

ON THE FINITE PRESENTATION OF SUBDIRECT PRODUCTS AND THE NATURE OF RESIDUALLY FREE GROUPS

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

ABSTRACT. We establish *virtual surjection to pairs* (VSP) as a general criterion for the finite presentability of subdirect products of groups: if $\Gamma_1, \dots, \Gamma_n$ are finitely presented and $S < \Gamma_1 \times \dots \times \Gamma_n$ projects to a subgroup of finite index in each $\Gamma_i \times \Gamma_j$, then S is finitely presentable, indeed there is an algorithm that will construct a finite presentation for S .

We use the VSP criterion to characterise the finitely presented residually free groups. We prove that the class of such groups is recursively enumerable. We describe an algorithm that, given a finite presentation of a residually free group, constructs a canonical embedding into a direct product of finitely many limit groups. We solve the (multiple) conjugacy problem and membership problem for finitely presentable subgroups of residually free groups. We also prove that there is an algorithm that, given a finite generating set for such a subgroup, will construct a finite presentation.

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.

1. INTRODUCTION

A very challenging problem is to determine which subgroups $S < G_1 \times \dots \times G_n$ of a direct product of finitely presented groups are themselves finitely presented. Indeed this problem is subtle even when the G_i are free groups.

Some terminology is useful for describing how a subgroup sits inside a direct product. 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. A subgroup $S < G_1 \times \dots \times G_n$ is said to be *virtually*

1991 *Mathematics Subject Classification*. Primary 20F65, 20E08, 20F67.

Key words and phrases. Virtual surjection to pairs, subdirect products, residually free groups, limit groups, finitely presentation, algorithms.

This work grew out of a project funded by l'Alliance Scientific grant # PN 05.004. Bridson was supported in part by an EPSRC Senior Fellowship and a Royal Society Nuffield Research Merit Award. Howie was supported in part by Leverhulme Trust grant F/00 276/J.

surjective on pairs (VSP) if for all $i \neq j \in \{1, \dots, n\}$, the projection $p_{ij}(S) \subset G_i \times G_j$ has finite index. (We implicitly assume that $n \geq 2$.)

In [14] we showed that the full subdirect products of non-abelian free and surface groups which are finitely presented must satisfy the VSP condition, from which it follows that they contain a term of the lower central series of a subgroup of finite index in the direct product. The VSP condition also played an important role in our previous work on subdirect products of limit groups [11]. The first purpose of this article is to establish VSP as a criterion for the finite presentability of subgroups of more general direct products of groups. We remind the reader that a subgroup $S < \Gamma$ is termed *separable* if for every $\gamma \in \Gamma \setminus S$ there exists a normal subgroup $K \triangleleft \Gamma$ of finite index such that $\gamma \notin SK$.

Theorem A (The VSP Criterion). *Let $S < G_1 \times \dots \times G_n$ be a subgroup of a direct product of finitely presented groups. If S is virtually surjective on pairs (VSP), then it is finitely presented and separable.*

Note that we do not assume, *a priori*, that the subgroup S is finitely generated. The converse of Theorem A is false in general; even finitely presented full subdirect products need not satisfy VSP. For example, if N is a finitely-generated torsion-free nilpotent group that is not cyclic, and if $\phi : N \times N \rightarrow \mathbb{Z}$ is a homomorphism whose restriction to each factor is non-trivial, then the kernel of ϕ is a finitely presented, separable full subdirect product without VSP.

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

Theorem B (Asymmetric 1-2-3 Theorem). *Let $f_1 : \Gamma_1 \rightarrow Q$ and $f_2 : \Gamma_2 \rightarrow Q$ 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.

We shall concentrate on the *effective* version of this result (Theorem 2.2) which yields an explicit finite presentation for P . (Proofs of the non-effective version can be found in [12] and [22].) In Theorem 3.7 we use Theorem 2.2 to prove an effective version of Theorem A: there is a uniform partial algorithm that, given finite presentations for the factors G_i and a finite generating set for S satisfying VSP, will output a finite presentation for S .

In this paper we describe a number of algorithmic processes. Often they are *partial algorithms* meaning that when applied to an object X which satisfies some condition \mathcal{C} , the process will halt with some appropriate information about X ; but if X does not satisfy \mathcal{C} either the process will halt saying $X \notin \mathcal{C}$ or the process will fail to halt.

For instance in Theorem 3.7, which is our effective version of Theorem A, we describe such a partial algorithm which, when given a direct product of finitely presented groups and a finite set of generators for a subgroup S , succeeds when S actually satisfies VSP and yields a finite presentation for S . The algorithm is *uniform* in the given data (a direct product and a finite generating set for S) and so can be started without knowing whether or not S satisfies VSP. If S does not satisfy VSP, the algorithm of Theorem 3.7 does not halt. (In a direct product of free groups for instance, there is no algorithm to determine whether or not the subgroup generated by a finite set satisfies VSP nor whether it is finitely presentable.)

Residually free groups. In the second half of this article we use Theorem A and more specialised results to advance the understanding of finitely presented residually free groups. 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 . The prototypes for these groups are the finitely presented subgroups of finite direct products of free and surface groups. In general a finitely-presented residually-free 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 C below). In our earlier studies [14], [9], [10], [11], 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 [35]. 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. Using Theorem A, 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.

By definition, a group G is residually free if it 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 [27, Corollary 2] and [37, Claim 7.5]). In [28] Kharlampovich and Myasnikov describe an algorithm to find such an embedding, based on the deep work of Makanin [34] and Razborov [36]. We shall describe a new algorithm that does not depend on [34] and [36]; the embedding that we construct is canonical in a strong sense (see Theorem C).

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 [37]. They have been much studied in recent years in connection with Tarski's problems on the first order logic of free groups [37], [27]. 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, it is difficult 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 [11].

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

- (1) $\exists\text{Env}(S) = \Gamma_{\text{ab}} \times \exists\text{Env}_0(S)$ where $\Gamma_{\text{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$ and $\exists\text{Env}_0(S) = \Gamma_1 \times \cdots \times \Gamma_n$ is a direct product of non-abelian limit groups Γ_i . The intersection of S with the kernel of the projection $\rho : \exists\text{Env}(S) \rightarrow \exists\text{Env}_0(S)$ 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 $i = 1, \dots, n$, and therefore $\text{Nilp}_{\exists}(S) := \exists\text{Env}(S)/(L_1 \times \cdots \times L_n)$ is virtually nilpotent.

- (3) [Universal Property] *For every homomorphism $\phi : S \rightarrow D = \Lambda_1 \times \cdots \times \Lambda_m$, with $\phi(S)$ subdirect and the Λ_i non-abelian limit groups, there exists a unique homomorphism $\hat{\phi} : \exists\text{Env}_0(S) \rightarrow D$ with $\hat{\phi} \circ \rho|_S = \phi$;*
- (4) [Uniqueness] *moreover, if $\phi : S \hookrightarrow D$ embeds S as a full subdirect product, then $\hat{\phi} : \exists\text{Env}_0(S) \rightarrow D$ is an isomorphism that respects the direct sum decomposition.*

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

Proceeding in the opposite direction, Guirardel and Levitt [26] prove that, given a subdirect product S of limit groups, one can algorithmically construct a finitely presented group whose maximal centreless residually free quotient is isomorphic to $S/Z(S)$. They also show that there is no algorithm to determine whether the maximal residually free quotient of a finitely presentable group is finitely presentable. In a similar vein, we note that there is no algorithm to determine whether or not a finitely generated subdirect product of limit groups is finitely presentable. Indeed, if F is a non-abelian free group, then there is no algorithm to determine which finitely generated full subdirect products of $F \times F$ are finitely presentable (cf. [14]).

The second of the drawbacks we identified in the discussion preceding Theorem C 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 $i = 1, \dots, n$.

Theorem D. *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 $\text{FP}_2(\mathbb{Q})$;*
- (3) *$\dim H_2(S_0; \mathbb{Q}) < \infty$ for all subgroups $S_0 \subset S$ of finite index;*
- (4) *[\exists neat VSP] 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 $1 \leq i < j \leq n$;*
- (5) *[neat \implies VSP] for every neat embedding $S \hookrightarrow \Lambda_0 \times \cdots \times \Lambda_n$ into a product of limit groups, the image of S under the projection to $\Lambda_i \times \Lambda_j$ has finite index for $1 \leq i < j \leq n$.*

Corollary E. *For all $k \in \mathbb{N}$, a residually free group S is of type F_k if and only if it is of type $\text{FP}_k(\mathbb{Q})$.*

In this context it is worth noting that D. Kochloukova [29] has obtained results concerning the question of which subdirect products of limit groups are FP_k for $2 < k < n$. In the pro- p category, the analogous question has been completely answered [30], but for discrete groups it remains open in general.

