Universität Stuttgart

Fakultät Informatik, Elektrotechnik und Informationstechnik

Inverse monoids: decidability and complexity of algebraic questions

Markus Lohrey and Nicole Ondrusch

Report Nr. 2005/2

Institut für Formale
Methoden der Informatik
Universitätsstraße 38
D–70569 Stuttgart

May 20, 2005

CR: F.4.1, F.4.2

Abstract

This paper investigates the word problem for inverse monoids generated by a set Γ subject to relations of the form $e = f$, where e and f are both idempotents in the free inverse monoid generated by Γ. It is shown that for every fixed monoid of this form the word problem can be solved in polynomial time which solves an open problem of Margolis and Meakin. For the uniform word problem, where the presentation is part of the input, EXPTIME-completeness is shown. For the Cayley-graphs of these monoids, it is shown that the first-order theory with regular path predicates is decidable. Regular path predicates allow to state that there is a path from a node x to a node y that is labeled with a word from some regular language. As a corollary, the decidability of the generalized word problem is deduced. Finally, it is shown that the Cayley-graph of the free inverse monoid has an undecidable monadic second-order theory.

1 Introduction

The decidability and complexity of algebraic questions in various kinds of structures is a classical topic at the borderline of computer science and mathematics. The most basic algorithmic question concerning algebraic structures is the word problem, which asks whether two given expressions denote the same element of the underlying structure. Markov [20] and Post [28] proved independently that the word problem for finitely presented monoids is undecidable in general. This result can be seen as one of the first undecidability results that touched real mathematics. Later, Novikov [25] and Boone [3] extended the result of Markov and Post to finitely presented groups.

In this paper, we are interested in a class of monoids that lies somewhere between groups and general monoids: inverse monoids [27]. In the same way as groups can be represented by sets of permutations, inverse monoids can be represented by sets of partial injections [27]. Algorithmic questions for inverse monoids received increasing attention in the past and inverse monoid theory found several applications in combinatorial group theory, see e.g. the survey [19]. In [18], Margolis and Meakin presented a large class of finitely presented inverse monoids with decidable word problems. An inverse monoid from that class is of the form $\mathrm{FIM}(\Gamma)/P$, where $\mathrm{FIM}(\Gamma)$ is the free inverse monoid generated by the set Γ and P is a presentation consisting of a finite number of identities between idempotents of $\mathrm{FIM}(\Gamma)$; we call such a presentation idempotent. In fact, in [18] it is shown that even the uniform word problem for idempotent presentations is decidable. In this problem, also the presentation is part of the input.

The decidability proof of Margolis and Meakin uses Rabin's seminal tree theorem [29], concerning the decidability of the monadic second-order theory of the complete binary tree. From the view point of complexity, the use of Rabin's tree theorem is somewhat unsatisfactory, because it leads to a nonelementary algorithm for the word problem. Therefore, in [1, 18] the question for a more efficient approach was asked. In Section 6 we show by using tree automata techniques that for every fixed idempotent presentation the word problem for $\mathrm{FIM}(\Gamma)/P$ can be solved in polynomial time. For the uniform word problem for idempotent presentations we prove completeness for EXPTIME (deterministic exponential time). Similarly to the method of Margolis and Meakin, we use results from logic for the upper bound. But instead of translating the uniform word problem into monadic second-order logic over the complete binary tree, we exploit a translation into the modal μ-calculus, which is a popular logic for the verification of reactive systems. Then, we can use a result from [12, 38] stating that the model-checking problem of the modal μ-calculus over context-free graphs [23] is EXPTIME-complete.

In Section 7 we will investigate Cayley-graphs of inverse monoids of the form $\mathrm{FIM}(\Gamma)/P$. The Cayley-graph of a finitely generated monoid $\mathcal{M}$ w.r.t. a finite generating set Γ is a Γ-labeled directed graph with node set $\mathcal{M}$ and an a-labeled edge from a node x to a node y if $y = xa$ in $\mathcal{M}$. Cayley-graphs of groups are a fundamental tool in combinatorial group theory [17] and serve as a link to other fields like topology, graph theory, and automata theory, see, e.g., [22, 23]. Here we consider Cayley-graphs from a logical point of view, see [13, 14] for previous results in this direction. More precisely, we consider an expansion of the Cayley-graph G that contains for every regular language L over the generators of $\mathcal{M}$ a binary predicate reach_L. Two nodes u and v of G are related by reach_L if there exists a path from u to v in the Cayley-graph, which is labeled with a word from the language L.

Our main result of Section 7 states that this structure has a decidable first-order theory, whenever the underlying monoid is of the form $\mathrm{FIM}(\Gamma)/P$ for an idempotent presentation P (Theorem 7.2). An immediate corollary of this result is that the generalized word problem of $\mathrm{FIM}(\Gamma)/P$ is decidable. The generalized word problem asks whether for given elements $w, w_1, \ldots, w_n \in \mathrm{FIM}(\Gamma)/P$, w belongs to the submonoid of $\mathrm{FIM}(\Gamma)/P$ generated by $w_1, \ldots, w_n$. Our decidability result for Cayley-graphs should be also compared with the undecidability result for the existential theory of the free inverse monoid $\mathrm{FIM}(\{a, b\})$ [31], which consists of all true statements over $\mathrm{FIM}(\{a, b\})$ of the form $\exists x_1 \cdots \exists x_m : \varphi$, where φ is a boolean combination of word equations (with constant).

It is not hard to see that an atomic proposition $\mathrm{reach}_L(x, y)$ can be expressed in monadic second-order logic over the Cayley-graph. Thus, one might ask, whether our decidability result for first-order logic with the reach_L-predicates can be extended to the full monadic second-order theory. Our final result states that already the Cayley-graph (without the reach_L-predicates) of the free inverse monoid generated by two elements has an undecidable monadic second-order theory (Theorem 7.8).

2 Preliminaries

The length of a word u is denoted by $|u|$. The empty word is ε. For a finite alphabet Γ, we denote with $\Gamma^{-1} = \{a^{-1} \mid a \in \Gamma\}$ a disjoint copy of Γ. For $a^{-1} \in \Gamma^{-1}$ we define $(a^{-1})^{-1} = a$; thus, $^{-1}$ becomes an involution on the alphabet $\Gamma \cup \Gamma^{-1}$. We extend this involution to words from $(\Gamma \cup \Gamma^{-1})^*$ by setting $(b_1 b_2 \cdots b_n)^{-1} = b_n^{-1} \cdots b_2^{-1} b_1^{-1}$, where $b_i \in \Gamma \cup \Gamma^{-1}$. The set of all regular languages over an alphabet Γ will be denoted by $\mathrm{REG}(\Gamma)$.

We assume that the reader has some basic background in complexity theory [26]. We will make use of alternating Turing-machines, see [4] for more details. Roughly speaking, an *alternating Turing-machine* $T = (Q, \Sigma, \delta, q_0, q_f)$ (where Q is the state set, Σ is the tape alphabet, δ is the transition relation, q_0 is the initial state, and q_f is the unique accepting state) is a nondeterministic Turing-machine, where the set of nonfinal states $Q \setminus \{q_f\}$ is partitioned into two sets: $Q_\exists$ (existential states) and $Q_\forall$ (universal states). We assume that T cannot make transitions out of the accepting state q_f. A configuration C with current state q is accepting, if

- $q = q_f$, or
- $q \in Q_\exists$ and there exists a successor configuration of C that is accepting, or
- $q \in Q_\forall$ and every successor configuration of C is accepting.

An input word w is accepted by T if the corresponding initial configuration is accepting. It is known that EXPTIME (deterministic exponential time) equals APSPACE (the class of all problems that can be accepted by an alternating Turing-machine in polynomial space) [4].

3 Relational Structures and Logic

See [8] for more details on the subject of this section. A signature is a countable set $\mathcal{S}$ of relational symbols, where each relational symbol $R \in \mathcal{S}$ has an associated arity n_R. A (relational) structure over the signature $\mathcal{S}$ is a tuple $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \mathcal{S}})$, where A is a set (the universe of $\mathcal{A}$) and $R^{\mathcal{A}}$ is a relation of arity n_R over the set A, which interprets the relational symbol R. We will assume that every signature contains the equality symbol $=$ and that $=^{\mathcal{A}}$ is the identity relation on the set A. As usual, a constant $c \in A$ can be encoded by the unary relation $\{c\}$. Usually, we denote the relation $R^{\mathcal{A}}$ also with R. With $\mathrm{RS}(\mathcal{S})$ we denote the class of all relational structures over the signature $\mathcal{S}$. For $B \subseteq A$ we define the restriction $\mathcal{A}{\restriction}B = (B, (R^{\mathcal{A}} \cap B^{n_R})_{R \in \mathcal{S}})$; it is again a structure over the signature $\mathcal{S}$.

Next, let us introduce *monadic second-order logic (MSO-logic)*. Let $\mathbb{V}_1$ (resp. $\mathbb{V}_2$) be a countably infinite set of *first-order variables* (resp. *second-order variables*) which range over elements (resp. subsets) of the universe A. First-order variables (resp. second-order variables) are denoted x, y, z, x', etc. (resp. X, Y, Z, X', etc.). *MSO-formulas* over the signature $\mathcal{S}$ are constructed from the atomic formulas $R(x_1, \ldots, x_{n_R})$ and $x \in X$ (where $R \in \mathcal{S}$, $x_1, \ldots, x_{n_R}, x \in \mathbb{V}_1$, and $X \in \mathbb{V}_2$) using the boolean connectives $\neg, \wedge$, and $\vee$, and quantifications over variables from $\mathbb{V}_1$ and $\mathbb{V}_2$. The notion of a free occurrence of a variable is defined as usual. A formula without free occurrences of variables is called an *MSO-sentence*. If $\varphi(x_1, \ldots, x_n, X_1, \ldots, X_m)$ is an MSO-formula such that at most the first-order variables among $x_1, \ldots, x_n$ and the second-order variables among $X_1, \ldots, X_m$ occur freely in φ, and $a_1, \ldots, a_n \in A$, $A_1, \ldots, A_m \subseteq A$, then $\mathcal{A} \models \varphi(a_1, \ldots, a_n, A_1, \ldots, A_m)$ means that φ evaluates to true in $\mathcal{A}$ if the free variable x_i (resp. X_j) evaluates to a_i (resp. A_j). The *MSO-theory* of $\mathcal{A}$, denoted by $\mathrm{MSOTh}(\mathcal{A})$, is the set of all MSO-sentences φ such that $\mathcal{A} \models \varphi$. For an MSO-formula $\varphi(x_1, \ldots, x_n, X_1, \ldots, X_m)$ and a variable $Y \in \mathbb{V}_2 \setminus \{X_1, \ldots, X_m\}$ we need the relativation $\varphi{\restriction}_Y(x_1, \ldots, x_n, X_1, \ldots, X_m, Y)$. It is inductively defined by restricting every quantifier in φ to the set Y. Then for all $B \subseteq A$ and all $a_1, \ldots, a_n \in B$, $A_1, \ldots, A_m \subseteq B$ we have $\mathcal{A}{\restriction}_B \models \varphi(a_1, \ldots, a_n, A_1, \ldots, A_m)$ if and only if $\mathcal{A} \models \varphi{\restriction}_Y(a_1, \ldots, a_n, A_1, \ldots, A_m, B)$.

