KNAPSACK IN HYPERBOLIC GROUPS

MARKUS LOHREY

Abstract. Recently knapsack problems have been generalized from the integers to arbitrary finitely generated groups. The knapsack problem for a finitely generated group G is the following decision problem: given a tuple $(g, g_1, \ldots, g_k)$ of elements of G, are there natural numbers $n_1, \ldots, n_k \in \mathbb{N}$ such that $g = g_1^{n_1} \cdots g_k^{n_k}$ holds in G? Myasnikov, Nikolaev, and Ushakov proved that for every (Gromov-)hyperbolic group, the knapsack problem can be solved in polynomial time. In this paper, the precise complexity of the knapsack problem for hyperbolic group is determined: for every hyperbolic group G, the knapsack problem belongs to the complexity class LogCFL, and it is LogCFL-complete if G contains a free group of rank two. Moreover, it is shown that for every hyperbolic group G and every tuple $(g, g_1, \ldots, g_k)$ of elements of G the set of all $(n_1, \ldots, n_k) \in \mathbb{N}^k$ such that $g = g_1^{n_1} \cdots g_k^{n_k}$ in G is semilinear and a semilinear representation where all integers are of size polynomial in the total geodesic length of the $g, g_1, \ldots, g_k$ can be computed. Groups with this property are also called knapsack-tame. This enables us to show that knapsack can be solved in LogCFL for every group that belongs to the closure of hyperbolic groups under free products and direct products with $\mathbb{Z}$.

1. Introduction

In [22], Myasnikov, Nikolaev, and Ushakov initiated the investigation of discrete optimization problems, which are usually formulated over the integers, for arbitrary (possibly non-commutative) groups. One of these problems is the *knapsack problem* for a finitely generated group G: The input is a sequence of group elements $g_1, \ldots, g_k, g \in G$ (specified by finite words over the generators of G) and it is asked whether there exists a tuple $(n_1, \ldots, n_k) \in \mathbb{N}^k$ such that $g_1^{n_1} \cdots g_k^{n_k} = g$ in G. For the particular case $G = \mathbb{Z}$ (where the additive notation $n_1 \cdot g_1 + \cdots + n_k \cdot g_k = g$ is usually preferred) this problem is NP-complete (resp., TC^0-complete) if the numbers $g_1, \ldots, g_k, g \in \mathbb{Z}$ are encoded in binary representation [12, 9] (resp., unary notation [2]).

In [22], Myasnikov et al. 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. There is also an encoding of words that corresponds to the binary encoding of integers, so called straight-line programs, and knapsack problems under this encoding have been studied in [18]. In this paper, we only consider the case where input words are explicitly represented. Here is a list of known results concerning the knapsack problem:

- Knapsack can be solved in polynomial time for every hyperbolic group [22]. In [4] this result was extended to free products of any finite number of hyperbolic groups and finitely generated abelian groups.
- There are nilpotent groups of class 2 for which knapsack is undecidable. Examples are direct products of sufficiently many copies of the discrete Heisenberg group $H_3(\mathbb{Z})$ [13], and free nilpotent groups of class 2 and sufficiently high rank [20].

This work has been supported by the DFG research project LO 748/13-1.

- Knapsack for $H_3(\mathbb{Z})$ is decidable [13]. In particular, together with the previous point it follows that decidability of knapsack is not preserved under direct products.
- Knapsack is decidable for every co-context-free group [13], i.e., groups where the set of all words over the generators that do not represent the identity is a context-free language. Lehnert and Schweitzer [15] have shown that the Higman-Thompson groups are co-context-free.
- Knapsack belongs to NP for all virtually special groups (finite extensions of subgroups of graph groups) [19]. The class of virtually special groups is very rich. It contains all Coxeter groups, one-relator groups with torsion, fully residually free groups, and fundamental groups of hyperbolic 3-manifolds. For graph groups (also known as right-angled Artin groups) a complete classification of the complexity of knapsack was obtained in [19]: If the underlying graph contains an induced path or cycle on 4 nodes, then knapsack is NP-complete; in all other cases knapsack can be solved in polynomial time (even in LogCFL).
- Decidability of knapsack is preserved under finite extensions, HNN-extensions over finite associated subgroups and amalgamated free products over finite subgroups [18].

In this paper we further investigate the knapsack problem in hyperbolic groups. The definition of hyperbolic groups requires that all geodesic triangles in the Cayley-graph are δ-slim for a constant δ; see Section 3 for details. The class of hyperbolic groups has several alternative characterizations (e.g., it is the class of finitely generated groups with a linear Dehn function), which gives hyperbolic groups a prominent role in geometric group theory. Moreover, in a certain probabilistic sense, almost all finitely presented groups are hyperbolic [8, 23]. Also from a computational viewpoint, hyperbolic groups have nice properties: it is known that the word problem and the conjugacy problem can be solved in linear time [3, 10]. As mentioned above, knapsack can be solved in polynomial time for every hyperbolic group [22]. Our first main result of this paper provides a precise characterization of the complexity of knapsack for hyperbolic groups: for every hyperbolic group, knapsack belongs to LogCFL, which is the class of all problems that are logspace-reducible to a context-free language. LogCFL has several alternative characterizations, see Section 4 for details. The LogCFL upper bound for knapsack in hyperbolic groups improves the polynomial upper bound shown in [22], and also generalizes a result from [16], stating that the word problem for a hyperbolic group is in LogCFL. For hyperbolic groups that contain a copy of a non-abelian free group (such hyperbolic groups are called non-elementary) it follows from [19] that knapsack is LogCFL-complete. Hyperbolic groups that contain no copy of a non-abelian free group (so called elementary hyperbolic groups) are known to be virtually cyclic, in which case knapsack belongs to nondeterministic logspace (NL), which is contained in LogCFL.

In Section 8 we prove our second main result: for every hyperbolic group G and every tuple $(g, g_1, \ldots, g_k)$ of elements of G the set of all $(n_1, \ldots, n_k) \in \mathbb{N}^k$ such that $g = g_1^{n_1} \cdots g_k^{n_k}$ in G is effectively semilinear. In other words: the set of all solutions of a knapsack instance in G is semilinear. Groups with this property are also called knapsack-semilinear. For the special case $G = \mathbb{Z}$ this is well-known (the set of solutions of a linear equation is Presburger definable and hence semilinear). Clearly, knapsack is decidable for every knapsack-semilinear group (due to the effectiveness assumption). In a series of recent papers it turned out that the class of knapsack-semilinear groups is surprisingly rich. It contains all virtually special groups [17] and all co-context-free group [13] and is closed under the following constructions:

- going to a finitely generated subgroup (this is trivial) and going to a finite group extension [18],
- HNN-extensions over finite associated subgroups and amalgamated free products over finite subgroups [18],
- direct products (in contrast, the class of groups with a decidable knapsack problem is not closed under direct products),
- restricted wreath products [5].

Our proof of the knapsack-semilinearity of a hyperbolic group shows an additional quantitative statement: If the group elements $g, g_1, \ldots, g_k$ are represented by words over the generators and the total length of these words is N, then the set $\{(n_1, \ldots, n_k) \in \mathbb{N}^k \mid g = g_1^{n_1} \cdots g_k^{n_k} \text{ in } G\}$ has a semilinear representation, where all vectors only contain integers of size at most $p(N)$. Here, $p(x)$ is a fixed polynomial that only depends on G. Groups with this property are called knapsack-tame in [19]. In [19], it is shown that the class of knapsack-tame groups is closed under free products and direct products with $\mathbb{Z}$. Using this, we can show in Section 9 that knapsack can be solved in LogCFL for every group that belongs to the closure of hyperbolic groups under free products and direct products with $\mathbb{Z}$.

Recently, it was shown that the compressed version of the knapsack problem, where input words are encoded by straight-line programs, is NP-complete for every infinite hyperbolic group [11].

2. General notations

We assume that the reader is familiar with basic concepts from group theory and formal languages. The empty word is denoted with ε. For a word $w = a_1 a_2 \cdots a_n$ let $|w| = n$ be the length of w, and for $1 \leq i \leq j \leq n$ let $w[i] = a_i$, $w[i:j] = a_i \cdots a_j$, $w[:i] = w[1:i]$ and $w[i:] = w[i:n]$. Moreover, let $w[i:j] = \varepsilon$ for $i > j$.

A set of vectors $A \subseteq \mathbb{N}^k$ is *linear* if there exist vectors $v_0, \ldots, v_n \in \mathbb{N}^k$ such that $A = \{v_0 + \lambda_1 \cdot v_1 + \cdots + \lambda_n \cdot v_n \mid \lambda_1, \ldots, \lambda_n \in \mathbb{N}\}$. The tuple of vectors $(v_0, \ldots, v_n)$ is a *linear represention* of A. Its *magnitude* is the largest number appearing in one the vectors $v_0, \ldots, v_n$. A set $A \subseteq \mathbb{N}^k$ is *semilinear* if it is a finite union of linear sets $A_1, \ldots, A_m$. A *semilinear representation* of A is a list of linear representations for the linear sets $A_1, \ldots, A_m$. Its *magnitude* is the maximal magnitude of the linear representations for the sets $A_1, \ldots, A_m$. The magnitude of a semilinear set A is the smallest magnitude among all semilinear representations of A.

In the context of knapsack problems, we will consider semilinear sets as sets of mappings $f : \{x_1, \ldots, x_k\} \to \mathbb{N}$ for a finite set of variables $X = \{x_1, \ldots, x_k\}$. Such a mapping f can be identified with the vector $(f(x_1), \ldots, f(x_k))$. This allows to use all vector operations (e.g. addition and scalar multiplication) on the set $\mathbb{N}^X$ of all mappings from X to $\mathbb{N}$. The pointwise product $f \cdot g$ of two mappings $f, g \in \mathbb{N}^X$ is defined by $(f \cdot g)(x) = f(x) \cdot g(x)$ for all $x \in X$. Moreover, for mappings $f \in \mathbb{N}^X$, $g \in \mathbb{N}^Y$ with $X \cap Y = \emptyset$ we define $f \oplus g : X \cup Y \to \mathbb{N}$ by $(f \oplus g)(x) = f(x)$ for $x \in X$ and $(f \oplus g)(y) = g(y)$ for $y \in Y$. All operations on $\mathbb{N}^X$ will be extended to subsets of $\mathbb{N}^X$ in the standard pointwise way.

It is well-known that the semilinear subsets of $\mathbb{N}^k$ are exactly the sets definable in *Presburger arithmetic*. These are those sets that can be defined with a first-order formula $\varphi(x_1, \ldots, x_k)$ over the structure $(\mathbb{N}, 0, +, \leq)$ [7]. Moreover, the transformations between such a first-order formula and an equivalent semilinear representation are effective. In particular, the semilinear sets are effectively closed under Boolean operations.

