Application of verification techniques to inverse monoids

Markus Lohrey*

Universität Leipzig, Institut für Informatik
Johannisgasse 26, 04103 Leipzig, Germany
lohrey@informatik.uni-leipzig.de

Abstract. 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, is investigated. 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, some results on free partially commutative inverse monoids are presented.

Keywords. inverse monoids, word problems, Cayley-graphs, complexity

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 and Post 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 and Boone extended the result of Markov and Post to finitely presented groups, see [1] for references.

In this paper, we are interested in a class of monoids that lies somewhere between groups and general monoids: inverse monoids [2]. In the same way as groups can be represented by sets of permutations, inverse monoids can be represented by sets of partial injections [2]. 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 [3]. In [4], 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

* This extended abstract is based on three papers, written with Volker Diekert, Nicole Ondrusch, and Alexander Miller. The work is supported by the DFG research project GELO.

Dagstuhl Seminar Proceedings 07441
Algorithmic-Logical Theory of Infinite Structures
http://drops.dagstuhl.de/opus/volltexte/2008/1410

identities between idempotents of $\mathrm{FIM}(\Gamma)$; we call such a presentation idempotent. In fact, in [4] 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 [5], 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, i.e., the running time is not bounded by an exponent tower of fixed height. Therefore, in [6,4] the question for a more efficient approach was asked. In this extended abstract we present solutions to this question based on techniques from the theory of automata and verification. In Section 6 we show by using tree automata 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 [7,8] stating that the model-checking problem of the modal μ-calculus over context-free graphs [9] is EXPTIME-complete.

In Section 7 we study 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 [1] and serve as a link to other fields like topology, graph theory, and automata theory, see, e.g., [10,9]. Here we consider Cayley-graphs from a logical point of view, see [11,12] 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 6). An immediate corollary of this result is that membership in rational subsets of $\mathrm{FIM}(\Gamma)/P$ is decidable.

Our decidability result for Cayley-graphs should be also compared with two undecidability results from the literature: (i) the monadic second-order theory of the Cayley-graph of $\mathrm{FIM}(\{a\})$ [13] as well as (ii) the existential theory of $\mathrm{FIM}(\{a,b\})$ (i.e. the set 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) [14] are undecidable.

In Section 8 we briefly consider free partially commutative inverse monoids and their quotients by idempotent presentation. A free partially commutative inverse monoid is the quotient of a free inverse monoids by a partial commutation relation. Hence, these monoids can be seen as analoges of free partially commutative monoids (also known as trace monoids, see e.g. [15]) and free partially commutative groups (also known as graph groups or right-angled Artin groups, see e.g. [16]). It turns out that for some restricted cases our algorithmic techniques for free inverse monoids modulo idempotent presentations can be generalzed to the partially commutative case.

This extended abstract is based on the papers [17,18,19].

2 Preliminaries

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 is denoted by $\mathrm{REG}(\Gamma)$. We assume that the reader has some basic background in complexity theory. An *alternating Turing-machine* [20] $T = (Q, \Sigma, \delta, q_0, q_f)$ is a nondeterministic Turing-machine (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), 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 final state q_f. A configuration C with current state q is accepting, if (i) $q = q_f$, or (ii) $q \in Q_\exists$ and there exists a successor configuration of C that is accepting, or (iii) $q \in Q_\forall$ 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) [20].

3 Relational structures and logic

See any text book on logic 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. 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 $\mathcal{A} \models \varphi$. 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 $\mathcal{A} \models \varphi$.

Several times, we will use implicitly the well-known fact that reachability in graphs can be expressed in MSO. More precisely, there exists an MSO-formula $\mathrm{reach}(x, y)$ (over the signature containing a binary relation symbol E) such that for every directed graph $G = (V, E)$ and all nodes $s, t \in V$ we have $G \models \mathrm{reach}(s, t)$ iff $(s, t) \in E^*$. Another important fact is that finiteness of a subset of a finitely-branching tree can be expressed in MSO: There is an MSO-formula $\mathrm{fin}(X)$ (over the signature containing a binary relation symbol E) such that for every finitely-branching (and downward-directed) tree $T = (V, E)$ and all subsets $U \subseteq V$ we have $T \models \mathrm{fin}(U)$ iff U is finite (by König's lemma, U is infinite iff the upward-closure of U contains an infinite path), see also [5, Lemma 1.8].