Remark 3.1. *Several times, we will use implicitly the well-known fact that reachability in graphs can be expressed in MSO. More precisely, if* $\mathrm{reach}(x, y)$ *is the formula*

$$\forall X : ((x \in X \ \wedge \ \forall u, v : (u \in X \wedge E(u, v) \ \Rightarrow \ v \in X)) \ \Rightarrow \ y \in X),$$

then for every directed graph $G = (V, E)$ *and all nodes* $s, t \in V$ *we have* $G \models \mathrm{reach}(s, t)$ *if and only if* $(s, t) \in E^*$. *Another important fact is that finiteness of a subset of a finitely-branching tree can be expressed in MSO, i.e., there is an MSO-formula* $\mathrm{fin}(X)$ *(over the signature containing a binary relation symbol* E*) such that for every (finitely-branching and undirected) tree* $T = (V, E)$ *and all subsets* $U \subseteq V$ *we have* $T \models \mathrm{fin}(U)$ *if and only if* U *is finite, see also [29, Lemma 1.8]. First, let us define two auxiliary formulas, where* $N(x)$ *denotes the set* $\{y \in V \mid (x, y) \in E\}$*:*

$$\begin{aligned}
\textit{ω-path}(x, X) \quad &= \quad x \in X \ \wedge \ |N(x) \cap X| = 1 \ \wedge \ \forall y \in X \setminus \{x\} : |N(y) \cap X| = 2 \ \wedge \\
&\qquad \forall y \in X : \mathrm{reach}(x, y){\restriction}_X \\
\textit{fin-path}(x, y, X) \quad &= \quad (X = \{x\} \wedge x = y) \ \vee \ (x \neq y \ \wedge \ x, y \in X \ \wedge \\
&\qquad\quad |N(x) \cap X| = |N(y) \cap X| = 1 \ \wedge \\
&\qquad\quad \forall z \in X \setminus \{x, y\} : |N(z) \cap X| = 2 \ \wedge \\
&\qquad\quad \forall z \in X : \mathrm{reach}{\restriction}_X(x, z, X))
\end{aligned}$$

Then we have $T \models \omega$-path(u, U) if and only if U is an ω-path starting in node u, whereas $T \models$ fin-path(u, v, U) if and only if U is a finite path with end points u and v. Now $U \subseteq V$ is finite if and only if the following holds:

$$\begin{aligned}\exists r\, \exists X : \forall x : (x \in X \Leftrightarrow \exists y \in U\, \exists Y : (\textit{fin-path}(r, y, Y) \wedge x \in Y)) \wedge \\ \neg \exists Z : (\omega\textit{-path}(r, Z) \wedge Z \subseteq X)\end{aligned}$$

We select first an arbitrary root r. Then the formula $\forall x : (x \in X \Leftrightarrow \exists y \in U\, \exists Y : (\textit{fin-path}(r, y, Y) \wedge x \in Y))$ says that X is the upward-closure of the set U, when r is the root of the tree. Finally, we say that there does not exist an infinite path Z that is contained in X. Since T is finitely-branching, by König's lemma this is equivalent to the fact that X (and hence U) is finite.

A *first-order formula* over the signature $\mathcal{S}$ is an MSO-formula that does not contain any occurrences of second-order variables. In particular, first-order formulas do not contain atomic subformulas of the form $x \in X$. The *first-order theory* $\mathrm{FOTh}(\mathcal{A})$ of $\mathcal{A}$ is the set of all first-order sentences φ such that $\mathcal{A} \models \varphi$.

Let $\mathcal{S}$ and $\mathcal{T}$ be two relational signatures. A mapping $f : \mathrm{RS}(\mathcal{S}) \to \mathrm{RS}(\mathcal{T})$ is an *MSO-transduction* if there exists $m \in \mathbb{N}$ (the copy number) and for every $P \in \mathcal{T}$ of rank $n = n_P$ and every tuple $\overline{a} = (i_1, \ldots, i_n) \in \{1, \ldots, m\}^n$ there exists an MSO-formula $\theta_{P,\overline{a}}(x_1, \ldots, x_n)$ over the signature $\mathcal{S}$ such that the following holds: If $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \mathcal{S}})$ is a structure from $\mathrm{RS}(\mathcal{S})$, then

$$f(\mathcal{A}) = (A \times \{1, \ldots, m\}, (P^{f(\mathcal{A})})_{P \in \mathcal{T}}),$$

where for every $P \in \mathcal{T}$ of rank n:

$$P^{f(\mathcal{A})} = \bigcup_{(i_1,\ldots,i_n) \in \{1,\ldots,m\}^n} \{((a_1, i_1), \ldots, (a_n, i_n)) \mid \mathcal{A} \models \theta_{P,(i_1,\ldots,i_n)}(a_1, \ldots, a_n)\}.$$ [1]

By a result of Courcelle [6] it is known that an MSO-transduction f as defined above is MSO-compatible. This means that there is a total recursive mapping $f^\sharp$ (also called the backwards translation of f) from the set of MSO-sentences over the signature $\mathcal{T}$ to the set of MSO-sentences over the signature $\mathcal{S}$ such that for every MSO-sentences φ over the signature $\mathcal{T}$ and every $\mathcal{A} \in \mathrm{RS}(\mathcal{S})$: $\mathcal{A} \models f^\sharp(\varphi)$ if and only if $f(\mathcal{A}) \models \varphi$.

Let us end this section with a brief introduction into the *modal μ-calculus*, which is a popular logic for the verification of reactive systems, see [37] for more details. Formulas of this logic are interpreted over edge-labeled directed graphs. Let Σ be a finite set of edge labels. The syntax of the modal μ-calculus is given by the following grammar:

$$\varphi ::= \text{true} \mid \text{false} \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu X.\varphi \mid \nu X.\varphi$$

Here $X \in \mathbb{V}_2$ is a second-order variable ranging over sets of nodes and $a \in \Sigma$. We require that in every subformula $\mu X.\varphi$ or $\nu X.\varphi$ of a given formula, the variable X occurs within φ only in an even number

[1] Here we only need a very restricted version of MSO-transductions. More generally one allows to specify MSO-formulas $\delta_i(x)$ for all $1 \leq i \leq m$ over the signature $\mathcal{S}$ and then defines the universe of $f(\mathcal{A})$ as $\{(a, i) \mid 1 \leq i \leq m, a \in A, \mathcal{A} \models \delta_i(a)\}$.

of negations. Variables from $\mathbb{V}_2$ are bounded by the μ- and ν-operator. We define the semantics of the modal μ-calculus w.r.t. an edge-labeled graph $G = (V, (E_a)_{a\in\Sigma})$ ($E_a \subseteq V \times V$ is the set of all a-labeled edges) and a valuation $\sigma : \mathbb{V}_2 \to 2^V$. To each formula φ we assign the set $\varphi^G(\sigma) \subseteq V$ of nodes where φ evaluates to true under the valuation σ. For a valuation σ, a variable $X \in \mathbb{V}_2$, and a set $U \subseteq V$ define $\sigma[U/X]$ as the valuation with $\sigma[U/X](X) = U$ and $\sigma[U/X](Y) = \sigma(Y)$ for $X \neq Y$. Now we can define $\varphi^G(\sigma)$ inductively as follows:

- $\text{true}^G(\sigma) = V$, $\text{false}^G(\sigma) = \emptyset$
- $X^G(\sigma) = \sigma(X)$ for every $X \in \mathbb{V}_2$
- $(\neg\varphi)^G(\sigma) = V \setminus \varphi^G(\sigma)$
- $(\varphi \vee \psi)^G(\sigma) = \varphi^G(\sigma) \cup \psi^G(\sigma)$, $(\varphi \wedge \psi)^G(\sigma) = \varphi^G(\sigma) \cap \psi^G(\sigma)$
- $(\langle a\rangle\varphi)^G(\sigma) = \{u \in V \mid \exists v \in V : (u, v) \in E_a \ \wedge \ v \in \varphi^G(\sigma)\}$
- $([a]\varphi)^G(\sigma) = \{u \in V \mid \forall v \in V : (u, v) \in E_a \ \Rightarrow \ v \in \varphi^G(\sigma)\}$
- $(\mu X.\varphi)^G(\sigma) = \bigcap\{U \subseteq V \mid \varphi^G(\sigma[U/X]) \subseteq U\}$
- $(\nu X.\varphi)^G(\sigma) = \bigcup\{U \subseteq V \mid U \subseteq \varphi^G(\sigma[U/X])\}$

The set $(\mu X.\varphi)^G(\sigma)$ is the smallest fixpoint of the monotonic mapping $U \mapsto \varphi^G(\sigma[U/X])$, whereas $(\nu X.\varphi)^G(\sigma)$ is the largest fixpoint of this mapping. Monotony follows from the restriction that X occurs within an even number of negations within φ. Note that only the values of the valuation σ for free variables is important. In particular, if φ is a sentence (i.e., a formula where all variables are bounded by fixpoint operators), then the valuation σ is not relevant and we can write φ^G instead of $\varphi^G(\sigma)$, where σ is an arbitrary valuation. For a sentence φ and a node $v \in V$ we write $(G, v) \models \varphi$ if $v \in \varphi^G$. It is known that for every sentence φ of the modal μ-calculus one can construct an MSO-formula $\psi(x)$ such that for every node $v \in V$: $(G, v) \models \varphi$ if and only if $G \models \psi(v)$.

A *context-free graph* [23] is the transition graph of a pushdown automaton, i.e., nodes are the configurations of a given pushdown automaton, and edges are given by the transitions of the automaton. A more formal definition is not necessary for the purpose of this paper. We will only need the following result:

Theorem 3.2 ([12, 38]). *The following problem is in EXPTIME:*

INPUT: A pushdown automaton A defining a context-free graph $G(A)$, a node v of $G(A)$, and a formula φ of the modal μ-calculus

QUESTION: $(G(A), v) \models \varphi$?

Moreover, there exists already a fixed formula φ for which this question becomes EXPTIME-complete.

4 Word problems and Cayley-graphs

Let $\mathcal{M} = (M, \circ, 1)$ be a finitely generated monoid with identity 1 and let Σ be a finite generating set for $\mathcal{M}$, i.e., there exists a surjective monoid homomorphism $h : \Sigma^* \to \mathcal{M}$. The *word problem* for $\mathcal{M}$ w.r.t. Σ is the following problem:

INPUT: Words $u, v \in \Sigma^*$

QUESTION: $h(u) = h(v)$?

The following fact is well-known:

Proposition 4.1. *Let $\mathcal{M}$ be a finitely generated monoid and let Σ_1 and Σ_2 be two finite generating sets for $\mathcal{M}$. Then the word problem for $\mathcal{M}$ w.r.t. Σ_1 is logspace reducible to the word problem for $\mathcal{M}$ w.r.t. Σ_2.*

Since we are only interested in the complexity (resp. decidability) status of word problems, we can just speak of the word problem for a given monoid.

Thus, the computational complexity of the word problem does not depend on the underlying set of generators.

The *Cayley-graph* of $\mathcal{M}$ w.r.t. Σ is the following relational structure:

$$\mathcal{C}(\mathcal{M}, \Sigma) = (M, (\{(u, v) \in M \times M \mid u \circ h(a) = v\})_{a \in \Sigma}, 1)$$

It is a rooted (1 is the root) directed graph, where every edge has a label from Σ and $\{(u, v) \mid u \circ h(a) = v\}$ is the set of a-labeled edges. Since Γ generates $\mathcal{M}$, every $u \in M$ is reachable from the root 1.

Cayley-graphs of groups play an important role in combinatorial group theory [17], see also the survey of Schupp [32]. On the other hand, only a few papers deal with Cayley-graphs of monoids. Combinatorial aspects of Cayley-graphs of monoids are studied in [9, 10, 11, 39]. In [34, 35], Cayley-graphs of automatic monoids are investigated.