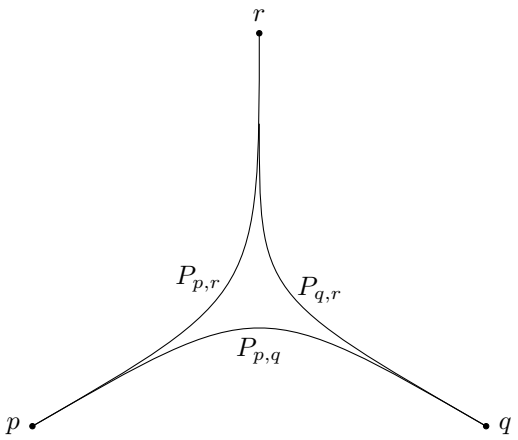

FIGURE 1. The shape of a geodesic triangle in a hyperbolic group

3. Hyperbolic groups

Let G be a finitely generated group with the finite symmetric generating set Σ, i.e., $a \in \Sigma$ implies that $a^{-1} \in \Sigma$. The *Cayley-graph* of G (with respect to Σ) is the undirected graph $\Gamma = \Gamma(G)$ with node set G and all edges (g, ga) for $g \in G$ and $a \in \Sigma$. We view Γ as a geodesic metric space, where every edge (g, ga) is identified with a unit-length interval. It is convenient to label the directed edge from g to ga with the generator a. The distance between two points p, q is denoted with $d_\Gamma(p, q)$. For $g \in G$ let $|g| = d_\Gamma(1, g)$. For $r \geq 0$, let $\mathcal{B}_r(1) = \{g \in G \mid d_\Gamma(1, g) \leq r\}$.

Paths can be defined in a very general way for metric spaces, but we only need paths that are induced by words over Σ. Given a word $w \in \Sigma^*$ of length n, one obtains a unique path $P[w] : [0, n] \to \Gamma$, which is a continuous mapping from the real interval $[0, n]$ to Γ. It maps the subinterval $[i, i+1] \subseteq [0, n]$ isometrically onto the edge (g_i, g_{i+1}) of Γ, where g_i (resp., g_{i+1}) is the group element represented by the word $w[: i]$ (resp., $w[: i+1]$). The path $P[w]$ starts in $1 = g_0$ and ends in g_n (the group element represented by w). We also say that $P[w]$ is the unique path that starts in 1 and is labelled with the word w. More generally, for $g \in G$ we denote with $g \cdot P[w]$ the path that starts in g and is labelled with w. When writing $u \cdot P[w]$ for a word $u \in \Sigma^*$, we mean the path $g \cdot P[w]$, where g is the group element represented by u. A path $P : [0, n] \to \Gamma$ of the above form is geodesic if $d_\Gamma(P(0), P(n)) = n$; it is a (λ, ϵ)-*quasigeodesic* if for all points $p = P(a)$ and $q = P(b)$ we have $|a - b| \leq \lambda \cdot d_\Gamma(p, q) + \varepsilon$; and it is ζ-*local* (λ, ϵ)-*quasigeodesic* if for all points $p = P(a)$ and $q = P(b)$ with $|a - b| \leq \zeta$ we have $|a - b| \leq \lambda \cdot d_\Gamma(p, q) + \varepsilon$.

A word $w \in \Sigma^*$ is geodesic if the path $P[w]$ is geodesic, which means that there is no shorter word representing the same group element from G. Similarly, we define the notion of (ζ-local) (λ, ϵ)-quasigeodesic words. A word $w \in \Sigma^*$ is *shortlex reduced* if it is the length-lexicographically smallest word that represents the same group element as w. For this, we have to fix an arbitrary linear order on Σ. Note that if $u = xy$ is shortlex reduced then x and y are shortlex reduced too. For a word $u \in \Sigma^*$ we denote with $\mathsf{shlex}(u)$ the unique shortlex reduced word that represents the same group element as u.

A *geodesic triangle* consists of three points $p, q, r \in G$ and geodesic paths $P_1 = P_{p,q}$, $P_2 = P_{p,r}$, $P_3 = P_{q,r}$ (the three sides of the triangle), where $P_{x,y}$ is a geodesic path from x to y. We call a geodesic triangle δ-*slim* for $\delta \geq 0$, if for all $i \in \{1, 2, 3\}$, every point on P_i has distance at most δ from a point on $P_j \cup P_k$, where $\{j, k\} = \{1, 2, 3\} \setminus \{i\}$. The group G is called δ-*hyperbolic*, if every geodesic

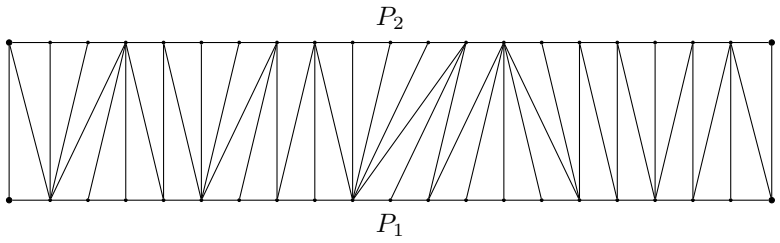

FIGURE 2. Paths that asynchronously K-fellow travel

triangle is δ-slim. Finally, G is hyperbolic, if it is δ-hyperbolic for some $\delta \geq 0$. Figure 1 shows the shape of a geodesic triangle in a hyperbolic group. Finitely generated free groups are for instance 0-hyperbolic. The property of being hyperbolic is independent of the chosen generating set Σ. The word problem for every hyperbolic group can be decided in real time [10].

Let us fix a δ-hyperbolic group G with the finite symmetric generating set Σ for the rest of the section, and let Γ be the corresponding geodesic metric space. We will apply a couple of well-known results for hyperbolic groups.

Lemma 3.1 (c.f. [6, 8.21])**.** *Let $g \in G$ be of infinite order and let $n \geq 0$. Let u be a geodesic word representing g. Then the word u^n is (λ, ϵ)-quasigeodesic, where $\lambda = N|g|$, $\epsilon = 2N^2|g|^2 + 2N|g|$ and $N = |\mathcal{B}_{2\delta}(1)|$.*

Consider two paths $P_1 : [0, n_1] \to \Gamma$, $P_2 : [0, n_2] \to \Gamma$ and let K be a positive real number. We say that P_1 and P_2 *asynchronously K-fellow travel* if there exist two continuous non-decreasing mappings $\varphi_1 : [0, 1] \to [0, n_1]$ and $\varphi_2 : [0, 1] \to [0, n_2]$ such that $\varphi_1(0) = \varphi_2(0) = 0$, $\varphi_1(1) = n_1$, $\varphi_2(1) = n_2$ and for all $0 \leq t \leq 1$, $d_\Gamma(P_1(\varphi_1(t)), P_2(\varphi_2(t))) \leq K$. Intuitively, this means that one can travel along the paths P_1 and P_2 asynchronously with variable speeds such that at any time instant the current points have distance at most K. By slightly increasing K one obtains a ladder graph of the form shown in Figure 2, where the edges connecting the horizontal P_1- and P_2-labelled paths represent paths of length at most K that connect elements from G.

Lemma 3.2 (c.f. [21])**.** *Let P_1 and P_2 be (λ, ϵ)-quasigeodesic paths in Γ_G and assume that P_i starts in g_i and ends in h_i. Assume that $d_\Gamma(g_1, g_2), d_\Gamma(h_1, h_2) \leq h$. Then there exists a computable bound $K = K(\delta, \lambda, \epsilon, h) \geq h$ such that P_1 and P_2 asynchronously K-fellow travel.*

Finally we need the following lemma for splitting quasigeodesic rectangles:

Lemma 3.3. *Fix constants λ, ϵ and let $\kappa = K(\delta, \lambda, \epsilon, 0)$ be taken from Lemma 3.2. Let $v_1, v_2 \in \Sigma^*$ be geodesic words and $u_1, u_2 \in \Sigma^*$ (λ, ϵ)-quasigeodesic words such that $v_1u_1 = u_2v_2$ in G. Consider a factorization $u_1 = x_1y_1$ with $|x_1| \geq \lambda(|v_1|+2\delta+\kappa)+\epsilon$ and $|y_1| \geq \lambda(|v_2| + 2\delta + \kappa) + \epsilon$ Then there exists a factorization $u_2 = x_2y_2$ and $c \in \mathcal{B}_{2\delta+2\kappa}(1)$ such that $v_1x_1 = x_2c$ and $cy_1 = y_2v_2$ in G.*

Proof. The construction is shown in Figure 3.3. Let t_1, t_2, x_1', y_1' be geodesic words with $t_1 = u_1$, $t_2 = u_2$, $x_1 = x_1'$ and $y_1 = y_1'$ in G. Since u_1 is (λ, ϵ)-quasigeodesic, we get $|x_1'| \geq (|x_1| - \epsilon)/\lambda \geq |v_1| + 2\delta + \kappa$ and $|y_1'| \geq (|y_1| - \epsilon)/\lambda \geq |v_2| + 2\delta + \kappa$. By Lemma 3.2 the paths $P[t_1]$ and $P[u_1]$ asynchronously κ-fellow travel. Hence, there exists a factorization $t_1 = r_1s_1$ and $c_1 \in \mathcal{B}_\kappa(1)$ such that $r_1c_1 = x_1 = x_1'$ and $c_1y_1' = c_1y_1 = s_1$ in G. This implies $|r_1| \geq |x_1'| - \kappa \geq |v_1| + 2\delta$ and $|s_1| \geq |y_1'| - \kappa \geq |v_2| + 2\delta$. Consider the geodesic rectangle with the paths $Q_1 = P[v_1]$, $P_1 = v_1 \cdot P[t_1]$, $P_2 = P[t_2]$, and $Q_2 = u_2 \cdot P[v_2]$. Since geodesic rectangles are 2δ-slim, there exists

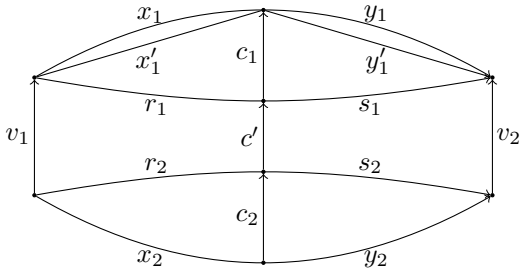

FIGURE 3. Splitting a quasigeodesic rectangle according to Lemma 3.3.

a point $p_2 \in P_2 \cup Q_1 \cup Q_2$ that has distance at most 2δ from $p_1 = P_1(|r_1|)$. By the triangle inequality we must have $p_2 \in P_2$. This yields a factorization $t_2 = r_2 s_2$ (where $p_2 = P_2(|r_2|)$) and $c' \in \mathcal{B}_{2\delta}(1)$ such that $v_1 r_1 = r_2 c'$ and $c' s_1 = s_2 v_2$ in G. Finally, since $P[t_2]$ and $P[u_2]$ asynchronously κ-fellow travel, we obtain a factorization $u_2 = x_2 y_2$ and $c_2 \in \mathcal{B}_\kappa(1)$ such that $x_2 c_2 = r_2$ and $c_2 s_2 = y_2$ in G. Let $c = c_2 c' c_1 \in \mathcal{B}_{2\delta+2\kappa}(1)$. We get $x_2 c = x_2 c_2 c' c_1 = r_2 c' c_1 = v_1 r_1 c_1 = v_1 x_1$ and $c y_1 = c_2 c' c_1 y_1 = c_2 c' s_1 = c_2 s_2 v_2 = y_2 v_2$. □

4. The complexity class LogCFL

The complexity class LogCFL consists of all computational problems that are logspace reducible to a context-free language. The class LogCFL is included in the parallel complexity class NC^2 and has several alternative characterizations (see e.g. [24, 26]):

- logspace bounded alternating Turing-machines with polynomial 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 AuxPDAs 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. The class of languages that are accepted by AuxPDAs is exactly LogCFL [24]. A *one-way* AuxPDA is an AuxPDA that never moves the input head to the left. Hence, in every step, the input head either does not move, or moves to the right.

