REACHABILITY PROBLEM FOR WEAK MULTI-PUSHDOWN AUTOMATA

WOJCIECH CZERWIŃSKI [a], PIOTR HOFMAN [b], AND SŁAWOMIR LASOTA [c]

[a] Universität Bayreuth
e-mail address: wczerwin@mimuw.edu.pl

[b,c] Institute of Informatics, University of Warsaw
e-mail address: {ph209519,sl}@mimuw.edu.pl

ABSTRACT. This paper is about reachability analysis in a restricted subclass of multi-pushdown automata. We assume that the control states of an automaton are partially ordered, and all transitions of an automaton go downwards with respect to the order. We prove decidability of the reachability problem, and computability of the backward reachability set. As the main contribution, we identify relevant subclasses where the reachability problem becomes NP-complete. This matches the complexity of the same problem for communication-free vector addition systems, a special case of stateless multi-pushdown automata.

1. INTRODUCTION

This paper is about reachability analysis of *multi-pushdown systems*, i.e., systems with a global control state and multiple stacks. The motivation for our work is twofold. On one side, a practical motivation coming from context-bounded analysis of recursive concurrent programs [28, 25, 3]. On the other side, a theoretical motivation coming from partially-commutative context-free grammars, developed recently in [14, 15, 17].

Context bounded analysis. Multi-pushdown systems may be used as an abstract model of concurrent programs with recursive procedures. As multi-pushdown systems are a Turing-complete model of computation, they are only applicable for verification under further tractable restrictions. One remarkably successful restriction is imposing a bound on the number of context switches; between consecutive context switches, the system may only perform operations on one stack (local operations). In [28], the context-bounded reachability has been shown decidable, by reduction to reachability of ordinary pushdown systems [7]. This line of research, with applications in formal verification, has been continued successfully, e.g., in [8, 25, 3].

Weak control states. As our starting point we observe that if the number of context switches is bounded, one may safely assume that the control state space is *weak*, in the sense that there is a partial order on control states such that transitions go only downwards with respect to the order. Indeed, the local state space of every thread may be eliminated using a stack, and the global control state essentially enumerates context switches. Roughly speaking, the model investigated in this paper extends the above one with respect to operations allowed between two context switches, namely, we do not restrict these operations to one stack only. Thus, if k is the number of stacks, we assume that transitions of a system are of the following form:

$$q, \ \mathrm{X} \ \xrightarrow{a} \ q', \ \alpha_1, \ldots, \alpha_k, \tag{1.1}$$

to mean that in control state q, symbol X is popped from one of the stacks, and sequences of symbols $\alpha_1, \ldots, \alpha_k$, respectively, are pushed on stacks. Wlog. one may assume that the symbols of different stacks are different.

Partially-commutative context-free grammars. A special case of the model investigated in this paper is *stateless* multi-pushdown systems. This is still a quite expressible model that subsumes, among the others, context-free graphs (so called Basic Process Algebra [10]) and communication-free Petri nets (so called Basic Parallel Processes [10], or commutative context-free graphs). In the stateless case, transitions (1.1) may be understood as productions of a grammar, with the nonterminal symbols on the right-hand side (stack symbols) subject to a commutativity law. More precisely, for any two symbols X and Y from different stacks, we impose the commutativity law

$$\mathrm{XY} \ = \ \mathrm{YX}.$$

One easily observes that this is a special case of *independence relation* over nonterminal symbols, as defined in trace theory [18][1]. In multi-pushdown systems, the *dependency relation* (the complement of independence relation) is always transitive. A general theory of context-free grammars modulo dependency relation that is not necessarily transitive, has been studied recently in [17]; complexity of bisimulation equivalence checking has been investigated in [14, 15]. The present paper complements these results by focusing on reachability analysis.

---

[1]Note however that the independence is imposed on nonterminal symbols, and not on input letters, as usually in trace theory.

Contributions. We investigate the reachability problem for weak multi-pushdown automata that asks, given two sets of configurations, namely the source and target set, whether there is a sequence of transitions leading from some source configuration to some target one. For finite representation of the source and target set, we restrict to regular sets of configurations only. Roughly speaking, in this paper we present an almost complete map of decidability and complexity results, with respect to natural restrictions of the source and target sets, and of the control states of multi-pushdown automata.

The paper contains two main results and a few accompanying ones. As our first main result, we prove decidability of reachability for weak multi-pushdown automata, in case when target set is a singleton. Our argument is based on a suitable well order on the set of configurations, that strongly depends on the assumption that the control states are weak.

As the second main result, we identify additional restrictions under which the problem is NP-complete. One such restriction is stateless multi-pushdown systems; another restriction is that in every control state, the stacks may be emptied without changing the state. Our result subsumes (and gives a simpler nondeterministic polynomial time algorithm for) the case of communication-free Petri nets, assuming that the source and target markings of a net are given in unary. Reachability of communication-free Petri nets is NP-complete as shown in [19], even when the source and target markings are presented in binary. NP-completeness of reachability for stateless multi-pushdown automata is similar to NP-completeness of the word problem for partially-commutative context-free grammars [21], where one asks if the given input word is accepted. The reachability question is more difficult to answer, as no input word is given in advance. In fact the main technical difficulty is to show the existence of a polynomial witness for reachability.

