Describing free groups

J. Carson, V. Harizanov, J. Knight, K. Lange, C. Maher,
C. McCoy, A. Morozov, S. Quinn, J. Wallbaum*

July 1, 2009

Abstract

We consider countable free groups of different ranks. For these groups, we investigate computability theoretic complexity of index sets within the class of free groups and within the class of all groups. For a computable free group of infinite rank, we consider the difficulty of finding a basis.

1 Introduction

Free groups play an important role in several branches of mathematics, including algebra, logic, and topology. Within logic, around 1945, Tarski asked whether free groups on different finite numbers of generators (greater than 1), are elementarily equivalent. Sela gave a positive answer to this question [10], [11], [12], [13], [14], [15], [16] (see also Kharlampovich and Myasnikov [4]).

In light of this result, we try to describe the different free groups, as simply as possible, using infinitary sentences. Formulas of $L_{\omega_1\omega}$ are infinitary formulas with countable disjunctions and conjunctions, but only finite strings of quantifiers. If we restrict the disjunctions and conjunctions to c.e. sets, then we have the *computable infinitary formulas* [1].

Scott [9] showed that if $\mathcal{A}$ is a countable structure for a countable language L, then there is a sentence of $L_{\omega_1\omega}$ whose countable models are exactly the isomorphic copies of $\mathcal{A}$. A sentence with this property is called a "Scott sentence" for $\mathcal{A}$. To describe specific countable free groups, we use computable infinitary sentences, and we aim for the simplest possible form.

For infinitary formulas, in particular, for computable infinitary formulas, we cannot bring the quantifiers "outside", but we can bring negations "inside". We have a kind of normal form, and we can classify formulas according to the number of alternations of infinite disjunction and $\exists$ with infinite conjunction and $\forall$.

*The authors acknowledge partial support under NSF Grant # DMS-0554841. The second author also received partial support under NSF Grant # DMS-0904101.

- $\varphi(\overline{x})$ is computable Π_0 and computable Σ_0 if it is finitary quantifier-free.
- For a computable ordinal $\alpha > 0$,
  - $\varphi(\overline{x})$ is computable Σ_α if it is a c.e. disjunction of formulas $\exists\overline{u}\,\psi(\overline{x},\overline{u})$, where ψ is computable Π_β for some $\beta < \alpha$,
  - $\varphi(\overline{x})$ is computable Π_α if it is a c.e. conjunction of formulas $\forall\overline{u}\,\psi(\overline{x},\overline{u})$, where ψ is computable Σ_β for some $\beta < \alpha$.

For a formula φ, in normal form, we write $neg(\varphi)$ for the dual formula, in normal form, that is logically equivalent to the negation of φ. For a discussion of computable infinitary formulas, see [1].

We fix a group language, including a binary operation symbol for the group operation, a unary operation symbol for inverse, and a constant for the identity. In this language, the axioms for groups are universal. A group G is *free* if there is a set B of elements such that B generates G and there are no non-trivial relations on elements of B. We call B a *basis* for G. If B and U are two bases for a free group G, then B and U have the same cardinality. For a free group G, the cardinality of a basis is called the *rank*. We write F_n for the free group of rank n, and F_∞ for the free group of rank $\aleph_0$. The groups F_n and F_∞ all have computable copies. If two computable structures satisfy the same computable infinitary sentences, then they are isomorphic. Thus, it is natural to look for descriptions using computable infinitary sentences.

To show that our descriptions is "optimal," we consider index sets.

Definition 1 (Computable index)**.** *A* computable index *for a structure $\mathcal{A}$ is a number e such that φ_e is the characteristic function of the atomic diagram of $\mathcal{A}$.*

Definition 2 (Index set)**.**

1. *For a structure $\mathcal{A}$, the* index set*, denoted by $I(\mathcal{A})$, is the set of computable indices for structures isomorphic to $\mathcal{A}$.*
2. *For a class K of structures, the* index set*, denoted by $I(K)$, is the set of computable indices for elements of K.*

In [2], there are results on index sets for some familiar kinds of structures, including the computable Abelian p-groups of computable lengths. These results support the thesis that for a computable structure, the complexity of the index set matches the complexity of an optimal description. If, for instance, we can describe $\mathcal{A}$, up to isomorphism, by a computable Π_3 sentence, then $I(\mathcal{A})$ is Π^0_3. If $I(\mathcal{A})$ is m-complete Π^0_3, then there can be no simpler description of $\mathcal{A}$.

Sometimes we are interested only in members of a class K, and we want to describe a particular $\mathcal{A}$ in K so as to distinguish it from other members of K, not from arbitrary structures. We define complexity of one class "within" a larger class.

Definition 3 (Complexity within a larger set)**.** *Let Γ be a complexity class (such as Π^0_3, or d-Σ^0_2) and let $A \subseteq B$.*

1. *We say that A is* Γ within B *if there is some $C \in \Gamma$ such that $A = C \cap B$.*

2. *We say that A is* Γ-hard within B *if for any set S in Γ, there is a computable function $f : \omega \to B$ such that $f(n) \in A$ iff $n \in S$.*

3. *We say that A is* m-complete Γ within B *if A is Γ within B and A is Γ-hard within B.*

For a structure $\mathcal{A}$ in a class K, where K is closed under isomorphism, we consider the complexity of $I(\mathcal{A})$ within $I(K)$. If $I(\mathcal{A})$ is Γ, or Γ-hard, within $I(K)$, we may simply say that it is Γ, or Γ-hard *within* K.

1.1 Summary of results

When we look at index sets for specific free groups within the class of free groups, we find that $I(F_1)$ is m-complete Π^0_1; $I(F_2)$ is m-complete Π^0_2; for $n > 2$, $I(F_n)$ is m-complete d-Σ^0_2; and $I(F_\infty)$ is m-complete Π^0_3. When we look at index sets for specific free groups within the class of all groups, we find that for all $n \geq 1$, $I(F_n)$ is m-complete d-Σ^0_2. For F_∞, we do not have a sharp result. We can show that $I(F_\infty)$ is Π^0_4, and it is Π^0_3-hard. These results are in Section 2. In Section 3, we consider the index sets of three classes of groups: finitely generated groups, locally free groups, and free groups. We show that for the class of finitely generated groups, the index set is m-complete Σ^0_3, and for the class of locally free groups, the index set is m-complete Π^0_2. We do not have a sharp result for the class of free groups. We can show that the index set is Π^0_4, and it is Π^0_3-hard.

