

THE COMPLEXITY OF KNAPSACK IN GRAPH GROUPS

MARKUS LOHREY AND GEORG ZETZSCHE

ABSTRACT. Myasnikov et al. have introduced the knapsack problem for arbitrary finitely generated groups. In [17] the authors proved that for each graph group, the knapsack problem can be solved in NP. Here, we determine the exact complexity of the problem for every graph group. While the problem is TC^0 -complete for complete graphs, it is LogCFL -complete for each (non-complete) transitive forest. For every remaining graph, the problem is NP-complete.

1. INTRODUCTION

In their paper [19], Myasnikov, Nikolaev, and Ushakov started the investigation of classical discrete optimization problems, which are classically formulated over the integers, for arbitrary (possibly non-commutative) groups. The general goal of this line of research is to study to what extent results from the classical commutative setting can be transferred to the non-commutative setting. Among other problems, Myasnikov et al. introduced for a finitely generated group G the *knapsack problem* and the *subset sum problem*. The input for the knapsack problem is a sequence of group elements $g_1, \dots, g_k, g \in G$ (specified by finite words over the generators of G) and it is asked whether there exists a solution $(x_1, \dots, x_k) \in \mathbb{N}^k$ of the equation $g_1^{x_1} \cdots g_k^{x_k} = g$. For the subset sum problem one restricts the solution to $\{0, 1\}^k$.

For the particular case $G = \mathbb{Z}$ (where the additive notation $x_1 \cdot g_1 + \cdots + x_k \cdot g_k = g$ is usually preferred) these problems are NP-complete if the numbers g_1, \dots, g_k, g are encoded in binary representation. For subset sum, this is a classical result from Karp's seminal paper [14] on NP-completeness. Knapsack for integers is usually formulated in a more general form in the literature; NP-completeness of the above form (for binary encoded integers) was shown in [10], where the problem was called MULTISUBSET SUM .¹ Interestingly, if we consider subset sum for the group $G = \mathbb{Z}$, but encode the input numbers g_1, \dots, g_k, g in unary notation, then the problem is in $\text{DLOGTIME-uniform TC}^0$ (a small subclass of polynomial time and even of logarithmic space that captures the complexity of multiplication of binary encoded numbers; see e.g. the book [23] for more details) [5], and the same holds for knapsack (see Theorem 4.1). Related results can be found in [12].

In [19] the authors encode elements of the finitely generated group G by words over the group generators and their inverses, which corresponds to the unary encoding of integers. Another, more succinct encoding of group elements uses *straight-line programs* (SLP). These are context-free grammars that produce a single word. Over a unary alphabet, one can achieve for every word exponential compression with

The second author is supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD).

¹Note that if we ask for a solution (x_1, \dots, x_k) in \mathbb{Z}^k , then knapsack can be solved in polynomial time (even for binary encoded integers) by checking whether $\text{gcd}(g_1, \dots, g_k)$ divides g .

SLPs: The word a^n can be produced by an SLP of size $O(\log n)$. This shows that knapsack and subset sum for the group \mathbb{Z} with SLP-compressed group elements correspond to the classical knapsack and subset sum problem with binary encoded numbers. To distinguish between the two variants, we will speak in this introduction of uncompressed knapsack (resp., subset sum) if the input group elements are given explicitly by words over the generators. On the other hand, if these words are represented by SLPs, we will speak of SLP-compressed knapsack (resp., subset sum). Later in this paper, we will only use the uncompressed versions, and denote these simply with knapsack and subset sum.

In our recent paper [17], we started to investigate knapsack and subset sum for graph groups, which are also known as right-angled Artin groups in group theory. A graph group is specified by a finite simple graph Γ and denoted with $\mathbb{G}(\Gamma)$. The vertices are the generators of the group, and two generators a and b are allowed to commute if and only if a and b are adjacent in Γ . Graph groups interpolate between free groups and free abelian groups and can be seen as a group counterpart of trace monoids (free partially commutative monoids), which have been used for the specification of concurrent behavior. In combinatorial group theory, graph groups are currently an active area of research, mainly because of their rich subgroup structure, see e.g. [3, 4, 7].

Contribution. In [17] the authors proved that for every graph group, SLP-compressed knapsack (resp., subset sum) is NP-complete. This result generalizes the classical result for knapsack with binary encoded integers. Moreover, we proved that uncompressed knapsack and subset sum are NP-complete for the group $F_2 \times F_2$ (F_2 is the free group on two generators). The group $F_2 \times F_2$ is the graph group $\mathbb{G}(\Gamma)$, where the graph Γ is a cycle on four nodes. This result leaves open the complexity of uncompressed knapsack (and subset sum) for graph groups, whose underlying graph does not contain an induced cycle on four nodes. In this paper, we completely settle this open problem by showing the following results:

- (i) Uncompressed knapsack and subset sum for $\mathbb{G}(\Gamma)$ are complete for TC^0 if Γ is a complete graph (and thus $\mathbb{G}(\Gamma)$ is a free abelian group).²
- (ii) Uncompressed knapsack and subset sum for $\mathbb{G}(\Gamma)$ are LogCFL-complete if Γ is not a complete graph and neither contains an induced cycle on four nodes (C4) nor an induced path on four nodes (P4).
- (iii) Uncompressed knapsack for $\mathbb{G}(\Gamma)$ is NP-complete if Γ contains an induced C4 or an induced P4.

Overview of the proofs. The result (i) is a straightforward extension of the corresponding result for \mathbb{Z} [5]. The statements in (ii) and (iii) are less obvious. Recall that LogCFL is the closure of the context-free languages under logspace reductions; it is contained in NC^2 .

To show the upper bound in (ii), we use the fact that the graph groups $\mathbb{G}(\Gamma)$, where Γ neither contains an induced C4 nor an induced P4 (these graphs are the so called transitive forests), are exactly those groups that can be built up from \mathbb{Z} using the operations of free product and direct product with \mathbb{Z} . We then construct inductively over these operations a logspace-bounded auxiliary pushdown automaton working in polynomial time (these machines accept exactly the languages in LogCFL) that checks whether an acyclic finite automaton accepts a word that is trivial in the

²In the following, TC^0 always refers to its DLOGTIME-uniform version.

graph group. In order to apply this result to knapsack, we finally show that every solvable knapsack instance over a graph group $\mathbb{G}(\Gamma)$ with Γ a transitive forest has a solution with polynomially bounded exponents. This result might be of independent interest.

For the lower bound in (ii), it suffices to consider the group F_2 (the free group on two generators). Our proof is based on the fact that the context-free languages are exactly those languages that can be accepted by valence automata over F_2 . This is a reinterpretation of the classical theorem of Chomsky and Schützenberger. To the authors' knowledge, the result (ii) is the first completeness result for LogCFL in the area of combinatorial group theory.

Finally, for the result (iii) it suffices to show NP-hardness of knapsack for the graph group $\mathbb{G}(\mathbf{P4})$ (the NP upper bound and the lower bound for C4 is shown in [17]). We apply a technique that was first used in a paper by Aalbersberg and Hoogeboom [1] to show that the intersection non-emptiness problem for regular trace languages is undecidable for P4.

2. KNAPSACK AND EXPONENT EQUATIONS

We assume that the reader has some basic knowledge concerning (finitely generated) groups (see e.g. [18] for further details). Let G be a finitely generated group, and let A be a finite generating set for G . Then, elements of G can be represented by finite words over the alphabet $A^{\pm 1} = A \cup A^{-1}$. An *exponent equation* over G is an equation of the form

$$h_0 g_1^{x_1} h_1 g_2^{x_2} h_2 \cdots g_k^{x_k} h_k = 1$$

where $g_1, g_2, \dots, g_k, h_0, h_1, \dots, h_k \in G$ are group elements that are given by finite words over the alphabet $A^{\pm 1}$ and x_1, x_2, \dots, x_k are not necessarily distinct variables. Such an exponent equation is *solvable* if there exists a mapping $\sigma: \{x_1, \dots, x_k\} \rightarrow \mathbb{N}$ such that $h_0 g_1^{\sigma(x_1)} h_1 g_1^{\sigma(x_2)} h_2 \cdots g_k^{\sigma(x_k)} h_k = 1$ in the group G . The *size* of an equation is $\sum_{i=0}^k |h_i| + \sum_{i=1}^k |g_i|$, where $|g|$ denotes the length of the shortest word $w \in (A^{\pm 1})^*$ representing g . *Solvability of exponent equations over G* is the following computational problem:

Input: An exponent equation E over G (where elements of G are specified by words over the group generators and their inverses).

Question: Is E solvable?

The *knapsack problem* for the group G is the restriction of solvability of exponent equations over G to exponent equations of the form $g_1^{x_1} g_2^{x_2} \cdots g_k^{x_k} g^{-1} = 1$ or, equivalently, $g_1^{x_1} g_2^{x_2} \cdots g_k^{x_k} = g$ where the exponent variables x_1, \dots, x_k have to be pairwise different. The *subset sum problem* for the group G is defined in the same way as the knapsack problem, but the exponent variables x_1, \dots, x_k have to take values in $\{0, 1\}$.

It is a simple observation that the decidability and complexity of solvability of exponent equations over G as well as the knapsack problem and subset sum problem for G does not depend on the chosen finite generating set for the group G . Therefore, we do not have to mention the generating set explicitly in these problems.

Remark 2.1. *Since we are dealing with a group, one might also allow solution mappings $\sigma: \{x_1, \dots, x_k\} \rightarrow \mathbb{Z}$ to the integers. This variant of solvability of (compressed) exponent equations (knapsack, respectively) can be reduced to the above*




FIGURE 1. P4 and C4

version, where σ maps to \mathbb{N} , by simply replacing a power $g_i^{x_i}$ by $g_i^{x_i}(g_i^{-1})^{y_i}$, where y_i is a fresh variable.

3. TRACES AND GRAPH GROUPS

Let (A, I) be a finite simple graph. In other words, the edge relation $I \subseteq A \times A$ is irreflexive and symmetric. It is also called the *independence relation*, and (A, I) is called an *independence alphabet*. Symbols $a, b \in A$ are called dependent if $(a, b) \in I$. We consider the monoid $\mathbb{M}(A, I) = A^*/\equiv_I$, where \equiv_I is the smallest congruence relation on the free monoid A^* that contains all pairs (ab, ba) with $a, b \in A$ and $(a, b) \in I$. This monoid is called a *trace monoid* or *partially commutative free monoid*. Elements of $\mathbb{M}(A, I)$ are called *Mazurkiewicz traces* or simply *traces*. The trace represented by the word u is denoted by $[u]_i$, or simply $[u]$ if no confusion can arise. The empty trace $[\varepsilon]_I$ is the identity element of the monoid $\mathbb{M}(A, I)$ and is denoted by 1. For a language $L \subseteq A^*$ we denote with $[L]_I = \{[u]_I \in \mathbb{M}(A, I) \mid u \in L\}$ the set of traces represented by L . The length of the trace $[u]_I$ is $|[u]_I| = |u|$.

Figure 1 shows two important independence alphabets that we denote with P4 (path on four nodes) and C4 (cycle on four nodes). Note that $\mathbb{M}(\text{C4}) = \{a, c\}^* \times \{b, d\}^*$.

With an independence alphabet (A, I) we associate the finitely presented group

$$\mathbb{G}(A, I) = \langle A \mid ab = ba \ ((a, b) \in I) \rangle.$$

More explicitly, this group can be defined as follows: Let $A^{-1} = \{a^{-1} \mid a \in A\}$ be a disjoint copy of the alphabet A . We extend the independence relation I to $A^{\pm 1} = A \cup A^{-1}$ by $(a^x, b^y) \in I$ for all $(a, b) \in I$ and $x, y \in \{-1, 1\}$. Then $\mathbb{G}(A, I)$ is the quotient monoid $(A^{\pm 1})^*/\sim_I$, where \sim_I is the smallest congruence relation that contains (i) all pairs (ab, ba) for $a, b \in A^{\pm 1}$ with $(a, b) \in I$ and (ii) all pairs (aa^{-1}, ε) and $(a^{-1}a, \varepsilon)$ for $a \in A$.

A group $\mathbb{G}(A, I)$ is called a *graph group*, or *right-angled Artin group*³, or *free partially commutative group*. Here, we use the term graph group. Graph groups received a lot of attention in group theory during the last years, mainly due to their rich subgroup structure [3, 4, 7], and their relationship to low dimensional topology (via so-called virtually special groups) [2, 11, 24].