For a finitely generated group G with the symmetric generating set Σ we define the word problem for G (with respect to Σ) as the set of all words $w \in \Sigma^*$ such that $w = 1$ in G. Let us say that a finitely generated group G belongs to the class OW-AuxPDA if the word problem for G is recognized by a one-way AuxPDA. It is easy to see that the latter property is independent of the generating set of G (this

holds, since the class of languages recognized by one-way AuxPDAs is closed under inverse homomorphisms).

Theorem 4.1. *Every hyperbolic group belongs to the class OW-AuxPDA.*

Proof. Let G be a hyperbolic group. In [16] it is shown that the word problem for G is a growing context-sensitive language, i.e., it can be generated by a grammar where all productions are strictly length-increasing (except for the start production $S \to \varepsilon$). In [1] it was shown that every growing context-sensitive language can be recognized by a one-way AuxPDA in logarithmic space and polynomial time. The result follows. □

Theorem 4.2. *If the groups G and H belong to OW-AuxPDA then also $G * H$ and $G \times \mathbb{Z}$ belong to OW-AuxPDA.*

Proof. The proof is essentially the same as in [19, Lemma 4.8], but is presented for completeness. Let us first consider the group $G \times \mathbb{Z}$. Let $\mathcal{P}(G)$ be a one-way AuxPDA for the word problem of G. The one-way AuxPDA $\mathcal{P}(G \times \mathbb{Z})$ for the word problem of G simulates $\mathcal{P}(G)$ 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, $\mathcal{P}(G \times \mathbb{Z})$ accepts if and only if $\mathcal{P}(G)$ accepts and the $\mathbb{Z}$-component on the work tape is zero.

Next, we consider the group $G * H$. We have one-way AuxPDAs $\mathcal{P}(G)$ and $\mathcal{P}(H)$ for the word problems of G and H, respectively. We can assume that $\mathcal{P}(G)$ (resp., $\mathcal{P}(H)$) accepts an input word w if after reading w the stack is empty and $\mathcal{P}(G)$ (resp., $\mathcal{P}(H)$) is in the unique final state q_G (resp., q_H). This can be achieved by doing ε-transitions at the end of the computation. In the following, we call q_G (resp., q_H) the 1-state of $\mathcal{P}(G)$ (resp., $\mathcal{P}(H)$).

Let Σ (resp., Γ) be the input alphabet of $\mathcal{P}(G)$ (resp., $\mathcal{P}(H)$), which is a symmetric generating set for G (resp., H). We assume that $\Sigma \cap \Gamma = \emptyset$. Consider now an input word $w \in (\Sigma \cup \Gamma)^*$. Let us assume that $w = u_1 v_1 u_2 v_2 \cdots u_k v_k$ with $u_i \in \Sigma^+$ and $v_i \in \Gamma^+$ (other cases can be treated analogously). The AuxPDA $\mathcal{P}(G * H)$ starts with empty stack and simulates the AuxPDA $\mathcal{P}(G)$ on the prefix u_1. If it turns out that $u_1 = 1$ in G (which means that $\mathcal{P}(G)$ is in its 1-state and the stack is empty) then the AuxPDA $\mathcal{P}(G * H)$ continues with simulating $\mathcal{P}(H)$ on v_1. On the other hand, if $u_1 \neq 1$ in G, then $\mathcal{P}(G * H)$ pushes the state together with the work tape content of $\mathcal{P}(G)$ reached after reading u_1 on the stack (on top of the final stack content of $\mathcal{P}(G)$). This allows $\mathcal{P}(G * H)$ to resume the computation of $\mathcal{P}(G)$ later. Then $\mathcal{P}(G * H)$ continues with simulating $\mathcal{P}(H)$ on v_1.

The computation of $\mathcal{P}(G * H)$ will continue in this way. More precisely, if after reading u_i (resp. v_i with $i < k$) the AuxPDA $\mathcal{P}(G)$ (resp. $\mathcal{P}(H)$) 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 $\mathcal{P}(H)$ (resp. $\mathcal{P}(G)$), q is a state of $\mathcal{P}(H)$ (resp. $\mathcal{P}(G)$) and t is a work tape content of $\mathcal{P}(H)$ (resp. $\mathcal{P}(G)$).

In case (i), $\mathcal{P}(G * H)$ continues with the simulation of $\mathcal{P}(H)$ (resp. $\mathcal{P}(G)$) on the word v_i (resp. u_{i+1}) in the initial configuration. In case (ii), $\mathcal{P}(G * H)$ continues with the simulation of $\mathcal{P}(H)$ (resp. $\mathcal{P}(G)$) 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 $\mathcal{P}(G)$ (resp. $\mathcal{P}(H)$) is not in its 1-state then $\mathcal{P}(G * H)$ pushes on the stack the state and work tape content of $\mathcal{P}(G)$ reached after its simulation on u_i. This concludes the description of the AuxPDA $\mathcal{P}(G * H)$. It is a one-way AuxPDA that accepts the word problem of $G * H$. □

5. Knapsack problems

Let G be a finitely generated group with the finite symmetric generating set Σ. Moreover, let X be a set of formal variables that take values from $\mathbb{N}$. For a subset $U \subseteq X$, we use $\mathbb{N}^U$ to denote the set of maps $\nu\colon U \to \mathbb{N}$, which we call *valuations*. An *exponent expression* over G is a formal expression of the form $E = u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k$ with $k \geq 1$ and words $u_i, v_i \in \Sigma^*$. Here, the variables do not have to be pairwise distinct. If every variable in an exponent expression occurs at most once, it is called a *knapsack expression*. Let $X_E = \{x_1, \ldots, x_k\}$ be the set of variables that occur in E. For a valuation $\nu \in \mathbb{N}^U$ such that $X_E \subseteq U$ (in which case we also say that ν is a valuation for E), we define $\nu(E) = u_1^{\nu(x_1)} v_1 u_2^{\nu(x_2)} v_2 \cdots u_k^{\nu(x_k)} v_k \in \Sigma^*$. We say that ν is a *solution* of the equation $E = 1$ if $\nu(E)$ evaluates to the identity element 1 of G. With $\mathsf{sol}(E)$ we denote the set of all solutions $\nu \in \mathbb{N}^{X_E}$ of E. We can view $\mathsf{sol}(E)$ as a subset of $\mathbb{N}^k$. The *length* of E is defined as $|E| = \sum_{i=1}^{k} |u_i| + |v_i|$, whereas k is its *depth*. We define *solvability of exponent equations over* G as the following decision problem:

Input: A finite list of exponent expressions $E_1, \ldots, E_n$ over G.
Question: Is $\bigcap_{i=1}^{n} \mathsf{sol}(E_i)$ non-empty?

The *knapsack problem for* G is the following decision problem:

Input: A single knapsack expression E over G.
Question: Is $\mathsf{sol}(E)$ non-empty?

It is easy to observe that the concrete choice of the generating set Σ has no influence on the decidability and complexity status of these problems. Later, we will also allow exponent expressions of the form $v_0 u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k$, which do not start with a power $u_1^{x_1}$. Such an exponent expression can be replaced by $u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k v_0$ without changing the set of solutions.

The group G is called *knapsack-semilinear* if for every knapsack expression E over G, the set $\mathsf{sol}(E)$ is a semilinear set of vectors and a semilinear representation can be effectively computed from E. Since the emptiness of the intersection of finitely many semilinear sets is decidable, solvability of exponent equations is decidable for every knapsack-semilinear group. As mentioned in the introduction, the class of knapsack-semilinear groups is very rich. An example of a group G, where knapsack is decidable but solvability of exponent equations is undecidable is the Heisenberg group $H_3(\mathbb{Z})$ (which consists of all upper triangular (3×3)-matrices over the integers, where all diagonal entries are 1), see [13]. In particular, $H_3(\mathbb{Z})$ is not knapsack-semilinear.

The group G is called *polynomially knapsack-bounded* if there is a fixed polynomial $p(n)$ such that for a given a knapsack expression E over G, one has $\mathsf{sol}(E) \neq \emptyset$ if and only if there exists $\nu \in \mathsf{sol}(E)$ with $\nu(x) \leq p(|E|)$ for all variables x in E.

The group G is called *knapsack-tame* if there is a fixed polynomial $p(n)$ such that for a given a knapsack expression E over G one can compute a semilinear representation for $\mathsf{sol}(E)$ of magnitude at most $p(|E|)$. Thus, every knapsack-tame group is knapsack-semilinear as well as polynomially knapsack-bounded. The following result was shown in [19]:

Proposition 5.1 ([19, Proposition 4.11 and 4.17])**.** *If G and H are knapsack-tame groups then also the free product $G * H$ and the direct product $G \times \mathbb{Z}$ are knapsack-tame.*

6. Membership for acyclic automata

An *acyclic NFA* is a nondeterministic finite automaton $\mathcal{A} = (Q, \Sigma, \Delta, q_0, F)$ (Q is a finite set of states, Σ is the input alphabet, $\Delta \subseteq Q \times \Sigma^* \times Q$ is the set of transition triples, $q_0 \in Q$ is the initial state, and $F \subseteq Q$ is the set of final states) such that the

relation $\{(p,q) \in Q \times Q \mid \exists w \in \Sigma^* : (p,w,q) \in \Delta\}$ is acyclic. Note that we allow transitions labelled with words, which will be convenient in the following.

Let G be a finitely generated group with the finite symmetric generating set Σ. The *membership problem for acyclic NFAs over G* is the following computational problem:

Input: an acyclic NFA $\mathcal{A}$ with input alphabet Σ.
Question: does $\mathcal{A}$ accept a word $w \in \Sigma^*$ such that $w = 1$ in G?

Again, the concrete choice of the generating set Σ has no influence on the decidability and complexity status of this problem.

Theorem 6.1. *If the group G belongs to the class OW-AuxPDA, then membership for acyclic NFAs over G belongs to* LogCFL.