The *free group* $\mathrm{FG}(\Gamma)$ generated by the set Γ is the quotient monoid $(\Gamma \cup \Gamma^{-1})^*/\delta$, where δ is the smallest congruence on $(\Gamma \cup \Gamma^{-1})^*$ that contains all pairs (bb^{-1}, ε) for $b \in \Gamma \cup \Gamma^{-1}$. Let $\gamma : (\Gamma \cup \Gamma^{-1})^* \to \mathrm{FG}(\Gamma)$ denote the canonical morphism mapping a word $u \in (\Gamma \cup \Gamma^{-1})^*$ to the group element represented by u. It is well known that for every $u \in (\Gamma \cup \Gamma^{-1})^*$ there exists a unique word $r(u) \in (\Gamma \cup \Gamma^{-1})^*$ (the reduced normalform of u) such that $\gamma(u) = \gamma(r(u))$ and $r(u)$ does not contain a factor of the form bb^{-1} for $b \in \Gamma \cup \Gamma^{-1}$. The word $r(u)$ can be calculated from u in linear time [2]. It holds $\gamma(u) = \gamma(v)$ if and only if $r(u) = r(v)$.

The Cayley-graph of $\mathrm{FG}(\Gamma)$ w.r.t. the standard generating set $\Gamma \cup \Gamma^{-1}$ will be denoted by $\mathcal{C}(\Gamma)$; it is a finitely-branching tree and a context-free graph [23]. Figure 1 shows a finite portion of $\mathcal{C}(\{a, b\})$. Here, and in the following, we only draw one directed edge between two points. Thus, for every drawn x-labeled edge we omit the x^{-1}-labeled reversed edge.

The concrete shape of a Cayley-graph $\mathcal{C}(\mathcal{M}, \Sigma)$ depends heavily on the chosen set of generators Σ. Nevertheless, and similarly to the word problem, the chosen generating set has no influence on the decidability (or complexity) of the first-order (resp. monadic second-order) theory of the Cayley-graph:

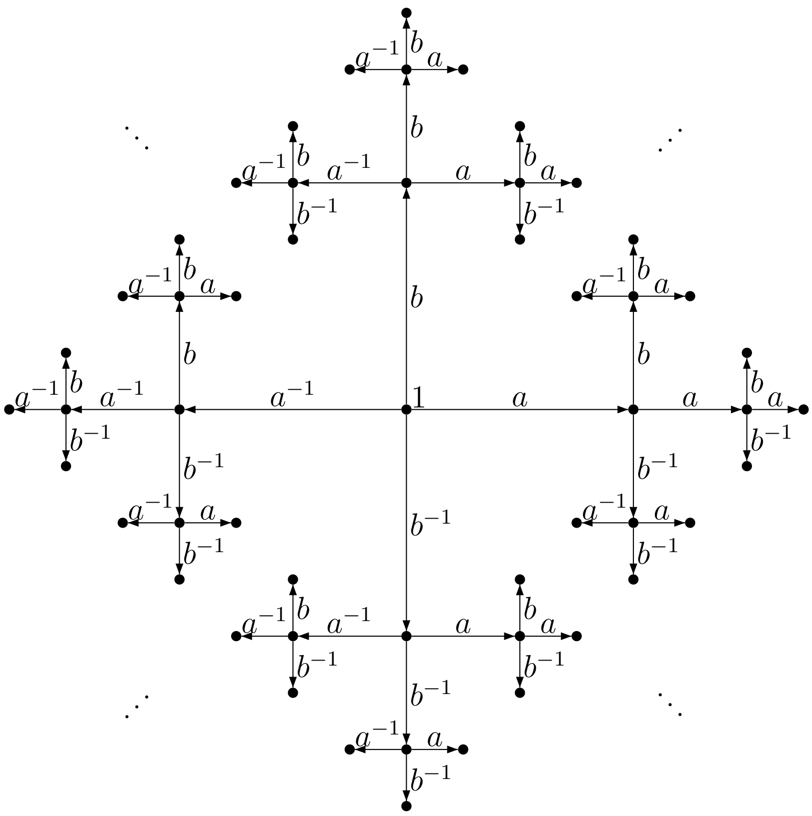

Figure 1: The Cayley-graph $\mathcal{C}(\{a,b\})$ of the free group $\mathrm{FG}(\{a,b\})$

Proposition 4.2 ([14]). *Let Σ_1 and Σ_2 be finite generating sets for the monoid $\mathcal{M}$. Then the first-order theory of $\mathcal{C}(\mathcal{M},\Sigma_1)$ is logspace reducible to the first-order theory of $\mathcal{C}(\mathcal{M},\Sigma_2)$ and the same holds for the MSO-theories.*

Thus, similarly to the word problem, we will just speak of the Cayley-graph of a monoid in statements concerning the complexity (resp. decidability) of the first-order (monadic second-order) theory of Cayley-graphs.

It is easy to see that the decidability of the first-order theory of the Cayley-graph implies the decidability of the word problem. On the other hand, there exists a finitely presented monoid for which the word problem is decidable, but the first-order theory of the Cayley-graph is undecidable, see [14]. When restricting to groups, the situation is different: The Cayley-graph of a finitely generated group has a decidable first-order theory if and only if the group has a decidable word problem [13]. Moreover, the Cayley-graph of a finitely generated group has a decidable monadic second-order theory if and only if the group is virtually free (i.e., has a free subgroup of finite index) [13, 23]. We will only need the latter result for the Cayley-graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$:

Theorem 4.3 ([23]). *For every finite set Γ, $\mathrm{MSOTh}(\mathcal{C}(\Gamma))$ is decidable.*

Remark 4.4. *It is known that already the complexity of the MSO-theory of $\mathbb{Z}$ with the successor function is nonelementary [21], i.e., the running time of every algorithm for deciding this theory cannot be bounded by an exponent tower of fixed height. It follows that also the complexity of* $\mathrm{MSOTh}(\mathcal{C}(\Gamma))$ *is nonelementary.*

5 Inverse Monoids

A monoid $\mathcal{M}$ is called an *inverse monoid* if for each $m \in \mathcal{M}$ there is a unique $m^{-1} \in \mathcal{M}$ such that $m = mm^{-1}m$ and $m^{-1} = m^{-1}mm^{-1}$. For detailed reference on inverse monoids see [27]; here we only recall the basic notions. Since the class of inverse monoids forms a variety it follows from universal algebra that *free inverse monoids* exist. The free inverse monoid generated by a set Γ is denoted by $\mathrm{FIM}(\Gamma)$; it is isomorphic to $(\Gamma \cup \Gamma^{-1})^*/\rho$, where ρ is the smallest congruence on the free monoid $(\Gamma \cup \Gamma^{-1})^*$ which contains for all words $v, w \in (\Gamma \cup \Gamma^{-1})^*$ the pairs $(w, ww^{-1}w)$ and $(ww^{-1}vv^{-1}, vv^{-1}ww^{-1})$ (which are also called the Vagner equations). Let $\alpha : (\Gamma \cup \Gamma^{-1})^* \to \mathrm{FIM}(\Gamma)$ denote the canonical morphism mapping a word $u \in (\Gamma \cup \Gamma^{-1})^*$ to the element of $\mathrm{FIM}(\Gamma)$ represented by u. Obviously, there exists a morphism $\beta : \mathrm{FIM}(\Gamma) \to \mathrm{FG}(\Gamma)$ such that $\gamma = \beta \circ \alpha$, where $\gamma : (\Gamma \cup \Gamma^{-1})^* \to \mathrm{FG}(\Gamma)$ is the canonical morphism from the previous section.

The free inverse monoid $\mathrm{FIM}(\Gamma)$ can be also represented via *Munn trees*: The Munn tree $\mathrm{MT}(u)$ of $u \in (\Gamma \cup \Gamma^{-1})^*$ is a finite and connected subset of the Cayley-graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$; it is defined by

$$\mathrm{MT}(u) = \{\gamma(v) \in \mathrm{FG}(\Gamma) \mid \exists w \in (\Gamma \cup \Gamma^{-1})^* : u = vw\}.$$

In other words, $\mathrm{MT}(u)$ is the set of all nodes along the unique path in $\mathcal{C}(\Gamma)$ that starts in 1 and that is labeled with the word u. We identify $\mathrm{MT}(u)$ with the subtree $\mathcal{C}(\Gamma)\restriction_{\mathrm{MT}(u)}$ of $\mathcal{C}(\Gamma)$. Munn's theorem [24] states that $\alpha(u) = \alpha(v)$ for $u, v \in (\Gamma \cup \Gamma^{-1})^*$ if and only if $r(u) = r(v)$ (i.e., $\gamma(u) = \gamma(v)$) and $\mathrm{MT}(u) = \mathrm{MT}(v)$. It is well known that for a word $u \in (\Gamma \cup \Gamma^{-1})^*$, the element $\alpha(u) \in \mathrm{FIM}(\Gamma)$ is an idempotent element, i.e., $\alpha(uu) = \alpha(u)$, if and only if $r(u) = \varepsilon$, i.e., $\gamma(u) = 1$.

For a finite set $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ define $\mathrm{FIM}(\Gamma)/P = (\Gamma \cup \Gamma^{-1})^*/\tau$ to be the inverse monoid with the set Γ of generators and the set P of relations, where τ is the smallest congruence on $(\Gamma \cup \Gamma^{-1})^*$ generated by $\rho \cup P$. Then the canonical morphism $\mu_P : (\Gamma \cup \Gamma^{-1})^* \to \mathrm{FIM}(\Gamma)/P$ factors as $\mu_P = \nu_P \circ \alpha$ with $\nu_P : \mathrm{FIM}(\Gamma) \to \mathrm{FIM}(\Gamma)/P$.

We say that $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ is an *idempotent presentation* if for all $(e, f) \in P$, $\alpha(e)$ and $\alpha(f)$ are both idempotents of $\mathrm{FIM}(\Gamma)$, i.e., $r(e) = r(f) = \varepsilon$. In this paper, we are concerned with inverse monoids of the form $\mathrm{FIM}(\Gamma)/P$ for a finite idempotent presentation P. In this case, since every identity $(e, f) \in P$ is true in $\mathrm{FG}(\Gamma)$ (we have $\gamma(e) = \gamma(f) = 1$), there also exists a canonical morphism $\beta_P : \mathrm{FIM}(\Gamma)/P \to \mathrm{FG}(\Gamma)$. The following commutative diagram summarizes all morphisms introduced so far.

$$\begin{tikzcd}
 & (\Gamma \cup \Gamma^{-1})^* \arrow[dl, "\mu_P"'] \arrow[d, "\alpha"] \arrow[dr, "\gamma"] & \\
\mathrm{FIM}(\Gamma)/P \arrow[rr, bend right, "\beta_P"'] & \mathrm{FIM}(\Gamma) \arrow[l, "\nu_P"'] \arrow[r, "\beta"] & \mathrm{FG}(\Gamma)
\end{tikzcd}$$

For the rest of this paper, the meaning of the morphisms $\alpha, \beta, \beta_P, \gamma, \mu_P$, and ν_P will be fixed.

To solve the word problem for $\mathrm{FIM}(\Gamma)/P$, Margolis and Meakin [18] used a closure operation for Munn trees, which is based on work of Stephen [36]. We shortly review the ideas here. As remarked in [18], every idempotent presentation P can be replaced by the idempotent presentation $P' = \{(e, ef), (f, ef) \mid (e, f) \in P\}$, i.e., $\mathrm{FIM}(\Gamma)/P \cong \mathrm{FIM}(\Gamma)/P'$. Since $\mathrm{MT}(e) \subseteq \mathrm{MT}(ef) \supseteq \mathrm{MT}(f)$ if $r(e) = r(f) = \varepsilon$, we can restrict in the following to idempotent presentations P such that $\mathrm{MT}(e) \subseteq \mathrm{MT}(f)$ for all $(e, f) \in P$. Let $V \subseteq \mathrm{FG}(\Gamma)$. Define sets $V_i \subseteq \mathrm{FG}(\Gamma)$ $(i \geq 1)$ inductively as follows: (i) $V_1 = V$ and (ii) for $n \geq 1$ let

$$V_{n+1} = V_n \cup \bigcup_{(e,f)\in P} \{u \circ v \mid u \in V_n, \forall w \in \mathrm{MT}(e) : u \circ w \in V_n, v \in \mathrm{MT}(f)\}.^2$$

Finally, define the closure of V w.r.t. the presentation P as $\mathrm{cl}_P(V) = \bigcup_{n\geq 1} V_n$.