In Section 6 we will make use of the *modal μ-calculus*, which is a popular logic for the verification of reactive systems. 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 (we only introduce those operators that are needed later; other operators like $\neg\varphi$ or $[a]\varphi$ are defined as usual): $\varphi ::= \text{true} \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid \mu X.\varphi$. Here $X \in \mathbb{V}_2$ is a second-order variable ranging over sets of nodes and $a \in \Sigma$. Variables from $\mathbb{V}_2$ are bounded by the μ-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:

- $\text{true}^G(\sigma) = V$, $X^G(\sigma) = \sigma(X)$ for every $X \in \mathbb{V}_2$,
- $(\varphi \vee \psi)^G(\sigma) = \varphi^G(\sigma) \cup \psi^G(\sigma)$, $(\varphi \wedge \psi)^G(\sigma) = \varphi^G(\sigma) \cap \psi^G(\sigma)$,
- $(\langle a \rangle \varphi)^G(\sigma) = \{u \in V \mid \exists v \in V : (u, v) \in E_a \ \wedge \ v \in \varphi^G(\sigma)\}$,
- $(\mu X.\varphi)^G(\sigma)$ is the smallest fixpoint of the monotonic function $U \mapsto \varphi^G(\sigma[U/X])$

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 μ-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 $v \in \varphi^G$.

A *context-free graph* [9] 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 1 ([7,8]). *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: $(G(A), v) \models \varphi$?

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 computational problem that asks for two given words $u, v \in \Sigma^*$, whether $h(u) = h(v)$. It is well-known that if Σ_1 and Σ_2 are 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. Thus, the computational complexity of the word problem does not depend on the underlying set of generators.

The *Cayley-graph* of the monoid $\mathcal{M}$ w.r.t. the generating set Σ is the 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 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 Σ 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 [1].

The *free group* $\mathrm{FG}(\Gamma)$ generated by the set Γ is the quotient $(\Gamma \cup \Gamma^{-1})^*/\delta$, 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. It holds $\gamma(u) = \gamma(v)$ iff $r(u) = r(v)$. 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 [9].

Similarly to the word problem, if Σ_1 and Σ_2 are finite generating sets for the same monoid $\mathcal{M}$, then $\mathrm{FOTh}(\mathcal{C}(\mathcal{M}, \Sigma_1))$ is logspace reducible to $\mathrm{FOTh}(\mathcal{C}(\mathcal{M}, \Sigma_2))$ and the same holds for the MSO-theories, see [12]. 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 [12]. When restricting to groups, the situation is different: The Cayley-graph of a finitely generated group has a decidable first-order theory iff the group has a decidable word problem [11]. Moreover, the Cayley-graph of a finitely generated group has a decidable MSO-theory iff the group is virtually free (i.e., has a free subgroup of finite index) [11,9]. We will only need this result for the Cayley-graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$:

Theorem 2 ([9]). *For every finite Γ, $\mathrm{MSOTh}(\mathcal{C}(\Gamma))$ is decidable but 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 [2]; 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 $\gamma = \beta \circ \alpha$. 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 $\mathrm{MT}(u) = \{\gamma(v) \in \mathrm{FG}(\Gamma) \mid \exists w \in (\Gamma \cup \Gamma^{-1})^* : u = vw\}$; it is a finite and connected subset of the Cayley-graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$. 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 [21] states that $\alpha(u) = \alpha(v)$ for $u, v \in (\Gamma \cup \Gamma^{-1})^*$ iff $r(u) = r(v)$ (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., $\alpha(uu) = \alpha(u)$, iff $r(u) = \varepsilon$, 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 $\nu_P : \mathrm{FIM}(\Gamma) \to \mathrm{FIM}(\Gamma)/P$. For the rest of the paper, the meaning of the morphisms α, γ, μ_P, and ν_P will be fixed. 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. To solve the word problem for such a monoid, Margolis and Meakin [4] constructed a closure operation for Munn trees. We shortly review the ideas here. As remarked in [4], every idempotent presentation P can be replaced by the 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

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

where $\circ$ refers to the multiplication in the free group $\mathrm{FG}(\Gamma)$. Finally, define the closure of V w.r.t. the presentation P as $\mathrm{cl}_P(V) = \bigcup_{n\geq 1} V_n$.