Proof. Let $\mathcal{P}$ be a one-way AuxPDA for the word problem of G. An AuxPDA for the membership problem for acyclic NFAs over G guesses a path in the acyclic input NFA $\mathcal{A}$ and thereby simulates the AuxPDA $\mathcal{P}$ on the word spelled by the guessed path. If the final state of the input NFA $\mathcal{A}$ is reached and the AuxPDA $\mathcal{P}$ accepts at the same time, then the overall AuxPDA accepts. It is important that the AuxPDA $\mathcal{P}$ works one-way since the guessed path in $\mathcal{A}$ cannot be stored in logspace. This implies that the AuxPDA cannot re-access the input symbols that have already been processed. Also note that the AuxPDA is logspace bounded and polynomially time bounded since $\mathcal{A}$ is acyclic. □

Theorem 6.2. *Let G be a polynomially knapsack-bounded group. Then there is a logspace reduction from the knapsack problem for G to membership for acyclic NFAs over G.*

Proof. Let G be a polynomially knapsack-bounded group with the symmetric generating set Σ. We present a logspace reduction from knapsack for G to the membership problem for acyclic NFAs. Consider a knapsack expression $E = u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k$ over G. Since G is polynomially knapsack-bounded, there exists a polynomial $p(x)$ such that $\mathsf{sol}(E) \neq \emptyset$ if and only if there exists a solution $\nu \in \mathsf{sol}(E)$ such that $\nu(x_i) \leq p(|E|)$ for all $1 \leq i \leq k$. We now construct an NFA $\mathcal{A}$ as follows: It has the state set $Q = [1, k+1] \times [0, p(n)]$ and the following transitions. 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 $(1,0)$ and the unique final state is $(k+1, 0)$.

It is clear that $\mathcal{A}$ accepts a word that represents 1 if and only if $\mathsf{sol}(E) \neq \emptyset$. Finally, the NFA can be clearly computed in logarithmic space from E. □

7. Complexity of knapsack in hyperbolic groups

In this section we consider the complexity of the knapsack problem for a hyperbolic group. In [22] it was shown that for every hyperbolic group, knapsack can be solved in polynomial time. Here, we improve the complexity to LogCFL. We need one more result from [22]:

Theorem 7.1 (c.f. [22])**.** *Every hyperbolic group is polynomially knapsack-bounded.*

This result is also a direct corollary of Theorem 8.1 from the next section, stating that every hyperbolic group is knapsack-tame.

We can now easily derive the following two results:

Corollary 7.2. *Membership for acyclic NFAs over a hyperbolic group belongs to* LogCFL.

Proof. This follows from Theorem 4.1 and 6.1. □

Corollary 7.3. *For every hyperbolic groups G, knapsack can be solved in* LogCFL*. Moreover, if G contains a copy of F_2 (the free group of rank 2) then knapsack for G is* LogCFL*-complete.*

Proof. The first statement follows from Theorems 6.2 and 7.1 and Corollary 7.2. The second statement follows from [19, Proposition 4.26], where it was shown that knapsack for F_2 is LogCFL-complete. □

8. Hyperbolic groups are knapsack-semilinear

In this section, we prove the following strengthening of Theorem 7.1:

Theorem 8.1. *Every hyperbolic group is knapsack-tame.*

Let us remark that the total number of vectors in a semilinear representation can be exponential, even for the simplest case $G = \mathbb{Z}$. Take the (additively written) knapsack expression $E = x_1 + x_2 + \cdots + x_n - n$. Then $\mathsf{sol}(E)$ is finite and consists of $\binom{2n-1}{n} \geq 2^n$ vectors.

Let us fix a δ-hyperbolic group G for the rest of Section 8 and let Σ be a finite symmetric generating set for G.

8.1. **Knapsack expressions of depth two.** We first consider knapsack expressions of depth 2 where all powers are quasigeodesic. It is well known that the semilinear sets are exactly the Parikh images of the regular languages. We need a quantitative version of this result that was independently discovered by Kopczynski and Lin:

Theorem 8.2 (c.f. [25, Theorem 4.1], see also [14])**.** *Let k be a fixed constant. Given an NFA $\mathcal{A}$ over an alphabet of size k with n states, one can compute in polynomial time a semilinear representation of the Parikh image of $L(\mathcal{A})$. Moreover, all numbers appearing in the semilinear representation are polynomially bounded in n (in other words: one can compute the semilinear representation with unary encoded numbers).*

Lemma 8.3. *Let λ and ϵ be fixed constants. For all geodesic words $u_1, v_1, u_2, v_2 \in \Sigma^*$ such that $u_1 \neq \varepsilon \neq u_2$ and u_1^n, u_2^n are (λ, ϵ)-quasigeodesic for all $n \geq 0$, the set $\{(x_1, x_2) \in \mathbb{N} \times \mathbb{N} \mid v_1 u_1^{x_1} = u_2^{x_2} v_2 \text{ in } G\}$ is semilinear. Moreover, one can compute a semi-linear representation whose magnitude is bounded by $p(|u_1|+|v_1|+|u_2|+|v_2|)$ for a fixed polynomial $p(n)$.*

Proof. Let $S := \{(x_1, x_2) \in \mathbb{N} \times \mathbb{N} \mid v_1 u_1^{x_1} = u_2^{x_2} v_2 \text{ in } G\}$. We will define an NFA $\mathcal{A}$ over the alphabet $\{a_1, a_2\}$ such that the Parikh image of $L(A)$ is S. Moreover, the number of states of $\mathcal{A}$ is polynomial in $|u_1| + |u_2| + |v_1| + |v_2|$. This allows us to apply Theorem 8.2. We will allow transitions that are labelled with words (having length polynomial in $|u_1| + |u_2| + |v_1| + |v_2|$). Moreover, instead of writing in the transitions these words, we write their Parikh images (so, for instance, a transition $p \xrightarrow{a_1^2 a_2^3} q$ is written as $p \xrightarrow{(2,3)} q$.

Let $\ell_i = |u_i|$ and $m_i = |v_i|$. Take the constant κ from Lemma 3.3 and define $N_1 = \lambda(m_1 + 2\delta + \kappa) + \epsilon$ and $N_2 = \lambda(m_2 + 2\delta + \kappa) + \epsilon$. We split the set S into two parts:

- $S_1 = S \cap \{(n_1, n_2) \in \mathbb{N} \times \mathbb{N} \mid n_1 < (N_1 + N_2)/\ell_1\}$
- $S_2 = S \cap \{(n_1, n_2) \in \mathbb{N} \times \mathbb{N} \mid n_1 \geq (N_1 + N_2)/\ell_1\}$

For all $(n_1, n_2) \in S_1$ we have $|u_1^{n_1}| = n_1 \ell_1 < N_1 + N_2$. Hence, $|\mathsf{shlex}(u_2^{n_2})| = |\mathsf{shlex}(v_1 u_1^{n_1} v_2^{-1})| < N_1 + N_2 + m_1 + m_2$. Since $u_2^{n_2}$ is (λ, ϵ)-quasigeodesic we get $|u_2^{n_2}| = n_2 \ell_2 < \lambda(N_1 + N_2 + m_1 + m_2) + \epsilon$, i.e., $n_2 < (\lambda(N_1 + N_2 + m_1 + m_2) + \epsilon)/\ell_2$.

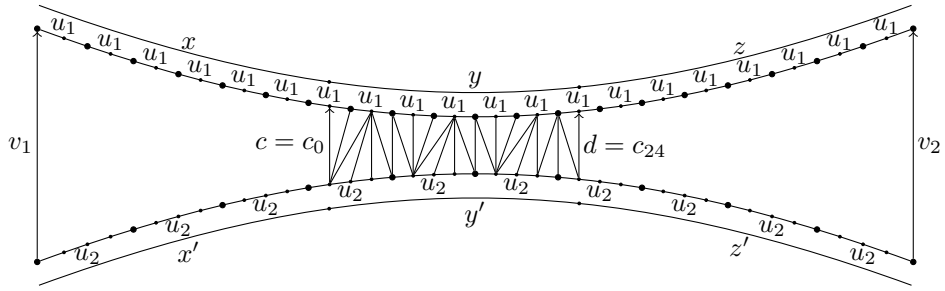

FIGURE 4. Example for the construction from the proof of Lemma 8.3.

Hence, the set S_1 is finite and has a semilinear representation where all numbers are bounded by $\mathcal{O}(m_1 + m_2)$.

We now deal with pairs $(n_1, n_2) \in S_2$, where $v_1 u_1^{n_1} = u_2^{n_2} v_2$ in G and $n_1 \geq (N_1 + N_2)/\ell_1$, i.e., $|u_1^{n_1}| \geq N_1 + N_2$. Consider such a pair (n_1, n_2) and the quasigeodesic rectangle consisting of the four paths $Q_1 = P[v_1]$, $P_1 = v_1 \cdot P[u_1^{n_1}]$, $P_2 = P[u_2^{n_2}]$, and $Q_2 = u_2^{n_2} \cdot P[v_2]$. We factorize the word $u_1^{n_1}$ as $u_1^{n_1} = xyz$ with $|x| = N_1$ and $|z| = N_2$. By Lemma 3.3 we can factorize $u_2^{n_2}$ as $u_2^{n_2} = x'y'z'$ such that there exist $c, d \in \mathcal{B}_{2\delta+2\kappa}(1)$ with $v_1 x = x'c$ and $dz = z'v_2$ in G, see Figure 4 (where $n_1 = 20$, $n_2 = 10$, $\ell_1 = 2$ and $\ell_2 = 4$). Since $u_2^{n_2}$ is (λ, ϵ)-quasigeodesic, we have

$$(1) \qquad |x'| \leq \lambda(m_1 + |x| + 2\delta + 2\kappa) + \epsilon = \lambda(m_1 + N_1 + 2\delta + 2\kappa) + \epsilon,$$

$$(2) \qquad |z'| \leq \lambda(m_2 + |z| + 2\delta + 2\kappa) + \epsilon = \lambda(m_2 + N_2 + 2\delta + 2\kappa) + \epsilon.$$

Consider now the subpath P_1' of P_1 from $P_1(|x|)$ to $P_1(n_1\ell_1 - |z|)$ and the subpath P_2' of P_2 from $P_2(|x'|)$ to $P_2(n_2\ell_2 - |z'|)$. These are the paths labelled with y and y', respectively, in Figure 4. By Lemma 3.2 these paths asynchronously γ-fellow travel, where $\gamma := K(\delta, \lambda, \epsilon, 2\delta + 2\kappa)$ is a constant. In Figure 4 this is visualized by the part between the c-labelled edge and the d-labelled edge. W.l.o.g. we assume that $\gamma \geq 2\delta + 2\kappa$.