When we specify a free group, we often specify a set of letters such that the group elements are the reduced words on these letters and their inverses. An automorphism of the group may take the original set of letters to another basis. Recall that a *basis* for a free group G is a generating set B with the feature that the identity cannot be expressed as a non-trivial word on elements of B. In trying to describe the different free groups, especially F_∞, we need to describe tuples that can be included in a basis. Finding formulas which describe basis elements is an old question of Mal'tsev [6], who produced a finitary formula with parameters that distinguished bases of F_2 from all other pairs. We may also ask how difficult it is to find a basis in a given computable free group. We have only partial results in this direction. We show that for any computable copy of F_∞, there is a Π^0_2 basis, but there is a computable copy of F_∞ with no infinite c.e. set that extends to a basis. These results are in Section 4. In the remainder of the introduction, we state some facts about free groups and their bases (see Lyndon and Schupp [5]).

1.2 Facts about free groups and their bases

Let U be a tuple of elements in a group G with identity e. If U is a finite tuple, say $(a_1, \ldots, a_n)$, then to denote the group generated by U we write $\langle a_1, \ldots, a_n \rangle$. Otherwise, if U consists of infinitely many elements, we write $Gp(U)$ for the group generated by U.

Definition 4. *Let $U = (a_1, \dots, a_n)$ be a tuple of elements in a group G with identity element e. The following operations on this tuple are called elementary Nielsen transformations:*

1. *for some i, replace a_i by a_i^{-1},*
2. *for some i and j, replace a_i by $(a_i)(a_j)$,*
3. *for some i such that $u_i = e$, delete u_i.*

A Nielsen transformation *is the result of a finite sequence of elementary Nielsen transformations.*

Nielsen transformations, and the following important facts about them (taken from [5]), will be used throughout this paper.

Fact 1 (2.1 of [5]). *If U is carried into V by a Nielsen transformation, then $Gp(U) = Gp(V)$.*

Definition 5. *Fixing a basis X for a free group G, let U be a set of elements, expressed as reduced words on X. Let $|u|$ be the length of u. We say that U is* N-reduced *with respect to X if for all $v_1, v_2, v_3 \in U$,*

N0 $v_1 \neq e$,

N1 $v_1 v_2 \neq e$ *implies* $|v_1 v_2| \geq |v_1|$ *and* $|v_1 v_2| \geq |v_2|$,

N2 $v_1 v_2 \neq e$ *and* $v_2 v_3 \neq e$ *implies* $|v_1 v_2 v_3| \geq |v_1| - |v_2| + |v_3|$.

Fact 2 (Proposition 2.2 of [5]). *Fix a basis X of a free group G. If $U = (u_1, u_2, \dots, u_n)$ is finite, then U can be carried by a Nielsen transformation into some V that is N-reduced with respect to X.*

Fact 3 (Proposition 2.5 of [5]). *Fix a basis X of a free group G. If U is N-reduced with respect to X, then U is a basis of $Gp(U)$.*

Definition 6. *If U is a tuple of elements of a free group, then let $U^{\pm 1}$ consist of u and u^{-1} for all $u \in U$.*

Fact 4 (Proposition 2.8 of [5]). *Let G be free with basis X and let U be N-reduced. Then $X^{\pm 1} \cap Gp(U) = X^{\pm 1} \cap U^{\pm 1}$.*

Fact 5 (Proposition 2.7 of [5]). *Let G be a free group of finite rank n. Then G cannot be generated by fewer than n elements, and if a set U of n elements generates G, then it is a basis for G (i.e., there are no non-trivial relations on the elements of U).*

Fact 6 (Proposition 2.26 of [5]). *There is an algorithm, uniform in n and $k \leq n$, to decide whether a k-tuple of words $(w_1, w_2, \dots, w_k)$ on a basis $(x_1, \dots, x_n)$ of the free group F_n is part of a basis of F_n.*

Fact 7 (Proposition 2.6 of [5]). *Every finitely generated subgroup of a free group is free of finite rank.*

Fact 8 (Proposition 2.11 of [5]). *Every subgroup of a free group is free.*

Definition 7. *For each $n \in \omega$, let $(x_1, \ldots, x_n)$ denote a basis of the free group F_n.*

1. *A k-tuple of words $(w_1, w_2, \ldots, w_k)$ on the basis $(x_1, \ldots, x_n)$ is called* primitive *if it forms part of a basis of F_n. (By Fact 5, it must be that $k \leq n$ if a k-tuple is primitive.) Otherwise, the tuple is called* non-primitive.

2. *For each n, let V_n, be the set of all primitive tuples of words on the generators $(x_1, \ldots, x_n)$ of F_n. (By Fact 6, the sets V_n are uniformly computable.)*

It is important to note that if $\overline{x}$ and $\overline{y}$ are *any* two bases of a free group F_n, then a tuple of words $(w_1(\overline{x}), \ldots, w_k(\overline{x}))$ extends to a basis iff the tuple $(w_1(\overline{y}), \ldots, w_k(\overline{y}))$ extends to a basis. Therefore, the set of primitive words should be thought of as a set of formal words on "dummy variables" rather than a set of words tied to any particular set of basis elements.

The following lemma is an easy consequence of the facts above.

Lemma 1.1. *If G is a countable free group, then for a tuple $\overline{x} = x_0, \ldots, x_n$ in G, the following are equivalent:*

1. *$\overline{x}$ is part of a basis,*

2. *for every finitely generated subgroup $H \subseteq G$ with $\overline{x}$ in H, $\overline{x}$ is part of a basis for H.*

Proof. To show $1 \Rightarrow 2$, assume that $\overline{x}$ is part of a basis X for G. Let H be a finitely generated subgroup with $\overline{x}$ in H. By Fact 7, H is free and finitely generated, with basis $(y_1, \ldots, y_k)$. Now, by Fact 2, $(y_1, \ldots, y_k)$ can be transformed, using Nielsen transformations, into N-reduced (with respect to X) set $(z_1, \ldots, z_\ell)$. By Fact 1, $(z_1, \ldots, z_\ell)$ generates the same group as $(y_1, \ldots, y_k)$, and by Fact 3, $(z_1, \ldots, z_\ell)$ is a basis of $\langle z_1, \ldots, z_\ell \rangle = \langle y_1, \ldots, y_k \rangle = H$. (So, in fact, $\ell = k$.) Then $\overline{x}$ is in $\langle z_1, \ldots, z_\ell \rangle$. By Fact 4, $\overline{x}$ is in $\{z_1, \ldots, z_\ell\}^{\pm 1}$. Then $\overline{x}$ is part of a basis for H.