Theorem 3 ([4]). *Let P be an idempotent presentation and let $u, v \in (\Gamma \cup \Gamma^{-1})^*$. Then $\mu_P(u) = \mu_P(v)$ iff $r(u) = r(v)$ (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.

Example 1. Let $\Gamma = \{a, b\}$, $u = aa^{-1}bb^{-1}$, and $P = \{(aa^{-1}, a^2a^{-2}), (bb^{-1}, b^2b^{-2})\}$. The Munn trees for the words in the presentation P and u are shown below; the bigger circle represents the 1 of $\mathrm{FG}(\Gamma)$.

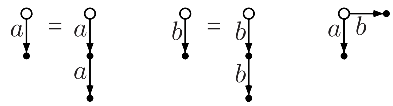

Then $\mathrm{cl}_P(\mathrm{MT}(u)) = \{a^n \mid n \geq 0\} \cup \{b^n \mid n \geq 0\} \subseteq \mathrm{FG}(\Gamma)$.

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., $r(w) = \varepsilon$) explicitly, we will only show its Munn tree, where as above 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; the resulting word represents the same element of $\mathrm{FIM}(\Gamma)$ as the original one.

Margolis and Meakin used Theorem 3 in order to decide the word problem for $\mathrm{FIM}(\Gamma)/P$. More precisely, they have shown that from a finite and 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 $A \subseteq \mathrm{FG}(\Gamma)$: $\mathcal{C}(\Gamma) \models \mathrm{CL}_P(\mathrm{MT}(u), A)$ iff $A = \mathrm{cl}_P(\mathrm{MT}(u))$. The decidability of the word problem for $\mathrm{FIM}(\Gamma)/P$ is an immediate consequence of Theorem 2 and Theorem 3. But the application of Theorem 2 results in a nonelementary algorithm.

6 Complexity of the word problem

Using the efficient translation of MSO-formulas over trees into tree automata, and the fact that the membership for a fixed tree automaton can be checked in (i) linear time on a RAM and (ii) logspace on a Turing machine [22], we can prove:

Theorem 4. *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 (i) linear time on a RAM and (ii) logspace on a Turing machine.*

An alternative proof of Theorem 4, which does not rely on the translation of MSO into tree automata, was given in [18]. Moreover, that proof works for a larger class of inverse monoids, where Munn trees over free groups are replaced by subgraphs of the Cayley-graph of a virtually-free group.

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

Theorem 5. *The following problem is EXPTIME-complete:*

INPUT: A finite alphabet Γ, 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: $\mu_P(u) = \mu_P(v)$?

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 assume that:

– 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 $Q_\exists \cup \{q_f\}$.
– The initial state q_0 belongs to $Q_\exists$.
– For each pair $(q, a) \in (Q \setminus \{q_f\}) \times \Sigma$, the machine T has precisely two choices according to 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: (i) $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*(Q_\exists \times \Sigma)\Sigma^*\{\varepsilon, \#\}$ and $c_1c_2c_3 \xrightarrow{j} d_j$ for $j \in \{1, 2\}$ or (ii) $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*\{\varepsilon, \#\}$ 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 $c_1c_2c_3 \in \{\varepsilon, \#\}\Sigma^*(Q_\forall \times \Sigma)\Sigma^*\{\varepsilon, \#\}$.

We encode a configuration $\#c_1c_2 \cdots c_m\#$, where the current state is from $Q_\exists$ by the following subtree of $\mathcal{C}(\Gamma)$, where $i = 1$ or $i = 2$.

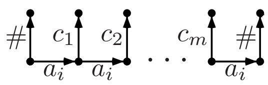

If the current state is from $Q_\forall$, then we take the same subgraph, except that b_i replaces a_i.

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 identity 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).

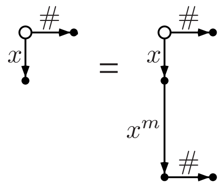

Successor configurations of the current configuration are generated by the equations below, where $i \in \{1,2\}$, $0 \le k \le m-1$, and $c_1c_2c_3 \overset{\exists}{\to} (d_1,d_2)$ (resp. $c_1c_2c_3 \overset{\forall}{\to} (d_1,d_2)$) for the left (resp. right) equation:

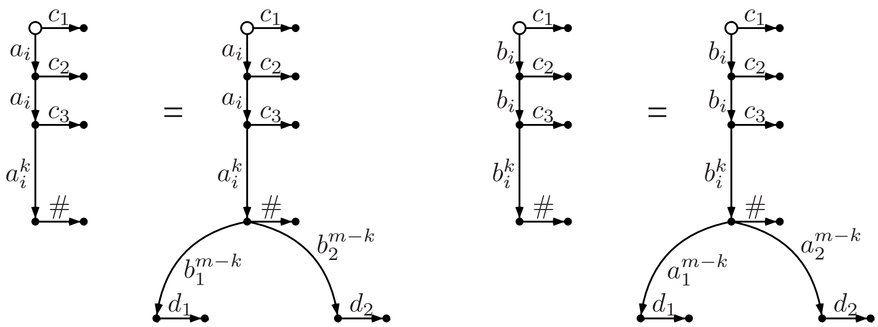

The remaining equations propagate acceptance information back to the initial Munn tree. Here the separation of the state set into existential and universal states becomes important. Let $f = (q_f, \$)$; recall that \$ is the symbol under the head of T when T terminates in state q_f. For all $i, j \in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

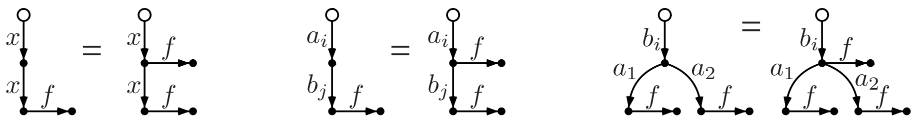

This concludes the description of the presentation P. Now define the words $u, v \in (\Gamma \cup \Gamma^{-1})^*$ as follows: Assume that the input word for our alternating Turing machine w is of the form $w = w_1 w_2 \cdots w_n$ with $w_i \in \Sigma$. For $n+1 \le i \le m$ define $w_i = \square$, where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

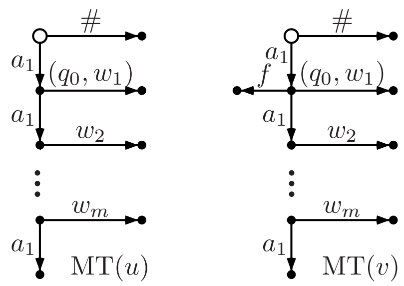

We claim that $\mu_P(u) = \mu_P(v)$ iff the machine T accepts the word w. From the construction of u, v, and P it follows that T accepts the word w iff $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$. Since $\mathrm{MT}(u) \subseteq \mathrm{MT}(v)$ this is equivalent to $\mathrm{cl}_P(\mathrm{MT}(v)) = \mathrm{cl}_P(\mathrm{MT}(u))$, i.e., $\mu_P(u) = \mu_P(v)$ due to Theorem 3 (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 $u, v \in (\Gamma \cup \Gamma^{-1})^*$. Since $r(u) = r(v)$ can be checked in linear time, it suffices by Theorem 3 to show that we can verify in EXPTIME whether $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$ (note that $\mathrm{cl}_P(\mathrm{MT}(v)) = \mathrm{cl}_P(\mathrm{MT}(u))$ iff $\mathrm{MT}(u) \subseteq \mathrm{cl}_P(\mathrm{MT}(v))$ and $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$). Let G be the graph that results from the Cayley-graph $\mathcal{C}(\Gamma)$ by taking a new edge label $\#$, 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. Since $\mathcal{C}(\Gamma)$ is a context-free graph, also G is context-free. By Theorem 1 it suffices to construct in polynomial time a modal μ-calculus formula $\varphi_{u,v,P}$ such that $\mathrm{MT}(v) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$ iff $(G, 1) \models \varphi_{u,v,P}$.

For $w = a_1 a_2 \cdots a_m$ ($a_i \in \Gamma \cup \Gamma^{-1}$) and two positions $i, j \in \{1, \ldots, m\}$, $i \leq j$, let $w[i,j] = a_i \cdots a_j$. If $i > j$, then set $w[i,j] = \varepsilon$. Moreover, we use $\langle w \rangle \phi$ as an abbreviation for $\langle a_1 \rangle \langle a_2 \rangle \cdots \langle a_m \rangle \phi$. 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