We now define the NFA $\mathcal{A}$ over the alphabet $\{a_1, a_2\}$ (recall the we replace edge labels from $\{a_1, a_2\}^*$ by their Parikh images). The state set of $\mathcal{A}$ is

$$Q = \{q_0, q_f\} \cup \{(i, b, j) \mid 0 \leq i < \ell_1, 0 \leq j < \ell_2, b \in \mathcal{B}_\gamma(1)\}.$$

The unique initial state is q_0 and the unique final state is q_f. To define the transitions of $\mathcal{A}$ set $p = \lfloor N_1/\ell_1 \rfloor = \lfloor |x|/|u_1| \rfloor$, $r = N_1 \bmod \ell_1 = |x| \bmod |u_1|$, $s = \lceil N_2/\ell_1 \rceil = \lceil |z|/|u_1| \rceil$, $t = -N_2 \bmod \ell_1 = -|z| \bmod |u_1|$. Thus, we have $x = u_1^p u_1[:r]$ and $z = u_1^s[t+1:]$. There are the following types of transitions (transitions without a label are implicitly labelled by the zero vector $(0,0)$), where $0 \leq i < \ell_1$, $0 \leq j < \ell_2$, $b, b' \in \mathcal{B}_\gamma(1)$.

(1) $q_0 \xrightarrow{(p,p')} (r, c, r')$ if there exists a number $0 \leq k \leq \lambda(m_1 + N_1 + 2\delta + 2\kappa) + \epsilon$ (this is the possible range for the length of x' in (1)) such that $p' = \lfloor k/\ell_2 \rfloor$, $r' = k \bmod \ell_2$, and $v_1 u_1^p u_1[:r] = u_2^{p'} u_2[:r']c$ in G.
(2) $(i, b, j) \to (i+1, b', j)$ if $i + 1 < \ell_1$ and $bu_1[i+1] = b'$ in G.
(3) $(\ell_1 - 1, b, j) \xrightarrow{(1,0)} (0, b', j)$ if $bu_1[\ell_1] = b'$ in G.
(4) $(i, b, j) \to (i, b', j+1)$ if $j + 1 < \ell_2$ and $b = u_2[j+1]b'$ in G.
(5) $(i, b, \ell_2 - 1) \xrightarrow{(0,1)} (i, b', 0)$ if $b = u_2[\ell_2]b'$ in G.
(6) $(t, d, t') \xrightarrow{(s,s')} q_f$ if there exists a number $0 \leq k \leq \lambda(m_2 + N_2 + 2\delta + 2\kappa) + \epsilon$ (this is the possible range for the length of z' in (2)) such that $s' = \lceil k/\ell_2 \rceil$, $t' = -k \bmod \ell_2$, and $du_1[t+1:]u_1^s = u_2[t'+1:]u_2^{s'}v_2$ in G.

The construction is best explained using the example in Figure 4. As mentioned above, the vertical lines between $c = c_0$ and $d = c_{24}$ represent the asynchronous γ-fellow travelling. The vertical lines are labelled with group elements $c_0, c_1, \ldots, c_{23}, c_{24} \in \mathcal{B}_\gamma(1)$ from left to right. In order to not overload the figure we only show c_0 and c_{24}. Note that $x = u_1^6 u_1[1]$, $x' = u_2^3 u_2[1]$, $z = u_1^8[2:]$, $z' = u_2^4[2:]$. Basically, the NFA $\mathcal{A}$ moves the vertical edges from left to right and thereby stores (i) the label c_i of the vertical edge, (ii) the position in the current u_2-factor where the vertical edge starts (position 0 means that we have just completed a u_2-factor), and (iii) the position in the current u_1-factor where the vertical edge ends. If a u_1-factor (resp., u_2-factor) is completed then the automaton makes a $(1,0)$-labelled (resp., $(0,1)$-labelled) transition. The automaton run corresponding to Figure 4 is:

$$\begin{aligned}
q_0 \xrightarrow{(6,3)} &(1, c_0, 1) \xrightarrow{(1,0)} (0, c_1, 1) \to (1, c_2, 1) \to (1, c_3, 2) \to (1, c_4, 3) \xrightarrow{(0,1)} \\
&(1, c_5, 0) \xrightarrow{(1,0)} (0, c_6, 0) \to (0, c_7, 1) \to (1, c_8, 1) \xrightarrow{(1,0)} (0, c_9, 1) \to \\
&(1, c_{10}, 1) \to (1, c_{11}, 2) \to (1, c_{12}, 3) \xrightarrow{(0,1)} (1, c_{13}, 0) \xrightarrow{(1,0)} (0, c_{14}, 0) \to \\
&(0, c_{15}, 1) \to (1, c_{16}, 1) \xrightarrow{(1,0)} (0, c_{17}, 1) \to (1, c_{18}, 1) \to (1, c_{19}, 2) \to \\
&(1, c_{20}, 3) \xrightarrow{(1,0)} (0, c_{21}, 3) \xrightarrow{(0,1)} (0, c_{22}, 0) \to (0, c_{23}, 1) \to (1, c_{24}, 1) \xrightarrow{(8,4)} q_f
\end{aligned}$$

With the above intuition it is straightforward to show that the Parikh image of $L(\mathcal{A})$ is indeed S_2. Also note that the number of states of $\mathcal{A}$ is bounded by $\mathcal{O}(\ell_1 \ell_2)$. The statement of the lemma then follows directly from Theorem 8.2. □

8.2. **Reduction to quasi-geodesic knapsack expressions.** Let us call a knapsack expression $E = u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k$ over G (λ, ϵ)-quasigeodesic if all words $u_1, \ldots, u_k, v_1, \ldots, v_k$ are geodesic and for all $1 \leq i \leq k$ and all $n \geq 0$ the word u_i^n is (λ, ϵ)-quasigeodesic. We say that E has *infinite order*, if all u_i represent group elements of infinite order. The goal of this section is to reduce a knapsack expression to a finite number (in fact, exponentially many) of (λ, ϵ)-quasigeodesic knapsack expressions of infinite order for certain constants λ, ϵ:

Proposition 8.4. *There exist fixed constants λ, ϵ such that from a given knapsack expression E over G one can compute a finite list of knapsack expressions E_i ($i \in I$) over G such that*

$$\mathsf{sol}(E) = \bigcup_{i \in I} \big((m_i \cdot \mathsf{sol}(E_i) + d_i) \oplus \mathcal{F}_i \big),$$

where the following additional properties hold:

- *every $\mathcal{F}_i$ is a semilinear subset of $\mathbb{N}^Y$ for a subset $Y \subseteq X_E$,*
- *the magnitude of every $\mathcal{F}_i$ is bounded by a constant that only depends on G,*
- *every E_i is a (λ, ϵ)-quasigeodesic knapsack expression of infinite order with variables from $Z := X_E \setminus Y$,*
- *the size of every E_i is bounded by $\mathcal{O}(|E|)$, and*
- *all m_i and d_i are vectors from $\mathbb{N}^Z$ where all entries are bounded by a constant that only depends on G (here, $m_i \cdot \mathsf{sol}(E_i) = \{m_i \cdot z \mid z \in \mathsf{sol}(E)\}$ and $m_i \cdot z$ is the pointwise multiplication of the vectors m_i and z).*

Once Proposition 8.4 is shown, we can conclude the proof of Theorem 8.1 by showing that all sets $\mathsf{sol}(E_i)$ are semilinear and that their magnitudes are bounded by $p(|E_i|)$ for a fixed polynomial $p(n)$. This will be achieved in the next section.

For the proof of Proposition 8.4 we mainly build on results from [3]. We fix the constants $L = 34\delta + 2$ and $K = |\mathcal{B}_{4\delta}(1)|^2$.

Lemma 8.5 (c.f. [3, Lemma 3.1]). *Let $u = u_1u_2$ be shortlex reduced, where $|u_1| \leq |u_2| \leq |u_1| + 1$. Let $\tilde{u} = \mathsf{shlex}(u_2u_1)$. If $|\tilde{u}| \geq 2L + 1$ then for every $n \geq 0$, the word $\tilde{u}^n$ is L-local $(1, 2\delta)$-quasigeodesic.*

The following lemma is not stated explicitly in [3] but is shown in Section 3.2 (where the main argument is attributed to Delzant).

Lemma 8.6 (c.f. [3]). *Let u be geodesic such that $|u| \geq 2L + 1$ and for every $n \geq 0$, the word u^n is L-local $(1, 2\delta)$-quasigeodesic. Then one can compute $c \in \mathcal{B}_{4\delta}(1)$ and an integer $1 \leq m \leq K$ such that $(\mathsf{shlex}(c^{-1}u^mc))^n$ is geodesic for all $n \geq 0$.*

Proof of Proposition 8.4. We set $\lambda = N(2L+1)$ and $\epsilon = 2N^2(2L+1)^2 + 2N(2L+1)$, where $N = |\mathcal{B}_{2\delta}(1)|$. Consider a knapsack expression $E = u_1^{x_1}v_1u_2^{x_2}v_2 \cdots u_k^{x_k}v_k$. We can assume that every u_i is shortlex reduced. Let $g_i \in G$ be the group element represented by the word u_i.

Step 1. In this first step we show how to reduce to the case where all g_i have infinite order. In a hyperbolic group G the order of torsion elements is bounded by a fixed constant that only depends on G, see also the proof of [22, Theorem 6.7]). This allows to check for each g_i whether it has finite order, and to compute the order in the positive case. Let $Y \subseteq \{x_1, \ldots, x_k\}$ be those variables x_i such that g_i has finite order. For $x_i \in Y$ let $o_i < \infty$ be the order of g_i. Let $\mathcal{F}$ be the set of mappings $f : Y \to \mathbb{N}$ such that $0 \leq f(x_i) < o_i$ for all $x_i \in Y$. For every such mapping $f \in \mathcal{F}$ let E_f be the knapsack expression that is obtained from E by replacing for every $x_i \in Y$ the power $u_i^{x_i}$ by $u_i^{f(x_i)}$ (which is merged with the word v_i). Moreover, let $\mathcal{F}_f$ be the set of all mappings $g : Y \to \mathbb{N}$ such that $g(x_i) \equiv f(x_i) \bmod o_i$ for every $x_i \in Y$. Then the set $\mathsf{sol}(E)$ can be written as

$$\mathsf{sol}(E) = \bigcup_{f \in \mathcal{F}} \mathsf{sol}(E_f) \oplus \mathcal{F}_f.$$

Note that $\mathcal{F}_f$ is a semilinar set of magnitude $\mathcal{O}(1)$.

Step 2. We now consider a knapsack expression from $\mathcal{F}_f$. To simplify notation, we denote this expression again with $E = u_1^{x_1}v_1u_2^{x_2}v_2 \cdots u_k^{x_k}v_k$. For every i, the group element g_i represented by u_i has infinite order. We factorize u_i uniquely as $u_i = u_{i,1}u_{i,2}$ where $|u_{i,1}| \leq |u_{i,2}| \leq |u_{i,1}| + 1$, and let $\tilde{u}_i = \mathsf{shlex}(u_{i,2}u_{i,1})$. Note that $|\tilde{u}_i| \leq |u_i|$. Let $\tilde{g}_i$ be the group element represented by $\tilde{u}_i$. Since $\tilde{g}_i$ is conjugated to g_i, also $\tilde{g}_i$ has infinite order. By Lemma 3.1, for every $n \geq 0$, the word $\tilde{u}_i^n$ is (λ_i, ϵ_i)-quasigeodesic for $\lambda_i = N|\tilde{u}_i|$, $\epsilon_i = 2N^2|\tilde{u}_i|^2 + 2N|\tilde{u}_i|$. If $|\tilde{u}_i| < 2L + 1$ then $\tilde{u}_i^n$ is (λ, ϵ)-quasigeodesic for the constants λ and ϵ defined at the beginning of the proof. We then replace $u_i^{x_i}$ by $u_{i,1}\tilde{u}_i^{x_i}u_{i,1}^{-1}$. Note that for every $n \geq 0$, $u_{i,1}\tilde{u}_i^nu_{i,1}^{-1} = u_{i,1}(u_{i,2}u_{i,1})^nu_{i,1}^{-1} = (u_{i,1}u_{i,2})^n = u_i^n$ in G.