4. TC⁰- AND LogCFL-COMPLETENESS

Let us first consider free abelian groups \mathbb{Z}^m . Note that \mathbb{Z}^m is isomorphic to the graph group $\mathbb{G}(A, I)$ where (A, I) is the complete graph on m nodes. Our first result is a simple combination of known results [20, 5].

³This term comes from the fact that right-angled Artin groups are exactly the Artin groups corresponding to right-angled Coxeter groups.

Theorem 4.1. *For every fixed $m \geq 1$, knapsack and subset sum for the free abelian group \mathbb{Z}^m are complete for TC^0 . Hence, knapsack and subset sum for $\mathbb{G}(A, I)$ are complete for TC^0 if (A, I) is a non-empty complete graph.*

Proof. Hardness for TC^0 follows from the fact that the word problem for \mathbb{Z} is TC^0 -complete: this is exactly the problem of checking whether for a given word $w \in \{a, b\}^*$, $|w|_a = |w|_b$ holds.

Let us now show that knapsack for \mathbb{Z}^m belongs to TC^0 . Let $A = \{a_1, \dots, a_m\}$ be the generating set for \mathbb{Z}^m . Given a word $w \in (A \cup A^{-1})^*$ we can compute the vector $(b_1, \dots, b_m) \in \mathbb{Z}^m$ with $b_i := |w|_{a_i} - |w|_{a_i^{-1}}$ represented in unary notation in TC^0 (counting the number of occurrences of a symbol in a string and subtraction can be done in TC^0). Hence, we can transform in TC^0 an instance of knapsack for \mathbb{Z}^m into a system of equations $Ax = b$, where $A \in \mathbb{Z}^{m \times n}$ is an integer matrix with unary encoded entries, $b \in \mathbb{Z}^m$ is an integer vector with unary encoded entries, and x is a vector of n variables ranging over \mathbb{N} . Let $t = n(ma)^{2m+1}$, where a is the maximal absolute value of an entry in $(A \mid b)$. By [20] the system $Ax = b$ has a solution if and only if it has a solution with all entries of x from the interval $[0, t]$. Since m is a constant, the unary encoding of the number t can be computed in TC^0 (iterated multiplication can be done in TC^0). However, the question whether the system $Ax = b$ has a solution from $[0, t]^n$ is an instance of the m -integer-linear-programming problem from [5], which was shown to be in TC^0 in [5]. For subset sum for \mathbb{Z}^m one can use the same argument with $t = 1$. \square

We now characterize those graph groups where knapsack for $\mathbb{G}(A, I)$ is LogCFL -complete. The class LogCFL consists of all problems that are logspace reducible to a context-free language. The *comparability graph* of a tree t is the simple graph with the same vertices as t , but has an edge between two vertices whenever one is a descendent of the other in t . A graph (A, I) is a *transitive forest* if it is a disjoint union of comparability graphs of trees.

Theorem 4.2. *If (A, I) is a transitive forest and not complete, then knapsack and subset sum for $\mathbb{G}(A, I)$ are LogCFL -complete.*

If the graph (A, I) is the disjoint union of graphs Γ_0 and Γ_1 , then by definition, we have $\mathbb{G}(A, I) \cong \mathbb{G}(\Gamma_0) * \mathbb{G}(\Gamma_1)$. If one vertex v of (A, I) is adjacent to every other vertex and removing v from (A, I) results in the graph Γ_0 , then $\mathbb{G}(A, I) \cong \mathbb{G}(\Gamma_0) \times \mathbb{Z}$. Therefore, we have the following *inductive characterization* of the graph groups $\mathbb{G}(A, I)$ for transitive forests (A, I) : It is the smallest class of groups containing the trivial group that is closed under taking (i) free products and (ii) direct products with \mathbb{Z} .

Acyclic Automata. In both the upper and the lower bound of our completeness result, we employ the membership problem for acyclic automata, which has already been studied in connection with the knapsack and subset sum problem [15, 6].

We define a *finite automaton* as a tuple $\mathcal{A} = (Q, \Sigma, \Delta, q_0, q_f)$, where Q is a finite set of states, Σ is the *input alphabet*, $q_0 \in Q$ is the *initial state*, $q_f \in Q$ is the *final state*, and $\Delta \subseteq Q \times \Sigma^* \times Q$ is a finite set of *transitions*. The language accepted by \mathcal{A} is denoted with $L(\mathcal{A})$. An *acyclic automaton* is a finite automaton $\mathcal{A} = (Q, \Sigma, \Delta, q_0, q_f)$ such that the relation $\{(p, q) \mid \exists w \in \Sigma^* : (p, w, q) \in \Delta\}$ is acyclic. For a graph group $\mathbb{G}(A, I)$ the *membership problem for acyclic automata* is the following computational problem:

Input: An acyclic automaton \mathcal{A} over the input alphabet $A \cup A^{-1}$.

Question: Is there a word $w \in L(\mathcal{A})$ such that $w = 1$ in $\mathbb{G}(A, I)$?

In order to show the upper bound in Theorem 4.2, we reduce knapsack for $\mathbb{G}(A, I)$ with (A, I) a transitive forest to the membership problem for acyclic automata for $\mathbb{G}(A, I)$ (note that for subset sum this reduction is obvious). Then, we apply the following proposition. From work of Frenkel, Nikolaev, and Ushakov [6], it follows that the membership problem for acyclic automata is in P. We strengthen this to LogCFL.

Proposition 4.3. *If (A, I) is a transitive forest, then the membership problem for acyclic automata over $\mathbb{G}(A, I)$ is in LogCFL.*

The class LogCFL is included in the parallel complexity class NC^2 and has several alternative characterizations (see e.g. [22, 23]):

- logspace bounded alternating Turing-machines with polynomial proof tree size,
- semi-unbounded Boolean circuits of polynomial size and logarithmic depth, and
- logspace bounded auxiliary pushdown automata with polynomial running time.

For our purposes, the last characterization is most suitable. An AuxPDA (for auxiliary pushdown automaton) is a nondeterministic pushdown automaton with a two-way input tape and an additional work tape. Here we only consider AuxPDA with the following two restrictions:

- The length of the work tape is restricted to $O(\log n)$ for an input of length n (logspace bounded).
- There is a polynomial $p(n)$, such that every computation path of the AuxPDA on an input of length n has length at most $p(n)$ (polynomially time bounded).

Whenever we speak of an AuxPDA in the following, we implicitly assume that the AuxPDA is logspace bounded and polynomially time bounded. Deterministic AuxPDA are defined in the obvious way. The class of languages that are accepted by AuxPDA (resp., deterministic AuxPDA) is denoted by LogCFL (resp., LogDCFL). The proof of Proposition 4.3 uses the following lemma:

Lemma 4.4. *For every transitive forest (A, I) with the associated graph group $G = \mathbb{G}(A, I)$ there is a deterministic AuxPDA $\mathcal{P}(G)$ with input alphabet $A^{\pm 1}$ and the following properties:*

- *In each step, the input head for $\mathcal{P}(G)$ either does not move, or moves one step to the right.*
- *If the input word is equal to 1 in G , then $\mathcal{P}(G)$ terminates in the distinguished state q_1 with empty stack. Let us call this state the 1-state of $\mathcal{P}(G)$.*
- *If the input word is not equal to 1 in G , then $\mathcal{P}(G)$ terminates in a state different from q_1 (and the stack is not necessarily empty).*

Proof. We construct the AuxPDA $\mathcal{P}(G)$ by induction over the structure of the group G . For this, we consider the three cases that $G = 1$, $G = G_1 * G_2$, and $G = \mathbb{Z} \times G'$. The case that $G = 1$ is of course trivial.

- Case $G = \mathbb{Z} \times G'$. We have already constructed the AuxPDA $\mathcal{P}(G')$. The AuxPDA $\mathcal{P}(G)$ simulates the AuxPDA $\mathcal{P}(G')$ on the generators of G' . Moreover, it stores the current value of the \mathbb{Z} -component in binary notation on the work tape. If the input word has length n , then $O(\log n)$ bits are sufficient for this. At the end, $\mathcal{P}(G)$ goes into its 1-state if and only if $\mathcal{P}(G')$ is in its 1-state (which implies that the stack will be empty) and the \mathbb{Z} -component is zero.
- Case $G = G_1 * G_2$. For $i \in \{1, 2\}$, we have already constructed the AuxPDA $\mathcal{P}_i = \mathcal{P}(G_i)$. Let $A_i^{\pm 1}$ be its input alphabet, which is a monoid generating set for G_i . Consider now an input word $w \in (A_1^{\pm 1} \cup A_2^{\pm 1})^*$. Let us assume that $w = u_1 v_1 u_2 v_2 \cdots u_k v_k$ with $u_i \in (A_1^{\pm 1})^+$ and $v_i \in (A_2^{\pm 1})^+$ (other cases can be treated analogously). The AuxPDA $\mathcal{P}(G)$ starts with empty stack and simulates the AuxPDA \mathcal{P}_1 on the prefix u_1 . If it turns out that $u_1 = 1$ in G_1 (which means that \mathcal{P}_1 is in its 1-state) then the stack will be empty and the AuxPDA $\mathcal{P}(G)$ continues with simulating \mathcal{P}_2 on v_1 . On the other hand, if $u_1 \neq 1$ in G_1 , then $\mathcal{P}(G)$ pushes the state together with the work tape content of \mathcal{P}_1 reached after reading u_1 on the stack (on top of the final stack content of \mathcal{P}_1). This allows $\mathcal{P}(G)$ to resume the computation of \mathcal{P}_1 later. Then $\mathcal{P}(G)$ continues with simulating \mathcal{P}_2 on v_1 .

The computation of $\mathcal{P}(G)$ will continue in the above way. More precisely, if after reading u_i (resp. v_i with $i < k$) the AuxPDA \mathcal{P}_1 (resp. \mathcal{P}_2) is in its 1-state then either

- (i) the stack is empty or
- (ii) the top part of the stack is of the form sqt (t is the top), where s is a stack content of \mathcal{P}_2 (resp. \mathcal{P}_1), q is a state of \mathcal{P}_2 (resp. \mathcal{P}_1) and t is a work tape content of \mathcal{P}_2 (resp. \mathcal{P}_1).

In case (i), $\mathcal{P}(G)$ continues with the simulation of \mathcal{P}_2 (resp. \mathcal{P}_1) on the word v_i (resp. u_{i+1}) in the initial configuration. In case (ii), $\mathcal{P}(G)$ continues with the simulation of \mathcal{P}_2 (resp. \mathcal{P}_1) on the word v_i (resp. u_{i+1}), where the simulation is started with stack content s , state q , and work tape content t . On the other hand, if after reading u_i (resp. v_i with $i < k$) the AuxPDA \mathcal{P}_1 (resp. \mathcal{P}_2) is not in its 1-state then $\mathcal{P}(G)$ pushes on the stack the state and work tape content of \mathcal{P}_1 reached after its simulation on u_i .

This concludes the description of the AuxPDA $\mathcal{P}(G)$. □

We can now prove Proposition 4.3:

Proof of Proposition 4.3. Fix the graph group $G = \mathbb{G}(A, I)$, where (A, I) is a transitive forest. An AuxPDA for the membership problem for acyclic automata guesses a path in the input automaton \mathcal{A} and thereby simulates the AuxPDA $\mathcal{P}(G)$ from Lemma 4.4. If the final state of the input automaton \mathcal{A} is reached while the AuxPDA $\mathcal{P}(G)$ is in the accepting state q_1 , then the overall AuxPDA accepts. It is important that the AuxPDA $\mathcal{P}(G)$ works one-way since the guessed path in \mathcal{A} cannot be stored in logspace. This implies that the AuxPDA cannot reaccess the input symbols that already have been processed. Also note that the AuxPDA is logspace bounded and polynomially time bounded since \mathcal{A} is acyclic. □

4.1. Bounds on knapsack solutions. As mentioned above, we reduce for graph groups $\mathbb{G}(A, I)$ with (A, I) a transitive forest the knapsack problem to the membership problem for acyclic automata. To this end, we show that every positive

knapsack instance has a polynomially bounded solution. The latter is the most involved proof in our paper.