To show $2 \Rightarrow 1$, let G have an infinite basis $B = \{b_0, b_1, \ldots\}$, and write $\overline{x}$ as a tuple of words over B. Assume, without loss of generality, that the first k elements are the only letters that appear in $\overline{x}$. Let $H = \langle b_0, \ldots, b_k \rangle$. Then there exists a tuple $\overline{y}$ in H so that $\overline{x} \cup \overline{y}$ is a basis for H. Then $\overline{x} \cup \overline{y} \cup \{b_{k+1}, \ldots\}$ is a basis for G. □

There is a computable sequence $(\gamma_k(\overline{x}))_{k \in \omega}$ of computable Π_2 formulas, where k is the length of the tuple $\overline{x}$. These formulas express, within free groups, Property 2 from Lemma 1.1. To express this property for a k-tuple $\overline{x}$, we need to say that there are no non-trivial relations on $\overline{x}$, and if $\overline{x}$ is in any finite subgroup generated by a tuple $\overline{y}$, then $\overline{x}$ must be generated as a set of primitive words on $\overline{y}$.

First, there is a computable sequence $(\varrho_k(\overline{x}))_{k\in\omega}$ of computable Π_1 formulas stating that there are no non-trivial relations on the k-tuple $\overline{x}$. Let R be the set consisting of non-empty reduced words on no more than n letters.

$$\varrho_k(\overline{x}) = \bigwedge\!\!\!\bigwedge_{w\in R} (w(\overline{x}) \neq e)$$

Next, recall the uniformly computable sets of primitive tuples V_n from Definition 7. Then the following formula, which we will call $\gamma_k(\overline{x})$, expresses the desired property from Lemma 1.1.

$$\varrho_k(\overline{x}) \wedge \bigwedge\!\!\!\bigwedge_{n\in\omega} \forall y_1, \ldots, y_n \left[\varrho_n(\overline{y}) \rightarrow \bigwedge\!\!\!\bigwedge_{(w_1,\ldots,w_k)\notin V_n} \neg\,(x_1 = w_1(\overline{y}) \wedge \ldots \wedge x_k = w_k(\overline{y}))\right]$$

Note that the formulas γ_k make sense in all groups, not just free ones. Note also that the only clause making this formula Π_2, rather than simply Π_1, is the antecedant of the implication, namely, $\varrho_n(\overline{y})$. We will refer to this sequence of formulas $(\gamma_k(x_1, \ldots, x_k))_{k\in\omega}$ throughout the rest of the paper.

2 Index sets for free groups

2.1 Working within the class of free groups

Let $FrGr$ be the class of free groups. Working within this class, we have the following results.

Proposition 2.1. *The set $I(F_1)$ is m-complete Π^0_1 within $FrGr$.*

Proof. We can describe F_1 within $FrGr$ by saying that it is Abelian. This implies that $I(F_1)$ is Π^0_1 within $FrGr$. For hardness, let S be a Π^0_1 set. We show that there is a uniformly computable sequence of free groups $(\mathcal{C}_n)_{n\in\omega}$ such that $\mathcal{C}_n \cong F_1$ iff $n \in S$. For each n, we enumerate the diagram of $\mathcal{C}_n$ in stages. We copy F_1 so long as we believe that $n \in S$. If we discover that $n \notin S$, then we add a second generator. □

For each $n \geq 1$, there is a natural computable Π_2 sentence φ_n saying that for any $(n+1)$-tuple of elements, there is an n-tuple that generates it. We let φ_n say that for any $x_1, \ldots, x_{n+1}$, there exist $y_1, \ldots, y_n$ such that for some $(n+1)$-tuple of words $\overline{w}$, we have $x_i = w_i(\overline{y})$. The group F_n satisfies φ_m iff $m \geq n$. The group F_∞ does not satisfy any φ_m. Throughout the rest of this paper, we will refer to these sentences φ_n for $n \in \omega$.

Proposition 2.2. *The set $I(F_2)$ is m-complete Π^0_2 within $FrGr$.*

Proof. We can describe F_2 within $FrGr$ by the conjunction of φ_2 and a finitary Σ_1 sentence saying that the group is not Abelian. The only free groups that satisfy φ_2 are F_1 and F_2, and F_1 is Abelian. It follows that $I(F_2)$ is Π^0_2. For hardness, let P be a Π^0_2 set. We show that there is a uniformly computable sequence of free groups $(\mathcal{C}_n)_{n\in\omega}$ such that $\mathcal{C}_n \cong F_2$ iff $n \in P$. When we guess that $n \in P$, then we build a group with generators a and b. If we guess that $n \notin P$, then we add a new generator c. If we later guess that $n \in P$, then we make the third generator into a word on a and b. The result of this is that if $n \notin P$, then we eventually always guess that $n \notin P$, and we get a copy of F_3. If $n \in P$, then infinitely often we guess $n \in P$, so we collapse all the extra generators, and we have a copy of F_2. □

Proposition 2.3. *For $n > 2$, $I(F_n)$ is m-complete d-Σ^0_2 within FrGr.*

Proof. Recall the computable Π_2 sentences φ_n describing the groups of rank less than or equal to n. The sentence

$$\varphi_n \ \wedge \ neg(\varphi_{n-1})$$

describes F_n, up to isomorphism, within the class $FrGr$. It follows that $I(F_n)$ is d-Σ^0_2 within $FrGr$. For hardness, let S_1 and S_2 be Σ^0_2 sets. We can produce a uniformly computable sequence of free groups $(H_n)_{n\in\omega}$ such that

$$H_n \cong \begin{cases} F_{n-1} & \text{if } n \notin S_1, \\ F_n & \text{if } n \in S_1 \ \& \ n \notin S_2, \\ F_{n+1} & \text{if } n \in S_1 \cap S_2. \end{cases}$$

We begin by building a free group with generators $(a_1, a_2, \ldots, a_{n-1})$ that we will never collapse, that is, we will never make any a_i into a word on the remaining generators. If we guess that $n \in S_1$, we add a new potential generator b. If, additionally, we guess that $n \in S_2$, we add a second new potential generator c. After this point, if we ever guess that $n \notin S_1$, we collapse both b and c by making them words on $(a_1, \ldots, a_{n-1})$. As long as we continue to guess that $n \in S_1$, we maintain b as a generator and concentrate on S_2. When we guess $n \in S_2$, we maintain c as a generator. If we guess $n \notin S_2$, we collapse c. We can then later add another potential generator if we think again that $n \in S_2$.