Now assume that $|\tilde{u}_i| \geq 2L + 1$. By Lemma 8.5, $\tilde{u}_i^n$ is L-local $(1, 2\delta)$-quasigeodesic for every $n \geq 0$. By Lemma 8.6, one can compute $c_i \in \mathcal{B}_{4\delta}(1)$ and an integer $1 \leq m_i \leq K$ such that $(\mathsf{shlex}(c_i^{-1}\tilde{u}_i^{m_i}c_i))^n$ is geodesic (and hence $(1, 0)$-quasigeodesic) for all $n \geq 0$. We then produce for every number $0 \leq d_i \leq m_i - 1$ a new knapsack instance by replacing $u_i^{x_i}$ by $u_{i,1}\tilde{u}_i^{d_i}c_i(\mathsf{shlex}(c_i^{-1}\tilde{u}_i^{m_i}c_i))^{x_i}c_i^{-1}u_{i,1}^{-1}$. To make the description of the resulting knapsack expression more uniform we set $m_i = 1$ and $c_i = 1$ in case $|\tilde{u}_i| < 2L + 1$. Then, the replacement of $u_i^{x_i}$ by $u_{i,1}\tilde{u}_i^{x_i}u_{i,1}^{-1}$ in case $|\tilde{u}_i| < 2L + 1$ is the same as the one for the case $|\tilde{u}_i| \geq 2L + 1$. Let $m : \{x_1, \ldots, x_k\} \to \mathbb{N}$ be the mapping with $m(x_i) = m_i$.

From the above discussion, we obtain a finite set of (λ, ϵ)-quasigeodesic knapsack expressions E_d that are parameterized by a mapping $d : \{x_1, \ldots, x_k\} \to \mathbb{N}$ with

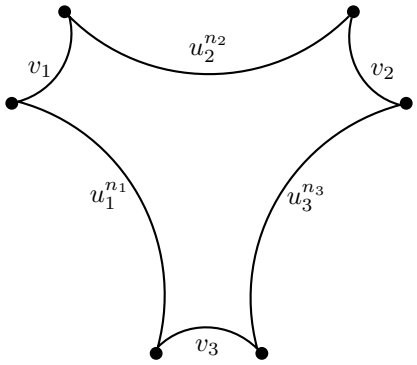

FIGURE 5. The $2k$-gon for $k = 3$ from the proof of Theorem 8.1

$0 \leq d(x_i) < m_i$ for all $1 \leq i \leq k$. Let $\mathcal{D}$ be the set of all such mappings. We then have

$$\mathsf{sol}(E) = \bigcup_{d \in \mathcal{D}} (m \cdot \mathsf{sol}(E_d) + d).$$

Note that the magnitude of every E_d is bounded linearly in the magnitude of E.

Finally, the statement of the proposition is directly obtained by combining the above steps 1 and 2. □

8.3. **Proof of Theorem 8.1.** We now come to the proof of Theorem 8.1. Consider a knapsack expression $E = u_1^{x_1} v_1 u_2^{x_2} v_2 \cdots u_k^{x_k} v_k$. We can assume that all u_i, v_i are geodesic. By Proposition 8.4 we can moreover assume that for all $1 \leq i \leq k$, u_i represents a group element of infinite order and that u_i^n is (λ, ϵ)-quasigeodesic for all $n \geq 0$, where λ, ϵ are fixed constants. We want to show that $\mathsf{sol}(E)$ is semilinear and has a magnitude that is polynomially bounded by $|E|$.

For the case $k = 1$ we have to consider all natural numbers n with $u_1^n = v_1^{-1}$ in G. Since u_1 represents a group element of infinite order there is at most one such n. Moreover, since u_i^n is (λ, ϵ)-quasigeodesic, such an n has to satisfy $|u_1| \cdot n \leq \lambda |v_1| + \epsilon$, which yields a linear bound on n.

For the case $k = 2$ we can directly use Proposition 8.3. Now assume that $k \geq 3$. We want to show that the set $\mathsf{sol}(E)$ is a semilinear subset of $\mathbb{N}^k$ (later we will consider the magnitude of $\mathsf{sol}(E)$). For this we construct a Presburger formula with free variables $x_1, \ldots, x_k$ that is equivalent to $E = 1$. We do this by induction on the depth k. Therefore, we can use in our Presburger formula also knapsack equations of the form $F = 1$, where F has depth at most $k - 1$.

It suffices to construct a Presburger formula for $\mathsf{sol}(E) \cap (\mathbb{N} \setminus \{0\})^k$. Note that $E = 1$ is equivalent to $\bigvee_{I \subseteq \{1,\ldots,k\}} (E_I = 1 \wedge \bigwedge_{i \in I} x_i > 0)$, where E_I is obtained from E by removing for every $i \notin I$ the power $u_i^{x_i}$.

Consider a tuple $(n_1, \ldots, n_k) \in \mathsf{sol}(E) \cap (\mathbb{N} \setminus \{0\})^k$ and the corresponding $2k$-gon that is defined by the (λ, ϵ)-quasigeodesic paths $P_i = (u_1^{n_1} v_1 \cdots u_{i-1}^{n_{i-1}} v_{i-1}) \cdot P[u_i^{n_i}]$ and the geodesic paths $Q_i = (u_1^{n_1} v_1 \cdots u_i^{n_i}) \cdot P[v_i]$, see Figure 5 for the case $k = 3$. Since all paths P_i and Q_i are (λ, ϵ)-quasigeodesic, we can apply [22, Lemma 6.4]: Every side of the $2k$-gon is contained in the h-neighborhoods of the other sides, where $h = \xi + \xi \log(2k)$ for a constant ξ that only depends on the constants $\delta, \lambda, \varepsilon$.

Let us now consider the side P_2 of the quasigeodesic $(2k)$-gon. It is labelled with $u_2^{x_2}$. Its neighboring sides are Q_1 and Q_2, which are labelled with v_1 and v_3, respectively. We distinguish several cases. In each case we cut the $2k$-gon into smaller pieces along paths of length $\leq 2h + 1$ (length h in some cases), and these smaller pieces will correspond to knapsack expressions of depth $< k$. This is done

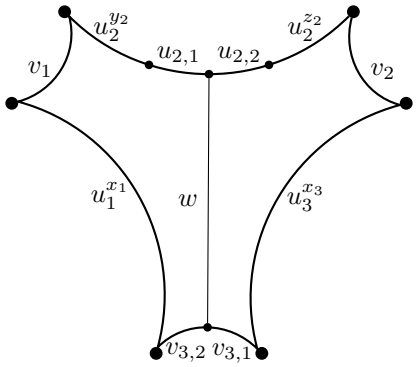

FIGURE 6. Case 1.1 from the proof of Theorem 8.1

until all knapsack expressions have depth at most two. When we speak of a point on the $2k$-gon, we mean a node of the Cayley graph (i.e., an element of the group G) and not a point in the interior of an edge. Moreover, when we speak of the successor point of a point p, we refer to the clockwise order on the $2k$-gon, where the sides are traversed in the order $P_1, Q_1, \ldots, P_k, Q_k$. We now distinguish the following cases:

Case 1: There is a point $p \in P_2$ that has distance at most h from a point q that does not belong to $P_1 \cup Q_1 \cup Q_2 \cup P_3$. Thus q must belong to one of the paths $Q_3, P_4, \ldots Q_{k-1}, P_k, Q_k$. Let w be a geodesic word of length at most h that labels a path from p to q. There are two subcases:

Case 1.1: q belongs to the paths Q_i, where $3 \leq i \leq k$. The situation is shown in Figure 6. We construct two new knapsack expressions F_t and G_t for all tuples $t = (w, u_{2,1}, u_{2,2}, v_{i,1}, v_{i,2})$ such that $w \in \Sigma^*$ is of length at most h, $u_2 = u_{2,1}u_{2,2}$ and $v_i = v_{i,1}v_{i,2}$:

$$\begin{aligned} F_t &= u_1^{x_1} v_1 u_2^{y_2} (u_{2,1} w v_{i,2}) u_{i+1}^{x_{i+1}} v_{i+1} \cdots u_k^{x_k} v_k \text{ and} \\ G_t &= u_{2,2} u_2^{z_2} v_2 u_3^{x_3} v_3 \cdots u_i^{x_i} (v_{i,1} w^{-1}) \end{aligned}$$

Here y_2 and z_2 are new variables. Note that F_t and G_t have depth at most $k-1$. Moreover, let $A_{1.1}$ be the following formula, where t ranges over all tuples of the above form:

$$A_{1.1} = \bigvee_t \exists y_2, z_2 : x_2 = y_2 + 1 + z_2 \wedge F_t = 1 \wedge G_t = 1$$

Case 1.2: q belongs to the path P_i, where $4 \leq i \leq k$ (this case can only occur if $k \geq 4$). This case is analogous to Case 1.1. We only have to split $u_i^{x_i}$ as $u_i^{y_i}(u_{i,1}u_{i,2})u_i^{z_i}$ (as we do for $u_2^{x_2}$). We construct two new knapsack expressions F_t and G_t for all tuples $t = (w, u_{2,1}, u_{2,2}, u_{i,1}, u_{i,2})$ such that $w \in \Sigma^*$ is of length at most h, $u_2 = u_{2,1}u_{2,2}$ and $u_i = u_{i,1}u_{i,2}$:

$$\begin{aligned} F_t &= u_1^{x_1} v_1 u_2^{y_2} (u_{2,1} w u_{i,2}) u_i^{z_i} v_i u_{i+1}^{x_{i+1}} v_{i+1} \cdots u_k^{x_k} v_k \text{ and} \\ G_t &= u_{2,2} u_2^{z_2} v_2 u_3^{x_3} v_3 \cdots u_{i-1}^{x_{i-1}} v_{i-1} u_i^{y_i} (u_{i,1} w^{-1}) \end{aligned}$$

Here y_2, z_2, y_i, z_i are new variables. Note that F_t and G_t have depth at most $k-1$. Moreover, let $A_{1.2}$ be the following formula, where t ranges over all tuples of the above form:

$$A_{1.2} = \bigvee_t \exists y_2, z_2, y_i, z_i : x_2 = y_2 + 1 + z_2 \wedge x_i = y_i + 1 + z_i \wedge F_t = 1 \wedge G_t = 1$$

Case 2: Every point on P_2 that has distance at most h from a point on $P_1 \cup Q_1 \cup Q_2 \cup P_3$.