Frenkel, Nikolaev, and Ushakov [6] call groups with this property *polynomially bounded knapsack groups* and show that this class is closed under taking free products. However, it is not clear if direct products with \mathbb{Z} also inherit this property and we leave this question open.

Hence, we are looking for a property that yields polynomial size solutions and is passed on to free products and to direct products with \mathbb{Z} . It is known that the solution sets are always semilinear. If (A, I) is a transitive forest, this follows from a more general semilinearity property of rational sets [16] and for arbitrary graph groups, this was shown in [17]. Note that it is not true that the solution sets always have polynomial size semilinear representations. This already fails in the case of \mathbb{Z} : The equation $x_1 + \dots + x_k = k$ has $\binom{2k-1}{k} \geq 2^k$ solutions. We will show here that the solution sets have semilinear representations where every occurring number is bounded by a polynomial.

For a vector $x = (x_1, \dots, x_k) \in \mathbb{Z}^k$, we define the norm

$$\|x\| = \max\{|x_i| \mid i \in [1, k]\}.$$

For a subset $T \subseteq \mathbb{N}^k$, we write T^\oplus for the smallest subset of \mathbb{N}^k that contains 0 and is closed under addition. A subset $S \subseteq \mathbb{N}^k$ is called *linear* if there is a vector $x \in \mathbb{N}^k$ and a finite set $F \subseteq \mathbb{N}^k$ such that $S = x + F^\oplus$. Note that a set is linear if and only if it can be written as $x + \mathbb{A}\mathbb{N}^t$ for some $x \in \mathbb{N}^k$ and some matrix $A \in \mathbb{N}^{k \times t}$. Here, $\mathbb{A}\mathbb{N}^t$ denotes the set of all vectors Ay for $y \in \mathbb{N}^t$. A *semilinear set* is a finite union of linear sets. If $S = \bigcup_{i=1}^n x_i + F_i^\oplus$ for $x_1, \dots, x_n \in \mathbb{N}^k$ and finite sets $F_1, \dots, F_n \subseteq \mathbb{N}^k$, then the tuple $(x_1, F_1, \dots, x_n, F_n)$ is a *semilinear representation* of S and the *magnitude* of this representation is defined as the maximum of $\|y\|$, where y ranges over all elements of $\bigcup_{i=1}^n \{x_i\} \cup F_i$. The *magnitude* of a semilinear set S is the smallest magnitude of a semilinear representation for S .

Definition 4.5. *A group G is called knapsack tame if there is a polynomial p such that for every exponent equation $h_0 g_1^{x_1} h_1 g_2^{x_2} h_2 \dots g_n^{x_n} h_n = 1$ of size n with pairwise distinct variables x_1, \dots, x_k , the set $S \subseteq \mathbb{N}^k$ of solutions is semilinear of magnitude at most $p(n)$.*

Observe that although the size of an exponent equation may depend on the chosen generating set of G , changing the generating set increases the size only by a constant factor. Thus, whether or not a group is knapsack tame is independent of the chosen generating set.

Theorem 4.6. *If (A, I) is a transitive forest, then $\mathbb{G}(A, I)$ is knapsack tame.*

Note that Theorem 4.6 implies in particular that every solvable exponent equation with pairwise distinct variables has a polynomially bounded solution. Theorem 4.6 and Proposition 4.3 easily yield the upper bound in Theorem 4.2, see Proposition 4.21.

We prove Theorem 4.6 by showing that knapsack tameness transfers from groups G to $G \times \mathbb{Z}$ (Proposition 4.7) and from G and H to $G * H$ (Proposition 4.13). Since the trivial group is obviously knapsack tame, the inductive characterization of groups $\mathbb{G}(A, I)$ for transitive forests (A, I) immediately yields Theorem 4.6.

4.2. Tameness of direct products with \mathbb{Z} . In this section, we show the following.

Proposition 4.7. *If G is knapsack tame, then so is $G \times \mathbb{Z}$.*

Linear Diophantine equations. Vectors will be column vectors in this section. For a vector $x \in \mathbb{Z}^k$ we denote with x^T the corresponding row vector. We employ a result of Pottier [21], which bounds the norm of minimal non-negative solutions to a linear Diophantine equation. Let $A \in \mathbb{Z}^{k \times m}$ be an integer matrix where a_{ij} is the entry of A at row i and column j and let $x \in \mathbb{Z}^m$ be a vector with entries x_i . We will use the following norms:

$$\begin{aligned} \|A\|_{1,\infty} &= \max_{i \in [1,k]} \left(\sum_{j \in [1,m]} |a_{ij}| \right), \\ \|A\|_{\infty,1} &= \max_{j \in [1,m]} \left(\sum_{i \in [1,k]} |a_{ij}| \right), \\ \|A\|_{\infty} &= \max_{i \in [1,k], j \in [1,m]} |a_{ij}|, \\ \|x\|_1 &= \sum_{i=1}^m |x_i|. \end{aligned}$$

Recall that $\|x\| = \max_{i \in [1,m]} |x_i|$. A non-trivial solution $x \in \mathbb{N}^m \setminus \{0\}$ to the equation $Ax = 0$ is *minimal* if there is no $y \in \mathbb{N}^m \setminus \{0\}$ with $Ay = 0$ and $y \leq x$, $y \neq x$. Here $y \leq x$ means that $y_i \leq x_i$ for all $i \in [1,m]$. The set of all solutions clearly forms a submonoid of \mathbb{N}^m . Let r be the rank of A .

Theorem 4.8 (Pottier [21]). *Each non-trivial minimal solution $x \in \mathbb{N}^m$ to $Ax = 0$ satisfies $\|x\|_1 \leq (1 + \|A\|_{1,\infty})^r$.*

We only need Theorem 4.8 for the case that A is a row vector u^T for $u \in \mathbb{Z}^k$.

Corollary 4.9. *Let $u \in \mathbb{Z}^k$. Each non-trivial minimal solution $x \in \mathbb{N}^k$ to $u^T x = 0$ satisfies $\|x\|_1 \leq 1 + \|u\|_1$.*

By applying Theorem 4.8 to the row vector $(u^T, -b)$ for $b \in \mathbb{Z}$, it is easy to deduce that for each $x \in \mathbb{N}^k$ with $u^T x = b$, there is a $y \in \mathbb{N}^k$ with $u^T y = b$, $y \leq x$, and $\|y\|_1 \leq 1 + \left\| \begin{pmatrix} u \\ b \end{pmatrix} \right\|_1 = 1 + \|u\|_1 + |b|$. We reformulate Corollary 4.9 as follows.

Lemma 4.10. *Let $u \in \mathbb{Z}^k$ and $b \in \mathbb{Z}$. Then the set $\{x \in \mathbb{N}^k \mid u^T x = b\}$ admits a decomposition $\{x \in \mathbb{N}^k \mid u^T x = b\} = \bigcup_{i=1}^s c_i + C\mathbb{N}^t$, where $c_i \in \mathbb{N}^k$ and $C \in \mathbb{N}^{k \times t}$ with $\|c_i\|_1$ and $\|C\|_{\infty,1}$ bounded by $1 + \|u\|_1 + |b|$.*

Proof. Let $\{c_1, \dots, c_s\}$ be the set of minimal solutions of $u^T x = b$. Then, as explained above, Corollary 4.9 yields $\|c_i\|_1 \leq 1 + \|u\|_1 + |b|$. Moreover, let $C \in \mathbb{N}^{k \times t}$ be the matrix whose columns are the non-trivial minimal solutions of $u^T x = 0$. Then we have $\|C\|_{\infty,1} \leq 1 + \|u\|_1$. This clearly yields the desired decomposition. \square

The problem is that we want to apply Lemma 4.10 in a situation where we have no bound on $\|u\|_1$, but only one on $\|u\|$. The following lemma yields such a bound.

Lemma 4.11. *If $u \in \mathbb{Z}^k$ and $b \in \mathbb{Z}$ with $\|u\|, |b| \leq M$, then we have a decomposition $\{x \in \mathbb{N}^k \mid u^T x = b\} = \bigcup_{i=1}^s c_i + C\mathbb{N}^t$ where $\|c_i\|_1, \|C\|_{\infty,1} \leq 1 + (M + 2)M$.*

Proof. Write $u^T = (b_1, \dots, b_k)$ and consider the row vector $v^T = (b'_1, \dots, b'_{2M+1})$, $v \in \mathbb{Z}^{2M+1}$, with entries $b'_i = i - (M + 1)$. Thus, we have

$$v^T = (b'_1, \dots, b'_{2M+1}) = (-M, -M + 1, \dots, -1, 0, 1, \dots, M).$$

Moreover, define the matrix $S = (s_{ij}) \in \mathbb{N}^{(2M+1) \times k}$ with

$$s_{ij} = \begin{cases} 1 & \text{if } b_j = b'_i, \\ 0 & \text{otherwise.} \end{cases}$$

Then clearly $u = v^T S$ and $\|v\|_1 = (M+1)M$. Furthermore, observe that we have $\|Sx\|_1 = \|x\|_1$ for every $x \in \mathbb{N}^k$ and that for each $y \in \mathbb{N}^{2M+1}$ the set

$$T_y = \{x \in \mathbb{N}^k \mid Sx = y\}$$

is finite. According to Lemma 4.10, we can write

$$(4.1) \quad \{x \in \mathbb{N}^{2M+1} \mid v^T x = b\} = \bigcup_{i=1}^{s'} c'_i + C' \mathbb{N}^{t'}$$

where $\|c'_i\|_1, \|C'\|_{\infty,1} \leq 1 + (M+1)M + M = 1 + (M+2)M$. Let $\{c_1, \dots, c_s\}$ be the union of all sets $T_{c'_i}$ for $i \in [1, s']$ and let $C \in \mathbb{N}^{k \times t}$ be the matrix whose columns comprise all T_v where $v \in \mathbb{N}^{2M+1}$ is a column of C' . Since we have $\|Sx\|_1 = \|x\|_1$ for $x \in \mathbb{N}^k$, the vectors c_i obey the same bound as the vectors c'_i , meaning $\|c_i\|_1 \leq 1 + (M+2)M$. By the same argument, we have $\|C\|_{\infty,1} \leq \|C'\|_{\infty,1} \leq 1 + (M+2)M$. It remains to be shown that the equality from the lemma holds.

Suppose $u^T x = b$. Then $v^T Sx = b$ and hence $Sx = c'_i + C'y$ for some $y \in \mathbb{N}^{t'}$. Observe that if $Sz = p + q$ for $p, q \in \mathbb{N}^{2M+1}$, then z decomposes as $z = p' + q'$, $p', q' \in \mathbb{N}^k$, so that $Sp' = p$ and $Sq' = q$. Therefore, we can write $x = x_0 + \dots + x_n$ with $Sx_0 = c'_i$ and Sx_j is some column of C' for each $j \in [1, n]$. Hence, $x_0 = c_r$ for some $r \in [1, s]$ and for each $j \in [1, n]$, x_j is a column of C . This proves $x \in c_r + C\mathbb{N}^t$.

On the other hand, the definition of c_1, \dots, c_s and C implies that for each column v of C , Sv is a column of C' . Moreover, for each $i \in [1, s]$, there is a $j \in [1, s']$ with $Sc_i = c'_j$ and thus $Sc_i + SC\mathbb{N}^t \subseteq c'_j + C'\mathbb{N}^{t'}$. Therefore

$$u^T(c_i + C\mathbb{N}^t) = v^T S(c_i + C\mathbb{N}^t) \subseteq v^T(c'_j + C'\mathbb{N}^{t'})$$

and the latter set contains only b because of (4.1). \square

Lemma 4.12. *Let $S \subseteq \mathbb{N}^k$ be a semilinear set of magnitude M and $u \in \mathbb{Z}^k$, $b \in \mathbb{Z}$ with $\|u\|, |b| \leq m$. Then $\{x \in S \mid u^T x = b\}$ is a semilinear set of magnitude at most $(kmM)^d$ for some constant d .*

Proof. Let $T = \{x \in \mathbb{N}^k \mid u^T x = b\}$. We may assume that S is linear of magnitude M , because if $S = L_1 \cup \dots \cup L_n$, then $S \cap T = (L_1 \cap T) \cup \dots \cup (L_n \cap T)$.