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

Then $(G, x) \models \varphi_{u,P}$ iff $x \in \mathrm{cl}_P(\mathrm{MT}(u))$. The disjunction $\bigvee_{i=0}^{|u|} \langle u[1,i]^{-1} \rangle \langle \# \rangle \mathrm{true}$ expresses $\mathrm{MT}(u) \subseteq \mathrm{cl}_P(\mathrm{MT}(u))$, whereas $\bigvee_{i=1}^{n} \bigvee_{j=0}^{|f_i|} \langle f_i[1,j]^{-1} \rangle (\bigwedge_{k=0}^{|e_i|} \langle e_i[1,k] \rangle X)$ defines all nodes such that via the inverse of 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. Finally, set $\varphi_{u,v,P} = \bigwedge_{i=0}^{|v|} \langle v[1,i] \rangle \varphi_{u,P}$. □

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 $\mathcal{C}(\mathcal{M}, \Sigma)$: $\mathcal{C}(\mathcal{M}, \Sigma)_{\mathrm{reg}} = (M, (\mathrm{reach}_L)_{L \in \mathrm{REG}(\Sigma)}, 1)$ with $\mathrm{reach}_L = \{(u, v) \in M \times M \mid \exists w \in L : u \circ h(w) = v\}$. Thus, $\mathcal{C}(\mathcal{M}, \Sigma) = (M, (\mathrm{reach}_{\{a\}})_{a \in \Sigma}, 1)$. The main result of this section is:

Theorem 6. *Let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be a finite and idempotent presentation. Then the first-order theory of the structure $\mathcal{C}(\mathrm{FIM}(\Gamma)/P, \Gamma \cup \Gamma^{-1})_{\mathrm{reg}}$ is decidable.*

The following undecidability result of Calbrix [13] contrasts Theorem 6. It is easy to see that for every Cayley-graph $\mathcal{C}(\mathcal{M}, \Gamma)$, if $\mathrm{MSOTh}(\mathcal{C}(\mathcal{M}, \Gamma))$ is decidable, then also $\mathrm{FOTh}(\mathcal{C}(\mathcal{M}, \Gamma)_{\mathrm{reg}})$ is decidable.

Theorem 7. $\mathrm{MSOTh}(\mathcal{C}(\mathrm{FIM}(\{a\}), \{a, a^{-1}\}))$ *is undecidable.*

Theorem 7 can be shown by identifying an infinite grid as a minor of the Cayley-graph $\mathcal{C}(\mathrm{FIM}(\{a\}), \{a, a^{-1}\})$.

Before we prove Theorem 6, let us first state a corollary. The *generalized word problem* for $\mathcal{M}$ asks whether for given words $u, u_1, \ldots, u_n \in \Sigma^*$ the monoid element $h(u)$ belongs to the submonoid of $\mathcal{M}$ that is generated by the elements $h(u_1), \ldots, h(u_n)$. Theorem 6 easily implies:

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

To prove Theorem 6 we first need some lemmas.

Lemma 1. *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: $G \models \varphi(s, t)$ iff there is a path in G with initial vertex s and terminal vertex t visiting all vertices from V.*

For the proof of Lemma 1 one defines a partial order $\prec$ on the set of strongly connected components of G: $U \prec V$ for two different strongly connected components U and V if and only if there is a (directed) path from a node of U to a node of V. Then there is a path in G with initial vertex s and terminal vertex t visiting all vertices from V iff $\prec$ is a total order and s (resp. t) belongs to the minimal (resp. maximal) strongly connected component of G. These conditions can be easily formalized in MSO-logic.

Lemma 2. *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 $G \models \psi_L$ iff 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 $1 \le i < n$, $a_1 a_2 \cdots a_{n-1} \in L$, and $V = \{v_1, v_2, \ldots, v_n\}$.*

Let us just give a brief sketch of the proof of Lemma 2. Let $A = (Q, \Sigma, \delta, q_0, F)$ be a deterministic finite automaton with $L(A) = L$. W.l.o.g. $Q = \{1, \ldots, m\}$. Define the structure $f_A(G)$ by $f_A(G) = (V \times Q, E, \Delta, I_s, F_t)$, where

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

Then one can show that f_A is an MSO-transduction in the sense of [23]. Thus, there exists a backwards translation $f_A^\sharp$ such that for every MSO-sentence ϕ over the signature of $f_A(G)$ we have: $f_A(G) \models \phi$ iff $G \models f_L^\sharp(\phi)$ [23]. Now, using Lemma 1 and the relation Δ it is easy to write down an MSO-sentence ϕ over the signature of $f_A(G)$ expressing that there exists a path from (s, q_0) to a node in F_t such that the set of first components of nodes along that path is precisely V. Then the sentence $f_A^\sharp(\phi)$ is the desired sentence.