Example 5.1. *Assume that* $\Gamma = \{a, b\}$*,* $P = \{(aa^{-1}, a^2a^{-2}), (bb^{-1}, b^2b^{-2})\}$ *and* $u = aa^{-1}bb^{-1}$*. The Munn trees for the words in the presentation* P *and* u *look as follows, the bigger circle represents the* 1 *of* $\mathrm{FG}(\Gamma)$*:*

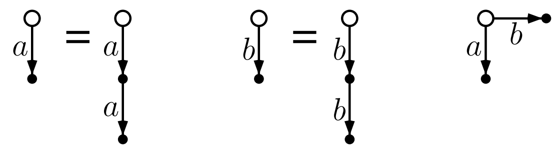

Then the closure $\mathrm{cl}_P(\mathrm{MT}(u))$ *is* $\{a^n \mid n \geq 0\} \cup \{b^n \mid n \geq 0\} \subseteq \mathrm{FG}(\Gamma)$.

In the next section, instead of specifying a word $w \in (\Gamma \cup \Gamma^{-1})^*$ (that represents an idempotent in $\mathrm{FIM}(\Gamma)$, i.e., $r(w) = 1$) explicitly, we will only show its Munn tree, where as in Example 5.1 the 1 of $\mathrm{FG}(\Gamma)$ is drawn as a bigger circle. In fact, one can replace w by any word that labels a path from the circle back to the circle and that visits all nodes in the diagram; the resulting word represents the same element of $\mathrm{FIM}(\Gamma)$ (and hence also of $\mathrm{FIM}(\Gamma)/P$) as the original word.

Theorem 5.2 ([18]). *Let* P *be an idempotent presentation and let* $u, v \in (\Gamma \cup \Gamma^{-1})^*$*. Then* $\mu_P(u) = \mu_P(v)$ *if and only if* $r(u) = r(v)$ *(i.e.,* $\gamma(u) = \gamma(v)$*) and* $\mathrm{cl}_P(\mathrm{MT}(u)) = \mathrm{cl}_P(\mathrm{MT}(v))$.

The result of Munn for $\mathrm{FIM}(\Gamma)$ mentioned above is a special case of this result for $P = \emptyset$.

Remark 5.3. *Note that* $\mathrm{cl}_P(\mathrm{MT}(u)) = \mathrm{cl}_P(\mathrm{MT}(v))$ *if and only if* $\mathrm{MT}(u) \subseteq \mathrm{cl}_P(\mathrm{MT}(v))$ *and* $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$.

Margolis and Meakin used Theorem 5.2 in order to give a solution for the word problem for the monoid $\mathrm{FIM}(\Gamma)/P$. More precisely, they have shown that from a finite idempotent presentation P one can effectively construct an MSO-formula $\mathrm{CL}_P(X, Y)$ over the signature of the Cayley-graph $\mathcal{C}(\Gamma)$ such that for all words $u \in (\Gamma \cup \Gamma^{-1})^*$ and all subsets $A \subseteq \mathrm{FG}(\Gamma)$: $\mathcal{C}(\Gamma) \models \mathrm{CL}_P(\mathrm{MT}(u), A)$ if and only if $A = \mathrm{cl}_P(\mathrm{MT}(u))$. The decidability of the word problem for $\mathrm{FIM}(\Gamma)/P$ is an immediate consequence of Theorem 4.3 and Theorem 5.2.

[2]Here, $\circ$ refers to the multiplication in the free group $\mathrm{FG}(\Gamma)$.

6 Complexity of the word problem

The direct use of Theorem 4.3 leads to a nonelementary algorithm for the word problem for the monoid $\mathrm{FIM}(\Gamma)/P$, see Remark 4.4. Using tree automata techniques we will show:

Theorem 6.1. *For every finite idempotent presentation $P \subseteq (\Gamma\cup\Gamma^{-1})^* \times (\Gamma\cup\Gamma^{-1})^*$ the word problem for $\mathrm{FIM}(\Gamma)/P$ can be solved in deterministic polynomial time.*

Proof. Let us fix a finite and idempotent presentation $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ and let $u, v \in (\Gamma \cup \Gamma^{-1})^*$. By Theorem 5.2 we have to check whether $r(u) = r(v)$ and $\mathrm{cl}_P(\mathrm{MT}(u)) = \mathrm{cl}_P(\mathrm{MT}(v))$. The first property $r(u) = r(v)$ can be checked in linear time [2]. By Remark 5.3, the property $\mathrm{cl}_P(\mathrm{MT}(u)) = \mathrm{cl}_P(\mathrm{MT}(v))$ is equivalent to

$$\forall w \in \mathrm{MT}(u) : w \in \mathrm{cl}_P(\mathrm{MT}(v)) \ \wedge\ \forall w \in \mathrm{MT}(v) : w \in \mathrm{cl}_P(\mathrm{MT}(u)).$$

Let us fix a prefix p of the word v. It suffices to show that we can check in polynomial time whether $\gamma(p) \in \mathrm{cl}_P(\mathrm{MT}(u)) \subseteq \mathrm{FG}(\Gamma)$.

Recall that there is an MSO-formula $\mathrm{CL}_P(X,Y)$ over the signature of the Cayley-graph $\mathcal{C}(\Gamma)$ such that for all subsets $A \subseteq \mathrm{FG}(\Gamma)$: $\mathcal{C}(\Gamma) \models \mathrm{CL}_P(\mathrm{MT}(u), A)$ if and only if $A = \mathrm{cl}_P(\mathrm{MT}(u))$. Let in-$\mathrm{cl}_P(x, X)$ be the formula $\exists Y : \mathrm{CL}_P(X, Y) \wedge x \in Y$. Thus, we have to check whether $\mathcal{C}(\Gamma) \models$ in-$\mathrm{cl}_P(\gamma(p), \mathrm{MT}(u))$. Here, it is important to note that since P is a fixed presentation, in-$\mathrm{cl}_P(x, X)$ is a fixed MSO-formula over the signature of the Cayley-graph $\mathcal{C}(\Gamma)$.

Let T_Γ be the $(2 \cdot |\Gamma|)$-ary tree

$$T_\Gamma = ((\Gamma \cup \Gamma^{-1})^*, (\mathrm{suc}_a)_{a\in\Gamma\cup\Gamma^{-1}}),$$

where $\mathrm{suc}_a = \{(w, wa) \mid w \in (\Gamma \cup \Gamma^{-1})^*\}$, and let $\mathrm{IRR}(\Gamma) = \{r(w) \mid w \in (\Gamma \cup \Gamma^{-1})^*\}$ be the set of all reduced words. In a next step, we translate the fixed MSO-formula in-$\mathrm{cl}_P(x, X)$ into a fixed MSO-formula $\psi_P(x, X)$ over the signature of T_Γ such that for every $t \in \mathrm{IRR}(\Gamma)$ and every $A \subseteq \mathrm{IRR}(\Gamma)$ we have $T_\Gamma \models \psi_P(t, A)$ if and only if $\mathcal{C}(\Gamma) \models$ in-$\mathrm{cl}_P(\gamma(t), \gamma(A))$. For this, one has to notice that $\mathcal{C}(\Gamma)$ is isomorphic to the structure

$$\begin{aligned}(\mathrm{IRR}(\Gamma), (\{(u, ua) \mid u \in \mathrm{IRR}(\Gamma) \setminus (\Gamma \cup \Gamma^{-1})^* a^{-1}\} \cup \\ \{(ua^{-1}, u) \mid u \in \mathrm{IRR}(\Gamma) \setminus (\Gamma \cup \Gamma^{-1})^* a\})_{a\in\Gamma\cup\Gamma^{-1}}, \varepsilon).\end{aligned}$$

Since $\mathrm{IRR}(\Gamma)$ is a regular subset of $(\Gamma \cup \Gamma^{-1})^*$ and hence MSO-definable in T_Γ, it follows that $\mathcal{C}(\Gamma)$ is MSO-definable in T_Γ, see also [18].

We now calculate the set

$$U = \{r(s) \mid \exists w \in (\Gamma \cup \Gamma^{-1})^* : u = sw\}$$

(which uniquely represents $\mathrm{MT}(u)$) and the word $s = r(p)$. Thus, it remains to check whether $T_\Gamma \models \psi_P(s, U)$.

Next, we translate the fixed MSO-formula $\psi_P(x, X)$ into a (top-down) ω-tree automaton $\mathcal{A}_P$. The automaton $\mathcal{A}_P$ runs on a labeled ω-tree $((\Gamma \cup \Gamma^{-1})^*, (\mathrm{suc}_a)_{a\in\Gamma\cup\Gamma^{-1}}, \lambda)$, where $\lambda : (\Gamma \cup \Gamma^{-1})^* \to$

$\{0,1\} \times \{0,1\}$ is the labeling function. The property of $\mathcal{A}_P$ is that $T_\Gamma \models \psi_P(s, U)$ if and only if $\mathcal{A}_P$ accepts the ω-tree

$$T_{s,U} = ((\Gamma \cup \Gamma^{-1})^*, (\mathrm{suc}_a)_{a \in \Gamma \cup \Gamma^{-1}}, \lambda),$$

where for all $w \in (\Gamma \cup \Gamma^{-1})^*$ with $\lambda(w) = (i, j)$ we have: $i = 1$ if and only if $w = s$ and $j = 1$ if and only if $w \in U$. Again, since $\psi_P(x, X)$ is a fixed MSO-formula, $\mathcal{A}_P$ is a fixed ω-tree automaton. The translation from $\psi_P(x, X)$ to $\mathcal{A}_P$ is the standard translation from MSO-formulas to automata, see [29, Theorem 1.7]. It remains to check whether $\mathcal{A}_P$ accepts the ω-tree $T_{s,U}$.

The final step translates $T_{s,U}$ into a finite tree $t^{\mathrm{fin}}_{s,U}$. Note that in $T_{s,U}$ almost all nodes are labeled with $(0, 0)$ (note that U is a finite set of words). Let B be the set of all words of the form wa, where $w \in (\Gamma \cup \Gamma^{-1})^*$, $a \in \Gamma \cup \Gamma^{-1}$, $\lambda(wat) = (0, 0)$ for every $t \in (\Gamma \cup \Gamma^{-1})^*$, but $\lambda(w) \neq (0, 0)$. We construct the tree $t^{\mathrm{fin}}_{s,U}$ by taking $T_{s,U}$ but making every node $w \in B$ to a leaf of $t^{\mathrm{fin}}_{s,U}$ that is labeled with the new symbol # (all proper prefixes of words from B are labeled as in $T_{s,U}$). Note that $t^{\mathrm{fin}}_{s,U}$ is a finite tree that can be constructed from s and U in polynomial time. Before we continue, let us give an example. Let $U = \{\varepsilon, a, aa, a^{-1}\}$ and $s = a^{-1}bb$. Then $t^{\mathrm{fin}}_{s,U}$ is the following tree.