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

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

The proof of Theorem D relies on our earlier work concerning the finiteness properties of subgroups of direct products of limit groups [11] as well as Theorem A (the VSP criterion).

In the final section of this paper we shall combine Theorem D with the effective form of Theorem A 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.

Subdirect products of free groups hold a particular historical interest, most notably in connection with Baumslag and Roseblade's groundbreaking work [4] on the (non)finite presentability of subgroups of $F \times F$, and the seminal constructions by Stallings [38] and Bieri [5] of finitely presentable groups that are not of type $\text{FP}_\infty(\mathbb{Q})$. Subdirect products of surface groups also have a special appeal: the work of Delzant and Gromov [20] 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. In the context of subdirect products of surface groups, Dimca, Papadima and Suciu [21] have constructed analogues of the Bieri-Stallings examples that are fundamental groups of smooth complex projective varieties (and hence Kähler). These are currently the only known examples of subdirect products of surface groups that are Kähler but not of type $\text{FP}_\infty(\mathbb{Q})$.

We construct the first examples of finitely presented subgroups of direct products of free groups that are neither $\text{FP}_\infty(\mathbb{Q})$ nor of Stallings-Bieri type, thus answering a question raised in [14]. (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 $i = 1, \dots, n$.*

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

Corollary I. *For all positive integers c and $n \geq c + 2$, there exists a finitely-presented residually-free group S for which $\text{Nilp}_{\exists}(S)$ 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 D describes the finitely-presented residually-free groups. A description of a quite different nature is given in Theorem 5.6: 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 C 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 subgroups [35].

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, \dots, u_n) and (v_1, \dots, 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 $i = 1, \dots, n$.*

As previously noted, in a direct product of non-abelian free groups there is no algorithm to determine whether a finitely generated subgroup can be finitely presented. Nevertheless in a finitely presented residually free group, if we are given a finite set of generators for a subgroup which is in fact finitely presentable, then we can effectively find a presentation it. The method is uniform in the given data. Here is a more formal statement.

Theorem K. *There is a uniform partial algorithm for finding presentations of finitely presentable subgroups of finitely-presented residually-free groups. More precisely, there is a partial algorithm that, given a*

finite presentation for a residually free group G and a finite set of words generating a subgroup H , will output a finite presentation for H if it exists.

In the direct product of non-abelian free groups there are finitely generated subgroups for which the membership problem is unsolvable, but the subgroups in question are not finitely presented. We establish a uniform solution to the membership problem for finitely-presentable subgroups of finitely-presented residually-free groups.

Theorem L. *There is a uniform partial algorithm that, given a finite presentation of a residually free group S , a finite generating set for a finitely presentable subgroup $H \subset S$ and a word w in the generators of S , can determine whether or not w defines an element of H .*

Following our work, alternative approaches to the conjugacy and membership problems were developed in [15] and [18].

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 $S \hookrightarrow \exists\text{Env}(S)$.

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 A. 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 D and Theorem 5.6. Section 6 is devoted to the proof of Theorem C and other aspects of the canonical embedding $S \hookrightarrow \exists\text{Env}(S)$. Finally, in Section 7, we turn our attention to decidability and enumeration problems, proving Theorems G, J, K and L.

Most of the results in this paper were proved in our preprint [12], posted in 2008. In the intervening period there have been a number of related developments by various authors [15, 18, 21, 26, 29, 30]. In presenting the current article, we have tried to strike a balance between the sometimes competing goals of presenting a single coherent account of our original material, and of taking these later developments into account.

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. We are also grateful to an anonymous referee for a careful reading of an earlier version of the paper, and insightful comments that led to significant improvements.

2. THE EFFECTIVE ASYMMETRIC 1-2-3 THEOREM

In this section we prove the following effective version of Theorem B. 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.

In order to gain a full understanding of the proof that we are going to present, the reader should be familiar with the original proof of the 1-2-3 Theorem as presented in [2, §§1.4–1.5]. In particular, we shall not rehearse the reasons why considerations of $\pi_2\mathcal{Q}$ enter naturally into the proof (cf. remark 2.4). We remind the reader that, 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 $Q = \pi_1K$. In the present context, though, it is better to regard elements of $\pi_2\mathcal{Q}$ as equivalence classes of identity sequences $[(w_1, r_1), \dots, (w_m, r_m)]$, where the w_i are elements of the free group $F(X)$, the $r_i \in R^{\pm 1}$, and where $\prod_{i=1}^m w_i^{-1}r_iw_i$ 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 [39].

We shall also need the following observations, which are addressed in a more pedestrian manner in [2].

Remark 2.1. Let $f : \Gamma \rightarrow Q$ be an epimorphism where $Q = \langle X \mid R \rangle$. Then

- (1) Γ can be presented as $\Gamma = \langle X, C \mid \hat{R}, S \rangle$, so that f is given by $f(x) = x$ for all $x \in X$ and $f(c) = 1$ for all $c \in C$; there is relation $\hat{r} \in \hat{R}$ of the form $r(X) = u(C)$ for each $r \in R$; and S consists of words that are products of subwords drawn from C^* , the set of conjugates of the symbols $c \in C$ by words in the free group on X .
- (2) If $\ker f$ is finitely generated, then one can further assume that C generates $\ker f$, and that S consists of two sets of relations: the first set expresses the fact that $\ker f$ is normal, with a relation of the form $x^{-\epsilon}cx^\epsilon = w(C)$ for each $x \in X, c \in C$ and $\epsilon = \pm 1$; the second set consists of words in the free group on C .

Theorem 2.2. *There exists a Turing machine that, given the following data describing group homomorphisms $f_i : \Gamma_i \rightarrow Q$ ($i = 1, 2$), 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 $\hat{a} \in F(X)$ such that $\hat{a} = f_i(a)$ in Q .

- (4) *A finite set of identity sequences that generates $\pi_2\mathcal{Q}$ as a $\mathbb{Z}\mathcal{Q}$ -module.*

Proof. We associate to a fixed finite group-presentation $\mathcal{Q} \equiv \langle X \mid R \rangle$ of a group Q the class $\mathcal{C}(\mathcal{Q})$ of finite group-presentations that have the form

$$\langle X \sqcup A \sqcup B \mid S_1, S_2, S_3, S_4, S_5 \rangle,$$

where

- S_1 consists of a relator $r(X)u_r(A)v_r(B^*)$ for each relator $r = r(X) \in R$, where B^* is the set of formal conjugates of the letters $b \in B$ by words in the free monoid on $X \cup X^{-1}$ and $v_r(B^*)$ is a word in the free group on this set, while $u_r(A)$ is a word in the free group on A ;
- S_2 consists of a relator $ax^\epsilon w_{a,x,\epsilon}(A)x^{-\epsilon}$ for each $a \in A$, $x \in X$ and $\epsilon = \pm 1$, with the $w_{a,x,\epsilon}(A)$ words in the free group on A ;
- $S_3 = \{aba^{-1}b^{-1} \mid a \in A, b \in B\}$;
- S_4 is a finite set of words in the free group on A ;
- S_5 is a finite set of words in the free group on B^* .

It is clear the class $\mathcal{C}(\mathcal{Q})$ is recursively enumerable. Moreover, each group G given by a presentation $\mathcal{P} \in \mathcal{C}(\mathcal{Q})$ comes naturally equipped with an epimorphism π onto Q , namely the map that at the level of generators restricts to the identity on X and sends each element of A and B to 1. We write G_A (resp. G_B) to denote the quotient of G by the normal closure N_A of A (resp. the normal closure N_B of B) and regard π as the canonical map $G \rightarrow G/N_A N_B \cong Q$. Note that the diagonal map $\Delta : G \rightarrow G_A \times G_B$ sends G onto the fibre-product of the epimorphisms $G_A \rightarrow Q$ and $G_B \rightarrow Q$. Note also that N_A is generated by A as a group (not just a normal subgroup) and that N_A commutes with N_B .