Lemma 2 easily implies the next lemma.

Lemma 3. *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 $G \models \theta_L(U)$ iff 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 $1 \le i < n$, $a_1 a_2 \cdots a_{n-1} \in L$, and $U \subseteq \{v_1, v_2, \ldots, v_n\}$.*

We are now able to finish the proof of Theorem 6. Let $P \subseteq (\Gamma \cup \Gamma^{-1})^* \times (\Gamma \cup \Gamma^{-1})^*$ be a finite and 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 3 and translate each first-order sentence φ over $\mathcal{A}$ into an MSO-sentence $\|\varphi\|$ over the Cayley graph $\mathcal{C}(\Gamma)$ of the free group $\mathrm{FG}(\Gamma)$ such that $\mathcal{A} \models \varphi$ iff $\mathcal{C}(\Gamma) \models \|\varphi\|$. Together with Theorem 2 this will complete the proof of Theorem 6.

To every variable x (ranging over $\mathrm{FIM}(\Gamma)/P$) in φ we associate two variables in $\|\varphi\|$: (i) an MSO-variable X' representing $\mathrm{cl}_P(\mathrm{MT}(u))$, where u is any word representing x, and (ii) an FO-variable x' representing $\gamma(u) \in \mathrm{FG}(\Gamma)$. 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 $\Theta(x', X, X') = (1, x' \in X \ \wedge\ X$ is connected and finite $\wedge$ $\mathrm{CL}_P(X, X'))$. Recall that finiteness and connectedness of a subset of the finitely-branching tree $\mathcal{C}(\Gamma)$ can be expressed in MSO, see the remarks in Section 3. Here $\mathrm{CL}_P(X, X')$ is the MSO-formula constructed by Margolis and Meakin in [4], see the remark at the end of Section 5. Next, note that by Lemma 3 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: $\mathcal{C}(\Gamma) \models \xi_L(u', U, v', V)$ iff $U \subseteq V$ and there is a path from u' to v' in $\mathcal{C}(\Gamma)\restriction_V$ 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 $\|\varphi\|$ inductively: $\|\mathrm{reach}_L(x, y)\| = \exists X, Y : \Theta(x', X, X') \ \wedge\ \Theta(y', Y, Y') \ \wedge\ \xi_L(x', X, y', Y)$, $\|\neg\psi\| = \neg\|\psi\|$, $\|\psi_1 \wedge \psi_2\| = \|\psi_1\| \wedge \|\psi_2\|$, and $\|\forall x : \psi\| = \forall x' \,\forall X' : \mathrm{MT}(x', X') \Rightarrow \|\psi\|$. It is straight-forward to verify that $\mathcal{A} \models \varphi$ iff $\mathcal{C}(\Gamma) \models \|\varphi\|$. This concludes the proof of Theorem 6. $\square$

8 Partially commutative inverse monoids

In this section we briefly consider free partially commutative inverse monoids and their quotients by idempotent presentations. These monoids generalize the inverse monoids that we have considered so far.

Let us fix again a finite alphabet Γ, and let $\Gamma^{-1} = \{a^{-1} \mid a \in \Gamma\}$ be a disjoint copy of Γ. By an *independence relation* we mean here an irreflexive and symmetric binary relation I on $\Gamma \cup \Gamma^{-1}$ such that $(a, b) \in I$ implies $(a^{-1}, b) \in I$ for all $a, b \in \Gamma$. Clearly, I is uniquely defined by its restriction $I \cap (\Gamma \times \Gamma)$. Let us define the *free partially commutative inverse monoid* generated by (Γ, I) as the quotient

$$\mathrm{FIM}(\Gamma, I) = \mathrm{FIM}(\Gamma)/\{ab = ba \mid (a, b) \in I\}.$$

Da Costa has studied $\mathrm{FIM}(\Gamma, I)$ in his thesis from a more general viewpoint of graph products [24]. As a consequence he showed that $\mathrm{FIM}(\Gamma, I)$ has a decidable word problem. In his construction he used the general approach via Schützenberger graphs and Stephen's iterative procedure [25]. The decidability of the word problem follows because da Costa can show that Stephen's procedure terminates. However, no complexity bounds are given in [24].