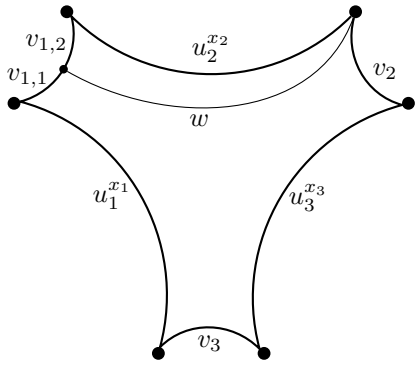

FIGURE 7. Case 2.1 from the proof of Theorem 8.1

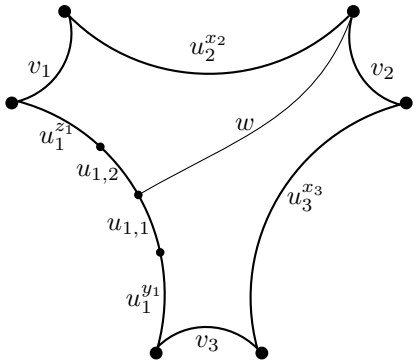

FIGURE 8. Case 2.2 from the proof of Theorem 8.1

Case 2.1: The end point of P_2 (i.e., the point connecting P_2 with Q_2) has distance at most h from a point on Q_1, see Figure 7. For all tuples $t = (w, v_{1,1}, v_{1,2})$ such that $w \in \Sigma^*$ is of length at most h and $v_1 = v_{1,1}v_{1,2}$ we construct two new knapsack expressions

$$F_t = u_2^{x_2}(wv_{1,2}) \text{ and } G_t = u_1^{x_1}(v_{1,1}w^{-1}v_2)u_3^{x_3}v_3 \cdots u_k^{x_k}v_k$$

and the formula

$$A_{2.1} = \bigvee_t F_t = 1 \wedge G_t = 1,$$

where t ranges over all tuples of the above form. Note that F_t has depth one and G_t has depth $k-1$.

Case 2.2: The end point of P_2 (i.e., the point connecting P_2 with Q_2) has distance at most h from a point on P_1, see Figure 8. For all tuples $t = (w, u_{1,1}, u_{1,2})$ such that $w \in \Sigma^*$ is of length at most h and $u_1 = u_{1,1}u_{1,2}$, we construct two new knapsack expressions

$$F_t = u_1^{z_1}v_1u_2^{x_2}(wu_{1,2}) \text{ and } G_t = u_1^{y_1}(u_{1,1}w^{-1}v_2)u_3^{x_3}v_3 \cdots u_k^{x_k}v_k$$

and the formula

$$A_{2.2} = \bigvee_t \exists y_1, z_1 : x_1 = y_1 + 1 + z_1 \wedge F_t = 1 \wedge G_t = 1,$$

where t ranges over all tuples of the above form. Note that F_t has depth two and G_t has depth $k-1$.

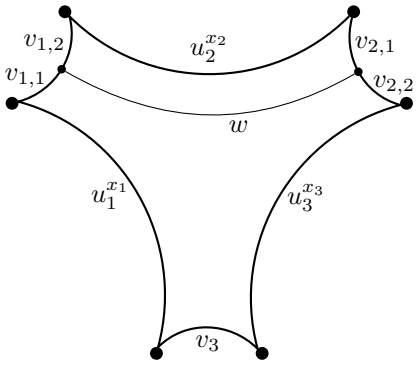

FIGURE 9. Case 2.3 from the proof of Theorem 8.1

If on the other hand the end point of P_2 has distance $> h$ from all points on $P_1 \cup Q_1$, then there must be two points p_1, p_2 on P_2 such that p_2 is the successor point of p_1 when travelling along P_2 (i.e., $d(p_1, p_2) = 1$), and p_1 has distance at most h from a point $q_1 \in P_1 \cup Q_1$, while p_2 has distance at most h from a point on $q_2 \in Q_2 \cup P_3$. Thus, the distance between q_1 and q_2 is at most $2h+1$. Let w be a word that labels a geodesic path from q_1 to q_2 (thus, $|w| \leq 2h+1$). This leads to the following four subcases.

Case 2.3: $q_1 \in Q_1$ and $q_2 \in Q_2$, see Figure 9. This case is very similar to Case 2.1. For every tuple $t = (w, v_{1,1}, v_{1,2}, v_{2,1}, v_{2,2})$ with $|w| \leq 2h+1$, $v_1 = v_{1,1}v_{1,2}$ and $v_2 = v_{2,1}v_{2,2}$ we obtain two new knapsack expressions

$$F_t = F_t = v_{1,2}u_2^{x_2}(v_{2,1}w) \text{ and } G_t = u_1^{x_1}(v_{1,1}w^{-1}v_{2,2})u_3^{x_3}v_3 \cdots u_k^{x_k}v_k$$

and the formula

$$A_{2.3} = \bigvee_t F_t = 1 \wedge G_t = 1,$$

where t ranges over all tuples of the above form.

Case 2.4: $q_1 \in P_1$ and $q_2 \in Q_2$, see Figure 10. This case is very similar to Case 2.2. For every tuple $t = (w, u_{1,1}, u_{1,2}, v_{2,1}, v_{2,2})$ such that $|w| \leq 2h+1$, $u_1 = u_{1,1}u_{1,2}$, and $v_2 = v_{2,1}v_{2,2}$ we obtain two new knapsack expressions

$$F_t = u_{1,2}u_1^{z_1}v_1u_2^{x_2}(v_{2,1}w) \text{ and } G_t = u_1^{y_1}(u_{1,1}w^{-1}v_{2,2})u_3^{x_3}v_3 \cdots u_k^{x_k}v_k$$

and the formula

$$A_{2.4} = \bigvee_t \exists y_1, z_1 : x_1 = y_1 + 1 + z_1 \wedge F_t = 1 \wedge G_t = 1,$$

where t ranges over all tuples of the above form.

Case 2.5: $q_1 \in Q_1$ and $q_2 \in P_3$. This case is analogous to Case 2.4.

Case 2.6: $q_1 \in P_1$ and $q_2 \in P_3$, see Figure 11. For every tuple

$$(w_1, w_2, w, u_{1,1}, u_{1,2}, u_{2,1}, u_{2,2}, u_{3,1}, u_{3,2})$$

such that $|w| \leq 2k+1$, $|w_1| \leq h$, $|w_2| \leq h+1$, $w = w_1^{-1}w_2$ in G, $u_1 = u_{1,1}u_{1,2}$, $u_2 = u_{2,1}u_{2,2}$, and $u_3 = u_{3,1}u_{3,2}$ we obtain three new knapsack expressions

$$\begin{aligned} F_t &= u_1^{z_1}v_1u_2^{y_2}(u_{2,1}w_1u_{1,2}), \\ G_t &= u_2^{z_2}v_2u_3^{y_3}(u_{3,1}w_2^{-1}u_{2,2}) \text{ and} \\ H_t &= u_3^{z_3}v_3u_4^{x_4}v_4 \cdots u_k^{x_k}v_ku_1^{y_1}(u_{1,1}wu_{3,2}). \end{aligned}$$

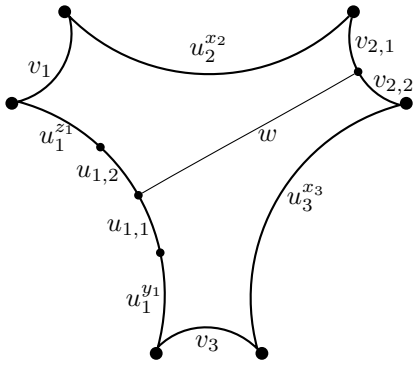

FIGURE 10. Case 2.4 from the proof of Theorem 8.1

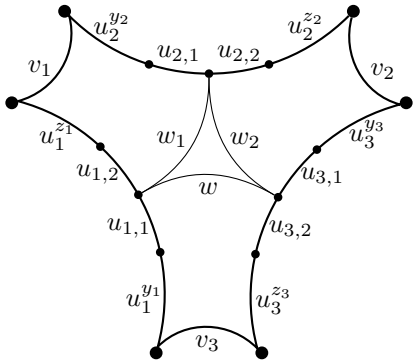

FIGURE 11. Case 2.6 from the proof of Theorem 8.1

and the formula

$$A_{2.6} = \bigwedge_t \exists y_1, z_1, y_2, z_2, y_3, z_3 : \bigwedge_{i=1}^{3} x_i = y_i + 1 + z_i \wedge F_t = 1 \wedge G_t = 1 \wedge H_t = 1,$$

where t ranges over all tuples of the above form. Note that F_t and G_t have depth 2 and that H_t has depth $k-1$.

This concludes the construction of a Presburger formula for the set $\mathsf{sol}(E)$ and shows the semilinearity of $\mathsf{sol}(E)$. It remains to argue that the magnitude of $\mathsf{sol}(E)$ is bounded polynomially in $|E|$. Iterating the above splitting procedure results in an exponentially large disjunction of conjunctive formulas of the form

$$\exists y_1, \ldots, y_m \bigwedge_{i \in I} E_i = 1 \bigwedge_{j \in J} z_j = z'_j + z''_j + 1 \tag{3}$$

where every E_i is a knapsack expression of depth at most two. Moreover, for $i \neq j$, E_i and E_j have no common variables. The existentially quantified variables $y_1, \ldots, y_m$ are the new variables that were introduced when splitting factors $u_i^{x_i}$ (e.g., y_2, z_2 in the formula $A_{1.1}$). The variables z_j, z'_j, z''_j in (3) are from $\{x_1, \ldots, x_k, y_1, \ldots, y_m\}$. The equations $z_j = z'_j + z''_j + 1$ in (3) result from the splitting of factors $u_i^{x_i}$. For instance, $x_2 = y_2 + 1 + z_2$ in $A_{1.1}$ is one such equation.

In order to bound the magnitude of $\mathsf{sol}(E)$ it suffices to consider a single conjunctive formula of the form (3), since disjunction corresponds to union of semilinear sets, which does not increase the magnitude. We can also ignore the existential quantifiers in (3), because existential quantification corresponds to projection onto

some of the coordinates, which cannot increase the magnitude. Hence, we have to consider the magnitude of the semilinear set A defined by

$$\bigwedge_{i\in I} E_i = 1 \bigwedge_{j\in J} z_j = z'_j + z''_j + 1. \tag{4}$$

The splitting process that finally produces formula (4) can be seen as a tree T, where every node v is labelled with a knapsack expression $E(v)$, the root is labelled with E, the leaves are labelled with the expressions E_i ($i \in I$) from (3) and the children of a node v are labelled with the expressions into which $E(v)$ is decomposed. The number of children of every node is at most three (three children are only produced in Case 2.6).