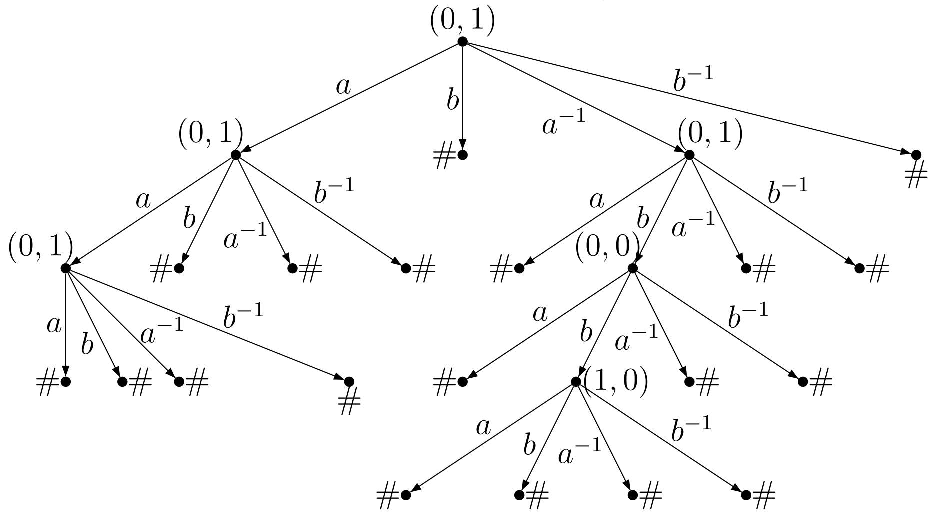

Now, from the fixed ω-tree automaton $\mathcal{A}_P$ it is easy to construct a fixed tree automaton $\mathcal{A}^{\mathrm{fin}}_P$ (working on finite trees) such that $\mathcal{A}_P$ accepts $T_{s,U}$ if and only if $\mathcal{A}^{\mathrm{fin}}_P$ accepts $t^{\mathrm{fin}}_{s,U}$. Basically, $\mathcal{A}^{\mathrm{fin}}_P$ has the same states and transitions as $\mathcal{A}_P$, except that $\mathcal{A}^{\mathrm{fin}}_P$ accepts in a #-labeled leaf in state q if and only if $\mathcal{A}_P$ accepts the full ω-tree with all nodes labeled $(0, 0)$ when starting in state q. Finally, whether $\mathcal{A}^{\mathrm{fin}}_P$ accepts $t^{\mathrm{fin}}_{s,U}$ can be checked in polynomial time. $\square$

A closer analysis of the algorithm in the previous proof shows that the word problem of $\mathrm{FIM}(\Gamma)/P$ even belongs to the class NC, which is the class of all problems that can be solved in poly-logarithmic time using polynomially many processors on a PRAM; roughly speaking, NC is the class of all problems in P that can be efficiently parallelized. To see that the word problem of $\mathrm{FIM}(\Gamma)/P$ belongs to NC, one has to notice that:

- The word problem of the free group $\mathrm{FG}(\Gamma)$ (i.e., the question, whether $r(u) = r(v)$) can be solved in NC; in fact it can be even solved in deterministic logspace [15].

- For every fixed tree automaton $\mathcal{A}$, the membership problem of $\mathcal{A}$ belongs to $\mathrm{NC}^1 \subseteq \mathrm{NC}$ [16].
- The tree $t^{\mathrm{fin}}_{s,U}$ can be build in NC from the word u and the prefix p of v, basically by using the NC-algorithm for the word problem of $\mathrm{FG}(\Gamma)$ in order to calculate in parallel which prefixes of the word u represent the same element of $\mathrm{FG}(\Gamma)$.

In the uniform case, where the presentation P is part of the input, the complexity increases considerably:

Theorem 6.2. *There exists a fixed alphabet Γ such that the following problem is EXPTIME-complete:*
INPUT: Words $u, v \in (\Gamma \cup \Gamma^{-1})^*$ and a finite idempotent presentation $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$
QUESTION: $\mu_P(u) = \mu_P(v)$?
The EXPTIME upper bound even holds if the alphabet Γ belongs to the input.

Proof. For the lower bound we use the fact that EXPTIME equals APSPACE. Thus, let

$$T = (Q, \Sigma, \delta, q_0, q_f)$$

be a fixed alternating Turing machine that accepts an EXPTIME-complete language. Assume that T works in space $p(n)$ for a polynomial p on an input of length n. W.l.o.g. we may assume the following:

- T alternates in each state, i.e., it either moves from a state of $Q_\exists$ to a state from $Q_\forall \cup \{q_f\}$ or from a state of $Q_\forall$ to a state from $Q_\exists \cup \{q_f\}$.
- $q_0 \in Q_\exists$
- For each pair $(q, a) \in (Q \setminus \{q_f\}) \times \Sigma$, the machine T has precisely two choices according to the transition relation δ, which we call choice 1 and choice 2.
- If T terminates in the final state q_f, then the symbol that is currently read by the head is some distinguished symbol $\# \in \Sigma$.

Define $\Gamma = \Sigma \cup (Q \times \Sigma) \cup \{a_1, a_2, b_1, b_2, \#\}$, where all unions are assumed to be disjoint. A configuration of T is encoded as a word from $\#\Sigma^*(Q \times \Sigma)\Sigma^*\# \subseteq \Gamma^*$. Now let $w \in \Sigma^*$ be an input of length n and let $m = p(n)$. Then a configuration of T is a word from $\bigcup_{i=0}^{m-1} \#\Sigma^i(Q \times \Sigma)\Sigma^{m-i-1}\# \subseteq \Gamma^{m+2}$. Clearly, the symbol at position $1 < i < m+2$ at time $t+1$ in a configuration only depends on the symbols at the positions $i-1$, i, and $i+1$ at time t. Assume that $c, c_1, c_2, c_3 \in \Sigma \cup (Q \times \Sigma) \cup \{\#\}$ such that $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*(Q \times \Sigma)\Sigma^*\{\varepsilon, \#\}$. We write $c_1c_2c_3 \xrightarrow{j} c$ for $j \in \{1, 2\}$ if the following holds: If three consecutive positions $i-1$, i, and $i+1$ of a configuration contain the symbol sequence $c_1c_2c_3$, then choice j of T results in the symbol c at position i. We write $c_1c_2c_3 \xrightarrow{\exists} (d_1, d_2)$ for $c_1, c_2, c_3, d_1, d_2 \in \Sigma \cup (Q \times \Sigma) \cup \{\#\}$ if one of the following two cases holds:

- $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*(Q_\exists \times \Sigma)\Sigma^*\{\varepsilon, \#\}$ and $c_1c_2c_3 \xrightarrow{j} d_j$ for $j \in \{1, 2\}$
- $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*\{\varepsilon, \#\}$ and $d_1 = d_2 = c_2$.

The notation $c_1c_2c_3 \xrightarrow{\forall} (d_1, d_2)$ is defined analogously, except that in the first case we require $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*(Q_\forall \times \Sigma)\Sigma^*\{\varepsilon, \#\}$.

Let us briefly describe the idea for the lower bound proof. We will encode a configuration $\#c_1c_2 \cdots c_m\#$, where the current state is from $Q_\exists$ by a subgraph of the Cayley-graph $\mathcal{C}(\Gamma)$ of the following form, where $i = 1$ or $i = 2$:

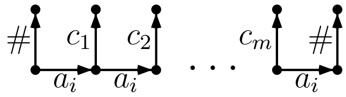

If the current state is from $Q_\forall$, then we take the same subgraph, except that a_i is replaced by b_i. The idempotent presentation $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ is constructed in such a way from the machine T that building the closure from a Munn tree that represents the initial configuration (in the above sense) corresponds to generating the whole computation tree of the Turing machine T starting from the initial configuration. We will describe each pair $(e, f) \in P$ by the Munn trees of e and f.

For all $x \in \{a_1, a_2, b_1, b_2\}$ put the following equation into P, which propagates the end-marker $\#$ along intervals of length $m + 2$ (here, the x^m-labeled edge abbreviates a path consisting of m many x-labeled edges):

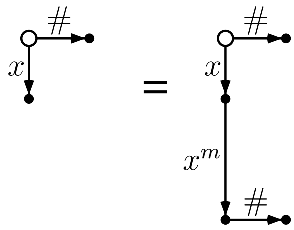

The next two equation types generate the two successor configurations of the current configuration. If $c_1c_2c_3 \xrightarrow{\exists} (d_1, d_2)$, then for every $0 \leq k \leq m - 1$ we include the following equation in P:

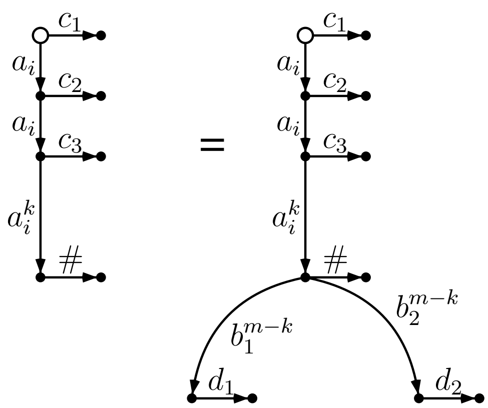

If $c_1c_2c_3 \xrightarrow{\forall} (d_1, d_2)$, then for every $0 \leq k \leq m - 1$ we take the following equation:

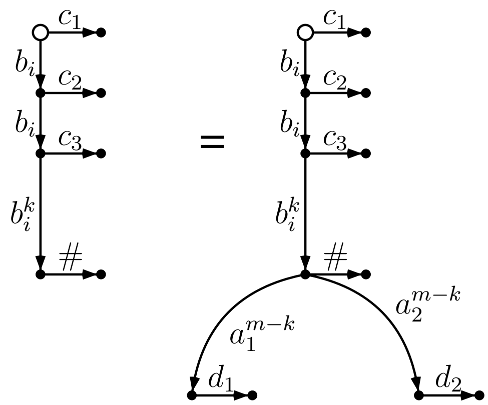

The remaining equations propagate acceptance information back to the initial Munn tree. Here the separation of the state set into existential and universal states becomes crucial. Let $c_f = (q_f, \#)$; recall that $\#$ is the symbol under the head of T when T terminates in state q_f. For all $x \in \{a_1, a_2, b_1, b_2\}$ and all $i, j \in \{1, 2\}$ we put the following equations into P:

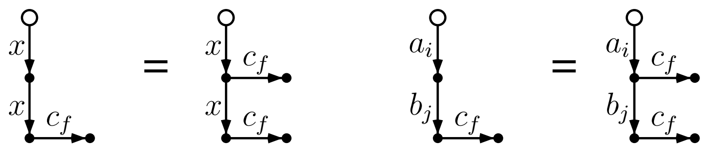

Here, the second equation expresses the fact that an existential configuration is accepting if and only if at least one successor configuration is accepting.

Finally, for $i \in \{1, 2\}$ we add the following equation to P, which reflects the fact that a universal configuration is accepting if and only if both successor configurations are accepting.

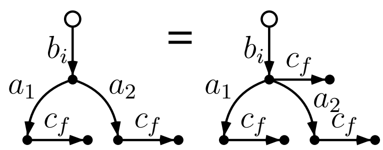

This concludes the description of the presentation P. Now define the words $u, v \in (\Gamma \cup \Gamma^{-1})^*$ as follows: Assume that the input word for our alternating Turing machine w is of the form $w = w_1 w_2 \cdots w_n$ with $w_i \in \Sigma$. For $n+1 \leq i \leq m$ define $w_i = \Box$, where $\Box$ is the blank symbol of T. Then the Munn trees of u and v look as follows (we assume $r(u) = r(v) = \varepsilon$):

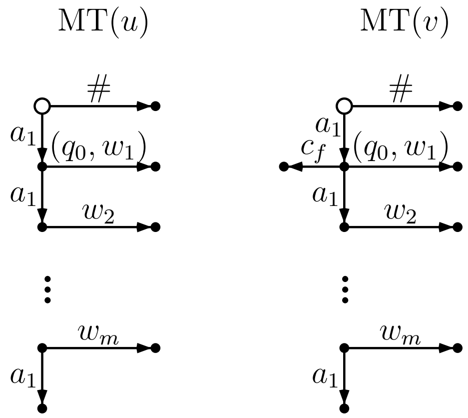

We claim that $\mu_P(u) = \mu_P(v)$ if and only if the machine T accepts the word w. From the construction of u, v, and P it follows easily that T accepts the word w if and only if $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$. Since $\mathrm{MT}(u) \subseteq \mathrm{MT}(v)$ this is equivalent to $\mathrm{cl}_P(\mathrm{MT}(v)) = \mathrm{cl}_P(\mathrm{MT}(u))$ (see Remark 5.3), i.e., $\mu_P(u) = \mu_P(v)$ due to Theorem 5.2 (note that $r(u) = r(v) = \varepsilon$). This proves the EXPTIME lower bound.