Write $S = a + A\mathbb{N}^n$ with $a \in \mathbb{N}^k$ and $A \in \mathbb{N}^{k \times n}$, where $\|a\| \leq M$ and $\|A\|_{\infty} \leq M$. Consider the set $U = \{x \in \mathbb{N}^n \mid u^T Ax = b - u^T a\}$. Note that $u^T A \in \mathbb{Z}^{1 \times n}$ and

$$\|u^T A\| \leq k \cdot \|u\| \cdot \|A\|_{\infty} \leq kmM,$$

$$|b - u^T a| \leq m + k \cdot \|u\| \cdot \|a\| \leq m + kmM.$$

According to Lemma 4.11, we can write $U = \bigcup_{i=1}^s c_i + C\mathbb{N}^t$ where $\|c_i\|_1$ and $\|C\|_{\infty,1}$ are at most $1 + (m + kmM)(m + kmM + 2)$. Observe that

$$a + AU = \bigcup_{i=1}^s a + Ac_i + AC\mathbb{N}^t$$

and

$$\|a + Ac_i\| \leq \|a\| + \|A\|_{\infty} \cdot \|c_i\|_1 \leq 2M + M(m + kmM)(m + kmM + 2),$$

$$\|AC\|_\infty \leq \|A\|_\infty \cdot \|C\|_{\infty,1} \leq M + M(m + kmM)(m + kmM + 2).$$

Finally, note that $S \cap T = a + AU$. \square

We are now ready to prove Proposition 4.7.

Proof of Proposition 4.7. Suppose G is knapsack tame with polynomial \bar{p} . Let

$$(4.2) \quad h_0 g_1^{x_1} h_1 g_2^{x_2} h_2 \cdots g_k^{x_k} h_k = 1$$

be an exponent equation of size n with pairwise distinct variables x_1, \dots, x_k and with $h_0, g_1, h_1, \dots, g_k, h_k \in G \times \mathbb{Z}$. Let $h_i = (\bar{h}_i, y_i)$ for $i \in [0, k]$ and $g_i = (\bar{g}_i, z_i)$ for $i \in [1, k]$.

The exponent equation $\bar{h}_0 \bar{g}_1^{x_1} \bar{h}_1 \bar{g}_2^{x_2} \bar{h}_2 \cdots \bar{g}_k^{x_k} \bar{h}_k = 1$ has a semilinear solution set $\bar{S} \subseteq \mathbb{N}^k$ of magnitude at most $\bar{p}(n)$. The solution set of (4.2) is

$$S = \{(x_1, \dots, x_k) \in \bar{S} \mid z_1 x_1 + \cdots + z_k x_k = y\},$$

where $y = -(y_0 + \cdots + y_k)$. Note that $|z_i| \leq n$ and $|y| \leq n$. By Lemma 4.12, S is semilinear of magnitude $(n^2 \bar{p}(n))^d$ for some constant d (recall that $k \leq n$). \square

4.3. Tameness of free products. This section is devoted to the proof of the following proposition.

Proposition 4.13. *If G_0 and G_1 are knapsack tame, then so is $G_0 * G_1$.*

Let $G = G_0 * G_1$. Suppose that for $i \in \{0, 1\}$, the group G_i is generated by A_i , where w.l.o.g. $A_i^{-1} = A_i$ and let $A = A_0 \uplus A_1$, which generates G . Recall that every $g \in G$ can be written uniquely as $g = g_1 \cdots g_n$ where $n \geq 0$, $g_i \in (G_0 \setminus \{1\}) \cup (G_1 \setminus \{1\})$ for each $i \in [1, n]$ and where $g_j \in G_t$ iff $g_{j+1} \in G_{1-t}$ for $j \in [1, n-1]$. We call g *cyclically reduced* if for some $t \in \{0, 1\}$, either $g_1 \in G_t$ and $g_n \in G_{1-t}$ or $g_1, g_n \in G_t$ and $g_n g_1 \neq 1$. Consider an exponent equation

$$(4.3) \quad h_0 g_1^{x_1} h_1 \cdots g_k^{x_k} h_k = 1,$$

of size n , where g_i is represented by $u_i \in A^*$ for $i \in [1, k]$ and h_i is represented by $v_i \in A^*$ for $i \in [0, k]$. Then clearly $\sum_{i=0}^k |v_i| + \sum_{i=1}^k |u_i| \leq n$. Let $S \subseteq \mathbb{N}^k$ be the set of all solutions to (4.3). Every word $w \in A^*$ has a (possibly empty) unique factorization into maximal factors from $A_0^+ \cup A_1^+$, which we call *syllables*. By $\|w\|$, we denote the number of syllables of w . The word w is *reduced* if none of its syllables represents 1 (in G_0 resp. G_1). We define the maps $\lambda, \rho: A^+ \rightarrow A^+$ ("rotate left/right"), where for each word $w \in A^+$ with its factorization $w = w_1 \cdots w_m$ into syllables, we set $\lambda(w) = w_2 \cdots w_m w_1$ and $\rho(w) = w_m w_1 w_2 \cdots w_{m-1}$.

Consider a word $w \in A^*$ and suppose $w = w_1 \cdots w_m$, $m \geq 0$, where for each $i \in [1, m]$, we have $w_i \in A_j^+$ for some $j \in \{0, 1\}$ (we allow that $w_i, w_{i+1} \in A_j^+$). A *cancellation* is a subset $C \subseteq 2^{[1, m]}$ that is

- *a partition*: $\bigcup_{I \in C} I = [1, m]$ and $I \cap J = \emptyset$ for any $I, J \in C$ with $I \neq J$.
- *consistent*: for each $I \in C$, there is an $i \in \{0, 1\}$ such that $w_j \in A_i^+$ for all $j \in I$.
- *cancelling*: if $\{i_1, \dots, i_\ell\} \in C$ with $i_1 < \cdots < i_\ell$, then $w_{i_1} \cdots w_{i_\ell}$ represents 1 in G .
- *well-nested*: there are no $I, J \in C$ with $i_1, i_2 \in I$ and $j_1, j_2 \in J$ such that $i_1 < j_1 < i_2 < j_2$.
- *maximal*: if $w_i, w_{i+1} \in A_j^+$ for $j \in \{0, 1\}$ then there is an $I \in C$ with $i, i+1 \in I$.

Since C can be regarded as a hypergraph on $[1, m]$, the elements of C will be called *edges*. We have the following simple fact:

Lemma 4.14. *Let $w = w_1 \cdots w_m$, $m \geq 0$, where for each $i \in [1, m]$, we have $w_i \in A_j^+$ for some $j \in \{0, 1\}$. Then w admits a cancellation if and only if it represents 1 in G .*

Proof. Assume that w represents 1 in the free product G . The case $w = \varepsilon$ is clear; hence assume that $w \neq \varepsilon$. Then there must exist a factor $w_i w_{i+1} \cdots w_j$ representing 1 in G such that (i) $w_i w_{i+1} \cdots w_j \in A_k^+$ for some $k \in \{0, 1\}$, (ii) either $i = 1$ or $w_{i-1} \in A_{1-k}^+$, and (iii) either $j = m$ or $w_{j+1} \in A_{1-k}^+$. The word $w' = w_1 \cdots w_{i-1} w_{j+1} \cdots w_m$ also represents 1 in G . By induction, w' admits a cancellation C' . Let C'' be obtained from C' by replacing every occurrence of an index $k \geq i$ in C' by $k + j - i + 1$. Then $C = C'' \cup \{[i, j]\}$ is a cancellation for w .

For the other direction let us call a partition $C \subseteq 2^{[1, m]}$ a *weak cancellation* if it is consistent, cancelling and well-nested (but not necessarily maximal). Then we show by induction that w represents 1 if it has a weak cancellation. So, let C be a weak cancellation of $w \neq \varepsilon$. Then there must exist an interval $[i, j] \in C$ (otherwise C would be not well-nested). Then $w_i w_{i+1} \cdots w_j$ represents 1 in G . Consider the word $w' = w_1 \cdots w_{i-1} w_{j+1} \cdots w_m$. Let C' be obtained from $C \setminus \{[i, j]\}$ by replacing every occurrence of an index $k \geq j + 1$ in $C \setminus \{[i, j]\}$ by $k - j + i - 1$. Then C' is a weak cancellation for w' . Hence, w' represents 1, which implies that w represents 1. \square

Of course, when showing that the solution set of (4.3) has a polynomial magnitude, we may assume that $g_i \neq 1$ for any $i \in [1, k]$. Moreover, we lose no generality by assuming that all words u_i , $i \in [1, k]$ and v_i , $i \in [0, k]$ are reduced. Furthermore, we may assume that each g_i is cyclically reduced. Indeed, if some g_i is not cyclically reduced, we can write $g_i = f^{-1} g f$ for some cyclically reduced g and replace h_{i-1} , g_i , and h_i by $h_{i-1} f^{-1}$, $g = f g_i f^{-1}$, and $f h_i$, respectively. This does not change the solution set because $h_{i-1} f^{-1} (f g_i f^{-1})^{x_i} f h_i = h_{i-1} g_i^{x_i} h_i$. Moreover, if we do this replacement for each g_i that is not cyclically reduced, we increase the size of the instance by at most $2|g_1| + \cdots + 2|g_k| \leq 2n$ (note that $|g| = |g_i|$). Applying this argument again, we may even assume that

$$(4.4) \quad u_i \in A_0^+ \cup A_1^+ \cup A_0^+ A^* A_1^+ \cup A_1^+ A^* A_0^+$$

for every $i \in [1, k]$. Note that λ and ρ are bijections on words of this form.

Consider a solution (x_1, \dots, x_k) to (4.3). Then the word

$$(4.5) \quad w = v_0 u_1^{x_1} v_1 \cdots u_k^{x_k} v_k$$

represents 1 in G . We factorize each v_i , $i \in [0, k]$, and each u_i , $i \in [1, k]$, into its syllables. These factorizations define a factorization $w = w_1 \cdots w_m$ and we call this the *block factorization* of w . This is the coarsest refinement of the factorization $w = v_0 u_1^{x_1} v_1 \cdots u_k^{x_k} v_k$ and of w 's factorization into syllables. The numbers $1, 2, \dots, m$ are the *blocks* of w . We fix this factorization $w = w_1 \cdots w_m$ for the rest of this section.

Cycles and certified solutions. In the representation $v_0 u_1^{x_1} v_1 \cdots u_k^{x_k} v_k = 1$ of (4.3), the words u_1, \dots, u_k are called the *cycles*. If $u_i \in A_0^+ \cup A_1^+$, the cycle u_i is said to be *simple* and otherwise *mixed* (note that $u_i = \varepsilon$ cannot happen because $g_i \neq 1$). Let p be a block of w . If w_p is contained in some $u_i^{x_i}$ for a cycle u_i , then p

is a u_i -blocks or block from u_i . If w_p is contained in some v_i , then p is a v_i -block or a block from v_i .

Note that if $C \subseteq 2^{[1,m]}$ is a cancellation for $w = w_1 \cdots w_m$ then by maximality, for each simple cycle u_i , all u_i -blocks are contained in the same edge of C .

A *certified solution* is a pair (x, C) , where x is a solution to (4.3) and C is a cancellation of the word w as in (4.5).

We will also need the following two auxiliary lemmas.

Lemma 4.15. *Let C be a cancellation. If i, j are two distinct blocks from the same mixed cycle, then there is no edge $I \in C$ with $i, j \in I$.*

Proof. Suppose there is such an $I \in C$. Furthermore, assume that i and j are chosen so that $|i - j|$ is minimal and $i < j$. Since $i, j \in I$, we have $w_i w_j \in A_0^+ \cup A_1^+$ by consistency of C . Hence, by (4.4), we cannot have $j = i + 1$. Therefore there is an $\ell \in [1, m]$ with $i < \ell < j$. This means there is a $J \in C$ with $\ell \in J$. By well-nestedness, $J \subseteq [i, j]$. Since every edge in C must contain at least two elements, we have $|J| \geq 2$ and thus a contradiction to the minimality of $|i - j|$. \square

An edge $I \in C$ is called *standard* if $|I| = 2$ and the two blocks in I are from mixed cycles. Intuitively, the following lemma tells us that in a cancellation, most edges are standard.

Lemma 4.16. *Let C be a cancellation and u_i be a mixed cycle. Then there are at most $n + 3k + 1$ non-standard edges $I \in C$ containing a u_i -block.*

Proof. Let $N \subseteq C$ be the set of all non-standard edges $I \in C$ that contain a u_i -block. Then, each edge $I \in N$ satisfies one of the following.