Let us first show that the size of this tree T is bounded by $\mathcal{O}(k^2)$. We assign to each node v of T the number $d(v) :=$ depth of the knapsack expression $E(v)$. Note that $d(v) \leq 2$ if and only if v is a leaf. If $E(v)$ is split according to one of the Cases 2.1–2.6 then v has $j \leq 3$ children $v_1, \ldots, v_j$, where $v_1, \ldots, v_{j-1}$ are leaves (their d-value is one or two) and $d(v_j) = d(v) - 1$. If $E(v)$ is split according to Case 1.1 or 1.2 then v has two children v_1 and v_2 such that (i) $d(v_1), d(v_2) < d(v)$, (ii) $d(v_1), d(v_2) \geq 2$ and $d(v_1)+d(v_2) = d(v)+1$ in Case 1.1, and (iii) $d(v_1), d(v_2) \geq 3$ and $d(v_1) + d(v_2) = d(v) + 2$ in Case 1.2. Let T' be the tree that is obtained by removing all leaves with d-value at most 2. It suffices to show that the size of T' is bounded by $\mathcal{O}(k^2)$. All leaves of T' have the d-value 3. Moreover, every non-leaf v of T' has either exactly one child v' with $d(v) > d(v')$ or two children v_1 and v_2 such that $d(v) \geq d(v_1) + d(v_2) - 2$. Let n_0 be the number of leaves of T' and n_2 be the number of nodes of T' with exactly two children. From the above equations, it follows that the root r of T' satisfies $k = d(r) \geq 3n_0 - 2n_2$. Moreover, $n_2 = n_0 - 1$. We get $k \geq n_0 + 2$, i.e., $n_0 \leq k-2$ and $n_2 \leq k-3$. Since every path from the root to a leaf can contain at most k nodes having a single child, we must have $n_1 \leq (k-2)k$. This shows that the size of T' and hence of T is bounded by $\mathcal{O}(k^2)$. Thus, we also have $|I| \leq \mathcal{O}(k^2)$ in (4).

Next, we show that for every $i \in I$, $|E_i|$ is bounded polynomially in $|E|$. To see this, consider a single splitting step. In each of the above Cases 1.1–2.6 the argument is similar. Consider for instance Case 2.6, where the knapsack expression E is replaced by three knapsack expressions F_t, G_t, H_t. We can bound the sizes of these expressions by $|F_t| \leq |E| + |u_{1,2}| + |u_{2,1}| + |w_1| \leq |E| + |u_1| + |u_2| + h$, $|G_t| \leq |E| + |u_{2,2}| + |u_{3,1}| + |w_2| \leq |E| + |u_2| + |u_3| + h + 1$, and $H_t| \leq |E| + |u_{1,1}| + |u_{3,2}| + |w| \leq |E| + |u_1| + |u_3| + 2h + 1$. The number of splitting steps that finally leads to an E_i is bounded by k (since the depth of the knapsack expressions is reduced in each step). Hence, the size of each knapsack expression E_i in (4) is bounded by $|E| + 2k|E| + k(2h+1) = (2k+1)|E| + k(2\xi + 2\xi \log(2k) + 1) \leq \mathcal{O}(|E|^2)$. Since every E_i has depth at most two, there is a fixed polynomial $p(n)$ such that the magnitude of every set $\mathsf{sol}(E_i)$ is bounded by $p(|E|)$. Hence, also $\bigoplus_{i\in I} \mathsf{sol}(E_i)$ is a semilinear set of magnitude at most $p(|E|)$ (the $\oplus$-operator on semilinear sets does not increase the magnitude). Note that $\bigoplus_{i\in I} \mathsf{sol}(E_i)$ is the semilinear set defined by the conjunction $\bigwedge_{i\in I} E_i = 1$.

To bound the magnitude of the semilinear set A defined by (4), one has to consider also the additional equations $z_j = z'_j + z''_j + 1$ for $j \in J$. Let U be the set of variables that appear in the knapsack expressions E_i ($i \in I$). Note that the dimension of $\bigoplus_{i\in I} \mathsf{sol}(E_i)$ is $|U|$. Since every knapsack expression E_i ($i \in I$) contains at most two variables, we can bound the dimension of $\bigoplus_{i\in I} \mathsf{sol}(E_i)$ by $2|I| \leq \mathcal{O}(k^2)$. Note that for each equation $z_j = z'_j + z''_j + 1$ there exists a node v in the tree T with children v', v'' such that z_j is a variable from $E(v)$, z'_j is a variable from $E(v')$, and z''_j is a variable from $E(v'')$. This implies that every variable z_j

is a sum of pairwise different variables from U plus a constant that is bounded by $|T| \leq \mathcal{O}(k^2)$. Therefore the magnitude of A is bounded by $\mathcal{O}(k^2 \cdot p(|E|))$, which is polynomial in $|E|$. This concludes the proof. □

9. More groups with knapsack in LogCFL

Let $\mathcal{C}$ be the smallest class of groups such that (i) every hyperbolic group belongs to $\mathcal{C}$, (ii) if $G \in \mathcal{C}$ then also $G \times \mathbb{Z} \in \mathcal{C}$, and (iii) if $G, H \in \mathcal{C}$ then also $G * H \in \mathcal{C}$ (where $G * H$ is the free product of G and H). The class $\mathcal{C}$ contains groups that are not hyperbolic (e.g., $\mathbb{Z} \times \mathbb{Z}$). From Theorem 8.1 and Proposition 5.1 we get:

Proposition 9.1. *Every group from the class $\mathcal{C}$ is knapsack-tame and hence polynomially knapsack-bounded.*

From Theorem 4.1 and 4.2 we get:

Proposition 9.2. *Every group from the class $\mathcal{C}$ belongs to OW-AuxPDA.*

Proposition 9.1 and 9.2 together with Theorem 6.1 and 6.2 yield:

Corollary 9.3. *For every group G from the class $\mathcal{C}$, membership for acyclic NFAs over G and knapsack for G both belong to* LogCFL.

Corollary 9.3 generalizes Corollaries 7.2 and 7.3 as well as [4, Corollary 22], where it was shown that knapsack can be solved in polynomial time for a free product of hyperbolic groups and finitely generated abelian groups.

10. Conclusion

In this paper, it is shown that every hyperbolic group is knapsack-tame and that the knapsack problem can be solved in LogCFL. Here is a list of open problems that one might consider for future work.

- For the following important groups, it is not known whether the knapsack problem is decidable: braid groups B_n (with $n \geq 3$), solvable Baumslag-Solitar groups $\mathsf{BS}_{1,p} = \langle a, t \mid t^{-1}at = a^p \rangle$ (with $p \geq 2$), and automatic groups which are not in any of the known classes with a decidable knapsack problem.
- In [13], it was shown that knapsack is decidable for every co-context-free group. The algorithm from [13] has an exponential running time. Is there a more efficient solution?
- Is there a polynomially knapsack-bounded group which is not knapsack-tame?

References

[1] Gerhard Buntrock and Friedrich Otto. Growing context-sensitive languages and Church-Rosser languages. *Information and Computation*, 141:1–36, 1998.

[2] Michael Elberfeld, Andreas Jakoby, and Till Tantau. Algorithmic meta theorems for circuit classes of constant and logarithmic depth. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:128, 2011.

[3] David B. A. Epstein and Derek F. Holt. The linearity of the conjugacy problem in word-hyperbolic groups. *International Journal of Algebra and Computation*, 16(2):287–306, 2006.

[4] Elizaveta Frenkel, Andrey Nikolaev, and Alexander Ushakov. Knapsack problems in products of groups. *Journal of Symbolic Computation*, 74:96–108, 2016.

[5] Moses Ganardi, Daniel König, Markus Lohrey, and Georg Zetzsche. Knapsack problems for wreath products. In *Proceedings of STACS 2018*, volume 96 of *LIPIcs*, pages 32:1–32:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[6] Etienne Ghys and Pierre de La Harpe. *Sur les groupes hyperboliques d'après Mikhael Gromov*. Progress in mathematics. Birkhäuser, 1990.

[7] Seymour Ginsburg and Edwin H. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.

[8] Mikhail Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory*, number 8 in MSRI Publ., pages 75–263. Springer, 1987.
[9] Christoph Haase. *On the complexity of model checking counter automata*. PhD thesis, University of Oxford, St Catherine's College, 2011.
[10] Derek F. Holt. Word-hyperbolic groups have real-time word problem. *International Journal of Algebra and Computation*, 10:221–228, 2000.
[11] Derek F. Holt, Markus Lohrey, and Saul Schleimer. Compressed decision problems in hyperbolic groups. In *Proceedings of STACS 2019*, volume 126 of *LIPIcs*, pages 37:1–37:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
[12] Richard 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.
[13] Daniel König, Markus Lohrey, and Georg Zetzsche. Knapsack and subset sum problems in nilpotent, polycyclic, and co-context-free groups. In *Algebra and Computer Science*, volume 677 of *Contemporary Mathematics*, pages 138–153. American Mathematical Society, 2016.
[14] Eryk Kopczynski and Anthony Widjaja To. Parikh images of grammars: Complexity and applications. In *Proceedings of LICS 2010*, pages 80–89. IEEE Computer Society, 2010.
[15] Jörg Lehnert and Pascal Schweitzer. The co-word problem for the Higman-Thompson group is context-free. *Bulletin of the London Mathematical Society*, 39(2):235–241, 2007.
[16] Markus Lohrey. Decidability and complexity in automatic monoids. *International Journal of Foundations of Computer Science*, 16(4):707–722, 2005.
[17] Markus Lohrey and Georg Zetzsche. Knapsack in graph groups, HNN-extensions and amalgamated products. *CoRR*, abs/1509.05957, 2015.
[18] Markus Lohrey and Georg Zetzsche. 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.
[19] Markus Lohrey and Georg Zetzsche. Knapsack in graph groups. *Theory of Computing Systems*, 62(1):192–246, 2018.
[20] Alexei Mishchenko and Alexander Treier. Knapsack problem for nilpotent groups. *Groups Complexity Cryptology*, 9(1):87–98, 2017.
[21] Alexei Myasnikov and Andrey Nikolaev. Verbal subgroups of hyperbolic groups have infinite width. *Journal of the London Mathematical Society*, 90(2):573–591, 2014.
[22] Alexei Myasnikov, Andrey Nikolaev, and Alexander Ushakov. Knapsack problems in groups. *Mathematics of Computation*, 84:987–1016, 2015.
[23] Alexander Yu. Ol'shanskii. Almost every group is hyperbolic. *International Journal of Algebra and Computation*, 2(1):1–17, 1992.
[24] Ivan H. Sudborough. On the tape complexity of deterministic context–free languages. *Journal of the ACM*, 25(3):405–414, 1978.
[25] Anthony Widjaja To. Parikh images of regular languages: Complexity and applications. *CoRR*, abs/1002.1464, 2010. URL: `http://arxiv.org/abs/1002.1464`.
[26] Heribert Vollmer. *Introduction to Circuit Complexity*. Springer-Verlag, 1999.

Universität Siegen, Germany
E-mail address: `lohrey@eti.uni-siegen.de`