism problem for finitely generated fully residually free groups.* J. Pure Appl. Algebra **208** (2007), 961–977.

[14] S. C. Chagas and P. A. Zalesskii, *Finite Index Subgroups of Conjugacy Separable Groups*, preprint, February 2008.

[15] F. Dahmani and D. Groves, The isomorphism problem for toral relatively hyperbolic groups, ArXiv:math:GR/0512605.

[16] T. Delzant and M. Gromov, *Cuts in Khler groups*, in *Infinite groups: geometric, combinatorial and dynamical aspects*, (L. Bartholdi and others, eds.), Progr. Math., 248, Birkhuser, Basel, 2005, pp. 3155. MR 2006j:32023

[17] W. Dison, *Isoperimetric functions for subdirect products and Bestvina-Brady groups*, Ph.D. thesis, Imperial College London, 2008.

[18] R. Geoghegan, *Topological Methods in Group Theory.* Graduate Texts in Mathematics # 243, Springer-Verlag, New York 2008.

[19] S. M. Gersten and H. B. Short, *Rational subgroups of biautomatic groups*, Ann. of Math. **134** (1991), 125–158.

[20] D. Groves and H. Wilton, *Enumerating limit groups*, arXiv:0704.0989.

[21] D. Kochloukova, *On subdirect products of type FP_m of limit groups*, preprint.

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

[23] O. Kharlampovich and A. Myasnikov, *Effective JSJ decompositions*, in: *Groups, languages, algorithms* (A. V. Borovik, ed.) Contemp. Math., 378, Amer. Math. Soc., Providence, RI, 2005, pp 87–212.

[24] E. H. Lo, *Finding intersections and normalizers in finitely generated nilpotent groups*, J. Symbolic Comput. **25** (1998), 45–59.

[25] W. Magnus, *Über Gruppen und zugeordnete Liesche Ringe*, J. Reine Angew. Math. **182** (1940), 142–149.

[26] W. Magnus, A. Karrass and D. Solitar, "Combinatorial Group Theory", Wiley, New Yok, 1966.

[27] G. S. Makanin, *Equations in a free group*, Math. USSR Izv. **21** (1983), 483-546.

[28] C. F. Miller III, "On group-theoretic decision problems and their classification", Annals of Mathematics Studies, No. 68, Princeton University Press (1971).

[29] A. A. Razborov, *On systems of equations in a free group*, Math. USSR Izv. **25** (1985), 115-162.

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

[31] J. H. C. Whitehead, *On Adding Relations to Homotopy Groups*, Annals of Math. **42** (1941), 409–428.

[32] H. Wilton. Hall's Theorem for limit groups. *Geom. Funct. Anal.*, to appear. arXiv:math/0605546.

MARTIN R. BRIDSON, MATHEMATICAL INSTITUTE, 24–29 ST GILES', OXFORD OX1 3LB, U.K.
  E-mail address: bridson@maths.ox.ac.uk

JAMES HOWIE, DEPARTMENT OF MATHEMATICS, HERIOT–WATT UNIVERSITY, EDINBURGH EH14 4AS
  E-mail address: jim@ma.hw.ac.uk

CHARLES F. MILLER III, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF MELBOURNE, PARKVILLE 3052, AUSTRALIA
  E-mail address: c.miller@ms.unimelb.edu.au

HAMISH SHORT, L.A.T.P., U.M.R. 6632, CENTRE DE MATHÉMATIQUES ET D'INFORMATIQUE, 39 RUE JOLIOT–CURIE, UNIVERSITÉ DE PROVENCE, F–13453, MARSEILLE CEDEX 13, FRANCE
  E-mail address: hamish@cmi.univ-mrs.fr