Now, we verify that we build the correct isomorphism types. If $n \notin S_1$, then infinitely often we guess that $n \notin S_1$ and we will collapse any potential generators we had added so that only $(a_1, a_2, \ldots, a_{n-1})$ will be true generators and $H_n \cong F_{n-1}$. If $n \in S_1 - S_2$, then since S_1 is Σ^0_2, we will eventually come to a stage after which we always guess $n \in S_1$. The final b that we add as a potential generator will never be collapsed, and therefore will be a true generator of H_n. However, for $n \notin S_2$, we will infinitely often guess $n \notin S_2$. When we guess $n \notin S_2$, we collapse any potential generator c we may have added. Then H_n will have true generators $(a_1, a_2, \ldots, a_{n-1}, b)$ and will be isomorphic to F_n. Finally, if $n \in S_1 \cap S_2$, then we will come to a stage after which we always guess that both $n \in S_1$ and $n \in S_2$. The two potential generators we add will never be collapsed and H_n will be isomorphic to F_{n+1}. □

Proposition 2.4. *The set $I(F_\infty)$ is m-complete Π^0_3 within FrGr.*

Proof. Consider the conjunction of the sentences $neg(\varphi_n)$. This is a computable Π_3 sentence that is true in F_∞ and false in F_n for any $n \in \omega$. For completeness, consider the m-complete Σ_3^0 set $\mathrm{Cof} = \{n : W_n \text{ is cofinite}\}$. We build a uniformly computable sequence of free groups $(H_n)_{n\in\omega}$ such that $H_n \cong F_\infty$ iff $n \notin \mathrm{Cof}$. To build H_n, we designate an infinite collection of potential generators, say g_e for each e. We then simultaneously begin to build our free group and enumerate W_n. Whenever we see e enter W_n, we collapse g_e as a potential generator by making it a word on the previous generators g_i for $i < e$. If $n \notin \mathrm{Cof}$, then there are infinitely many e that will never enter W_n, and we will maintain g_e as a generator for each of those values, so we will have $H_n \cong F_\infty$. If $n \in \mathrm{Cof}$, then we will collapse all but finitely many of the potential generators, so H_n will be isomorphic to F_k, where k is the cardinality of the complement of W_n. □

2.2 Working within the class of all groups

In this section, we give optimal descriptions of the groups F_n, $n \in \omega$, within the class of all groups. In each case, the "natural" or "obvious" description is not optimal, from a computability theoretic standpoint. For the free group F_n, the "obvious" definition says that there exists an n-tuple, with no non-trivial relations among its elements, so that every other group element can be written as a word on this tuple. This sentence is computable infinitary Σ_3, while we shall see that every F_n has, in fact, a d-Σ_2 definition. In our discovery of the optimal definition, the hardness argument actually led to the definition. That is, we were unable to establish Σ_3^0 hardness, and our examination of the reasons for failure suggested the correct level of complexity for the optimal definition.

Proposition 2.5. *The set $I(F_1)$ is m-complete d-Σ_2^0.*

Proof. We first show that $I(F_1)$ is d-Σ_2^0. We have a computable d-Σ_2 sentence saying the following:

1. the group is Abelian and torsion free,
2. there is a non-zero element not divisible by any prime,
3. for any pair of elements, there is a single element that generates both elements in the pair.

For this proof, the groups that we consider are Abelian, so we use additive notation. For any group satisfying the above sentence, take a non-zero element a not divisible by any prime. This must actually be a generator. For any other element b, we have c generating both. If $k \cdot c = a$, then k must be ± 1. It follows that a generates both c and b. Therefore, the group is isomorphic to $(\mathbb{Z}, +)$.

For hardness, let S_1 and S_2 be Σ_2^0 sets. In this (and the next) proof, it will be easier to give the construction by explicitly mentioning the approximations $S_{1,s}$ and $S_{2,s}$ such that $n \in S_1$ ($n \in S_2$, respectively) iff there is a stage t so that for all $s \geq t$, $n \in S_{1,s}$ ($n \in S_{2,s}$, respectively). We produce a uniformly computable sequence of Abelian groups $(H_n)_{n\in\omega}$ such that H_n will have a summand that is divisible if $n \notin S_1$, $H_n \cong \mathbb{Z}$ if $n \in S_1 - S_2$, and $H_n \cong \mathbb{Z} \oplus \mathbb{Z}$ if $n \in S_1 \cap S_2$. Recall that there are computable approximations. We start with two possible generating elements a_0 and b_0. If $n \in S_{1,s+1}$, then $a_{s+1} = a_s$, and if $n \notin S_{1,s+1}$, then a_{s+1} is new, with $2a_{s+1} = a_s$. To describe how we treat the other generator, we consider the following two cases.

Case 1. The element b_s is not expressed in terms of a_s. If $n \in S_{2,s+1}$, then we let $b_{s+1} = b_s$, and we continue not expressing b_s in terms of a_s. If $n \notin S_{2,s+1}$, then we let $b_{s+1} = b_s$, but now we express $b_s = m \cdot a_{s+1}$, where m is larger than the product of all numbers we have considered up to this point. (This ensures that in making b_{s+1} a part of the subgroup generated by a_{s+1}, we will not contradict any quantifier-free statements to which we have already committed.)

Case 2. The element b_s has been expressed in terms of a_s. If $n \notin S_{2,s+1}$, then, again, $b_{s+1} = b_s$. If $n \in S_{2,s+1}$, then b_{s+1} is new, and it is not expressed in terms of a_{s+1}. □

Definition 8. *A group is* locally free *if every finitely generated subgroup is a free group.*

There exist locally free groups that are not free. A trivial example is the Abelian group generated by $\{b_n : n \in \omega\}$, where for all n, $b_{n+1}^2 = b_n$.

Proposition 2.6. *For finite $n > 1$, $I(F_n)$ is m-complete d-Σ_2^0.*

Proof. Fix n, and recall the set V_n defined in Definition 7. Let N be the subset of V_n consisting precisely of the n-tuples of words in V_n. So an n-tuple of formal words $(w_1(\overline{z}), \ldots, w_n(\overline{z}))$ on "dummy" variables $\overline{z}$ belongs to N iff for a basis $\overline{a} = (a_1, \ldots, a_n)$ of the free group F_n, the tuple $(w_1(\overline{a}), \ldots, w_n(\overline{a}))$ is also a basis of F_n. (Recall, by the comment after Definition 7, that if this property holds for a tuple of words over one basis, then it holds for that same tuple of words over any basis.) Of course, N is computable, since V_n is computable. We can describe F_n, up to isomorphism, by the conjunction of sentences saying the following.