As further results, we investigate forward and backward reachability sets, and prove that the backward reachability set of a regular set of configurations is regular and computable, while the forward reachability set needs not be regular in general. We also identify the decidability border for reachability of weak multi-pushdown systems. Roughly speaking, the problem becomes undecidable when the target set of configurations is an unrestricted regular set, but is decidable if the target set is a singleton.

Finally, we prove a slightly surprising fact: reachability is decidable for multi-pushdown automata, with unrestricted control states, under the assumption that in every state, the stacks may be emptied without changing state.

The standard techniques useful for analysis of pushdown systems, such as pumping or the automaton-based approach of [7], do not extend to the multi-pushdown setting. This is why the proofs of our results are based on new insights. The NP-membership proofs are, roughly speaking, based on polynomial witnesses obtained by careful elimination of 'irrelevant' transitions. On the other hand, the decidability results are based on a suitable well order on configurations.

Related research. Multi-pushdown systems are a fundamental model of recursive multi-threaded programs. This is why different instantiations of the multi-pushdown paradigm have been appearing in the literature recently, most often in the context of formal verification. We only mention here a few relevant positions we are aware of, without claiming completeness. All the papers cited below bring some restricted decidability results for reachability or model checking.

Most often, a model has global control states, with stack operations subject to some restriction. For instance, the author of [1] assumes that the stacks are ordered, and pop

operation can only be performed on the first nonempty stack. Another example is the model introduced in [9] and then further investigated e.g. in [8, 2, 6], that allows for unbounded creation of new stacks; on the other hand, operations on each stack are local, thus no communication between threads is allowed.

Another possible approach is to replace global state space with some communication mechanism between threads. Some successful results on analysis of multi-threaded programs communicating via locks, in a restricted way, has been reported in [23, 22, 11].

In [26] the algorithm for reachability over PA [10] graphs has been provided. The PA class is a generalization of both BPA and BPP that allows, similarly like multi-pushdown systems, both for sequential and concurrent (interleaved) behavior. Finally, in [24] the reachability problem has been shown decidable for Process Rewrite Systems [27] extended with weak control states.

Outline. In Section 2 we define the model we work with. Then in Section 3 we make explicit the variants the reachability problem we investigate in the paper, and in Section 4 we state all our results. To make the presentation more concise and easier to assimilate, the proofs of all the results outlined in Section 4 are moved to remaining sections, namely Sections 5–8.

In Section 5 we prove the fundamental property that the source set of configurations may be assumed to be a singleton without affecting complexity of the reachability problem. This section introduces also some basic terminology used in the following sections. The following three sections contain the proofs of the remaining results, grouped with respect to complexity. In Section 6 we prove membership in NP, in Section 7 we show decidability, and Section 8 contains undecidability proofs. In the last section we briefly discuss the cases where the complexity of the reachabiltiy problem is still open.

A preliminary version of this paper has appeared in [16]. This paper is an improvement of [16], extended with full proofs omitted there, and with few new results (for instance Theorem 4.1 and Theorem 4.7). The content of this paper is included in the PhD thesis of the first author [13].

2. MULTI-PUSHDOWN AUTOMATA

A multi-pushdown automaton (MPDA) is like a single-pushdown one. In a single step one symbol is popped from one of stacks[2], and a number of symbols are pushed on the stacks. Assume there are k stacks. A transition of an automaton is thus of the form:

$$q, \text{X} \xrightarrow{a} q', \alpha_1, \ldots, \alpha_k, \tag{2.1}$$

to mean that when an automaton reads a in state q, it pops X from one of the stacks, pushes the sequence of symbols α_i on the ith stack, for $i = 1 \ldots k$, and goes to state q'. We allow for silent transitions with $a = \varepsilon$. Observe that wlog. one may assume that stack alphabets are disjoint.

Formally, the ingredients of an MPDA are: a finite set of states Q, the number of stacks k, pairwise-disjoint finite stack alphabets $\mathsf{S}_1 \ldots \mathsf{S}_k$, an input alphabet A, and a finite set of transition rules:

$$\rightarrow \ \subseteq \ \mathsf{Q} \times (\bigcup_{i \leq k} \mathsf{S}_i) \times (\mathsf{A} \cup \{\varepsilon\}) \times \mathsf{Q} \times \mathsf{S}_1{}^* \times \ldots \times \mathsf{S}_k{}^* \tag{2.2}$$

---

[2] If we allowed for popping from more than one stack at a time, the model would clearly become Turing-complete, even with one state only.

written as in (2.1). A configuration of an MPDA is a tuple $\langle q, \beta_1, \ldots, \beta_k \rangle \in \mathsf{Q} \times \mathsf{S_1}^* \times \ldots \times \mathsf{S_k}^*$. The transition rules (2.1) induce the transition relation over all configurations in a standard way:

$$\frac{q,\ \mathrm{X}\ \overset{a}{\to}\ q',\ \alpha_1, \ldots, \alpha_k \qquad \mathrm{X} \in \mathsf{S_i} \qquad \beta_i = \mathrm{X}\beta}{\langle q, \beta_1, \ldots \beta_i \ldots, \beta_k \rangle \overset{a}{\to} \langle q', \alpha_1 \beta_1, \ldots \alpha_i \beta \ldots, \alpha_k \beta_k \rangle}$$

thus defining the configuration graph of an MPDA. For a configuration $\langle q, \alpha_1, \ldots, \alpha_k \rangle$, its *size* is defined as the sum of lengths of the words α_i, $i \leq k$. The same applies to a right-hand side of any transition rule $q\ \mathrm{X}\ \overset{a}{\to}\ q'\ \alpha_1 \ldots \alpha_k$.