For the upper bound let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be an idempotent presentation and let $u, v \in (\Gamma \cup \Gamma^{-1})^*$. Since $r(u) = r(v)$ can be checked in linear time, it suffices by Theorem 5.2 to verify in EXPTIME whether $\mathrm{cl}_P(\mathrm{MT}(v)) = \mathrm{cl}_P(\mathrm{MT}(u))$. By Remark 5.3, it is enough to show that we can check in EXPTIME, whether $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$.

Let G be the graph that results from the Cayley-graph $\mathcal{C}(\Gamma)$ by adding a new node v_0 and adding a $\#$-labeled edge from node 1 (i.e., the origin) of $\mathcal{C}(\Gamma)$ to the new node v_0 . Here, the edge label $\#$ is assumed to be not in $\Gamma \cup \Gamma^{-1}$ (the label set of $\mathcal{C}(\Gamma)$). Since $\mathcal{C}(\Gamma)$ is a context-free graph, also G is context-free. We decide $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$ by constructing from u, v and P in polynomial time a formula $\varphi_{u,v,P}$ of the modal μ-calculus such that $(G, 1) \models \varphi_{u,v,P}$ if and only if $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$. Then the EXPTIME upper bound follows from Theorem 3.2.

In the following, we define for a word $w = a_1 a_2 \cdots a_m$ ($a_i \in \Gamma \cup \Gamma^{-1}$) and two positions $i, j \in \{1, \ldots, m\}$, $i \leq j$, the word $w[i, j] = a_i \cdots a_j$. If $i > j$, then set $w[i, j] = \varepsilon$. Moreover, we use $\langle w \rangle \phi$ as an abbreviation for $\langle a_1 \rangle \langle a_2 \rangle \cdots \langle a_m \rangle \phi$. Now assume that $P = \{(e_i, f_i) \mid 1 \leq i \leq n\}$, where $\mathrm{MT}(e_i) \subseteq \mathrm{MT}(f_i)$. First, let $\varphi_{u,P}$ be the following sentence:

$$\mu X. \left(\bigvee_{i=0}^{|u|} \langle u[1, i]^{-1} \rangle \langle \# \rangle \text{true} \ \vee \ \bigvee_{i=1}^{n} \bigvee_{j=0}^{|f_i|} \langle f_i[1, j]^{-1} \rangle (\bigwedge_{k=0}^{|e_i|} \langle e_i[1, k] \rangle X) \right)$$

Then $(G, x) \models \varphi_{u,P}$ if and only if the node x belongs to $\mathrm{cl}_P(\mathrm{MT}(u))$. In the formula $\varphi_{u,P}$, the disjunction $\bigvee_{i=0}^{|u|} \langle u[1, i]^{-1} \rangle \langle \# \rangle \text{true}$ expresses $\mathrm{MT}(u) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$. The disjunction

$$\bigvee_{i=1}^{n} \bigvee_{j=0}^{|f_i|} \langle f_i[1, j]^{-1} \rangle (\bigwedge_{k=0}^{|e_i|} \langle e_i[1, k] \rangle X)$$

defines all nodes such that via some prefix of some word f_i a node x can be reached such that the whole path starting in x and labeled with e_i already belongs to X. For the correctness, it is important to note that $\mathcal{C}(\Gamma)$ is a deterministic graph, i.e., for every $a \in \Gamma \cup \Gamma^{-1}$, every node x has exactly one a-labeled outgoing edge. Thus, it is not relevant, whether the $[a]$- or $\langle a \rangle$-modality is used. Moreover, for every node x the path starting in x and labeled with the word e_i also ends in x ($r(e_i) = \varepsilon$). Finally, we can take for $\varphi_{u,v,P}$ the sentence $\bigwedge_{i=0}^{|v|} \langle v[1,i] \rangle \varphi_{u,P}$. $\square$

The following result was conjectured in [38].

Corollary 6.3. *There exists a fixed context-free graph, for which the model-checking problem of the modal μ-calculus (restricted to formulas of nesting depth* 1*) is EXPTIME-complete.*

Proof. We can reuse the constructions from the previous proof. Note that the generating set Γ from the lower bound proof is a fixed set; thus, the Cayley-graph $\mathcal{C}(\Gamma)$ is a fixed context-free graph. Hence, also the graph G constructed in the upper bound proof by adding a $\#$-labeled edge that leaves the origin 1 is a fixed context-free graph. For the input word w for the Turing machine T let u, v, and P be the data constructed in the lower bound proof. Then w is accepted by T if and only if $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$ if and only if $(G, 1) \models \varphi_{u,v,P}$. This proves the corollary. $\square$

7 Cayley-graphs of Inverse Monoids

Let $\mathcal{M} = (M, \circ, 1)$ be a monoid with a finite generating set Σ and let $h : \Sigma^* \to \mathcal{M}$ be the canonical morphism. We define the following expansion $\mathcal{C}(\mathcal{M}, \Sigma)_{\mathrm{reg}}$ of the Cayley-graph $\mathcal{C}(\mathcal{M}, \Sigma)$:

$$\begin{aligned} \mathcal{C}(\mathcal{M}, \Sigma)_{\mathrm{reg}} &= (M, (\mathrm{reach}_L)_{L \in \mathrm{REG}(\Sigma)}, 1), \text{ where} \\ \mathrm{reach}_L &= \{(u, v) \in M \times M \mid \exists w \in L : u \circ h(w) = v\}. \end{aligned}$$

Thus, $\mathcal{C}(\mathcal{M}, \Sigma) = (M, (\mathrm{reach}_{\{a\}})_{a \in \Sigma}, 1)$. Again, the decidability (resp. complexity) of $\mathcal{C}(\mathcal{M}, \Sigma)_{\mathrm{reg}}$ does not depend on the generating set Σ:

Proposition 7.1. *Let Σ_1 and Σ_2 be finite generating sets for the monoid $\mathcal{M}$. Then the first-order theory of $\mathcal{C}(\mathcal{M}, \Sigma_1)_{\mathrm{reg}}$ is logspace reducible to the first-order theory of $\mathcal{C}(\mathcal{M}, \Sigma_2)_{\mathrm{reg}}$.*

Proof. There exists a morphism $f : \Sigma_1^* \to \Sigma_2^*$ such that for every word $w \in \Sigma_1^*$, $f(w)$ represents the same monoid element of $\mathcal{M}$ as w. Then, for a given sentence φ_1 over the signature of $\mathcal{C}(\mathcal{M}, \Sigma_1)_{\mathrm{reg}}$ we just have to replace every atomic predicate $\mathrm{reach}_L(x, y)$ by $\mathrm{reach}_{f(L)}(x, y)$. If φ_2 is the resulting sentence then $\mathcal{C}(\mathcal{M}, \Sigma_1)_{\mathrm{reg}} \models \varphi_1$ if and only if $\mathcal{C}(\mathcal{M}, \Sigma_2)_{\mathrm{reg}} \models \varphi_2$. $\square$

The main result of this section is:

Theorem 7.2. *Let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be a finite idempotent presentation. Then the first-order theory of $\mathcal{C}(\mathrm{FIM}(\Gamma)/P, \Gamma \cup \Gamma^{-1})_{\mathrm{reg}}$ is decidable.*

Remark 7.3. *It is easy to show that already for the free inverse monoid $\mathcal{M} = \mathrm{FIM}(\{a, b\})$ the complexity of $\mathrm{FOTh}(\mathcal{C}(\mathcal{M}, \{a, b, a^{-1}, b^{-1}\})_{\mathrm{reg}})$ is nonelementary: It is known that the first-order theory of the structure $\mathcal{A} = (\{a, b\}^*, (\{(w, wc) \mid w \in \{a, b\}^*\})_{c \in \{a,b\}}, \preceq)$, where $\preceq$ is the prefix relation on $\{a, b\}^*$, is nonelementary decidable, see e.g. [5]. It is straight-forward to define $\mathcal{A}$ in $\mathcal{C}(\mathcal{M}, \{a, b, a^{-1}, b^{-1}\})_{\mathrm{reg}}$ using first-order logic.*

Before we prove Theorem 7.2, let us first state a corollary. The *generalized word problem* for $\mathcal{M}$ is the following computational problem:

INPUT: Words $u, u_1, \ldots, u_n \in \Sigma^*$

QUESTION: Does $h(u)$ belong to the submonoid of $\mathcal{M}$ that is generated by $h(u_1), \ldots, h(u_n)$?

Corollary 7.4. *Let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be a finite idempotent presentation. Then the generalized word problem for $\mathrm{FIM}(\Gamma)/P$ is decidable.*

Proof. Let $u, u_1, \ldots, u_n \in (\Gamma \cup \Gamma^{-1})^*$. Then $\mu_P(u)$ belongs to the submonoid generated by the elements $\mu_P(u_1), \ldots, \mu_P(u_n)$ if and only if

$$\mathcal{C}(\mathrm{FIM}(\Gamma)/P, \Gamma \cup \Gamma^{-1})_{\mathrm{reg}} \models \exists x : \mathrm{reach}_K(1, x) \wedge \mathrm{reach}_L(1, x),$$

where $K = \{u_1, \ldots, u_n\}^*$ and $L = \{u\}$. By Theorem 7.2 this proves the corollary. $\square$

To prove Theorem 7.2 we first need some lemmas that are shown in the next section.

7.1 Some auxiliary MSO formulas

Since later on we are dealing with Munn trees, which are finite node sets, we restrict to finite graphs in this section.

Lemma 7.5. *There exists a fixed MSO-formula $\varphi(x, y)$ (over the signature consisting of a binary relation symbol E) such that for every finite directed graph $G = (V, E)$ and all nodes $s, t \in V$ we have: $G \models \varphi(s, t)$ if and only if there is a path in G with initial vertex s and terminal vertex t visiting all vertices from V.*

Proof. Let $\mathcal{V}$ be the set of all strongly connected components of G; this set forms a partition of V. We define a partial order $\prec$ on $\mathcal{V}$ by setting $S \prec T$ for $S, T \in \mathcal{V}$ with $S \neq T$ if $\exists u \in S\, \exists v \in T : (u, v) \in E^*$. Note that this implies $\forall u \in S\, \forall v \in T : (u, v) \in E^* \wedge (v, u) \notin E^*$. We claim that there is a path p in G from s to t visiting all vertices from V if and only if

(1) $\prec$ is a total order,

(2) s belongs to the minimal (w.r.t. $\prec$) strongly connected component, and

(3) t belongs to the maximal (w.r.t. $\prec$) strongly connected component.

To prove this claim, first assume that p is a path in G from s to t visiting all vertices from V. Let $S, T \in \mathcal{V}$. Since p visits all vertices of S and T, either $\exists u \in S\, \exists v \in T : (u, v) \in E^*$ (and thus $S \prec T$) or $\exists u \in T\, \exists v \in S : (u, v) \in E^*$ (and thus $T \prec S$). Thus $\prec$ is a total order. Now assume that $s \in S \in \mathcal{V}$ and that there exists $T \in \mathcal{V}$ with $T \prec S$. Thus, $\forall v \in T : (s, v) \notin E^*$, contradicting the fact that p starts in s and visits all nodes of T. Similarly, we can show that t belongs to the maximal (w.r.t. $\prec$) strongly connected component.