Remark 2.1(1) assures us that any finitely presented group Γ_1 admitting an epimorphism $f_1 : \Gamma_1 \rightarrow Q$ is isomorphic to G_A for some (indeed infinitely many) G given by a presentation from $\mathcal{C}(\mathcal{Q})$, via an isomorphism $\phi_A : G_A \rightarrow \Gamma_1$ such that $f_1 \circ \phi_A$ is the natural epimorphism

$$\pi_A : G_A = G/N_A \rightarrow G/N_A N_B.$$

Moreover, if one has a finitely presented group Γ_2 admitting an epimorphism $f_2 : \Gamma_2 \rightarrow Q$ with $\ker f_2$ finitely generated, then Remark 2.1(2) assures us Γ_2 is also isomorphic to G_B for some G given by a presentation in $\mathcal{C}(\mathcal{Q})$, via an isomorphism $\phi_B : G_B \rightarrow \Gamma_2$ such that

$$f_2 \circ \phi_B = \pi_B : G_B = G/N_B \rightarrow G/N_A N_B.$$

When one constructs a presentation in $\mathcal{C}(\mathcal{Q})$, the choice of words involving A and those involving B can be made entirely independently of one another; so the construction of maps ϕ_A and ϕ_B in the preceding paragraph can be carried out simultaneously, i.e. using the same G . In other words, for any pair of finitely presented groups Γ_1 and Γ_2 and

epimorphisms $f_1 : \Gamma_1 \rightarrow Q$, $f_2 : \Gamma_2 \rightarrow Q$ with $\ker f_2$ finitely generated, there is a group G given by a presentation from $\mathcal{C}(\mathcal{Q})$ and isomorphisms $\phi_A : G_A \rightarrow \Gamma_1$, $\phi_B : G_B \rightarrow \Gamma_2$ such that $f_1 \circ \phi_A = \pi_A$ and $f_2 \circ \phi_B = \pi_B$.

We now have sufficient notation to describe the algorithm that we seek. Given input data (1), (2) and (3), the algorithm works systematically through an enumeration of $\mathcal{C}(\mathcal{Q})$, searching diagonally for pairs $\phi_A : G_A \rightarrow \Gamma_i$, $\phi_B : G_B \rightarrow \Gamma_j$ as above, with $\{i, j\} = \{1, 2\}$. By hypothesis Γ_1 and Γ_2 are finitely presented, and at least one of the kernels of $\Gamma_k \rightarrow Q$ ($k = 1, 2$) is finitely generated. So our search procedure will eventually terminate successfully.

Let $\mathcal{P} \in \mathcal{C}(\mathcal{Q})$ be the presentation found by the procedure, let G be the group given by \mathcal{P} , and let N_A, N_B, G_A, G_B be as defined above. Then, $(\phi_A \times \phi_B) \circ \Delta$ maps G onto the fibre product $P < \Gamma_1 \times \Gamma_2$ of f_1 and f_2 , and the kernel of this map is $N_A \cap N_B$. Since N_A and N_B commute, it follows that $N_A \cap N_B$ is central in $N_A N_B$, and that conjugation in G gives it the structure of a $\mathbb{Z}Q$ -module. To complete our proof, it suffices to show that we can find algorithmically a finite generating set Z for $N_A \cap N_B$ as a $\mathbb{Z}Q$ -module; adjoining Z to the relators of \mathcal{P} will then present P , as required.

To obtain Z , we follow the construction of [2, §§1.4,1.5]. Killing the generators B in \mathcal{P} gives a presentation of G_B of the form $\langle X \sqcup A \mid S'_1 \cup S_2 \cup S_4 \rangle$ with $S'_1 = \{ru_r(A) : r \in R\}$, as in [2, §1.4]. By [2, Theorem 1.2], the normal closure $N_A/(N_A \cap N_B)$ of A in G_B has a presentation on the generators A with relators all the $F(X, A)$ -conjugates of $S_2 \cup S_4 \cup S_6 \cup Z$, where $S_6 = \{[ru_r, a] : r \in R, a \in A\}$ and where Z is a finite set of words in $A^{\pm 1}$ derived by a simple algorithm from a finite set of identity sequences that generate $\pi_2(\mathcal{Q})$ as a $\mathbb{Z}Q$ -module. (Note: this is where input datum (4) enters.) Now the relators S_2 and S_4 are already relators of \mathcal{P} , while the relators S_6 can readily be derived from the relators $S_1 \cup S_2 \cup S_3$ of \mathcal{P} . It follows that the elements of N_A represented by the words in Z generate $N_A \cap N_B$ as a normal subgroup of G (in other words as a $\mathbb{Z}Q$ -module), as required. \square

Remark 2.3. The algorithm in the preceding proof does not just produce a finite presentation $\mathcal{P} \equiv \langle T \mid \Sigma \rangle$ of the fibre product $P < \Gamma_1 \times \Gamma_2$, it also produces an explicit isomorphism $\Phi : |\mathcal{P}| \rightarrow P$ (induced by $(\phi_A \times \phi_B) \circ \Delta$ in the notation of the proof). It follows that if one has a preferred finite generating set Y for P , then one can construct a finite presentation for P with generators Y . Indeed, a naive search will identify, for each generator $t \in T$, a word u_t so that $\Phi(t) = u_t(Y)$ in P . Then $P = \langle T \sqcup Y \mid \Sigma, t^{-1}u_t(Y) (t \in T) \rangle$, and obvious Tietze moves remove the generators T .

Remark 2.4. There does not exist an algorithm that, on input a finite presentation of a group of type FP_∞ can output a finite set of module generators for π_2 of the presentation, so the last piece of input

data in the above theorem cannot easily be dispensed with. In fact, Bridson and Wilton [16] have proved that it is essential: there exists a recursive sequence of maps $\phi_n : \Gamma_n \rightarrow Q_n$, with Γ_n and Q_n given by finite presentations, such that each Q_n is of type FP_∞ and each kernel $\ker \phi_n$ is finitely generated, but the first Betti number of the associated (finitely presentable) fibre product $P_n < \Gamma_n \times \Gamma_n$ cannot be determined algorithmically (whereas it could be if one had a finite presentation in hand).

3. SUBDIRECT PRODUCTS AND VSP

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 $p_i : D \rightarrow G_i$ and $p_{ij} : D \rightarrow G_i \times G_j$ ($i, j = 1, \dots, n$). We implicitly assume that $n \geq 2$.

We remind the reader that a subgroup $S < D$ is said to be VSP (*virtually surjective on pairs*) if for all $i, j \in \{1, \dots, n\}, i \neq j$, the projection $p_{ij}(S) < G_i \times G_j$ has finite index.

Theorem 3.1 (= Theorem A). *Let $S < G_1 \times \cdots \times G_n$ be a subgroup of a direct product of finitely presented groups. If S is VSP, 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 [14] and [11].

Proposition 3.2. *Let G_1, \dots, G_n be groups and let $S < G_1 \times \cdots \times G_n$ be a subgroup. If S is VSP 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

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

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 $i = 1, \dots, n$.

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 $x_1, \dots, x_{n-1} \in G_1^0$ the commutator $([x_1, x_2, \dots, x_{n-1}], 1, \dots, 1)$ can be expressed as the commutator of elements from the subgroups $N_j \subset S$; explicitly it is

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

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

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

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

Remarks 3.3. (1) A slight variation on the above proof of (3) shows that if G_1, \dots, G_n are finitely generated groups and $H < G_1 \times \dots \times G_n$ is a subdirect product whose intersection with each of the factors G_i contains some term of the lower central series of a subgroup of finite index in G_i , then H is finitely generated.

(2) 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 [23] Theorem 7.1.10).

3.1. Proof of Theorem A. 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 $q : G_1 \times \cdots \times G_n \rightarrow G_1 \times \cdots \times G_{n-1}$ 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 $N_n = T \cap S$ and noting that

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

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

3.2. Separability. It remains to prove the assertion in the last phrase of Theorem A.

Lemma 3.4. *Let $D = G_1 \times \cdots \times G_n$ where the G_i are finitely generated. If a subgroup $H < D$ is such that, for each i , $H \cap G_i$ contains a subgroup N_i that is normal in G_i with virtually nilpotent quotient, then H is separable in D .*

Proof. Let $N = N_1 \times \cdots \times N_n$. Then D/N is virtually nilpotent, hence subgroup separable. So given $g \in D \setminus H$ we can find a finite-index subgroup K in D/N that contains H/N but not gN (noting that $N \subset H$). Then KN has finite index in D and separates H from g . \square

The following corollary completes the proof of Theorem A.

Corollary 3.5. *If $H < D$ has the VSP property, then it is separable.*

Proof. We replace the subgroup $G_i^0 < G_i$ of Proposition 3.2(1) with the intersection of all of its conjugates; it is then normal in G_i , as is each term of its lower central series. Thus it suffices to take $N_i = \gamma_{n-1}(G_i^0)$. \square

In our solution to the membership problem for finitely presented subgroups of residually free groups, we shall need a further consequence of Lemma 3.4.

Corollary 3.6. *If G_1, \dots, G_n are limit groups, then every finitely presented full subdirect product of $D = G_1 \times \dots \times G_n$ is separable in D .*

Proof. We proved in Theorem 4.2 of [11] that finitely presented full subdirect products of non-abelian limit groups have the VSP property, so it only remains to deal with the abelian factors. By collecting these factors, we may assume that precisely one factor (G_n say) is abelian. If $S < D$ is a finitely presented, full subdirect product, then its projection $q_n(S)$ in $G_1 \times \dots \times G_{n-1}$ is also finitely presentable, full and subdirect, and hence satisfies VSP.