- (i) I contains a block from some simple cycle. There are at most k such I .
- (ii) I contains a block from some v_j , $j \in [0, k]$. Since $\|v_0\| + \cdots + \|v_k\| \leq n$, there are at most n such I .
- (iii) I contains only blocks from mixed cycles and $|I| > 2$.

Let $M \subseteq C$ be the set of edges of type (iii). If we can show that $|M| \leq 2k + 1$, then the lemma is proven. Consider the sets

$$\begin{aligned} M_- &= \{I \in M \mid I \text{ contains a block from a mixed cycle } u_j, j < i\}, \\ M_+ &= \{I \in M \mid I \text{ contains a block from a mixed cycle } u_j, j > i\}. \end{aligned}$$

We shall prove that $|M_- \cap M_+| \leq 1$ and that $|M_+ \setminus M_-| \leq k$. By symmetry, this also means $|M_- \setminus M_+| \leq k$ and thus $|M| = |M_- \cup M_+| \leq 2k + 1$.

Suppose $I_1, I_2 \in M_- \cap M_+$, $I_1 \neq I_2$. Let $r \in I_1$ and $s \in I_2$ such that r and s are blocks from u_i , say with $r < s$. Since $I_1 \in M_+$, I_1 contains a block r' from a mixed cycle u_j , $j > i$. This means in particular $s < r'$. By well-nestedness, this implies $I_2 \subseteq [r, r']$, so that I_2 cannot contain a block from a mixed cycle u_ℓ with $\ell < i$, contradicting $I_2 \in M_-$. Thus, $|M_- \cap M_+| \leq 1$.

In order to prove $|M_+ \setminus M_-| \leq k$, we need another concept. For each $I \in M_+$, there is a maximal $j \in [1, k]$ such that u_j is a mixed cycle and I contains a block from u_j . Let $\mu(I) = j$. We will show $\mu(I_1) \neq \mu(I_2)$ for all $I_1, I_2 \in M_+ \setminus M_-$, $I_1 \neq I_2$. This clearly implies $|M_+ \setminus M_-| \leq k$.

Suppose $I_1, I_2 \in M_+ \setminus M_-$, $I_1 \neq I_2$, with $\mu(I_1) = \mu(I_2)$. Let $j = \mu(I_1) = \mu(I_2)$. Let r be a block from u_i contained in I_1 and let r' be a block from u_i contained in I_2 . (Recall that those exist because $I_1, I_2 \in M$.) Without loss of generality, assume

$r < r'$. Moreover, let s be a block from u_j contained in I_1 and let s' be a block from u_j contained in I_2 . Thus, we have $r < r' < s'$.

However, we have $|I_1| > 2$, meaning I_1 contains a block p other than r and s . Since an edge cannot contain two blocks of one mixed cycle (Lemma 4.15), p has to belong to a mixed cycle u_t other than u_i and u_j . By the maximality of j , we have $i < t < j$. This implies, however, $r < r' < p < s'$, which contradicts well-nestedness. \square

Mixed periods. From now on, for each $i \in [1, k]$, we use e_i to denote the i -th unit vector in \mathbb{N}^k , i.e. the vector with 1 in the i -th coordinate and 0 otherwise. A *mixed period* is a vector $\pi \in \mathbb{N}^k$ of the form $\|u_j\| \cdot e_i + \|u_i\| \cdot e_j$, where u_i and u_j are mixed cycles. Let $\mathbb{P} \subseteq \mathbb{N}^k$ be the set of mixed periods. Note that $|\mathbb{P}| \leq k^2$.

We will need a condition that guarantees that a given period $\pi \in \mathbb{P}$ can be added to a solution x to obtain another solution. Suppose we have two blocks p and q for which we know that if we insert a string f_1 to the left of w_p and a string f_2 to the right of w_q and $f_1 f_2$ cancels to 1 in G , then the whole word cancels to 1. Which string would we insert to the left of w_p and to the right of w_q if we build the solution $x + \pi$?

Suppose p is a u_i -block and q is a u_j -block. Moreover, let r be the first (left-most) u_i -block and let s be the last (right-most) u_j -block. If we add $\|u_j\| \cdot e_i$ to x , this inserts $\lambda^{p-r}(u_i^{\|u_j\|})$ to the left of w_p . Indeed, in the case $p = r$, we insert $u_i^{\|u_j\|}$; and when p moves one position to the right, the inserted string is rotated once to the left. Similarly, if we add $\|u_i\| \cdot e_j$ to x , we insert $\rho^{s-q}(u_j^{\|u_i\|})$ to the right of w_q . This is clear for $q = s$ and decrementing q means rotating the inserted string to the right. This motivates the following definition.

Let (x, C) be a certified solution and let u_i and u_j be mixed cycles with $i < j$. Moreover, let $r \in [1, m]$ be the left-most u_i -block and let $s \in [1, m]$ be the right-most u_j -block. Then the mixed period $\pi = \|u_j\| \cdot e_i + \|u_i\| \cdot e_j$ is *compatible with* (x, C) if there are a u_i -block p and a u_j -block q such that

$$(4.6) \quad \{p, q\} \in C \text{ and } \lambda^{p-r}(u_i^{\|u_j\|})\rho^{s-q}(u_j^{\|u_i\|}) \text{ represents 1 in } G.$$

With $\mathbb{P}(x, C)$, we denote the set of mixed periods that are compatible with (x, C) . One might wonder why we require an edge $\{p, q\} \in C$. In order to guarantee that $\lambda^{p-r}(u_i^{\|u_j\|})$ and $\rho^{s-q}(u_j^{\|u_i\|})$ can cancel, it would be sufficient to merely forbid edges $I \in C$ that intersect $[p, q]$ and contain a block outside of $[p - 1, q + 1]$. However, this weaker condition can become false when we insert other mixed periods. Our stronger condition is preserved, which implies:

Lemma 4.17. *Let (x, C) be a certified solution. Then every $x' \in x + \mathbb{P}(x, C)^\oplus$ is a solution.*

Proof. It suffices to show that if (x, C) is a certified solution and $\pi \in \mathbb{P}(x, C)$, then there is a certified solution (x', C') such that $x' = x + \pi$ and $\mathbb{P}(x, C) \subseteq \mathbb{P}(x', C')$. Suppose $\pi = \|u_j\| \cdot e_i + \|u_i\| \cdot e_j \in \mathbb{P}(x, C)$. Without loss of generality, assume $i < j$. Let $r \in [1, m]$ be the left-most u_i -block and $s \in [1, m]$ be the right-most u_j -block in w . Since $\pi \in \mathbb{P}(x, C)$, there is a u_i -block p and a u_j -block q such that (4.6) holds. As explained above, we can insert $\lambda^{p-r}(u_i^{\|u_j\|})$ on the left of w_p and $\rho^{s-q}(u_j^{\|u_i\|})$ on the right of w_q and thus obtain a word w' that corresponds to the vector $x' = x + \pi$.

Both inserted words consist of $\|u_j\| \cdot \|u_i\|$ many blocks and they cancel to 1, which means we can construct a cancellation C' from C as follows. Between the two sequences of inserted blocks, we add two-element edges so that the left-most inserted u_i -block is connected to the right-most inserted u_j -block, and so forth. The blocks that existed before are connected by edges as in C . It is clear that then, C' is a partition that is consistent, cancelling and maximal. Moreover, since there is an edge $\{p, q\}$, the new edges between the inserted blocks do not violate well-nestedness: If there were a crossing edge, then there would have been one that crosses $\{p, q\}$.

It remains to verify $\mathbb{P}(x, C) \subseteq \mathbb{P}(x', C')$. A mixed period $\pi \in \mathbb{P}(x, C) \setminus \{\pi\}$ is clearly contained in $\mathbb{P}(x', C')$ too. Hence, it remains to show $\pi \in \mathbb{P}(x', C')$. This, however, follows from the fact that instead of the edge $\{p, q\}$ that witnesses compatibility of π with (x, C) , we can use its counterpart in C' ; let us call this edge $\{p', q'\}$: If r' is the left-most u_i -block in w' and s' is the right-most u_j -block in w' , then $p' - r' = p - r + \|u_i\| \cdot \|u_j\|$ and $s' - q' = s - q + \|u_i\| \cdot \|u_j\|$. This implies $\lambda^{p-r}(u_i^{\|u_j\|}) = \lambda^{p'-r'}(u_i^{\|u_j\|})$ and $\rho^{s-q}(u_j^{\|u_i\|}) = \rho^{s'-q'}(u_j^{\|u_i\|})$ which implies that (4.6) holds for the edge $\{p', q'\}$. This completes the proof of the lemma. \square

We shall need another auxiliary lemma.

Lemma 4.18. *Let C be a cancellation for w . Let u_i and u_j be distinct mixed cycles. Let $D \subseteq C$ be the set of standard edges $I \in C$ that contain one block from u_i and one block from u_j . Then the set $B = \{p \in [1, m] \mid p \text{ is a } u_i\text{-block}, \exists I \in D : p \in I\}$ is an interval.*

Proof. We prove the case $i < j$, the other follows by symmetry. Suppose there are $r_1, r_2 \in B$ such that $r_1 < r_2$ and there is no $t \in B$ with $r_1 < t < r_2$.

Towards a contradiction, suppose $r_2 - r_1 > 1$. Since $r_1, r_2 \in B$, there are $I_1, I_2 \in D$ with $r_1 \in I_1$ and $r_2 \in I_2$. Let $I_1 = \{r_1, s_1\}$ and $I_2 = \{r_2, s_2\}$. Then s_1 and s_2 are u_j -blocks and by well-nestedness, we have $r_1 < r_2 < s_2 < s_1$. Since $r_2 - r_1 > 1$, there is a t with $r_1 < t < r_2$ and therefore some $J \in C$ with $t \in J$. Since $|J| \geq 2$, there has to be a $t' \in J$, $t' \neq t$. However, well-nestedness dictates that $t' \in [r_1, s_1] \setminus [r_2, s_2]$. Since J cannot contain another block from u_i (Lemma 4.15), we cannot have $t' \in [r_1, r_2]$, which only leaves $t' \in [s_2, s_1]$. Hence, t' is from u_j . By the same argument, any block $t'' \in J \setminus \{t, t'\}$ must be from u_i or u_j , contradicting Lemma 4.15. This means $|J| = 2$ and thus $t \in B$, in contradiction to the choice of r_1 and r_2 . \square

Let $M \subseteq [1, k]$ be the set of $i \in [1, k]$ such that u_i is a mixed cycle. We define a new norm on vectors $x \in \mathbb{N}^k$ by setting $\|x\|_m = \max_{i \in M} x_i$.

Lemma 4.19. *There is a polynomial q such that the following holds. For every certified solution (x, C) with $\|x\|_m > q(n)$, there exists a mixed period $\pi \in \mathbb{P}(x, C)$ and a certified solution (x', C') such that $x' = x - \pi$ and $\mathbb{P}(x, C) \subseteq \mathbb{P}(x', C')$.*

Proof. We show that the lemma holds if $q(n) \geq (n + 3k + 1) + kn^2$. (Recall that $k \leq n$.) Let (x, C) be a certified solution with $\|x\|_m > q(n)$. Then there is a mixed cycle u_i such that $x_i > q(n)$ and hence $u_i^{x_i}$ consists of more than $q(n)$ blocks. Let $D \subseteq C$ be the set of all edges $I \in C$ that contain a block from u_i . Since an edge can contain at most one block per mixed cycle (Lemma 4.15), we have $|D| > q(n)$. Hence, Lemma 4.16 tells us that D contains more than kn^2 standard edges. Hence, there exists a mixed cycle u_j such that the set $E \subseteq D$ of standard edges $I \in D$ that

consist of one block from u_i and one block from u_j satisfies $|E| > n^2$. If B_i (resp., B_j) denotes the set of blocks from u_i (resp., u_j) contained in some edge $I \in E$, then each of the sets B_i and B_j has to be an interval (Lemma 4.18) of size more than n^2 .

We only deal with the case $i < j$, the case $i > j$ can be done similarly. Let us take a subinterval $[p', p]$ of B_i such that $p - p' = \|u_i\| \cdot \|u_j\| \leq n^2$. By well-nestedness and since B_j is an interval, the neighbors (with respect to the edges from E) of $[p', p]$ form an interval $[q, q'] \subseteq B_j$ as well, and we have $p - p' = q' - q = \|u_i\| \cdot \|u_j\|$. Moreover, we have an edge $\{p - \ell, q + \ell\} \in E$ for each $\ell \in [0, p - p']$. In particular, $w_{p'} w_{p'+1} \cdots w_{p-1} w_{q+1} \cdots w_{q'}$ represents 1 in G .