In [17], a concrete realization of free partially commutative inverse monoid, which is based on closed subsets of free partially commutative groups (which are also known as graph groups [16]), is presented. Using this realization, the following results are proved:

Theorem 8. *For a fixed free partially commutative inverse monoid* $\mathrm{FIM}(\Gamma, I)$*, the word problem can be solved in time* $O(n \log(n))$ *on a RAM.*

Theorem 9. *For every fixed free partially commutative inverse monoid* $\mathrm{FIM}(\Sigma, I)$*, the membership problem for rational subsets of* $\mathrm{FIM}(\Sigma, I)$ *belongs to NP. Moreover, already membership in a finitly generated submonoid of* $\mathrm{FIM}(\{a, b\})$ *is NP-hard.*

An *idempotent presentation* over (Γ, I) is a finite set of identities $P = \{(e_i, f_i) \mid 1 \leq i \leq n\}$, where every e_i and f_i is an idempotent element in $\mathrm{FIM}(\Gamma, I)$.

In case the complement of the independence relation I is transitive, one obtains the same results for the word problem of $\mathrm{FIM}(\Gamma, I)$ modulo an idempotent presentation as in the non-commutative case. The proof of the EXPTIME upper bound in the following theorem is again based on a closure operation. But this time, the closure is not computed in the free group $\mathrm{FG}(\Gamma)$ but in the graph group $\mathrm{GG}(\Gamma, I)$. In case $(\Gamma \times \Gamma) \setminus I$ is transitive, this group is a direct product of free groups, which allows to express the closure w.r.t. an idempotent presentation P in the modal μ-calculus with *simultaneous fixpoint definitions* over the Cayley-graph of a free group (of suitable rank).

Theorem 10. *The following problem is EXPTIME-complete:*

INPUT: An independence relation I *over* $\Gamma \cup \Gamma^{-1}$ *such that* $(\Gamma \times \Gamma) \setminus I$ *is transitive, an idempotent presentation* P *over* (Γ, I) *and words* $u, v \in (\Gamma \cup \Gamma^{-1})^*$.

QUESTION: $u = v$ *in* $\mathrm{FIM}(\Gamma, I)/P$*?*

Theorem 11. *If* I *is an independence relation over* $\Gamma \cup \Gamma^{-1}$ *with* $(\Gamma \times \Gamma) \setminus I$ *transitive and* P *is an idempotent presentation over* (Γ, I)*, then the word problem for* $\mathrm{FIM}(\Gamma, I)/P$ *can be solved in (i) linear time on a RAM and (ii) logspace on a Turing machine.*

In the non-transitive case, we can encode the acceptance problem for a Turing-machine T in the word problem for $\mathrm{FIM}(\Gamma, I)/P$. Let $\Gamma = \{a, b, c\}$ and assume that $(a, b) \in I$ but $(a, c), (b, c) \notin I$. Then a and b generate in the Cayley graph of the graph group $\mathrm{GG}(\{a, b, c\}, I)$ a two dimensional grid. Using the letter c, which is dependent from both a and b, we can encode a labelling of the grid-points with tape symbols and states of T. By computing the closure w.r.t. a suitable idempotent presentation P, we generate a labelling consistent with the transition function of T and the input of T. Hence, we have:

Theorem 12. *Let* I *be an independence relation over* $\Gamma \cup \Gamma^{-1}$ *with* $(\Gamma \times \Gamma) \setminus I$ not *transitive. Then there exists an idempotent presentation* P *over* (Γ, I) *such that the word problem for* $\mathrm{FIM}(\Gamma, I)/P$ *is undecidable.*

9 Open problems

We plan to investigate for which monoids $\mathcal{M}$ the structure $\mathcal{C}(\mathcal{M},\Gamma)_{\mathrm{reg}}$ has a decidable first-order theory. 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 $\mathcal{C}(\mathcal{M},\Gamma)_{\mathrm{reg}}$. Thus, the class of groups $\mathcal{G}$ for which $\mathcal{C}(\mathcal{G},\Gamma)_{\mathrm{reg}}$ is decidable lies somewhere between the virtually-free groups (i.e., those groups for which the MSO-theory of the Cayley-graph is decidable) and the groups with a decidable word problem (i.e., those groups for which the first-order theory of the Cayley-graph is decidable).