The proof of Proposition 3.2(1) applies to $S < D$, even if it is not VSP: given $a_1, \dots, a_{n-2} \in G_i^0$, that proof gives $s_1, \dots, s_{n-2} \in S$ such that $p_i([s_1, \dots, s_{n-2}]) = [a_1, \dots, a_{n-2}]$ in G_i , and $p_j([s_1, \dots, s_{n-2}]) = 1$ in G_j for $1 \leq j \leq n-1, j \neq i$. But it is also true that $p_n([s_1, \dots, s_{n-2}]) = 1$ in G_n , since G_n is abelian. Hence $\gamma_{n-2}(G_i^0) \subset \gamma_{n-2}(S)$.

As in the proof of Corollary 3.5 above, we arrange that G_i^0 is normal and define $N_i = \gamma_{n-2}(G_i^0)$. (When $i = n$ this gives $N_n = 1$.) Lemma 3.4 now completes the proof. \square

3.3. The effective version. The following theorem will play a key part in our proof that the class of finitely-presentable residually-free groups is recursively enumerable.

Theorem 3.7. *There exists a Turing machine that, given a finite collection G_1, \dots, G_n of finitely presentable groups (each given by an explicit finite presentation) and a finite subset $Y \subset G_1 \times \dots \times G_n$ (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 $G_i \times G_j$ ($1 \leq i < j \leq n$), will output a finite presentation $\langle Y \mid R \rangle$ 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 A. As in Theorem A 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 case $n = 2$ is 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 $n \geq 3$.

By Theorem 2.2 and Remark 2.3 it is sufficient to find finite presentations for

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

$$(3) Q = G_n / (G_n \cap S),$$

together with

(4) explicit epimorphisms $T \rightarrow Q$ and $G_n \rightarrow Q$, and

(5) a finite set of generators for π_2 of the presentation for Q , as a $\mathbb{Z}Q$ -module.

A finite presentation $G_n = \langle Y \mid R_n \rangle$ 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 $\langle Y \mid r_1(Y), \dots, r_m(Y) \rangle$.

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

The epimorphism $T \rightarrow Q$ 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 $Q' \rightarrow P$, 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 B of generators of $\pi_2(X^{(2)})$ as a $\mathbb{Z}P$ -module (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 $P \rightarrow Q'$, by a naive search procedure.

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

We then have an exact homotopy sequence

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

(since the map $P \rightarrow Q$ is injective by hypothesis), together with a finite set B of generators for $\pi_2(X^{(2)})$ as a $\mathbb{Z}Q'$ -module (via our isomorphism $Q' \rightarrow P$). But $\pi_2(\widehat{K}, Z) \cong H_2(\widehat{K}/Z)$, since the quotient complex \widehat{K}/Z is simply connected. Hence $\pi_2(K) = \pi_2(\widehat{K})$ 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 $H_2(\widehat{K}/Z)$. Such a set C can

be found by a naive search over finite sets of identity sequences over \mathcal{Q} . \square

Addendum 3.8. *The Turing machine in Theorem 3.7 will not halt on input of a subset Y such that $p_{ij}(Y)$ generates a subgroup of infinite index in $G_i \times G_j$ $i \neq j$.*

Proof. The Todd-Coxeter process invoked in the preceding proof will not halt. \square

4. NOVEL EXAMPLES

From [11] (or [14] 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 [14] 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 [11] and [14] are optimal.

4.1. The groups $S(E, c)$. The first part of our discussion applies to groups of a rather general nature, but since our main interest lies with subgroups of direct products of free groups, we fix the suggestive notation F for a finitely presentable group with a fixed generating set $\{a, b\}$. (The restriction to the 2-generator case is just for notational convenience.) 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 $f : \mathbb{Z} \rightarrow F$ endowed with pointwise multiplication.

Let $\Gamma = \langle w, x, y, z \rangle$ be a free group of rank 4, and define a homomorphism $\phi : \Gamma \rightarrow \Phi$ 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 $E \rightarrow F$. We then obtain a projection $p_E : \Phi \rightarrow F^E$ by restriction: $p_E(f) = f|_E : E \rightarrow F$.

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 $p_{\{n\}} : \Phi \rightarrow \Phi_n$, 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 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 $1 \leq i \leq 4$.

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 $F^{E'} \rightarrow F^E$ induces an epimorphism $S(E', c) \rightarrow S(E, c)$.

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 A, $S(E, c)$ is finitely presentable.

(3) It is clear that $S(-E, c) \cong S(E, c)$ via the isomorphism $F^E \rightarrow F^{-E}$ 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 $\theta : F^E \rightarrow F^{E'}$ defined by $a_n \mapsto a_{n+1}$, $b_n \mapsto b_{n+1}$ is induced by the shift automorphism $\bar{\theta} : \Phi \rightarrow \Phi$ defined by $\bar{\theta}(f)(k) := f(k-1)$, in the sense that $p_{E'} \circ \bar{\theta} = \theta \circ p_E$.

Similarly, $\bar{\theta}$ commutes with the automorphism $\hat{\theta}$ of Γ defined by $w \mapsto w$, $x \mapsto x$, $y \mapsto yw^{-1}$, $z \mapsto zx^{-1}$, in the sense that $\bar{\theta} \circ \phi = \phi \circ \hat{\theta}$.

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.

Henceforth we assume that $F = \langle a, b \rangle$ is a free group of rank 2. (As above, this is purely for notational convenience; analogues of our examples can be constructed in the same way using non-abelian free groups of arbitrary rank.)

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 $\eta_n : \Phi_n \rightarrow U(A)$ 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 [32] (or [33, 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 $\eta : \Gamma \rightarrow U(\mathbb{Q}[t] \otimes_{\mathbb{Q}} A)$ 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 $\psi_n : \mathbb{Q}[t] \otimes_{\mathbb{Q}} A \rightarrow A$ 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

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

where Ω is the free monoid on $\{\alpha, \beta\}$ and $\pi_W(t) \in \mathbb{Q}[t]$ 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 $m \in E \setminus \{n\}$, we have

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

It follows that, in the expression $\eta(g) = \sum_{W \in \Omega} \pi_W(t) \cdot W(\alpha, \beta)$ for $\eta(g)$, the elements of $E \setminus \{n\}$, of which there are at least c , are roots of all the polynomials $\pi_W(t)$. In particular, for words W of length less than c , the polynomials π_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

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

□

4.2. Sample calculations. We use the explicit form of the map $\eta : \Gamma \rightarrow U(\mathbb{Q}[t] \otimes_{\mathbb{Q}} A)$ from the proof of Theorem 4.2 to make some calculations that illuminate the preceding proof. Recall that $\Gamma = \langle w, x, y, z \rangle$ is free of rank 4.

Remark 4.3. Suppose that $U, V \in \Gamma$, $k, \ell \geq 1$ 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 $\eta(U^{-1}V^{-1}) = 1 \bmod J$, 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

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

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 D 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 D). *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 $\text{FP}_2(\mathbb{Q})$;
- (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 of S as a full subdirect of finitely many limit groups so that the image is VSP;
- (5) the image of every neat embedding of S as a full subdirect of finitely many limit groups is VSP.

Proof. The implications (1) implies (2) implies (3) are clear. Theorem A 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. Now [3, Corollary 19] 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 , and 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).

Given a neat embedding $S \hookrightarrow \Lambda_0 \times \cdots \times \Lambda_m$, the image of the projection of S to $\Lambda_0 \times \Lambda_i$ has finite index for any $i > 0$, and the quotient \bar{S} of S by $Z(S) = S \cap \Lambda_0$ is a full subdirect product of the non-abelian limit groups $\Lambda_1, \dots, \Lambda_m$. Moreover, since $S \cap \Lambda_0$ is finitely generated, (3) implies that $H_2(\bar{S}_0; \mathbb{Q})$ is finite dimensional for all subgroups $\bar{S}_0 < \bar{S}$ of finite index in \bar{S} . It then follows from Theorem 4.2 of [11] that the image of the projection of S to $\Lambda_i \times \Lambda_j$ has finite index for any i, j with $0 < i < j \leq m$. Thus (3) implies (5). \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 F). *Let $k \geq 2$ be an integer, let $S \subset D := \Gamma_1 \times \cdots \times \Gamma_n$ be a full subdirect product of limit groups, and let $T \subset D$ be a subgroup that contains S . If S is of type $\text{FP}_k(\mathbb{Q})$ then so is T .*

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

In particular, S is of type $\text{FP}_2(\mathbb{Q})$, so by [11, Theorem 4.2] the quotient group D/L is virtually nilpotent, where $L = (S \cap \Gamma_1) \times \cdots \times (S \cap \Gamma_n)$.