1. There exists an n-tuple $\overline{x}$ such that there are no non-trivial relations on $\overline{x}$, and for all n-tuples $\overline{y}$ and all n-tuples of words $\overline{w}$ not in N, it is not the case that for all $1 \leq i \leq n$, $x_i = w_i(\overline{y})$.
2. For every tuple $\overline{y}$, there exists an n-tuple $\overline{x}$ that generates $\overline{y}$.

We can take the first sentence to be computable Σ_2, and we can take the second sentence to be computable Π_2. To show that $I(F_n)$ is d-Σ_2^0, we must show that the conjunction describes F_n up to isomorphism.

First, we show that F_n satisfies the conjunction of the first and second sentence. If the n-tuple $\overline{x}$ is a basis, and some other n-tuple $\overline{y}$ generates $\overline{x}$, then, by Fact 5, the tuple $\overline{y}$ must also be a basis. Therefore, by the definition of N, the n-tuple of words in $\overline{y}$ used to express $\overline{x}$ must belong to N. Conversely, if G

is any group satisfying the first sentence, then it has a free subgroup H of rank n generated by $\overline{x}$. (We are not assuming that $\overline{x}$ is a basis of G, or even that G is free—that is what we must show.) Furthermore, if $\overline{y}$ is any n-tuple that generates $\overline{x}$, then the generating words form an n-tuple of words from N. By Fact 4, a sequence of Nielsen transformations formally converts this n-tuple of words on $\overline{y}$ into the elements $\overline{y}$. And by Fact 1, it must be true that $\overline{x}$ generates $\overline{y}$.

Note that in the argument above, we do not assume that G is a free group. Consequently, no subgroup of G generated by n elements *properly* includes H. Now, let g be any element of G. Consider the tuple $(\overline{x}, g)$. By the second sentence, this tuple is generated by an n-tuple $\overline{y}$. However, by what we just concluded, the subgroup generated by $\overline{y}$ is identical to H, so $g \in H$. Since g was arbitrary, $H = G$. That is, $G = H \cong F_n$.

If $n > 2$, the proof showing that $I(F_n)$ is m-complete d-Σ^0_2 within $FrGr$ shows hardness as well. If $n = 2$, let $S = S_1 - S_2$, where S_1 and S_2 are Σ^0_2 sets with computable approximations as in the previous proof. We produce a uniformly computable sequence $(H_n)_{n\in\omega}$ such that if $n \notin S_1$, then H_n is locally free but not free; if $n \in S_1 - S_2$, then $H_n \cong F_2$; and if $n \in S_1 \cap S_2$, then $H_n \cong F_3$. We consider possible generators a_s, b, and c_s.

When we guess that $n \notin S_1$, we replace a_s by a_{s+1}, where $a^2_{s+1} = a_s$, so a_s cannot be part of a basis of the group H_n. When we guess that $n \in S_1$, we define $a_{s+1} = a_s$, so we are guessing that it is, in fact, a genuine basis element. If infinitely often we guess that $n \notin S_1$, then the subgroup generated by $(b, a_0, a_1, \ldots, a_s, \ldots)$ is not free, so, by Fact 8, H_n is not free—it is locally free. If eventually we always guess that $n \in S_1$, then this subgroup is generated by b and a single a_s.

When we guess that $n \notin S_2$, we collapse the current c_s, making it equal to some word on a_{s+1} and b. When we guess that $n \in S_2$, after having collapsed the previous c_s, we add a new generator c_{s+1} that is not expressed as a word on a_{s+1} and b. If infinitely often we guess that $n \notin S_2$, then all of the c_s elements are included in the subgroup generated by $(b, a_0, a_1, \ldots, a_s, \ldots)$ (which may or may not be free, depending on S_1). If eventually we always guess that $n \in S_2$, then some c_s cannot be generated by $(b, a_0, a_1, \ldots, a_s, \ldots)$. Then, if $n \in S_1$, we have $H_n \cong F_3$. □

For F_∞, we can show that $I(F_\infty)$ is Π^0_4 and it is Π^0_3-hard. The hardness result follows from Proposition 2.4, in which we showed that $I(F_\infty)$ is Π^0_3 m-complete within $FrGr$.

Proposition 2.7. *The set* $I(F_\infty)$ *is* Π^0_4.

Proof. Recall the sequence of formulas $(\gamma_k(x_1, \ldots, x_k))_{k\in\omega}$, which express, within a free group, that $\overline{x}$ is part of a basis. Now, to describe F_∞, we say the following.

1. There exists x_1 satisfying γ_1.

2. For each k, each $(x_1, \ldots, x_k)$ satisfying γ_k, and each y, there exist $\ell \geq k+1$ and $(x_{k+1}, \ldots, x_\ell)$ such that $(x_1, \ldots, x_k, x_{k+1}, \ldots, x_\ell)$ satisfies γ_ℓ and $y \in \langle x_1, \ldots, x_\ell \rangle$.

This description is computable Π_4, since the γ_k are uniformly Π_2. It is easy to see that F_∞ satisfies this sentence, since it has an infinite basis $(a_1, a_2, \ldots, a_n, \ldots)$. Let H be any other group (not assumed to be free) which satisfies this sentence. A straightforward back-and-forth argument shows that $H \cong F_\infty$. $\square$

Recall that each formula $\gamma_k(\overline{x})$ is Π_2 only because we have to say that for $\overline{y}$ *with no trivial relations*, it is not the case that $\overline{x}$ can be expressed as a non-primitive k-tuple of words on $\overline{y}$. Is it the case that some k-tuple of genuine basis elements in F_∞ can be written as a non-primitive k-tuple of words on some n-tuple $\overline{y}$ that *does have* a non-trivial relation? It turns out that the answer is positive ([3]). Therefore, there is no obvious way to simplify the formulas γ_k to Π_1 formulas in order to give a computable Π_3 description of F_∞.