Now assume that the properties (1)–(3) above are true. Let $\mathcal{V} = \{S_1, \ldots S_m\}$ with $S_1 \prec S_2 \prec \cdots \prec S_m$. We construct a path p from s to t that visits all nodes of G as follows. The path p starts in the node $s \in S_1$, then it visits all nodes of S_1 followed by a path from a node of S_1 to a node of S_2. Then p visits all nodes of S_2 and so on. This proves the claim.

The lemma follows, because the properties (1)–(3) above are easily expressible in MSO, using the fact that reachability is MSO-expressible. $\square$

Lemma 7.6. *Let Σ be a finite alphabet and let $L \in \mathrm{REG}(\Sigma)$. Then one can construct an MSO-sentence ψ_L (over a signature consisting of binary relation symbols E_a ($a \in \Sigma$) and two constants s and t) such that for every finite structure $G = (V, (E_a)_{a\in\Sigma}, s, t)$ we have $G \models \psi_L$ if and only if there exists a path $p = (v_1, a_1, v_2, a_2, \dots, v_n)$ ($v_i \in V$, $a_i \in \Sigma$) such that: $v_1 = s$, $v_n = t$, $(v_i, v_{i+1}) \in E_{a_i}$ for all $1 \leq i < n$, $a_1 a_2 \cdots a_{n-1} \in L$, and $V = \{v_1, v_2, \dots, v_n\}$.*

Proof. Let G and L be as in the lemma. Let $A = (Q, \Sigma, \delta, q_0, F)$ be a deterministic finite automaton with $L(A) = L$, where w.l.o.g. $Q = \{1, \dots, m\}$. Define the structure $f_A(G) = (V \times Q, E, \Delta, I_s, F_t)$, where

$$\begin{aligned}
E &= \{((u,i),(v,j)) \mid \exists a \in \Sigma : (u,v) \in E_a \wedge \delta(i,a) = j\}, \\
\Delta &= \{((v,1),\dots,(v,m)) \mid v \in V\}, \\
I_s &= \{(s,q_0)\}, \text{ and} \\
F_t &= \{t\} \times F.
\end{aligned}$$

We claim that f_A is an MSO-transduction. For this, we have to construct the defining MSO-formulas $\theta_{P,(i_1,\dots,i_k)}(x_1,\dots,x_k)$ for every (k-ary) relation P of $f_A(G)$ and every tuple $(i_1,\dots,i_k) \in Q^k = \{1,\dots,m\}^k$:

$$\begin{aligned}
\theta_{E,(i,j)}(x_1,x_2) &= \bigvee_{a\in\Sigma,\delta(i,a)=j} E_a(x_1,x_2) \\
\theta_{\Delta,(i_1,\dots,i_m)}(x_1,\dots,x_m) &= \begin{cases} x_1 = x_2 = \cdots = x_m & \text{if } i_j = j \text{ for } 1 \leq j \leq m \\ \text{false} & \text{else} \end{cases} \\
\theta_{I_s,(i)}(x) &= \begin{cases} x = s & \text{if } i = q_0 \\ \text{false} & \text{else} \end{cases} \\
\theta_{F_t,(i)}(x) &= \begin{cases} x = t & \text{if } i \in F \\ \text{false} & \text{else} \end{cases}
\end{aligned}$$

It is easy to see that these formulas indeed define the transduction f_A. Thus, f_A is MSO-compatible; so there exists a backwards translation $f_A^\sharp$ such that for every MSO-sentence ϕ over the signature of $f_A(G)$ we have: $f_A(G) \models \phi$ if and only if $G \models f_A^\sharp(\phi)$. Now consider the following MSO-sentence ϕ over the signature of $f_A(G)$:

$$\phi = \exists X \left\{ \begin{array}{l} \forall x_1 \cdots \forall x_m : \left(\Delta(x_1,\dots,x_m) \Rightarrow \bigvee_{i=1}^{m} x_i \in X \right) \wedge \\ \exists x \in I_s \, \exists y \in F_t : x, y \in X \wedge \varphi\restriction_X(x,y,X) \end{array} \right\}$$

Here φ is the formula from Lemma 7.5. Thus, $\varphi\restriction_X(x, y, X)$ expresses that there exists a path from x to y in the graph $(V, E)\restriction_X$ visiting all nodes of X. Now, we have $f_A(G) \models \phi$ if and only if there exists a path in G from s to t that visits all nodes of V and that is labeled with a word from the language L. Thus, $f_A^\sharp(\phi)$ is the desired sentence ψ_L. This completes the proof of the lemma. $\square$

Lemma 7.7. *Let Σ be a finite alphabet and let $L \in \mathrm{REG}(\Sigma)$. Then one can construct an MSO-formula $\theta_L(X)$ (over a signature consisting of binary relation symbols E_a ($a \in \Sigma$) and two constants s and t) such that for every finite structure $G = (V, (E_a)_{a\in\Sigma}, s, t)$ and every finite set $U \subseteq V$ we have $G \models \theta_L(U)$ if and only if there exists a path $p = (v_1, a_1, v_2, a_2, \ldots, v_n)$ ($v_i \in V$, $a_i \in \Sigma$) such that: $v_1 = s$, $v_n = t$, $(v_i, v_{i+1}) \in E_{a_i}$ for all $1 \leq i < n$, $a_1 a_2 \cdots a_{n-1} \in L$, and $U \subseteq \{v_1, v_2, \ldots, v_n\}$.*

Proof. For $\theta_L(X)$ we can take the formula $\exists Y : X \subseteq Y \;\wedge\; s, t \in Y \;\wedge\; \psi_L\restriction_Y(Y)$, where ψ_L is the sentence from Lemma 7.6. $\square$

7.2 Proof of Theorem 7.2

Let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be a finite idempotent presentation. We want to show that the first-order theory of the structure $\mathcal{A} = \mathcal{C}(\mathrm{FIM}(\Gamma)/P, \Gamma \cup \Gamma^{-1})_{\mathrm{reg}}$ is decidable. For this, we use Theorem 5.2 and translate each first-order sentence φ over $\mathcal{A}$ into an MSO-sentence $\widehat{\varphi}$ over the Cayley graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$ such that for a sentence φ over $\mathcal{A}$ we have: $\mathcal{A} \models \varphi$ if and only if $\mathcal{C}(\Gamma) \models \widehat{\varphi}$. Together with Theorem 4.3 this will complete the proof of Theorem 7.2.

To every variable x (ranging over $\mathrm{FIM}(\Gamma)/P$) in φ we associate two variables in $\widehat{\varphi}$:

- an MSO-variable X' representing $\mathrm{cl}_P(\mathrm{MT}(u))$, where $u \in (\Gamma \cup \Gamma^{-1})^*$ is any word with $\mu_P(u) = x$, and
- a first-order variable x', representing $\beta_P(x) \in \mathrm{FG}(\Gamma)$ (recall that $\beta_P : \mathrm{FIM}(\Gamma)/P \to \mathrm{FG}(\Gamma)$ is the canonical morphism).

Thus, by Theorem 5.2, $x = y$ if and only if $x' = y'$ and $X' = Y'$. The relationship between x' and X' is expressed by the MSO-formula (over the signature of $\mathcal{C}(\Gamma)$) $\mathrm{MT}(x', X') = \exists X : \Theta(x', X, X')$, where:

$$\Theta(x', X, X') \;=\; (1, x' \in X \;\wedge\; X \text{ is connected and finite} \;\wedge\; \mathrm{CL}_P(X, X'))$$

Recall that by Remark 3.1, finiteness and connectedness of a subset of the finitely-branching tree $\mathcal{C}(\Gamma)$ can be expressed in MSO. Here $\mathrm{CL}_P(X, X')$ is the MSO-formula constructed by Margolis and Meakin in [18], see the remark at the end of Section 5.

Next, note that by Lemma 7.7 for every language $L \in \mathrm{REG}(\Gamma \cup \Gamma^{-1})$ there exists an MSO-formula $\xi_L(x', X, y', Y)$ over the signature of $\mathcal{C}(\Gamma)$ such that for all finite sets $U, V \subseteq \mathrm{FG}(\Gamma)$ and all nodes $u', v' \in \mathrm{FG}(\Gamma)$ we have: $\mathcal{C}(\Gamma) \models \xi_L(u', U, v', V)$ if and only if $U \subseteq V$ and there is a path from u' to v' in $\mathrm{FG}(\Gamma)\restriction_V$ that visits all vertices of $V \backslash U$ and which is labeled with a word from the language L.

Now let φ be an FO-formula over the signature of $\mathcal{A}$. We define $\widehat{\varphi}$ inductively as follows:

- for $\varphi = \mathrm{reach}_L(x, y)$ define $\widehat{\varphi} = \exists X, Y : \Theta(x', X, X') \;\wedge\; \Theta(y', Y, Y') \;\wedge\; \xi_L(x', X, y', Y)$

- for $\varphi = \neg\psi$ define $\widehat{\varphi} = \neg\widehat{\psi}$
- for $\varphi = \psi_1 \wedge \psi_2$ define $\widehat{\varphi} = \widehat{\psi_1} \wedge \widehat{\psi_2}$
- for $\varphi = \forall x : \psi$ define $\widehat{\varphi} = \forall x' \, \forall X' : \mathrm{MT}(x', X') \Rightarrow \widehat{\psi}$

The intuition behind the first formula $\exists X, Y : \Theta(x', X, X') \wedge \Theta(y', Y, Y') \wedge \xi_L(x', X, y', Y)$ is the following: We express that starting from the node $x' \in \mathrm{FG}(\Gamma)$ we traverse a path p in $\mathcal{C}(\Gamma)$ labeled with a word from the language L that ends in the node $y' \in \mathrm{FG}(G)$. Moreover, Y is the union of X and the nodes along the path p, and the closure of X (resp. Y) is X' (resp. Y'). Thus, $Y = \mathrm{MT}(uv)$ for some word uv such that $X = \mathrm{MT}(u)$, $\gamma(u) = x'$, $\gamma(uv) = y'$, and $v \in L$. Hence, the word u (resp. uv) represents $x \in \mathrm{FIM}(\Gamma)/P$ (resp. $y \in \mathrm{FIM}(\Gamma)/P$) and there is a path from x to y in the Cayley-graph of $\mathrm{FIM}(\Gamma)/P$ that is labeled with the word $v \in L$. Now it is straight-forward to verify that $\mathcal{A} \models \varphi$ if and only if $\mathcal{C}(\Gamma) \models \widehat{\varphi}$. This concludes the proof of Theorem 7.2.

7.3 MSO-theory of the Cayley-graph of $\mathrm{FIM}(\Gamma)$

Theorem 7.8. *For every finite alphabet Γ with $|\Gamma| > 1$, the MSO-theory of the Cayley-graph of $\mathrm{FIM}(\Gamma)$ is undecidable.*

Proof. It suffices to prove the theorem for $\Gamma = \{a, b\}$. For the poof of the theorem we will detect an infinite grid as a minor of $\mathcal{C} = \mathcal{C}(\mathrm{FIM}(\{a, b\}), \{a, b, a^{-1}, b^{-1}\})$. Let $G = (V, E)$ be an undirected graph. For a relation $R \subseteq V \times V$ on the set V of nodes define G/R to be the graph, which results when identifying all nodes $u, v \in V$ with $(u, v) \in R$ and removing all resulting loops and multiple edges. A minor of G is a graph of the form H/R, where H is a subgraph of G. The infinite grid is the graph $(\mathbb{N} \times \mathbb{N}, \{((n, m), (n+1, m)) \mid n, m \in \mathbb{N}\} \cup \{((n, m), (n, m+1)) \mid n, m \in \mathbb{N}\}$. It is known that the MSO-theory of an undirected graph G is undecidable if the infinite grid is isomorphic to a minor of G [33].