By [11, Corollary 8.2], applied to T/L , there is a finite index subgroup $S_0 < S$, and a subnormal chain $S_0 \triangleleft S_1 \triangleleft \cdots \triangleleft S_\ell = T$ such that each quotient S_{i+1}/S_i is either finite or infinite cyclic.

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

Note that the condition $k \geq 2$ in Theorem 5.2 is essential. For example, if $G = \langle x, y | r_1, r_2, \dots \rangle$ is a 2-generator group that is not finitely presentable, then the subgroup T of $F(x, y) \times F(x, y)$ generated by $\{(x, x), (y, y), (1, r_1), (1, r_2), \dots\}$ 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 $F(x, y)$). 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. Theorem D 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 A of [14], 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 $\phi_i : \Lambda_i \rightarrow Q$, such that*

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

where

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

Proof. First we argue that the criterion in the statement is sufficient. Each ϕ_i is an epimorphism, so 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 $(\lambda_1, \lambda_2, \lambda_3) \in \ker(\phi)$, and the projection $p_{12} : \Gamma_1 \times \Gamma_2 \times \Gamma_3 \rightarrow \Gamma_1 \times \Gamma_2$ 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 A, and hence S is also finitely presentable.

Conversely, suppose that S is finitely presentable. By [11, Theorem 4.2] the image of each of the projections $p_{ij} : S \rightarrow \Gamma_i \times \Gamma_j$ ($1 \leq i < j \leq 3$) 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 $\phi : K_1 \times K_2 \times K_3 \rightarrow A$ 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 $\phi : \Lambda_1 \times \Lambda_2 \times \Lambda_3 \rightarrow Q$, and each $\phi_i : \Lambda_i \rightarrow Q$ is an epimorphism. \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. Our proof makes use of the following observation on subgroups of finitely generated nilpotent groups.

Lemma 5.4. *Let G be a finitely generated nilpotent group and let H be a subgroup of G . If the image of H in the abelianisation of G has finite index, then H has finite index in G .*

Proof. We argue by induction on the nilpotency class c of G . If $c = 1$ then G is abelian and there is nothing to prove. By induction, we may assume that the image of H in $G/\gamma_c G$ has finite index. Thus it suffices to show that $H \cap \gamma_c G$ has finite index in $\gamma_c G$. Now $\gamma_c G$ is an abelian group generated by finitely many commutators $w = [x_1, \dots, x_c]$ of weight c , so it in fact suffices to show that each has a power that lies in $\gamma_c G \cap H$. If $|G : H\gamma_2 G| = n$ then we can write $x_i^n = h_i y_i$ for some $h_i \in H$ and some $y_i \in \gamma_2 G$. An elementary commutator calculation using the fact that $\gamma_c G$ is central in G yields the equality

$$w^{n^c} = [x_1^n, \dots, x_c^n] = [h_1, \dots, h_c] \in H,$$

which completes the proof. \square

Definition 5.5. Let $\mathcal{G} = \{\Gamma_1, \dots, \Gamma_n\}$ be a finite collection of 2 or more limit groups, let $c \geq 2$ be an integer, and let $\underline{g} = \{(g_{k,1}, \dots, g_{k,n}), 1 \leq k \leq m\}$ be a finite subset of $\Gamma := \Gamma_1 \times \dots \times \Gamma_n$.

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 5.6. *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 A to know that the projection of T to $\Gamma_i \times \Gamma_j$ is virtually surjective for each $i < j$, and this follows from Lemma 5.4

Conversely, suppose that S is a finitely presentable residually free group. If S is not itself a limit group, then Theorem D tells us that S may be expressed as a full subdirect product of limit groups $\Delta_1, \dots, \Delta_n$ such that the projection of S to $\Delta_i \times \Delta_j$ is virtually surjective for each $i < j$. By Proposition 3.2(1), each Δ_i contains a finite-index subgroup Γ_i such that $\gamma_{n-1}(\Gamma_i) \subset S$. Set $\mathcal{G} = \{\Gamma_1, \dots, \Gamma_n\}$, and $c = n - 1$.

We choose any finite set $\underline{g} = \{(g_{1,1}, \dots, g_{1,n}), \dots, (g_{m,1}, \dots, g_{m,n})\}$ in the direct product $D := \Gamma_1 \times \dots \times \Gamma_n$ whose image in $D/\gamma_{n-1}(D)$ generates $(S \cap D)/\gamma_{n-1}(D)$. Then $T = T(\mathcal{G}, \underline{g}, n-1) = S \cap D$ is a finite-index subgroup of S . \square

6. THE CANONICAL EMBEDDING THEOREM

The purpose of this section is to prove Theorem C: we shall describe an effective construction for $\exists\text{Env}(S)$, hence $\exists\text{Env}_0(S)$, then establish the universal property of the latter. We shall see that the direct factors of $\exists\text{Env}(S)$ 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 $\exists\text{Env}_0(S)$. At the end of the section we shall discuss how $\exists\text{Env}(S)$ is related to the Makanin-Razborov diagram for S .

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

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

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

where $\Gamma_{\text{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$ and $\exists\text{Env}_0(S) = \Gamma_1 \times \dots \times \Gamma_n$ with each Γ_i ($i \geq 1$) a non-abelian limit group. The intersection of S with the kernel of the projection $\rho : \exists\text{Env}(S) \rightarrow \exists\text{Env}_0(S)$ 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* — which 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 [11]. In Lemma 6.9 we establish the existence of a simple algorithm that, given a maximal centralizer system, will construct $S \hookrightarrow \exists\text{Env}(S)$. 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 $S \rightarrow H_1(S, \mathbb{Z})/(\text{torsion})$ to $Z(S)$ is injective.*
- (2) *If Γ is a non-abelian limit group and $\psi : S \rightarrow \Gamma$ 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. \square

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 $\{L_1, \dots, L_n\}$ and $\{M_1, \dots, M_n\}$ 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. [13]). \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 $(Y_i; Z_i) = (Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ 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 $S_0/\langle\langle Y_1, \dots, Y_n \rangle\rangle$ 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 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 D, H can be embedded as a full subdirect product in some $D = \Gamma_1 \times \dots \times \Gamma_n$ where the Γ_i are non-abelian limit groups. Let $p_i : D \rightarrow \Gamma_i$ 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 $H < D$ is VSP and $\Gamma_i/(H \cap \Gamma_i)$ is virtually nilpotent, by [11, Theorem 4.2], 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 $(\hat{Y}_i; \hat{Z}_i)$ 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 $(\hat{Y}_i; \hat{Z}_i)$ satisfies MCS(1), note that $Z(S) \cap [S, S] = 1$. Modulo this observation, it is clear that $(\hat{Y}_i; \hat{Z}_i)$ inherits the properties MCS(1-6) from $(Y_i; Z_i)$. \square

6.2. Two useful lemmata. The following are the two principal lemmata used in the proof of Theorem C. 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 . Moreover, since $K/N \neq 1$, it has finite index in every (cyclic) edge stabiliser. Therefore, the action of $\Gamma = G/K$ on T has finite cyclic edge stabilisers of the form $Stab_G(e)/K$. But Γ is torsion-free so these stabilisers are all trivial. \square

Lemma 6.9. *Suppose S is a finitely presented residually free group and that $(Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ 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 $S \rightarrow S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$ 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 $S \rightarrow \Gamma_{\text{ab}} \times S/\langle Z_1 \rangle \times \cdots \times S/\langle Z_n \rangle$ is an embedding, where $\Gamma_{\text{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$.*

Definition 6.10. To obtain the *reduced existential envelope* of S we fix an MCS $(Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ and define $\exists\text{Env}_0(S) := S/\langle Z_1 \rangle \times \cdots \times$

$S/\langle Z_n \rangle$. The existential envelope of S is then defined to be $\exists \text{Env}(S) = \Gamma_{\text{ab}} \times \exists \text{Env}_0(S)$, where $\Gamma_{\text{ab}} = H_1(S, \mathbb{Z})/(\text{torsion})$.

Remark 6.11. The above definition makes sense in the light of Lemma 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 C(4-5) makes this assertion precise.

Proof of Lemma 6.9 Suppose that $(Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ 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 $\Gamma_1, \dots, \Gamma_n$ 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 $1 \neq u \in \langle Z_1 \rangle \cap \dots \cap \langle Z_n \rangle$ 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 $\langle Z_1 \rangle \cap \dots \cap \langle Z_n \rangle$ which centralizes each $\langle \langle Y_i \rangle \rangle$. So their images in $S/\langle \langle Y_1, \dots, Y_n \rangle \rangle$ freely generate a free subgroup which contradicts MCS(6). Thus $\langle Z_1 \rangle \cap \dots \cap \langle Z_n \rangle = Z(S)$. 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 C(1) and C(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 in 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 $\Delta = (Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ 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

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

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 be 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 $\Delta = (Y_1, \dots, Y_n; Z_1, \dots, Z_n)$ 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 $\langle A' \mid R', Y_1, \dots, Y_n \rangle$ 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 C follows immediately from part 1 in the light of Proposition 3.2. \square

6.4. Proof of Theorem C(3) [the universal property of $\exists \text{Env}_0(S)$].

We first record the following general result which is also used implicitly in our discussion of how $\exists \text{Env}(S)$ is related to the Makanin-Razborov diagram of S .