References

1. Lyndon, R.C., Schupp, P.E.: Combinatorial Group Theory. Springer (1977)
2. Petrich, M.: Inverse semigroups. Wiley (1984)
3. Margolis, S., Meakin, J., Sapir, M.: Algorithmic problems in groups, semigroups and inverse semigroups. In Fountain, J., ed.: Semigroups, Formal Languages and Groups, Kluwer (1995) 147–214
4. Margolis, S., Meakin, J.: Inverse monoids, trees, and context-free languages. Transactions of the American Mathematical Society **335** (1993) 259–276
5. Rabin, M.O.: Decidability of second-order theories and automata on infinite trees. Transactions of the American Mathematical Society **141** (1969) 1–35
6. Birget, J.C., Margolis, S.W., Meakin, J.: The word problem for inverse monoids presented by one idempotent relator. Theoretical Computer Science **123(2)** (1994) 273–289
7. Kupferman, O., Vardi, M.Y.: An automata-theoretic approach to reasoning about infinite-state systems. In Emerson, E.A., Sistla, A.P., eds.: Proceedings of the 12th International Conference on Computer Aided Verification (CAV 2000), Chiacago (USA). Number 1855 in Lecture Notes in Computer Science, Springer (2000) 36–52
8. Walukiewicz, I.: Pushdown Processes: Games and Model-Checking. Information and Computation **164** (2001) 234–263
9. Muller, D.E., Schupp, P.E.: The theory of ends, pushdown automata, and second-order logic. Theoretical Computer Science **37** (1985) 51–75
10. Muller, D.E., Schupp, P.E.: Groups, the theory of ends, and context-free languages. Journal of Computer and System Sciences **26** (1983) 295–310
11. Kuske, D., Lohrey, M.: Logical aspects of Cayley-graphs: the group case. Annals of Pure and Applied Logic **131** (2005) 263–286
12. Kuske, D., Lohrey, M.: Logical aspects of Cayley-graphs: the monoid case. International Journal of Algebra and Computation **16** (2006) 307–340
13. Calbrix, H.: La théorie monadique du second ordre du monoïde inversif libre est indécidable (The second-order monadic theory of the free inverse monoid is undecidable). Bulletin of the Belgian Mathematical Society **4** (1997) 53–65
14. Rozenblat, B.V.: Diophantine theories of free inverse semigroups. Siberian Mathematical Journal **26** (1985) 860–865 English translation.
15. Diekert, V., Rozenberg, G., eds.: The Book of Traces. World Scientific (1995)
16. Droms, C.: Graph groups, coherence and three-manifolds. Journal of Algebra **106** (1985) 484–489
17. Diekert, V., Lohrey, M., Miller, A.: Partially commutative inverse monoids. In Kralovic, R., Urzyczyn, P., eds.: Proceedings of the 31th International Symposium on Mathematical Foundations of Computer Science (MFCS 2006), Bratislave (Slovakia). Number 4162 in Lecture Notes in Computer Science, Springer (2006) 292–304 long version in preparation.

18. Diekert, V., Lohrey, M., Ondrusch, N.: Algorithmic problems on inverse monoids over virtually-free groups. International Journal of Algebra and Computation (2008) to appear.
19. Lohrey, M., Ondrusch, N.: Inverse monoids: decidability and complexity of algebraic questions. Information and Computation **205** (2007) 1212–1234
20. Chandra, A.K., Kozen, D.C., Stockmeyer, L.J.: Alternation. Journal of the Association for Computing Machinery **28** (1981) 114–133
21. Munn, W.: Free inverse semigroups. Proceedings of the London Mathematical Society, 3rd Series **30** (1974) 385–404
22. Lohrey, M.: On the parallel complexity of tree automata. In Middeldorp, A., ed.: Proceedings of the 12th International Conference on Rewrite Techniques and Applications (RTA 2001), Utrecht (The Netherlands). Number 2051 in Lecture Notes in Computer Science, Springer (2001) 201–215
23. Courcelle, B.: The expression of graph properties and graph transformations in monadic second-order logic. In Rozenberg, G., ed.: Handbook of graph grammars and computing by graph transformation, Volume 1 Foundations. World Scientific (1997) 313–400
24. Veloso da Costa, A.: Γ-Productos de Monóides e Semigrupos. PhD thesis, Universidade do Porto (2003)
25. Stephen, J.: Presentations of inverse monoids. Journal of Pure and Applied Algebra **63** (1990) 81–112