Let r be the left-most u_i -block and let s be the right-most u_j -block. Then, as shown before the definition of compatibility (p. 14), we have

$$\lambda^{p-r}(u_i^{\|u_j\|}) = w_{p'} w_{p'+1} \cdots w_{p-1}, \quad \rho^{s-q}(u_j^{\|u_i\|}) = w_{q+1} w_{q+1} \cdots w_{q'}.$$

Therefore, $\lambda^{p-r}(u_i^{\|u_j\|}) \rho^{s-q}(u_j^{\|u_i\|})$ represents 1 in G and $\{p, q\}$ witnesses compatibility of $\pi = \|u_j\| \cdot e_i + \|u_i\| \cdot e_j$ with (x, C) . Hence, $\pi \in \mathbb{P}(x, C)$.

Let $x' = x - \pi$. We remove the factors $w_{p'} \cdots w_{p-1}$ and $w_{q+1} \cdots w_{q'}$ from w . Then, the remaining blocks spell $w' = v_0 u_1^{x'_1} v_1 \cdots u_k^{x'_k} v_k$. Indeed, recall that removing from a word y^t any factor of length $\ell \cdot |y|$ will result in the word $y^{t-\ell}$. Moreover, let C' be the set of edges that agree with C on the remaining blocks. By the choice of the removed blocks, it is clear that C' is a cancellation for w' . Hence, (x', C') is a certified solution.

It remains to verify $\mathbb{P}(x, C) \subseteq \mathbb{P}(x', C')$. First note that for every mixed cycle u_ℓ , all u_ℓ -blocks that remain in w' change their position relative to the left-most and the right-most u_ℓ -block by a difference that is divisible by $\|u_\ell\|$ (if $i \neq \ell \neq j$ then these relative positions do not change at all). Note that the expression $\lambda^{p-r}(u_i^{\|u_j\|})$ is not altered when $p - r$ changes by a difference divisible by $\|u_i\|$, and an analogous fact holds for $\rho^{s-q}(u_j^{\|u_i\|})$. Hence, the edge in C' that corresponds to the C -edge $\{p, q\}$ is a witness for $\pi \in \mathbb{P}(x', C')$. Moreover, for all other mixed periods $\pi' \in \mathbb{P}(x, C) \setminus \{\pi\}$ that are witnessed by an edge $\{t, u\} \in C$, the blocks t and u do not belong to $[p', p - 1] \cup [q + 1, q']$. Therefore, the corresponding edge in C' exists and serves as a witness for $\pi' \in \mathbb{P}(x', C')$. \square

Repeated application of Lemma 4.17 now yields:

Lemma 4.20. *There exists a polynomial q such that the following holds. For every solution $x \in \mathbb{N}^k$, there exists a certified solution (x', C') such that $\|x'\|_m \leq q(n)$ and $x \in x' + \mathbb{P}(x', C')^\oplus$.*

Proof. Let q be the polynomial provided by Lemma 4.19. Since x is a solution, there is a certified solution (x, C) . Repeated application of Lemma 4.19 yields certified solutions $(x_0, C_0), \dots, (x_m, C_m)$ and mixed periods π_1, \dots, π_m such that $(x_0, C_0) = (x, C)$, $\pi_i \in \mathbb{P}(x_{i-1}, C_{i-1}) \subseteq \mathbb{P}(x_i, C_i)$, $x_i = x_{i-1} - \pi_i$, and $\|x_m\|_m \leq q(n)$. In particular, $\mathbb{P}(x_m, C_m)$ contains each π_i and hence

$$x = x_m + \pi_1 + \cdots + \pi_m \in x_m + \mathbb{P}(x_m, C_m)^\oplus.$$

Thus, $(x', C') = (x_m, C_m)$ is the desired certified solution. \square

We are now ready to prove Proposition 4.13 and thus Theorem 4.6.

Proof of Proposition 4.13. Suppose that p_0 and p_1 are the polynomials guaranteed by the knapsack tameness of G_0 and G_1 , respectively. Recall that $S \subseteq \mathbb{N}^k$ is the set of solutions to (4.3). We prove that there exists a polynomial p such that for every $x \in S$ there is a semilinear set $S' \subseteq \mathbb{N}^k$ of magnitude at most $p(n)$ such that $x \in S' \subseteq S$. This clearly implies that S has magnitude at most $p(n)$. First, we apply Lemma 4.20. It yields a polynomial q and a certified solution (x', C') with $\|x'\|_m \leq q(n)$ such that $x \in x' + \mathbb{P}(x', C')^\oplus$. Let $w' = v_0 u_1^{x'_1} v_1 \cdots u_k^{x'_k} v_k$ and consider w' decomposed into blocks as we did above with w .

Let $T \subseteq [1, k]$ be the set of all $i \in [1, k]$ for which the cycle u_i is simple. Since C' is maximal, for each $i \in T$, all u_i -blocks are contained in one edge $I_i \in C'$. Note that it is allowed that one edge contains the blocks of multiple simple cycles. We partition T into sets $T = T_1 \uplus \cdots \uplus T_t$ so that $i \in T$ and $j \in T$ belong to the same part if and only if the u_i -blocks and the u_j -blocks belong to the same edge of C , i.e. $I_i = I_j$.

For a moment, let us fix an $\ell \in [1, t]$ and let $I \in C'$ be the edge containing all u_i -blocks for all the $i \in T_\ell$. Moreover, let $T_\ell = \{i_1, \dots, i_r\}$. The words \bar{v}_j for $j \in [0, r]$ will collect those blocks that belong to I but are not u_{i_s} -blocks for any $s \in [1, r]$. Formally, \bar{v}_0 consists of all blocks that belong to I that are to the left of all u_{i_1} -blocks. Similarly, \bar{v}_r is the concatenation of all blocks belonging to I that are to the right of all u_{i_r} -blocks. Finally, for $j \in [1, r-1]$, \bar{v}_j consists of all blocks that belong to I and are to the right of all u_{i_j} -blocks and to the left of all $u_{i_{j+1}}$ -blocks. By consistency of C' , for some $s \in \{0, 1\}$, all the words \bar{v}_j for $j \in [0, r]$ and the words u_{i_j} for $j \in [1, r]$ belong to A_s^* and thus represent elements of G_s . Since G_s is knapsack tame, we know that the set

$$S_\ell = \{z \in \mathbb{N}^k \mid \bar{v}_0 u_{i_1}^{z_{i_1}} \bar{v}_1 u_{i_2}^{z_{i_2}} \bar{v}_2 \cdots u_{i_r}^{z_{i_r}} \bar{v}_r \text{ represents } 1 \text{ in } G_s, z_j = 0 \text{ for } j \notin T_\ell\}$$

has magnitude at most $p_s(n)$. Consider the vector $y \in \mathbb{N}^k$ with $y_i = 0$ for $i \in T$ and $y_i = x'_i$ for $i \in [1, k] \setminus T$ (i.e. when u_i is a mixed cycle). We claim that $S' = y + S_1 + \cdots + S_t + \mathbb{P}(x', C')^\oplus$ has magnitude at most $q(n) + p_0(n) + p_1(n) + n$ and satisfies $x \in S' \subseteq S$.

First, since y and the members of S_1, \dots, S_t are non-zero on pairwise disjoint coordinates, the magnitude of $y + S_1 + \cdots + S_t$ is the maximum of $\|y\|$ and the maximal magnitude of S_1, \dots, S_t . Hence, it is bounded by $q(n) + p_0(n) + p_1(n)$. The summand $\mathbb{P}(x', C')^\oplus$ contributes only periods, and their magnitude is bounded by n (recall that they are mixed periods). Thus, the magnitude of S' is at most $p(n) = q(n) + p_0(n) + p_1(n) + n$.

The cancelling property of (x', C') tells us that $x' - y$ is contained in $S_1 + \cdots + S_t$. By the choice of (x', C') , we have $x \in x' + \mathbb{P}(x', C')^\oplus$. Together, this means $x \in S'$. Hence, it remains to show $S' \subseteq S$. To this end, consider a vector $x'' \in y + S_1 + \cdots + S_t$. It differs from x' only in the exponents at simple cycles. Therefore, we can apply essentially the same cancellation to x'' as to x' : we just need to adjust the edges containing the blocks of simple cycles. It is therefore clear that the resulting cancellation C'' has the same compatible mixed periods as C' : $\mathbb{P}(x'', C'') = \mathbb{P}(x', C')$. Thus, by Lemma 4.17, we have $x'' + \mathbb{P}(x', C')^\oplus \subseteq S$. This proves $S' = y + S_1 + \cdots + S_t + \mathbb{P}(x', C')^\oplus \subseteq S$ and hence the proposition. \square

4.4. LogCFL-membership. In this section, we prove the upper bound in Theorem 4.2:

Proposition 4.21. *Let $G = \mathbb{G}(A, I)$ be a graph group where (A, I) is a transitive forest. Then the knapsack problem and the subset sum problem for G belong to LogCFL.*

Proof. Let us consider knapsack. According to Proposition 4.3, it suffices to provide a logspace reduction from the knapsack problem over G to the membership problem for acyclic automata over G . Suppose we have an instance $h_0 g_1^{x_1} h_1 \cdots g_k^{x_k} h_k = 1$ of the knapsack problem over G of size n . Moreover, let h_i be represented by $v_i \in A^*$ for each $i \in [0, k]$ and let g_i be represented by $u_i \in A^*$ for $i \in [1, k]$.

By Theorem 4.6, there is a polynomial p such that the above instance has a solution if and only if it has a solution $x \in \mathbb{N}^k$ with $\|x\| \leq p(n)$. We construct an acyclic automaton $\mathcal{A} = (Q, A, \Delta, q_0, q_f)$ as follows. It has the state set $Q = [0, k+1] \times [0, p(n)]$ and the following transitions. From $(0, 0)$, there is one transition labeled v_0 to $(1, 0)$. For each $i \in [1, k]$ and $j \in [0, p(n) - 1]$, there are two transitions from (i, j) to $(i, j+1)$; one labeled by u_i and one labeled by ε . Furthermore, there is a transition from $(i, p(n))$ to $(i+1, 0)$ labeled v_i for each $i \in [1, k]$. The initial state is $q_0 = (0, 0)$ and the final state is $q_f = (k+1, 0)$.

It is clear that \mathcal{A} accepts a word that represents 1 if and only if the exponent equation has a solution. Finally, the reduction can clearly be carried out in logarithmic space.

For subset sum the same reduction as above works but the polynomial bound on solutions is for free. \square

4.5. LogCFL-hardness. It remains to show the lower bound in Theorem 4.2. If (A, I) is not complete, then (A, I) contains two non-adjacent vertices and thus $\mathbb{G}(A, I)$ contains an isomorphic copy of F_2 , the free group of rank two. Hence, we will show that knapsack and subset sum for F_2 are LogCFL-hard. Let $\{a, b\}$ be a generating set for F_2 . Let $\theta: \{a, b, a^{-1}, b^{-1}\}^* \rightarrow F_2$ be the morphism that maps a word w to the group element represented by w .

A *valence automaton* over a group G is a tuple $\mathcal{A} = (Q, \Sigma, \Delta, q_0, q_f)$ where Q, Σ, q_0, q_f are as in a finite automaton and Δ is a finite subset of $Q \times \Sigma^* \times G \times Q$. The *language accepted by \mathcal{A}* is denoted $L(\mathcal{A})$ and consists of all words $w_1 \cdots w_n$ such that there is a computation $p_0 \xrightarrow{w_1, g_1} p_1 \rightarrow \cdots \rightarrow p_{n-1} \xrightarrow{w_n, g_n} p_n$ such that $(p_{i-1}, w_i, g_i, p_i) \in \Delta$ for $i \in [1, n]$, $p_0 = q_0$, $p_n = q_f$, and $g_1 \cdots g_n = 1$ in G . We call this computation also an *accepting run* of \mathcal{A} for w (of length n). Note that we allow ε -transitions of the form $(p, \varepsilon, g, q) \in \Delta$. This implies that an accepting run for a word w can be of length greater than $|w|$.