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

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 $\phi : G \rightarrow \Gamma$ 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 . \square

To prove Theorem C(3), let S be a finitely presented, non-abelian, residually free group with MCS $(Y_1, \dots, Y_n; Z_1, \dots, Z_n)$. We have $\rho : S \rightarrow \exists \text{Env}_0(S) = S/\langle Z_1 \rangle \times \dots \times S/\langle Z_n \rangle$, and we are given a homomorphism $\phi : S \rightarrow D = \Lambda_1 \times \dots \times \Lambda_m$ with the Λ_i non-abelian limit groups and $\phi(S)$ subdirect. We must prove that there is a unique homomorphism $\hat{\phi} : \exists \text{Env}_0(S) \rightarrow D$ with $\hat{\phi} \circ \rho = \phi$.

For $k = 1, \dots, m$ let ϕ_k denote the composition of ϕ with the projection $D \rightarrow \Lambda_k$. Since Λ_k is a non-abelian limit group, Proposition 6.12 says that the surjective map $\phi_k : S \rightarrow \Lambda_k$ factors through the projection $S \rightarrow S/\langle Z_i \rangle$ 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 $\zeta_k : S/\langle Z_{i(k)} \rangle \rightarrow \Lambda_k$. Putting all these maps together produces the required $\hat{\phi} : \exists\text{Env}_0(S) \rightarrow \Lambda_1 \times \dots \times \Lambda_m$. \square

6.5. Proof of Theorem C(4) [the uniqueness of $\exists\text{Env}_0(S)$]. We are assuming that $\phi : S \hookrightarrow D = \Lambda_1 \times \dots \times \Lambda_m$ is a full subdirect product of non-abelian limit groups, and we must prove that $\hat{\phi} : \exists\text{Env}_0(S) \rightarrow D$ 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 $(Y'_1, \dots, Y'_m; Z'_1, \dots, Z'_m)$. 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 $\bar{\phi}_i : S/\langle Z'_i \rangle \rightarrow \Lambda_i$ for $i = 1, \dots, m$.

By using $(Y'_i; Z'_i)$ in place of $(Y_i; Z_i)$ in Definition 6.10 we obtain an alternative model $\exists\text{Env}_0(S)' = S/\langle Z'_1 \rangle \times \dots \times S/\langle Z'_m \rangle$ for $\exists\text{Env}_0(S)$, and we have an isomorphism $\Phi = (\bar{\phi}_1, \dots, \bar{\phi}_m) : \exists\text{Env}_0(S)' \rightarrow D$ that restricts to ϕ on the canonical image of S in $\exists\text{Env}_0(S)'$.

In proving Theorem C(3) we established the universal property for $\exists\text{Env}_0(S)'$. We apply this to obtain a unique homomorphism $\alpha : \exists\text{Env}_0(S)' \rightarrow \exists\text{Env}_0(S)$ extending the inclusion $S \hookrightarrow \exists\text{Env}_0(S)$. Thus we obtain a homomorphism $\alpha \circ \Phi^{-1} : D \rightarrow \exists\text{Env}_0(S)$ such that $\alpha \circ \Phi^{-1} \circ \phi$ is the identity on S . But this means that $\alpha \circ \Phi^{-1} \circ \hat{\phi} : \exists\text{Env}_0(S) \rightarrow \exists\text{Env}_0(S)$ extends $\text{id} : S \rightarrow S$. The identity map of $\exists\text{Env}_0(S)$ is also such an extension, so by the uniqueness assertion in C(3) we have that $\alpha \circ \Phi^{-1}$ is a left-inverse to $\hat{\phi}$. By reversing the roles of $\exists\text{Env}_0(S)$ and $\exists\text{Env}_0(S)'$ 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 [34] and later Razborov [36] 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 labelling $e = (u, v)$ mapping the group labeling u onto the group labelling 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 labelling 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 [37, Section 7] and, in different language, [28, Section 8].

An immediate observation is that any epimorphism from G onto a free group factors through the canonical quotient $G/\text{FR}(G)$, where $\text{FR}(G)$ 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 $G/\text{FR}(G)$ are identical.

Observe that $\text{FR}(G/\text{FR}(G)) = 1$; in other words $G/\text{FR}(G)$ 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 $\exists\text{Env}(G)$, namely the maximal free abelian quotient $\Gamma_{\text{ab}}(G)$ and the non-abelian quotients $\Gamma_1, \dots, \Gamma_n$. The fact that one can construct this effectively was proved by Kharlampovich and Myasnikov in [28, Corollary 3.3]. Indeed, their construction will construct for *any* finitely presented G , the embedding of $G/\text{FR}(G)$ into its envelope. Our construction of the embedding $G \hookrightarrow \exists\text{Env}(G)$ is of a quite different nature, and it works only when G is residually free. Nevertheless we feel that there is considerable benefit in its explicit description. It is also worth noting 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 C 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 [35].

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 $x = (x_1, \dots, x_l)$ and $y = (y_1, \dots, y_l)$, can determine if there exists $g \in G$ such that $gx_i g^{-1} = y_i$ in G , for $i = 1, \dots, l$. There exist groups in which the conjugacy problem is solvable but the multiple conjugacy problem is not [8].

The scheme of our solution to the conjugacy problem uses an argument from [14] that is based on Theorem 3.1 of [6]. This is phrased in terms of bicomvable groups. Recall that a group G with finite generating set A is said to be *bicomvable* 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 bicomvable groups. First, the fundamental groups of compact non-positively curved spaces are the prototypical bicomvable 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 [6] using an argument from [24].) Finally, we need the fact that the multiple conjugacy problem is solvable in bicomvable groups. The proof of this is a mild variation on the standard proof that bicomvable 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, \dots, u_k) and (v_1, \dots, 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. 446 of [7]).

Proposition 7.1. *Let Γ be a bicomvable 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 $i = 1, \dots, l$. 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

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

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 [14] 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 [31]) 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 $\langle q \rangle$ intersects $\langle T \rangle$ non-trivially.* \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 C 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. \square

7.2. The finite presentation problem. We will need the following technical observation. This was first proved in [28], Theorem 3.21. It admits a short proof based on Wilton's theorem that finitely generated subgroups of limit groups are virtual retracts (see [25] Theorem 2.4).

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

Unlike limit groups, finitely generated subgroups of a finitely-presented residually-free group need not be finitely presentable. Our next result says that if such a finitely generated subgroup is finitely presentable, then we can effectively find a presentation.

Theorem 7.6 (=Theorem K). *There is a uniform partial algorithm for finding presentations of finitely presentable subgroups of finitely-presented residually-free groups. More precisely, there is a partial algorithm that, given a finite presentation for a residually free group G and a finite set of words generating a subgroup H , will output a finite presentation for H if it exists.*

To prove this theorem we begin by applying the algorithm of Theorem C to find the existential envelope D of G and the images of the generators of H in D . So it suffices to consider the case in which H is given by a finite set of generators in a specified direct product of limit groups. The theorem then follows from Proposition 7.9 below.

Remark 7.7. We pause to record the following observation. The algorithm in [25] provides an enumeration of non-abelian limit groups. An obvious modification of this enumeration produces a recursively enumerable sequence of finite presentations for direct products $\Gamma_1 \times \cdots \times \Gamma_n$ of limit groups. So if we are given any finite presentation of a group D and told that D is a direct product limit groups, then a naive search identifies a presentation for D on this list. Thus we may effectively replace the given presentation for D with one in which such a direct product decomposition and the coordinate projections $p_i : D \rightarrow \Gamma_i$ and $q_i : D \rightarrow D/\Gamma_i$ are manifest.

Lemma 7.8. *Suppose that H is a subgroup of a direct product $D = \Gamma_1 \times \cdots \times \Gamma_n$ of limit groups Γ_i . If H is finitely presented and*

$$\rho : D \rightarrow \Gamma_{i_1} \times \cdots \times \Gamma_{i_k} \text{ where } i_1 < \cdots < i_k$$

is the projection onto the product of any subset of k of the factors, then $\rho(H)$ is finitely presented.