Finally, note that if we would try to modify the sentence in the above proof by replacing the formulas γ_k with simpler formulas stating only that the tuple $(x_1, \ldots, x_k)$ has no non-trivial relations, then this sentence is no longer true in F_∞. Indeed, if a_1 is a basis element, then $(a_1)^2$, as a singleton, has no non-trivial relations; however, given the element a_1, there is no way to extend the singleton $(a_1)^2$ to an ℓ-tuple that generates a_1 and also has no non-trivial relations.

3 Index sets for some classes of groups

Let *FinGen* denote the class of all finitely generated groups. Based on the results above, we can quickly establish the complexity of $I(FinGen)$ within the class of free groups and within the class of all groups.

Proposition 3.1. *The set $I(FinGen)$ is m-complete Σ^0_3 within the class of free groups.*

Proof. Recall the computable Π_2 sentences φ_n saying that any $(n+1)$-tuple is generated by an n-tuple. Let ψ be the disjunction of these sentences. This is a computable Σ_3 sentence, and among free groups, it is satisfied exactly by those that are finitely generated. For completeness, recall that in the proof of Proposition 2.4, we defined a uniformly computable sequence of free groups $(H_n)_{n \in \omega}$ such that $n \in \text{Cof}$ if and only if H_n is finitely generated. $\square$

Proposition 3.2. *The set $I(FinGen)$ is m-complete Σ^0_3 within the class of all groups.*

Proof. We have a computable Σ_3 sentence saying that for some n, there is an n-tuple $\overline{x}$ such that for every element y, we can express y as a word $w(\overline{x})$. This sentence characterizes the finitely generated groups within the class of all groups. Again, the proof of Proposition 2.4 establishes completeness. $\square$

Let *LocFr* denote the class of all locally free groups.

Proposition 3.3. *The set $I(LocFr)$ is m-complete Π_2^0 within the class of all groups.*

Proof. Consider the computable Π_2 sentence saying that the group is torsion free, and that for any $n \in \omega$ and any n-tuple $\overline{y}$, if $\overline{y}$ has a non-trivial relation, then there is an $(n-1)$-tuple $\overline{x}$ so that $\overline{x}$ generates $\overline{y}$. We claim that a group G is a locally free group iff it satisfies this sentence.

($\Rightarrow$) Let $\overline{y}$ be an n-tuple in G. By definition, the subgroup generated by $\overline{y}$ is free, and, by Fact 5, it has a basis $(x_1, \ldots, x_m)$, where $m \leq n$. If $m < n$, then $\overline{y}$ is generated by fewer than n elements, and hence by $(n-1)$ elements. If $m = n$, then $\overline{y}$ generates a free group of rank n, so, by Fact 5, $\overline{y}$ is a basis, and hence has no non-trivial relations.

($\Leftarrow$) Let H be a non-trivial, finitely generated subgroup of G, generated by an n-tuple $(y_1, \ldots, y_n)$. If this tuple has no non-trivial relations, then it is a basis for H. Otherwise, there is an $(n-1)$-tuple $(x_1, \ldots, x_{n-1})$ generating H. If this $(n-1)$-tuple has no non-trivial relations, then it is a basis for H. Otherwise, we continue in this fashion until we either reach a k-tuple that is a basis for H, or we come down to a single element g that generates H. In the latter case, since G is torsion free, $H \cong F_1$.

To show hardness, we use the "S_2 half" of the hardness argument from Proposition 2.5. That is, let P be a Π_2^0 set. We construct a computable sequence $(H_n)_{n\in\omega}$ of groups so that if $n \in P$, then $H_n \cong \mathbb{Z}$ and if $n \notin P$, then $H_n \cong \mathbb{Z}\oplus\mathbb{Z}$. As usual, P has a computable approximation P_s so that for all n, we have $n \notin P$ iff there is some s so that for all $t \geq s$, $n \notin P_t$. We construct H_n as an Abelian group with two possible generators a and b_s. At stage s, if $n \notin P_s$, then we let $b_{s+1} = b_s$, and we continue not to express b_s as any multiple of a. Assume that $n \in P_s$. Then we declare $b_s = m \cdot a$, where m is greater than the product of all numbers we have considered up to this point. We define b_{s+1} to be a new number not expressed as a multiple of a. □

We now give a hardness result on the class of free groups (within the class of all groups).

Proposition 3.4. *The set $I(FrGr)$ is Π_3^0-hard within the class of groups.*

Proof. First, we show that $I(FrGr)$ is Σ_2^0-hard. Consider the m-complete Σ_2^0 set $\text{Fin} = \{n : W_n \text{ is finite}\}$. We produce a uniformly computable sequence of groups $(H_n)_{n\in\omega}$ so that H_n is free iff $n \in \text{Fin}$. Each H_n has $\mathbb{N}$ as its universe. Divide $\mathbb{N}$ into columns $(\mathbb{N}_i)_{i\in\omega}$. We begin constructing H_n so that, if we never see an element enumerated into W_n, then H_n will be isomorphic to F_∞, and the column $\mathbb{N}_0$ will be a set of independent generators. If we see a first element enumerated into W_n, then we choose two new elements a, b from the column $\mathbb{N}_j$, where j is the least number so that none of the elements of $\mathbb{N}_k$, $k \geq j$, have been mentioned so far in our enumeration of the atomic diagram of H_n. We then guarantee that the construction of H_n will continue as follows.

1. The group $Gp(\mathbb{N}_0)$ has universe $\bigcup_{k<j} \mathbb{N}_k$, and the column $\mathbb{N}_0$ freely generates this group.

2. All of the elements of the first column are words on the elements a, b. (This is possible because the group F_2 contains an isomorphic copy of F_∞ as a subgroup, with the elements of the subgroup isomorphic to F_∞ given as a computable set of words on a pair of generators of F_2.)

We then continue to construct H_n so that, if no other elements enter into W_n, then H_n will be isomorphic to F_∞, and the column $\mathbb{N}_j$ will be a set of independent generators. If we see a new element enumerated into W_n, then we proceed as above. It is evident that if W_n is finite, then H_n is indeed isomorphic to F_∞. However, if W_n is infinite, then this infinite "nesting" prevents any single element from being part of a basis.