An analysis of a proof (in this case [13]) of the Chomsky-Schützenberger theorem yields:

Lemma 4.22. *For every language $L \subseteq \Sigma^*$ the following statements are equivalent:*

- (i) L is context-free.
- (ii) There is a valence automaton \mathcal{A} over F_2 such that $L = L(\mathcal{A})$.
- (iii) There is a valence automaton \mathcal{A} over F_2 such that $L = L(\mathcal{A})$ and for every $w \in L$ there exists an accepting run of \mathcal{A} for w of length at most $c \cdot |w|$.

Proof. The equivalence of (i) and (ii) is well known (see [13]) and the implication from (iii) to (ii) is trivial. We show that (i) implies (iii). For this, we shall use the concept of rational transductions. If Σ and Γ are alphabets, subsets $T \subseteq \Gamma^* \times \Sigma^*$ are called *transductions*. Given a language $L \subseteq \Sigma^*$ and a transduction $T \subseteq \Gamma^* \times \Sigma^*$,




FIGURE 2. Transducer used in proof of Lemma 4.22

we define

$$TL = \{u \in \Gamma^* \mid (u, v) \in T \text{ for some } v \in L\}.$$

A *finite-state transducer* is a tuple $\mathcal{A} = (Q, \Sigma, \Gamma, \Delta, q_0, q_f)$, where Q is a finite set of *states*, Σ is its *input alphabet*, Γ is its *output alphabet*, Δ is a finite subset of $Q \times \Gamma^* \times \Sigma^* \times Q$, $q_0 \in Q$ is its *initial state*, and $q_f \in Q$ is its *final state*. The elements of Δ are called *transitions*. We say that a pair $(u, v) \in \Gamma^* \times \Sigma^*$ is *accepted by* \mathcal{A} if there is a sequence $(p_0, u_1, v_1, p_1), (p_1, u_2, v_2, p_2), \dots, (p_{n-1}, u_n, v_n, p_n)$ of transitions where $n \geq 1$, $p_0 = q_0$, $p_n = q_f$, $u = u_1 \cdots u_n$, and $v = v_1 \cdots v_n$. The set of all pairs $(u, v) \in \Gamma^* \times \Sigma^*$ that are accepted by \mathcal{A} is denoted by $T(\mathcal{A})$. A transduction $T \subseteq \Gamma^* \times \Sigma^*$ is called *rational* if there is a finite-state transducer \mathcal{A} with $T(\mathcal{A}) = T$.

Let $W_2 \subseteq \{a, b, a^{-1}, b^{-1}\}^*$ be the word problem of F_2 , i.e.

$$W_2 = \{w \in \{a, b, a^{-1}, b^{-1}\}^* \mid \theta(w) = 1\}.$$

For languages $K \subseteq \Sigma^*$ and $L \subseteq \Gamma^*$, we write $K \rightsquigarrow L$ if there is a rational transduction T and a constant c such that $K = TL$ and for each $u \in K$, there is a $v \in L$ with $|v| \leq c|u|$ and $(u, v) \in T$. Observe that the relation \rightsquigarrow is transitive, meaning that it suffices to show $L \rightsquigarrow W_2$ for every context-free language L .

Let D_2 be the one-sided Dyck language over two pairs of parentheses, in other words: D_2 is the smallest language $D_2 \subseteq \{x, \bar{x}, y, \bar{y}\}^*$ such that $\varepsilon \in D_2$ and whenever $uv \in D_2$, we also have $uwv \in D_2$ for $w \in \{x\bar{x}, y\bar{y}\}$.

It is easy to see that $L \rightsquigarrow D_2$ for every context-free language L . Indeed, an ε -free pushdown automaton (which exists for every context-free language [9]) for L can be converted into a transducer witnessing $L \rightsquigarrow D_2$. Therefore, it remains to show that $D_2 \rightsquigarrow W_2$.

Let F_3 be the free group of rank 3 and let $\{a, b, \#\}$ be a free generating set for F_3 . As above, let

$$W_3 = \{w \in \{a, b, \#, a^{-1}, b^{-1}, \#^{-1}\}^* \mid w \text{ represents } 1 \text{ in } F_3\}$$

be the word problem of F_2 . Since F_3 can be embedded into F_2 [18, Proposition 3.1], we clearly have $W_3 \rightsquigarrow W_2$. It therefore suffices to show $D_2 \rightsquigarrow W_3$.

For this, we use a construction of Kambites [13]. He proves that if \mathcal{A} is the transducer in Fig. 2 and $T = T(\mathcal{A})$, then $D_2 = TW_3$. Thus, for every $u \in D_2$, we have $(u, v) \in T$ for some $v \in W_3$. An inspection of \mathcal{A} yields that $|v| = 2|u| + |v|_{\#^{-1}}$ and $|v|_{\#} = |u|$. Since $v \in W_3$, we have $|v|_{\#^{-1}} = |v|_{\#}$ and thus $|v| = 3|u|$. Hence, the transduction T witnesses $D_2 \rightsquigarrow W_3$. We have thus shown $L \rightsquigarrow D_2 \rightsquigarrow W_3 \rightsquigarrow W_2$ and hence the lemma. \square

Given w , it is easy to convert the valence automaton \mathcal{A} from Lemma 4.22 into an acyclic automaton that exhausts all computations of \mathcal{A} of length $c \cdot |w|$. This yields the following.

Proposition 4.23. *For F_2 , the membership problem for acyclic automata is LogCFL-hard.*

Proof. Fix a context-free language $L \subseteq \Sigma^*$ with a LogCFL-complete membership problem; such languages exist [8]. Fix a valence automaton \mathcal{A} over F_2 and a constant $c \in \mathbb{N}$ such that the statement of Lemma 4.22 holds for L , \mathcal{A} , and c . Consider a word $w \in \Sigma^*$. From w we construct an acyclic automaton \mathcal{B} over the input alphabet $\{a, b, a^{-1}, b^{-1}\}$ such that $1 \in \theta(L(\mathcal{B}))$ if and only if $w \in L$. Let $m = |w|$, $w = a_1 a_2 \cdots a_m$ and $n = c \cdot m$. The set of states of \mathcal{B} is $[0, m] \times [0, n] \times Q$, where Q is the state set of \mathcal{A} . The transitions of \mathcal{B} are defined as follows:

- $(i-1, j-1, p) \xrightarrow{x} (i, j, q)$ if $p \xrightarrow{(a_i, x)} q$ for all $i \in [1, m]$, $j \in [1, n]$, and $x \in \{a, b, a^{-1}, b^{-1}\}^*$
- $(i, j-1, p) \xrightarrow{x} (i, j, q)$ if $p \xrightarrow{(\varepsilon, x)} q$ for all $i \in [0, m]$, $j \in [1, n]$, and $x \in \{a, b, a^{-1}, b^{-1}\}^*$

The initial state of \mathcal{B} is $(0, 0, q_0)$, where q_0 is the initial state of \mathcal{A} and all states (m, j, q) with $j \in [0, n]$ and q final in \mathcal{A} are final in \mathcal{B} . It is then straightforward to show that $1 \in h(L(\mathcal{B}))$ if and only if $w \in L$. The intuitive idea is that in state of \mathcal{B} we store in the first component the current position in the word w . In this way we enforce to the simulation of a run of \mathcal{A} on input w . In the second component of the state we store the total number of simulated \mathcal{A} -transitions. In this way we make \mathcal{B} acyclic. Finally, the third state component of \mathcal{B} stores the current \mathcal{A} -state. \square

Proposition 4.24. *For F_2 , knapsack and subset sum are LogCFL-hard.*

Proof. Let $\mathcal{A} = (Q, \{a, b, a^{-1}, b^{-1}\}, \Delta, q_0, q_f)$ be an acyclic automaton. We construct words $w, w_1, \dots, w_m \in \{a, b, a^{-1}, b^{-1}\}$ such that the following three statements are equivalent:

- (i) $1 \in \theta(L(\mathcal{A}))$.
- (ii) $\theta(w) \in \theta(w_1^* w_2^* \cdots w_m^*)$.
- (iii) $\theta(w) \in \theta(w_1^{e_1} w_2^{e_2} \cdots w_m^{e_m})$ for some $e_1, e_2, \dots, e_m \in \{0, 1\}$.

W.l.o.g. assume that $Q = \{1, \dots, n\}$, where 1 is the initial state and n is the unique final state of \mathcal{A} .

Let $\alpha_i = a^i b a^{-i}$ for $i \in [1, n+2]$. It is well known that the α_i generate a free subgroup of rank $n+2$ in F_2 [18, Proposition 3.1]. Define the embedding $\varphi: F_2 \rightarrow F_2$ by $\varphi(a) = \alpha_{n+1}$ and $\varphi(b) = \alpha_{n+2}$. For a transition $t = (p, w, q) \in \Delta$ let $\tilde{t} = \alpha_p \varphi(w) \alpha_q^{-1}$. Let $\tilde{\Delta} = \{\tilde{t}_1, \dots, \tilde{t}_m\}$ such that $\tilde{t}_i = (p, a, q)$ and $\tilde{t}_j = (q, b, r)$ implies $i < j$. Since \mathcal{A} is acyclic, such an enumeration must exist. Together with the fact that the α_i generate a free group, it follows that the following three statements are equivalent:

- (i) $1 \in \theta(L(\mathcal{A}))$.
- (ii) $\theta(\alpha_1 \alpha_n^{-1}) \in \theta(\tilde{t}_1^* \tilde{t}_2^* \cdots \tilde{t}_m^*)$.
- (iii) $\theta(\alpha_1 \alpha_n^{-1}) \in \theta(\tilde{t}_1^{e_1} \tilde{t}_2^{e_2} \cdots \tilde{t}_m^{e_m})$ for some $e_1, e_2, \dots, e_m \in \{0, 1\}$.

This shows the proposition. \square

5. NP-COMPLETENESS

In [17], the authors proved that knapsack for the graph group $\mathbb{G}(\mathbf{C4}) \cong F_2 \times F_2$ is NP-complete. Here we extend this result to all graph groups $\mathbb{G}(A, I)$ where (A, I) is not a transitive forest. An *acyclic loop automaton* is a finite automaton $\mathcal{A} = (Q, \Sigma, \Delta, q_0, q_f)$ such that there exists a linear order \preceq on Δ having the property that for all $(p, u, q), (q, v, r) \in \Delta$ it holds $(p, u, q) \preceq (q, v, r)$. Thus, an acyclic loop automaton is obtained from an acyclic automaton by attaching to some of the states a unique loop. For a trace monoid $\mathbb{M}(A, I)$, the *intersection nonemptiness problem for acyclic loop automata* is the following computational problem:

Input: Two acyclic loop automata $\mathcal{A}_1, \mathcal{A}_2$ over the input alphabet A .

Question: Does $[L(\mathcal{A}_1)]_I \cap [L(\mathcal{A}_2)]_I \neq \emptyset$ hold?

Aalbersberg and Hoogeboom [1] proved that for the trace monoid $\mathbb{M}(\mathbf{P4})$ the intersection nonemptiness problem for arbitrary finite automata is undecidable. We use their technique to show:

Lemma 5.1. *For $\mathbb{M}(\mathbf{P4})$, intersection nonemptiness for acyclic loop automata is NP-hard.*

Proof. We give a reduction from 3SAT. Let $\varphi = \bigwedge_{i=1}^m C_i$ where for every $i \in [1, m]$, $C_i = (L_{i,1} \vee L_{i,2} \vee L_{i,3})$ is a clause consisting of three literals. Let x_1, \dots, x_n be the boolean variables that occur in φ . In particular, every literal $L_{i,j}$ belongs to the set $\{x_1, \dots, x_n, \neg x_1, \dots, \neg x_n\}$.

Let p_1, p_2, \dots, p_n be a list of the first n prime numbers. So, for each boolean variable x_i we have the corresponding prime number p_i . We encode a valuation $\beta: \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ by any natural number N such that $N \equiv 0 \pmod{p_i}$ if and only if $\beta(x_i) = 1$. For a positive literal x_i let $S(x_i) = \{p_i \cdot n \mid n \in \mathbb{N}\}$ and for a negative literal $\neg x_i$ let $S(\neg x_i) = \{p_i \cdot n + r \mid n \in \mathbb{N}, r \in [1, p_i - 1]\}$. Moreover, for every $i \in [1, m]$ let $S_i = S(L_{i,1}) \cup S(L_{i,2}) \cup S(L_{i,3})$. Thus, S_i is the set of all numbers that encode a valuation, which makes the clause C_i true. Hence, the set $S = \bigcap_{i=1}^m S_i$ encodes the set of all valuations that make φ true.