Proof. If p_i denotes the projection onto the factor Γ_i , then each $p_i(H)$ is a finitely generated subgroup of a limit group and so is again a limit group. Hence we may assume from the outset that $\Gamma_i = p_i(H)$, that is, H is a subdirect product. Then every such $\rho(H)$ is also subdirect in the direct product $\rho(D)$.

If $n = 1$ or $n = k$ there is nothing to prove. Suppose that $H \cap \Gamma_j = 1$ where $\Gamma_j \subseteq \ker \rho$. Then the projection

$$q_j : D \rightarrow \Gamma_1 \times \cdots \times \Gamma_{j-1} \times \Gamma_{j+1} \times \cdots \times \Gamma_n$$

is injective on H . So by induction on n , it follows that $\rho(H) = \rho(q_j(H))$ is finitely presented as required.

Next suppose that $\rho(H) \cap \Gamma_{i_j} = 1$. Then $H \cap \Gamma_{i_j} = 1$ and the projection

$$q_{i_j} : D \rightarrow \Gamma_1 \times \cdots \times \Gamma_{i_j-1} \times \Gamma_{i_j+1} \times \cdots \times \Gamma_n$$

is injective on both H and $\rho(H)$. Now $q_{i_j}(D)$ is a product of $n - 1$ limit groups and $q_{i_j} \circ \rho = \rho \circ q_{i_j}$ is projection onto $k - 1$ factors, and hence

by induction $q_{i_j}(\rho(H))$ is finitely presented. But since $\rho(H) \cap \Gamma_{i_j} = 1$, it follows that $\rho(H)$ is finitely presented.

So we may now assume that both H and $\rho(H)$ are full subdirect products. Next we observe generally that if K is a full subdirect product of D and Γ_j is an abelian limit group, then $q_j(K) \subseteq D/\Gamma_j$ is finitely presented if and only if K is finitely presented, because Γ_j is a finitely-generated central subgroup. Hence it suffices to consider the case all of the Γ_i are non-abelian.

But a full subdirect product of non-abelian limit groups is finitely presented if and only if it satisfies VSP. Since H has VSP, either $\rho(H)$ is a limit group or it also has VSP. Hence $\rho(H)$ is finitely presented as desired. \square

Making use of this lemma we can now complete the proof of Theorem 7.6 by showing the following.

Proposition 7.9. *There is a partial algorithm that, given a direct product D of limit groups presented as $D = \Gamma_1 \times \cdots \times \Gamma_n$ and a finite set of words X generating a subgroup $H < D$,*

- (1) *in case H is finitely presentable, will output a finite presentation for H on the given generating set X ; or*
- (2) *in case H is not finitely presentable, will either halt saying that H is not finitely presentable, or will fail to halt.*

Moreover, in the case H is finitely presentable, the algorithm will also

- *determine for each i whether or not $H \cap \Gamma_i = 1$; and*
- *determine a finite generating set for the centre of H .*

Proof. Note that there is a uniform solution to the word problem for all such Γ_i and D since they are residually finite and given by finite presentations. Hence, as a finitely generated subgroup of D , H also has a solvable word problem.

If $n = 1$ then we apply the algorithm of Lemma 7.5 (in this case H is finitely presentable). We now proceed by induction on n . Let q_j be the projection of D onto all the factors other than Γ_j , so $q_j : D \rightarrow D/\Gamma_j$. By Lemma 7.8, if H is finitely presented then each $Q_j = q_j(H)$ is also finitely presented and lies in a direct product of fewer limit groups, so we have (by our inductive assumption) a partial algorithm to find a presentation for Q_j .

So we now launch n versions of this partial algorithm attempting to find a presentation for each $q_i(H)$ in case it is finitely presented ($i = 1, \dots, n$). If some process finds a Q_j is not finitely presented, we are done since H cannot be finitely presented and we halt with this information.

So we may assume each of our n processes finds that its corresponding $Q_i = \langle q_i(X) \rangle$ is finitely presented, say as $\langle x_1, \dots, x_s \mid r_1 = 1, \dots, r_t = 1 \rangle$ where $X = \{x_1, \dots, x_s\}$ and the normal subgroup generated by the r_j

is the kernel of the surjection $x \mapsto q_i(x)$ to Q from the free group on X . Then q_i is an isomorphism from H onto Q_i if and only if all of the $r_j =_H 1$. And these equalities can be tested using the solution to the word problem for D . Of course $H \cap \Gamma_i = 1$ if and only if q_i is an isomorphism.

So if H is finitely presentable, each of these n processes will (eventually) have halted and we have decided whether or not each $H \cap \Gamma_i = 1$. If some $H \cap \Gamma_i = 1$, then $Q_i \cong H$ and we have a presentation for H and our inductive algorithm also provides a set of generators for the centre of H and we are done.

So we may now assume that we have a finite presentation for each Q_i and that $H \cap \Gamma_i \neq 1$, thus H is a full subdirect product.

Observe that using the solution to the word problem we can easily determine which of the Γ_i are abelian (and then they must be free abelian). So next suppose that none of the Γ_i is abelian. In this case, we know H is finitely presentable if and only if it satisfies VSP. Also, the centre of H is trivial in this case.

If $n = 2$, this is true if and only if H has finite index in $\Gamma_1 \times \Gamma_2$. So if $n = 2$, we start trying to find a presentation for H using the classic Todd-Coxeter algorithm. If H is finitely presentable, it will terminate and give us a presentation for H . If H is not finitely presentable, the process will fail to terminate. (Note: whether this process will terminate is an unsolvable problem.)

If $n > 2$, then since each Q_i is full and finitely presented, each Q_i has VSP and hence H has VSP. Thus applying the algorithm of Theorem 3.7 we can effectively find a finite presentation \mathcal{P} of H on the generators X .

Now suppose that at least one of the Γ_i is abelian, and recall it is a finitely generated free abelian group. Group all the abelian factors together as one factor and revise the notation so that Γ_n is the only abelian direct factor. Then, applying our algorithm inductively, we can assume we have a finite presentation for $q_n(H) = Q_n = \langle q_n(X) \rangle$, say as $\langle x_1, \dots, x_s \mid r_1 = 1, \dots, r_t = 1 \rangle$. We also know that $q_n(H) \cong H/Z(H)$. Lifting the generators of this presentation back up to the corresponding generators of H , the same words r_j , which are relators of Q_i , are elements in D such that $r_j \in \Gamma_n$. Further, they generate the centre $Z(H)$ of H . Using standard algorithms for finitely generated abelian groups, we can find a set of defining relations, say $z_1 = 1, \dots, z_k = 1$ for the subgroup of Γ_n generated by the r_i . Finally, expressing the z_j as words \hat{z}_j in the x_i , we can then write down a presentation for H as

$$H = \langle x_1, \dots, x_s \mid [x_i, r_j] = 1 \text{ for } i = 1, \dots, s, j = 1, \dots, t, \\ \hat{z}_1 = 1, \dots, \hat{z}_k = 1 \rangle.$$

This complete the proof. □

As previously explained this also completes the proof of Theorem 7.6.

7.3. The membership problem. Let G be a recursively presented group, and H the subgroup generated by a set of words in the generators of G . The *membership problem* for H in G is the algorithmic problem of deciding, given a word w in the generators of G , whether or not the element $g \in G$ represented by w belongs to H .

Remark 7.10. There is an obvious (uniform) algorithm to solve the membership problem for separable subgroups of finitely presented groups: one runs a naive search to express the given element as a word in the generators of the subgroup while, in parallel, enumerating the finite quotients of the ambient group, checking to see if the element is separated from the subgroup in any of them.

We prove that there is a uniform partial algorithm to solve the membership problem for finitely presentable subgroups of finitely presented residually free groups.

Theorem 7.11 (=Theorem L). *There is a uniform partial algorithm that, given a finite presentation of a residually free group G , a finite generating set for a subgroup $H \subset G$ and a word g in the generators defining G , will determine whether or not g lies in H , provided that H is finitely presentable.*

Proof. The algorithm given by Theorem C embeds G in a direct product Δ of limit groups. A solution to the membership problem for $H \subset \Delta$ 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 [40].

We first employ the algorithm of Proposition 7.9: if H is not finitely presentable then this algorithm will either fail to halt or else halt and inform us that H is not finitely presentable; if H is finitely presentable then it will halt and list the indices i such that $L_i := H \cap \Lambda_i = 1$.

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 $g = (g_1, \dots, g_n)$. 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, \dots, 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 using the uniform solution to the word in residually free groups.

It remains to consider the case where H intersects each factor non-trivially. Again we are given $g = (g_1, \dots, g_n)$. The projection H_i of H to Λ_i is finitely generated and Wilton's theorem [40] 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 $G(= D)$. Such subgroups are separable, by Corollary 3.6, so remark 7.10 completes the proof. \square

Remark 7.12. Following our work, Bridson and Wilton [15] proved that in the profinite topology on any finitely generated residually free group, all finitely presentable subgroups are closed. Using the results of [15] and [11], Chagas and Zalesski [18] proved that all finitely presented residually free groups are conjugacy separable.