Let G be the undirected graph that results from the Cayley-graph $\mathcal{C}$ by forgetting edge-labels and the direction of edges. Clearly if the MSO-theory of G is undecidable then also the MSO-theory of $\mathcal{C}$ is undecidable. Hence, by the above remarks it suffices to show that the infinite grid is a minor of G. The grid-point (n, m) will be represented by the element $a^n a^{-n} b^m b^{-m} = b^m b^{-m} a^n a^{-n} \in \mathrm{FIM}(\{a, b\})$. Let H be the subgraph of G that is induced by all nodes of the form $b^m b^{-m} a^n a^{-j}$ ($0 \leq j \leq n, m \geq 0$) and $a^n a^{-n} b^m b^{-j}$ ($0 \leq j \leq m, n \geq 0$). Let $R \subseteq \mathrm{FIM}(\{a, b\}) \times \mathrm{FIM}(\{a, b\})$ be the following relation on these nodes:

$$\begin{aligned} R \;=\; & \{(b^m b^{-m} a^n a^{-j}, b^m b^{-m} a^n a^{-n}) \mid j, n, m \in \mathbb{N}, 0 \leq j \leq n\} \cup \\ & \{(a^n a^{-n} b^m b^{-j}, a^n a^{-n} b^m b^{-m}) \mid j, n, m \in \mathbb{N}, 0 \leq j \leq m\} \end{aligned}$$

The following diagram shows a part of the graph H (with directions and labels of edges as in $\mathcal{C}$). All nodes in one shaded area are identified when forming the quotient H/R.

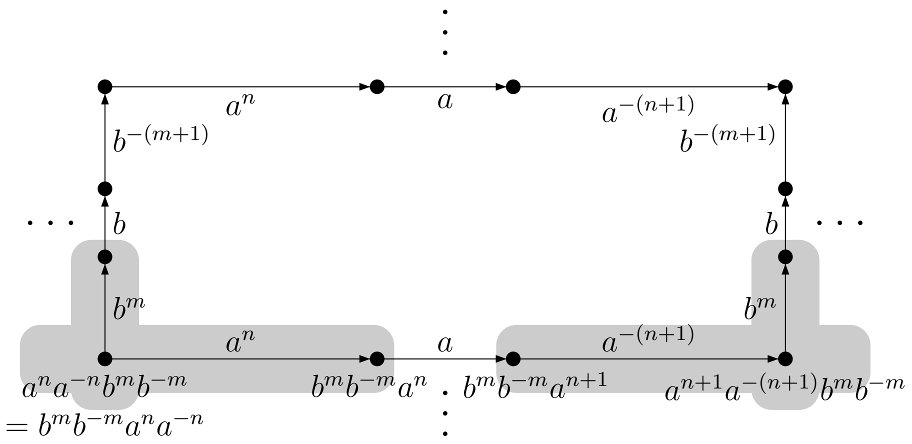

Note that for all natural numbers n, m we have

$$a^n a^{-n} b^m b^{-m} a^{n+1} a^{-(n+1)} = b^m b^{-m} a^n a^{-n} a^n a a^{-(n+1)} = b^m b^{-m} a^n a a^{-(n+1)} = a^{n+1} a^{-(n+1)} b^m b^{-m}.$$

Thus, there is a path from the node $a^n a^{-n} b^m b^{-m}$ of H to the node $a^{n+1} a^{-(n+1)} b^m b^{-m}$ of H that is labeled (in $\mathcal{C}$) with the word $a^{n+1} a^{-(n+1)}$. Similarly, there is a path from the node $a^n a^{-n} b^m b^{-m}$ to the node $a^n a^{-n} b^{m+1} b^{-(m+1)}$ labeled with $b^{m+1} b^{-(m+1)}$. Thus, in H/R, there are edges from the node $a^n a^{-n} b^m b^{-m}$ to both $a^{n+1} a^{-(n+1)} b^m b^{-m}$ and $a^n a^{-n} b^{m+1} b^{-(m+1)}$. Hence, these nodes define an infinite grid. $\square$

8 Open Problems

A promising research direction might be to investigate for which monoids $\mathcal{M}$ the structure $\mathcal{C}(\mathcal{M}, \Gamma)_{\text{reg}}$ has a decidable first-order theory. Here, in particular the group case is interesting. It is easy to see that the decidability of the MSO-theory of $\mathcal{C}(\mathcal{M}, \Gamma)$ implies the decidability of the first-order theory of $\mathcal{C}(\mathcal{M}, \Gamma)_{\text{reg}}$. The class of groups for which the first-order (resp. MSO-) theory of the Cayley-graph is decidable is precisely the class of groups with a decidable word problem (resp. the class of virtually free groups). Hence, the class of groups $\mathcal{G}$ for which $\mathcal{C}(\mathcal{G}, \Gamma)_{\text{reg}}$ is decidable lies somewhere between the virtually-free groups and the groups with a decidable word problem. Moreover, these inclusions are strict: By a reduction to Presburger's arithmetic it can be easily shown that for $\mathcal{G} = \mathbb{Z} \times \mathbb{Z}$ the first-order theory of $\mathcal{C}(\mathcal{G}, \Gamma)_{\text{reg}}$ is decidable, but since $\mathcal{C}(\mathcal{G}, \Gamma)$ is an infinite grid, $\mathrm{MSOTh}(\mathcal{C}(\mathcal{G}, \Gamma))$ is undecidable. Furthermore, there exists a hyperbolic group $\mathcal{G}$ [7], for which the generalized word problem is undecidable [30]. Thus, the first-order theory of $\mathcal{C}(\mathcal{G}, \Gamma)_{\text{reg}}$ is undecidable. On the other hand, every hyperbolic group has a decidable word problem [7].

References

[1] J.-C. Birget, S. W. Margolis, and J. Meakin. The word problem for inverse monoids presented by one idempotent relator. *Theoretical Computer Science*, 123(2):273–289, 1994.

[2] R. V. Book. Confluent and other types of Thue systems. *Journal of the Association for Computing Machinery*, 29(1):171–182, 1982.

[3] W. W. Boone. The word problem. *Annals of Mathematics (2)*, 70:207–265, 1959.

[4] A. K. Chandra, D. C. Kozen, and L. J. Stockmeyer. Alternation. *Journal of the Association for Computing Machinery*, 28(1):114–133, 1981.

[5] K. J. Compton and C. W. Henson. A uniform method for proving lower bounds on the computational complexity of logical theories. *Annals of Pure and Applied Logic*, 48:1–79, 1990.

[6] B. Courcelle. The expression of graph properties and graph transformations in monadic second-order logic. In G. Rozenberg, editor, *Handbook of graph grammars and computing by graph transformation, Volume 1 Foundations*, pages 313–400. World Scientific, 1997.

[7] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory*, number 8 in MSRI Publ., pages 75–263. Springer, 1987.

[8] W. Hodges. *Model Theory*. Cambridge University Press, 1993.

[9] A. V. Kelarev. On undirected Cayley graphs. *Australasian Journal of Combinatorics*, 25:73–78, 2002.

[10] A. V. Kelarev and C. E. Praeger. On transitive Cayley graphs of groups and semigroups. *European Journal of Combinatorics*, 24(1):59–72, 2003.

[11] A. V. Kelarev and S. J. Quinn. A combinatorial property and Cayley graphs of semigroups. *Semigroup Forum*, 66(1):89–96, 2003.

[12] O. Kupferman and M. Y. Vardi. An automata-theoretic approach to reasoning about infinite-state systems. In E. A. Emerson and A. P. Sistla, editors, *Proceedings of the 12th International Conference on Computer Aided Verification (CAV 2000), Chiacago (USA)*, number 1855 in Lecture Notes in Computer Science, pages 36–52. Springer, 2000.

[13] D. Kuske and M. Lohrey. Logical aspects of Cayley-graphs: the group case. *Annals of Pure and Applied Logic*, 131(1–3):263–286, 2005.

[14] D. Kuske and M. Lohrey. Logical aspects of Cayley-graphs: the monoid case. *International Journal of Algebra and Computation*, 2005. to appear.

[15] R. J. Lipton and Y. Zalcstein. Word problems solvable in logspace. *Journal of the Association for Computing Machinery*, 24(3):522–526, 1977.

[16] M. Lohrey. On the parallel complexity of tree automata. In A. Middeldorp, editor, *Proceedings of the 12th International Conference on Rewrite Techniques and Applications (RTA 2001), Utrecht (The Netherlands)*, number 2051 in Lecture Notes in Computer Science, pages 201–215. Springer, 2001.

[17] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer, 1977.

[18] S. Margolis and J. Meakin. Inverse monoids, trees, and context-free languages. *Trans. Amer. Math. Soc.*, 335(1):259–276, 1993.

[19] S. Margolis, J. Meakin, and M. Sapir. Algorithmic problems in groups, semigroups and inverse semigroups. In J. Fountain, editor, *Semigroups, Formal Languages and Groups*, pages 147–214. Kluwer, 1995.

[20] A. Markov. On the impossibility of certain algorithms in the theory of associative systems. *Doklady Akademii Nauk SSSR*, 55, 58:587–590, 353–356, 1947.

[21] A. R. Meyer. Weak monadic second order theory of one successor is not elementary recursive. In *Proceedings of the Logic Colloquium (Boston 1972–73)*, number 453 in Lecture Notes in Mathematics, pages 132–154. Springer, 1975.

[22] D. E. Muller and P. E. Schupp. Groups, the theory of ends, and context-free languages. *Journal of Computer and System Sciences*, 26:295–310, 1983.

[23] D. E. Muller and P. E. Schupp. The theory of ends, pushdown automata, and second-order logic. *Theoretical Computer Science*, 37(1):51–75, 1985.

[24] W. Munn. Free inverse semigroups. *Proc. London Math. Soc.*, 30:385–404, 1974.

[25] P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *American Mathematical Society, Translations, II. Series*, 9:1–122, 1958.

[26] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.

[27] M. Petrich. *Inverse semigroups*. Wiley, 1984.

[28] E. Post. Recursive unsolvability of a problem of Thue. *Journal of Symbolic Logic*, 12(1):1–11, 1947.

[29] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.

[30] E. Rips. Subgroups of small cancellation groups. *Bulletin of the London Mathematical Society*, 14:45–47, 1982.

[31] B. V. Rozenblat. Diophantine theories of free inverse semigroups. *Siberian Mathematical Journal*, 26:860–865, 1985. English translation.

[32] P. E. Schupp. Groups and graphs: Groups acting on trees, ends, and cancellation diagrams. *Mathematical Intelligencer*, 1:205–222, 1979.

[33] D. Seese. The structure of models of decidable monadic theories of graphs. *Annals of Pure and Applied Logic*, 53:169–195, 1991.

[34] P. V. Silva and B. Steinberg. Extensions and submonoids of automatic monoids. *Theoretical Computer Science*, 289:727–754, 2002.

[35] P. V. Silva and B. Steinberg. A geometric characterization of automatic monoids. *The Quarterly Journal of Mathematics*, 55:333–356, 2004.

[36] J. Stephen. Presentations of inverse monoids. *Journal of Pure and Applied Algebra*, 63:81–112, 1990.

[37] C. Stirling. *Modal and Temporal Properties of Processes*. Springer, 2001.

[38] I. Walukiewicz. Pushdown processes: Games and model-checking. *Information and Computation*, 164(2):234–263, 2001.

[39] B. Zelinka. Graphs of semigroups. *Casopis. Pest. Mat.*, 27:407–408, 1981.