We first construct an acyclic loop automaton \mathcal{A}_1 with

$$L(\mathcal{A}_1) = \prod_{i=1}^m \{a(bc)^{N_i}d \mid N_i \in S_i\}.$$

Note that φ is satisfiable iff $[L(\mathcal{A}_1)]_I$ contains a trace from $[\{(a(bc)^N d)^m \mid N \in \mathbb{N}\}]_I$. We will ensure this property with a second acyclic loop automaton \mathcal{A}_2 that satisfies the equality $L(\mathcal{A}_2) = b^*(ad(bc)^*)^{m-1}adc^*$.

We claim that $[L(\mathcal{A}_1)]_I \cap [L(\mathcal{A}_2)]_I = [\{(a(bc)^N d)^m \mid N \in S\}]_I$. First assume that $w \equiv_I (a(bc)^N d)^m$ for some $N \in S$. We have

$$w \equiv_I (a(bc)^N d)^m \equiv_I b^N (ad(bc)^N)^{m-1} adc^N$$

and thus $[w]_I \in [L(\mathcal{A}_2)]_I$. Moreover, since $N \in S$ we get $[w]_I \in [L(\mathcal{A}_1)]_I$. For the other direction, let $[w]_I \in [L(\mathcal{A}_1)]_I \cap [L(\mathcal{A}_2)]_I$. Thus

$$w \equiv_I \prod_{i=1}^m (a(bc)^{N_i}d) \equiv_I b^{N_1} \left(\prod_{i=1}^{m-1} (adc^{N_i} b^{N_{i+1}}) \right) adc^{N_m}$$

where $N_i \in S_i$ for $i \in [1, m]$. Moreover, the fact that $[w]_I \in [L(\mathcal{A}_2)]_I$ means that there are $k_0, k_1, \dots, k_{m-1}, k_m \geq 0$ with

$$\begin{aligned} b^{N_1} \left(\prod_{i=1}^{m-1} (adc^{N_i} b^{N_{i+1}}) \right) adc^{N_m} &\equiv_I b^{k_0} \left(\prod_{i=1}^{m-1} (ad(bc)^{k_i}) \right) adc^{k_m} \\ &\equiv_I b^{k_0} \left(\prod_{i=1}^{m-1} (adb^{k_i} c^{k_i}) \right) adc^{k_m} \end{aligned}$$

Since every symbol is dependent from a or d , this identity implies $N_i = N_{i+1}$ for $i \in [1, m-1]$. Thus, $[w]_I \in [\{(a(bc)^N d)^m \mid N \in S\}]_I$. \square

For a graph group $\mathbb{G}(A, I)$ the *membership problem for acyclic loop automata* is the following computational problem:

Input: An acyclic loop automaton \mathcal{A} over the input alphabet $A \cup A^{-1}$.

Question: Is there a word $w \in L(\mathcal{A})$ such that $w = 1$ in $\mathbb{G}(A, I)$?

It is straightforward to reduce the intersection emptiness problem for acyclic loop automata over $\mathbb{M}(A, I)$ to the membership problem for acyclic loop automata over $\mathbb{G}(A, I)$. For the rest of this section let $\Sigma = \{a, b, c, d, a^{-1}, b^{-1}, c^{-1}, d^{-1}\}$ and let $\theta: \Sigma^* \rightarrow \mathbb{G}(P_4)$ be the canonical homomorphism that maps a word over Σ to the corresponding group element.

Lemma 5.2. *For $\mathbb{G}(P_4)$, the membership problem for acyclic loop automata is NP-hard.*

Proof. The lemma follows easily from Lemma 5.1. Note that $[L(\mathcal{A}_1)]_I \cap [L(\mathcal{A}_2)]_I \neq \emptyset$ if and only if $1 \in \theta(L(\mathcal{A}_1)L(\mathcal{A}_2)^{-1})$ in the graph group $\mathbb{G}(P_4)$. Moreover, it is straightforward to construct from acyclic loop automata \mathcal{A}_1 and \mathcal{A}_2 an acyclic loop automaton for $L(\mathcal{A}_1)L(\mathcal{A}_2)^{-1}$. We only have to replace every transition label w in \mathcal{A}_2 by w^{-1} , then reverse all transitions in \mathcal{A}_2 and concatenate the resulting automaton with \mathcal{A}_1 on the left. \square

We can now use a construction from [16] to reduce knapsack to membership for acyclic loop automata.

Lemma 5.3. *Knapsack for the graph group $\mathbb{G}(P_4)$ is NP-hard.*

Proof. By Lemma 5.2 it suffices to reduce for $\mathbb{G}(P_4)$ the membership problem for acyclic loop automata to knapsack. Let $\mathcal{A} = (Q, \Sigma, \Delta, q_0, q_f)$ be an acyclic loop automaton with transitions $\Delta \subseteq Q \times \Sigma^* \times Q$. W.l.o.g. assume that $Q = \{1, \dots, n\}$.

We reuse a construction from [16], where the rational subset membership problem for $\mathbb{G}(P_4)$ was reduced to the submonoid membership problem for $\mathbb{G}(P_4)$. For a state $q \in Q$ let $\tilde{q} = (ada)^q d (ada)^{-q} \in \Sigma^*$. Let us fix the morphism $\varphi: \Sigma^* \rightarrow \Sigma^*$ with $\varphi(x) = xx$ for $x \in \Sigma$. For a transition $t = (p, w, q) \in \Delta$ let $\tilde{t} = \tilde{p} \varphi(w) \tilde{q}^{-1}$ and define $S = \{\tilde{t} \mid t \in \Delta\}^*$. In [16] it was shown that $1 \in \theta(L(\mathcal{A}))$ if and only if $\theta(\tilde{q}_0 \tilde{q}_f^{-1}) \in \theta(S)$.

We construct in polynomial time a knapsack instance over $\mathbb{G}(P_4)$ from the automaton \mathcal{A} as follows: Let us choose an enumeration t_1, t_2, \dots, t_m of the transitions of \mathcal{A} such that the following holds, where $t_i = (p_i, w_i, q_i)$: If $q_j = p_k$ then $j \leq k$. Since \mathcal{A} is an acyclic loop automaton such an enumeration exists. The following claim proves the theorem.

Claim: $1 \in \theta(L(\mathcal{A}))$ if and only if $\theta(\tilde{q}_0 \tilde{q}_f^{-1}) \in \theta(\tilde{t}_1^* \tilde{t}_2^* \dots \tilde{t}_m^*)$.

One direction is clear: If $\theta(\tilde{q}_0 \tilde{q}_f^{-1}) \in \theta(\tilde{t}_1^* \tilde{t}_2^* \cdots \tilde{t}_m^*)$, then $h(\tilde{q}_0 \tilde{q}_f^{-1}) \in h(S)$. Hence, by [16] we have $1 \in \theta(L(\mathcal{A}))$. On the other hand, if $1 \in \theta(L(\mathcal{A}))$, then there exists a path in A of the form

$$q_0 = s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \cdots s_{k-1} \xrightarrow{a_k} s_k = q_f$$

such that $\theta(a_1 a_2 \cdots a_k) = 1$. Let $(s_{j-1}, a_j, s_j) = t_{i_j}$, where we refer to the above enumeration of all transitions. Then, we must have $i_1 \leq i_2 \leq \cdots \leq i_k$. Moreover, we have

$$\theta(\tilde{q}_0 \tilde{q}_f^{-1}) = \theta(\tilde{q}_0 a_1 a_2 \cdots a_k \tilde{q}_f^{-1}) = \theta(\tilde{t}_{i_1} \tilde{t}_{i_2} \cdots \tilde{t}_{i_k}) \in \theta(\tilde{t}_1^* \tilde{t}_2^* \cdots \tilde{t}_m^*).$$

This proves the claim and hence the theorem. \square

Theorem 5.4. *Let (A, I) be an independence alphabet, which is not a transitive forest. Then, knapsack for the graph group $\mathbb{G}(A, I)$ is NP-complete.*

Proof. If (A, I) is not a transitive forest, then P4 or C4 is an induced subgraph of (A, I) [25]. Thus, $\mathbb{G}(P4)$ or $\mathbb{G}(C4) \cong F_2 \times F_2$ is a subgroup of $\mathbb{G}(A, I)$. Hence, NP-hardness of knapsack for $\mathbb{G}(A, I)$ follows from [17] or Lemma 5.3. \square

REFERENCES

- [1] I. J. Aalbersberg and H. J. Hoogeboom. Characterizations of the decidability of some problems for regular trace languages. *Mathematical Systems Theory*, 22:1–19, 1989.
- [2] I. Agol. The virtual Haken conjecture. Technical report, arXiv.org, 2012. <http://arxiv.org/abs/1204.2810>.
- [3] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. *Inventiones Mathematicae*, 129(3):445–470, 1997.
- [4] J. Crisp and B. Wiest. Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebraic & Geometric Topology*, 4:439–472, 2004.
- [5] M. Elberfeld, A. Jakoby, and T. Tantau. Algorithmic meta theorems for circuit classes of constant and logarithmic depth. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:128, 2011.
- [6] E. Frenkel, A. Nikolaev, and A. Ushakov. Knapsack problems in products of groups. *Journal of Symbolic Computation*, 74:96–108, 2016.
- [7] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Advances in Applied Mathematics*, 38(3):302–323, 2007.
- [8] S. Greibach. The hardest context-free language. *SIAM Journal on Computing*, 2(4):304–310, 1973.
- [9] S. A. Greibach. A new normal-form theorem for context-free phrase structure grammars. *Journal of the ACM*, 12(1):42–52, 1965.
- [10] C. Haase. *On the complexity of model checking counter automata*. PhD thesis, University of Oxford, St Catherine’s College, 2011.
- [11] F. Haglund and D. T. Wise. Coxeter groups are virtually special. *Advances in Mathematics*, 224(5):1890–1903, 2010.
- [12] B. Jenner. Knapsack problems for NL. *Information Processing Letters*, 54(3):169–174, 1995.
- [13] M. Kambites. Formal languages and groups as memory. *Communications in Algebra*, 37:193–208, 2009.
- [14] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [15] D. König, M. Lohrey, and G. Zetsche. Knapsack and subset sum problems in nilpotent, polycyclic, and co-context-free groups. Technical report, arXiv.org, 2015. <http://arxiv.org/abs/1507.05145>.
- [16] M. Lohrey and B. Steinberg. The submonoid and rational subset membership problems for graph groups. *Journal of Algebra*, 320(2):728–755, 2008.
- [17] M. Lohrey and G. Zetsche. Knapsack in graph groups, HNN-extensions and amalgamated products. In *Proceedings of STACS 2016*, volume 47 of *LIPICs*, pages 50:1–50:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

- [18] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer, 1977.
- [19] A. Myasnikov, A. Nikolaev, and A. Ushakov. Knapsack problems in groups. *Mathematics of Computation*, 84:987–1016, 2015.
- [20] C. H. Papadimitriou. On the complexity of integer programming. *Journal of the Association for Computing Machinery*, 28(4):765–768, 1981.
- [21] L. Pottier. Minimal solutions of linear diophantine systems : bounds and algorithms. In *Proc. of 4th International Conference on Rewriting Techniques and Applications (RTA 1991)*, pages 162–173, Berlin, Heidelberg, 1991. Springer.
- [22] I. H. Sudborough. On the tape complexity of deterministic context-free languages. *Journal of the ACM*, 25(3):405–414, 1978.
- [23] H. Vollmer. *Introduction to Circuit Complexity*. Springer, 1999.
- [24] D. T. Wise. Research announcement: the structure of groups with a quasiconvex hierarchy. *Electronic Research Announcements in Mathematical Sciences*, 16:44–55, 2009.
- [25] E. S. Wolk. A note on “The comparability graph of a tree”. *Proceedings of the American Mathematical Society*, 16:17–20, 1965.

UNIVERSITÄT SIEGEN, GERMANY
Email address: lohrey@eti.uni-siegen.de

LSV, CNRS & ENS CACHAN, UNIVERSITÉ PARIS-SACLAY, FRANCE
Email address: zetzsche@lsv.fr