Next, knowing that $I(FrGr)$ is Σ^0_2-hard, we show that it is actually Π^0_3-hard. Let P be a given Π^0_3 set, so that there is a Σ^0_2 binary relation $S(n,m)$ satisfying $n \in P \Leftrightarrow \forall m(S(n,m))$. Since $I(FrGr)$ is Σ^0_2-hard, there is a uniformly computable sequence of groups $(H_{n,m})_{n,m\in\omega}$ so that $H_{n,m}$ is free iff $S(n,m)$. Let $(B_n)_{n\in\omega}$ be the uniformly computable sequence of groups defined such that B_n is the free product of the groups $\{H_{n,m} : m \in \omega\}$. If $n \in P$, then $\forall m(S(n,m))$, so for all m, $H_{n,m}$ is free. That is, B_n is the free product of free groups, so B_n is free. Now, assume that $n \notin P$. Then $\exists m(\neg S(n,m))$, so there is m for which $H_{n,m}$ is not free. That is, B_n is a free product of groups, and one of the factors is not free, so B_n is not free. Consequently, $n \in P$ iff B_n is free. Therefore, the index set $I(FrGr)$ is Π^0_3-hard. □

4 Bases for free groups

We first consider bases for free groups of infinite rank. There is an old result of Metakides and Nerode in [7] on $\mathbb{Q}$-vector spaces, saying that there is a computable vector space of infinite dimension with no infinite c.e. linearly independent set. We have the analogous result in the setting of free groups.

Proposition 4.1. *There is a computable copy of F_∞ so that no infinite c.e. set can be extended to a basis.*

Proof. We build a computable group $G \cong F_\infty$ satisfying the following requirements.

R_e**:** If W_e is infinite, then W_e is not part of a basis.

We describe the strategy for a single requirement. Having enumerated finitely much of the diagram of G, and decided to put $b_1, \ldots, b_e$ into the basis of G, we continue adding basis elements $b_{e+1}, \ldots, b_r$, extending the diagram of G, and watching W_e until we see at least $e+1$ elements appear, all words on $b_1, \ldots, b_r$. (For if an element appeared in W_e that had not yet been defined as a word on $b_1, \ldots, b_r$, then we could simply define it to be a non-trivial power of b_1, thus ensuring that it was not a basis element.) At this point, we make all of the elements $b_{e+1}, \ldots, b_r$ into words on $b_1, \ldots, b_e$ by employing very large words that we have not yet considered in our construction of the diagram of G. This implies that the $e+1$ elements we have seen enumerated in W_e cannot be part of a basis. □

Proposition 4.2. *Every computable copy of F_∞ has a Π_2^0 basis.*

Proof. Let G be a computable copy of F_∞, and assume G has universe $\mathbb{N}$. First, we will show that Δ_3^0 can enumerate a basis in increasing order—hence, G has a Δ_3^0 basis. Recall that there is a computable sequence of computable Π_2 formulas $\gamma_k(x_1,\ldots,x_k)$ so that for a k-tuple $(a_1,\ldots,a_k) \in G$, $(a_1,\ldots,a_k)$ is part of a basis iff $G \models \gamma_k(a_1,\ldots,a_k)$. Using Δ_3^0, we search for the first (according to the ordering on $\mathbb{N}$) b_0 such that $G \models \gamma_1(b_0)$. Once we have found this b_0, we search, using Δ_3^0, for the first b_1 such that $G \models \gamma_2(b_0,b_1)$. (Notice that $b_0 < b_1$, since $G \models \gamma_1(b_1)$.) We continue in this way. Let this Δ_3^0 basis be $B = \{b_0, b_1, \ldots\}$.

Now we use this Δ_3^0 basis B to produce a Π_2^0 basis U. We give a Δ_2^0 enumeration of the complement $\overline{U}$. We use the fact that given (x,y) a basis for a free group of rank 2, we can apply Nielsen transformations to obtain infinitely many further bases, all disjoint. Starting with (x,y), we get (xy,y) and then (xy,xy^2), disjoint from (x,y). Each new basis is obtained from the current one by two steps just like these.

Relativizing the Limit Lemma, we obtain a binary Δ_2^0 function f, $f(i,s) = b_{i,s}$, so that for every $i \in \mathbb{N}$, $\lim_s b_{i,s} = b_i$. Moreover, we can assume that at any stage s, $b_{0,s} < b_{1,s} < \ldots < b_{s,s}$, and all these elements have no non-trivial relations among them, because otherwise, Δ_2^0 would see this aberration, and would keep re-approximating up to the $(s+1)$-st element.

The idea of the construction is as follows. To enumerate the complement of U, Δ_2^0 guesses the first few pairs from B, and enumerates elements into the complement of U based on the current guess. If the guess changes, and some pair has been put into the complement by mistake, Δ_2^0 chooses a new equivalent pair to preserve in U. Here is the formal construction.

Stage 0. Compute $b_{0,0}$ and $b_{1,0}$. Let $c_{0,0} = b_{0,0}$ and $c_{1,0} = b_{1,0}$. Enumerate into $\overline{U}$ all elements smaller than $c_{0,0}$, and all elements between $c_{0,0}$ and $c_{1,0}$. Declare $c_{k,0}$ undefined for all $k > 1$.

Stage $s+1$. Assume that at stage s, we have enumerated only finitely many elements into $\overline{U}$, so U_s is cofinite. Compute the elements $b_{0,s+1}$ and $b_{1,s+1}$.

Step 1.

a) If $b_{0,s+1} = b_{0,s}$ and $b_{1,s+1} = b_{1,s}$, then let $c_{0,s+1} = c_{0,s}$ and $c_{1,s+1} = c_{0,s}$. Proceed to Step 2.

b) If $b_{0,s+1} \neq b_{0,s}$ or $b_{1,s+1} \neq b_{1,s}$, and $b_{0,s+1}$ and $b_{1,s+1}$ belong to U_s, then enumerate all elements smaller than $b_{0,s+1}$ and all elements between $b_{0,s+1}$ and $b_{1,s+1}$ into $\overline{U}$. Define $c_{0,s+1} = b_{0,s+1}$, define $c_{1,s+1} = b_{1,s+1}$, and declare $c_{k,s+1}$ undefined for all $k > 1$. Otherwise, if it is not the case that both $b_{0,s+1}$ and $b_{1,s+1}$ belong to U_s, then systematically apply Nielsen transformations to the pair $(b_{0,s+1}, b_{1,s+1})$ until a new pair $(c_{0,s+1}, c_{1,s+1})$ is produced so that both $c_{0,s+1}$ and $c_{1,s+1}$ belong to U_s, and $c_{0,s+1} < c_{1,s+1}$ in the usual ordering on $\mathbb{N}$. (This can be done because Nielsen transformations on independent elements can produce arbitrarily long words on these elements, and the set U_s is cofinite.) Enumerate into $\overline{U}$ all elements smaller than $c_{0,s+1}$ and all elements between $c_{0,s+1}$ and $c_{1,s+1}$. Declare $c_{k,s+1}$ undefined for all $k > 1$. Proceed to stage $s+2$.