7.4. Recursive enumerability. 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 7.13 (= Theorem G). *The class of finitely presentable residually free groups is recursively enumerable. More precisely, there is a Turing machine that will output a list of finite group-presentations $\mathcal{P}_1, \mathcal{P}_2, \dots$ 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 [25]. This leads in a standard way to an enumeration of finite subsets Y of finite direct products thereof: $Y \subset D := \Gamma_1 \times \dots \times \Gamma_n$.

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 $\Gamma_1, \dots, \Gamma_n$ 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 3.7.

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 D this list contains (at least one isomorphic copy of) every finitely presentable residually free group. \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: 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 groups that are finitely presentable is recursively enumerable (Theorem 7.13).
- (4) The subsequence of (3) consisting of those finitely presented residually free groups with trivial (resp. non-trivial) centre is recursively enumerable.

7.5. 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 ([17, 19]), 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 $\Gamma_1, \dots, \Gamma_n$.

Moreover, by Proposition 7.9 we can effectively determine whether or not $Z(G) \cong Z(H)$. So we may assume that $Z(G) \cong Z(H) = Z$, say, and that 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 C(4) we have:

Proposition 7.14. *Any isomorphism $\theta : G \rightarrow H$ 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 $\Gamma_i \rightarrow \Gamma_{\sigma(i)}$ ($i = 1, \dots, n$) for some permutation σ of $\{1, \dots, n\}$.*

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 $\Gamma_1, \dots, \Gamma_n$ (at most one of which is abelian), can we find automorphisms θ_i of Γ_i for each i , such that

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

Recall that the automorphism groups of limit groups can be effectively described [17]; in particular one can find finite generating sets X_i for each $\text{Aut}(\Gamma_i)$.

Proposition 7.15. *There is a solution to the isomorphism problem in the case when at most 2 of the Γ_i are non-abelian.*

Proof. By Proposition 7.9 we can effectively find neat embeddings for G and H . Hence there is no loss of generality in assuming that the given embeddings $G, H \hookrightarrow D := \Gamma_1 \times \dots \times \Gamma_n$ are neat. In particular, at most two of the Γ_i are non-abelian, at most one is abelian, and G, H intersect any abelian direct factor in a subgroup that has finite index in that factor. The VSP property then ensures that each of G, H has finite index in D .

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 [17] we can find a finite set $X = X_1 \times \dots \times X_n$ of generators for $\Theta = \text{Aut}(\Gamma_1) \times \dots \times \text{Aut}(\Gamma_n)$.

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 and H lie in the same component of this graph. This happens if and only if G is isomorphic to H via an automorphism of D that preserves the direct factors. By Proposition 7.14, this suffices to solve the problem. □

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] G. Baumslag and J. Roseblade, *Subgroups of direct products of free groups*, J. London Math. Soc. (2) **30** (1984) 44–52.
- [5] R. Bieri, *Homological dimension of discrete groups*, Queen Mary College Mathematics Notes (1976).
- [6] 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.
- [7] M. R. Bridson and A. Haefliger, “Metric Spaces of Non-Positive Curvature”, Grundlehren Math. Wiss. **319**, Springer-Verlag, Heidelberg-Berlin, 1999.
- [8] M. R. Bridson and J. Howie, *Conjugacy of finite subsets in hyperbolic groups*. Internat. J. Algebra Comput. **15** (2005), no. 4, 725–756.
- [9] M. R. Bridson and J. Howie, *Subgroups of direct products of elementarily free groups*, Geom. Funct. Anal. (GAFA) **17** (2007), 385–403.
- [10] 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.
- [11] M. R. Bridson, J. Howie, C. F. Miller III and H. Short, *Subgroups of direct products of limit groups*. Ann. of Math. **170** (2009), 1447–1467.
- [12] M. R. Bridson, J. Howie, C. F. Miller III and H. Short, *Finitely presented residually free groups*. ArXiv:0809.3704 (2008).
- [13] 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.
- [14] M. R. Bridson and C. F. Miller III, *Structure and finiteness properties of subdirect products of groups*. Proc. London Math. Soc. (3) **98** (2009), 631–651.
- [15] M. R. Bridson and H. Wilton, *Subgroup separability in residually free groups*, Math. Z. **260** (2008), 25–30.
- [16] M. R. Bridson and H. Wilton, *On the difficulty of presenting finitely presentable groups*, Groups Geom. Dyn. **5** (2011), 301–325.
- [17] 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.
- [18] S. C. Chagas and P. A. Zalesskii, *Finite index subgroups of conjugacy separable groups*, Forum Math. **21** (2009), 347–353.
- [19] F. Dahmani and D. Groves, *The isomorphism problem for toral relatively hyperbolic groups*, Publ. Math. Inst. Hautes Études Sci. **107** (2008), 211–290.
- [20] T. Delzant and M. Gromov, *Cuts in Kähler groups*, in *Infinite groups: geometric, combinatorial and dynamical aspects*, (L. Bartholdi and others, eds.), Progr. Math., 248, Birkhäuser, Basel, 2005, pp. 3155. MR 2006j:32023
- [21] A. Dimca, S. Papadima and A. I. Suciuc, *Non-finiteness properties of fundamental groups of smooth projective varieties*, J. Reine Angew. Math. **629** (2009), 405–457.
- [22] W. Dison, *Isoperimetric functions for subdirect products and Bestvina-Brady groups*, Ph.D. thesis, Imperial College London, 2008.
- [23] R. Geoghegan, *Topological Methods in Group Theory*. Graduate Texts in Mathematics vol. 243, Springer-Verlag, New York 2008.
- [24] S. M. Gersten and H. B. Short, *Rational subgroups of biautomatic groups*, Ann. of Math. **134** (1991), 125–158.
- [25] D. Groves and H. Wilton, *Enumerating limit groups*, Groups Geom. Dyn. **3** (2009), 389–399.
- [26] V. Guirardel and G. Levitt, *Computing equations for residually free groups*, Illinois J. Math. **54** (2010), 129–135.

- [27] 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.
- [28] 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.
- [29] D. Kochloukova, *On subdirect products of type FP_m of limit groups*, J. Group Theory **13** (2010), 1–19.
- [30] D. Kochloukova and H. Short, *On subdirect products of free pro- p groups and Demushkin groups of infinite depth*, J. Algebra **343** (2011), 160–172.
- [31] E. H. Lo, *Finding intersections and normalizers in finitely generated nilpotent groups*, J. Symbolic Comput. **25** (1998), 45–59.
- [32] W. Magnus, *Über Gruppen und zugeordnete Liesche Ringe*, J. Reine Angew. Math. **182** (1940), 142–149.
- [33] W. Magnus, A. Karrass and D. Solitar, “Combinatorial Group Theory”, Wiley, New York, 1966.
- [34] G. S. Makanin, *Equations in a free group*, Math. USSR Izv. **21** (1983), 483–546.
- [35] C. F. Miller III, “On group-theoretic decision problems and their classification”, Annals of Mathematics Studies, No. 68, Princeton University Press (1971).
- [36] A. A. Razborov, *On systems of equations in a free group*, Math. USSR Izv. **25** (1985), 115–162.
- [37] Z. Sela, *Diophantine geometry over groups. I. Makanin-Razborov diagrams*, Publ. Math. Inst. Hautes Études Sci., pages 31–105, 2001.
- [38] J. R. Stallings, *A finitely generated group whose 3-dimensional homology group is not finitely generated*, Amer. J. Math., **85** (1963), 541–543.
- [39] J. H. C. Whitehead, *On Adding Relations to Homotopy Groups*, Ann. of Math. **42** (1941), 409–428.
- [40] H. Wilton, *Hall’s Theorem for limit groups*, *Geom. Funct. Anal.* **18** (2008), 25–30.

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 AND THE MAXWELL INSTITUTE FOR MATHEMATICAL SCIENCES, HERIOT-WATT UNIVERSITY, EDINBURGH EH14 4AS

E-mail address: `jin@ma.hw.ac.uk`

CHARLES F. MILLER III, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF MELBOURNE, MELBOURNE 3010, AUSTRALIA

E-mail address: `c.miller@ms.unimelb.edu.au`

HAMISH SHORT, L.A.T.P., U.M.R. 7353, CENTRE DE MATHÉMATIQUES ET D’INFORMATIQUE, 39 RUE JOLIOT-CURIE, UNIVERSITÉ D’AIX-MARSEILLE, F-13453, MARSEILLE CEDEX 13, FRANCE

E-mail address: `hamish@cmi.univ-mrs.fr`