Step 2. Compute $b_{2,s+1}, \ldots, b_{s+1,s+1}$. Let j be the first number so that $2 \leq j \leq s+1$ and either $b_{j,s+1} \neq b_{j,s}$ or $c_{j,s}$ is undefined. For all k so that $2 \leq k < j$, let $c_{k,s+1} = c_{k,s}$.

To complete Stage $s+1$, we find the least p so that $(c_{0,s+1})^p \cdot b_{j,s+1}$ belongs to U_s, and $(c_{0,s+1})^p \cdot b_{j,s+1}$ is not equal to $c_{k,s+1}$ for any $0 \leq k < j$. Call this element $c_{j,s+1}$. Enumerate all elements between $c_{j-1,s+1}$ and $c_{j,s+1}$ into $\overline{U}$. Declare $c_{k,s+1}$ undefined for all $k > j$. This completes Stage $s+1$.

We have described the whole construction. Given $i \in \omega$, there is a stage s so that for all $t \geq s$, we have $b_{0,t} = b_0, \ldots, b_{i,t} = b_i$. Therefore, by the construction, for each i, $\lim_{s\in\omega}(c_{i,s}) = c_i$. Moreover, the sequence of elements $(c_i)_{i\in\omega}$ has the following important properties.

1. The set $\{c_0, c_1, b_2, b_3, \ldots\}$ is a basis of G, because c_0, c_1 is derived from b_0, b_1 by Nielsen transformations.

2. For each $i \geq 2$, there is a k so that $c_i = (c_0)^k \cdot b_i$.

It can then be easily shown that the set $C = \{c_i : i \in \omega\}$ is a basis of G. By construction, C is Π^0_1 relative to Δ^0_2, and hence C is Π^0_2. □

Proposition 4.3. *For every $n \geq 2$, there is a computable Π_1 formula $\theta_n(x_1, \ldots, x_n)$ saying, in F_n, that $\{x_1, \ldots, x_n\}$ is a basis.*

Proof. The formula $\theta_n(x_1, \ldots, x_n)$ says that for all $(y_1, \ldots, y_n)$ and all non-primitive n-tuples of words $w_1(\overline{y}), \ldots, w_n(\overline{y})$, it is not the case that $x_i = w_i(\overline{y})$ for $i = 1, \ldots, n$. Suppose that $\overline{x}$ is a basis. If $\overline{y}$ is a basis, then $\overline{x}$ is not obtained by a non-primitive tuple of words. If $\overline{y}$ is not a basis, then, by Fact 5, $\overline{x}$ is not obtained by any tuple of words. Therefore, the formula is satisfied. Suppose that $\overline{x}$ is not a basis. Let $\overline{y}$ be a basis. Then $\overline{x}$ is obtained using a non-primitive tuple of words, so the formula is not satisfied. □

Recall that for $k < n$, the Π_2 formula $\gamma_k(x_1, \ldots, x_k)$ says that $(x_1, \ldots, x_k)$ forms part of a basis for F_n. By the previous proposition, there is also a computable Σ_2 formula expressing this; namely, $\exists z_{k+1} \ldots \exists z_n[\theta(x_1, \ldots, x_k, z_{k+1}, z_n)]$.

The problem of identifying the tuples that form part of a basis is interesting from several points of view. Sela showed that the common theory of the non-Abelian free groups is unstable. Pillay [8], has investigated the "generic" type, and showed that it has infinite "weight". Our proof of Proposition 4.2, using

Nielsen transformations, shows that for a given basis B for F_∞, we can produce infinitely many more bases, all disjoint. Pillay's proof uses forking calculus to do the same thing.

References

[1] C. J. Ash and J. F. Knight, *Computable Structures and the Hyperarithmetical Hierarchy*, Elsevier, 2000.

[2] W. Calvert, V. S. Harizanov, J. F. Knight, and S. Miller, "Index sets for computable structures", *Algebra and Logic*, vol. 45 (2006), pp. 306–325.

[3] D. Grove and M. Culler, personal correspondence.

[4] O. Kharlampovich and A. Myasnikov, "Elementary theory of free nonabelian groups", *Journal of Algebra*, vol. 302 (2006), pp. 451–552.

[5] R. C. Lyndon and P. E. Schupp, *Combinatorial Group Theory*, Springer, 2001.

[6] A. I. Mal'tsev, "On the equation $zxyx^{-1}y^{-1}z^{-1} = aba^{-1}b^{-1}$ in a free group", (Russian) *Algebra i Logika* Sem. 1, no. 5 (1962), pp. 45–50.

[7] G. Metakides and A. Nerode, "Effective content of field theory", *Annals of Math. Logic*, vol. 17 (1979), pp. 289–320.

[8] A. Pillay, "On genericity and weight in the free group", pre-print.

[9] D. Scott, "Logic with denumerably long formulas and finite strings of quantifiers", in *The Theory of Models*, ed. by J. Addison, L. Henkin, and A. Tarski, North-Holland, 1965, pp. 329–341.

[10] Z. Sela, "Diophantine geometry over groups I: Makanin-Razborov diagrams", *Publications Mathématiques*, Institute des Hautes Études Scientifiques, vol. 93 (2001), pp. 31–105.

[11] Z. Sela, "Diophantine geometry over groups II: Completions, closures, and formal solutions", *Israel J. of Math.*, vol. 134 (2003), pp. 173–254.

[12] Z. Sela, "Diophantine geometry over groups III: Rigid and solid solutions", *Israel J. of Math.*, vol. 147 (2005), pp. 1–73.

[13] Z. Sela, "Diophantine geometry over groups IV: An iterative procedure for validation of a sentence", *Israel J. of Math.*, vol. 143 (2004), pp. 1–130.

[14] Z. Sela, "Diophantine geometry over groups V_1: Quantifier elimination I", *Israel J. of Math.*, vol. 150 (2005), pp. 1–197.

[15] Z. Sela, "Diophantine geometry over groups V_2: Quantifier elimination II", *Geometric and Functional Analysis*, vol. 16 (2006), pp. 537–706.

[16] Z. Sela, "Diophantine geometry over groups VI: The elementary theory of a free group", *Geometric and Functional Analysis*, vol. 16 (2006), pp. 707–730.