An MPDA is *stateless* if there is just one state (or equivalently no states). Transition rules of an automaton are then of the form:

$$\mathrm{X}\ \overset{a}{\to}\ \alpha_1, \ldots, \alpha_k \tag{2.3}$$

and configurations are of the form $\langle \beta_1, \ldots, \beta_k \rangle$.

A less severe restriction on control states is the following one. We say that an automaton is *weak* if there is a partial order $\leq$ on its states such that every transition (2.1) satisfies $q' \leq q$. Clearly, every stateless automaton is weak.

Remark 2.1. Note that stateless one-stack automata are essentially context-free grammars in Greibach normal form. Thus the configuration graphs are precisely context-free graphs, called also BPA graphs [27, 10]. Another special case is many stacks with singleton alphabets. The stacks are thus essentially counters without zero tests. In this subclass, stateless automata corresponds to communication-free Petri nets [19], called also BPP [12], or commutative context-free graphs [15]. The BPA and BPP classes are members of the Process Rewrite Systems hierarchy of [27] that contains, among the others, unrestricted pushdown systems and Petri nets.

Example 2.2. Assuming a distinguished initial state and acceptance by all stacks empty, weak MPDAs can recognize non-context-free languages. For instance, the language

$$\{a^n b^n c^n : n \geq 0\} \tag{2.4}$$

is recognized by an automaton described below. The automaton has two states q_1, q_2 and two stacks. The alphabets of the stacks are $\{\mathrm{X}, \mathrm{B}, \mathrm{D}\}$ and $\{\mathrm{C}\}$, respectively. The starting configuration is $(q_1, \mathrm{XD}, \varepsilon)$. Besides the transition rules, we also present the automaton in a diagram, using *push* and *pop* operations with natural meaning.

$$q_1,\ \mathrm{X}\ \overset{a}{\to}\ q_1,\ \mathrm{XB},\ \mathrm{C}$$
$$q_1,\ \mathrm{X}\ \overset{\varepsilon}{\to}\ q_1,\ \varepsilon,\ \varepsilon$$
$$q_1,\ \mathrm{B}\ \overset{b}{\to}\ q_1,\ \varepsilon,\ \varepsilon$$
$$q_1,\ \mathrm{D}\ \overset{\varepsilon}{\to}\ q_2,\ \varepsilon,\ \varepsilon$$
$$q_2,\ \mathrm{C}\ \overset{c}{\to}\ q_2,\ \varepsilon,\ \varepsilon$$

The automaton is weak and uses ε-transitions, which may be however easily eliminated. Acceptance by empty stacks may be easily simulated using acceptance by states. The language (2.4) is not recognized by a stateless automaton, as shown in [17].

Example 2.3. Non-context-free languages are recognized even by stateless MPDAs with singleton stack alphabets. The class of languages recognized by this subclass is called *commutative context-free* languages [21], see also [17]. One example is the commutative closure of the language of the previous example: the set of all words with the same number of occurrences of a, b and c.

In the sequel we do not care about initial states nor about acceptance condition, as we will focus on the configuration graph of an automaton. Furthermore, as we only consider reachability problems, the labeling of transitions with input alphabet letters will be irrelevant, thus we write $\to$ instead of $\xrightarrow{a}$ from now on.

Using a standard terminology, we say that an MPDA is *normed* if for any state q and any configuration $\langle q, \alpha_1, \dots, \alpha_k \rangle$, there is a path to the empty configuration

$$\langle q, \alpha_1, \dots, \alpha_k \rangle \to \dots \to \langle p, \varepsilon, \dots, \varepsilon \rangle$$

for some state p. In general, whenever an MPDA is not assumed to be normed we call it *unnormed* for clarity. Note that in all examples above the automata were normed. In fact normedness is not a restriction as far as languages are considered. In the sequel we will however analyze the configuration graphs, and then normedness will play a role.

Further, we say that an MPDA is *strongly normed* if for any state q and any configuration $\langle q, \alpha_1, \dots, \alpha_k \rangle$, there is a path to the empty configuration

$$\langle q, \alpha_1, \dots, \alpha_k \rangle \to \dots \to \langle q, \varepsilon, \dots, \varepsilon \rangle$$

containing only transitions that do not change state. Intuitively, whatever is the state q we start in, any top-most symbol X in any stack may „disappear". For stateless automata, strong normedness is the same as normedness.

3. Regular sets and reachability problem

Regular sets. We will consider various reachability problems in the configuration graph of a given MPDA. Therefore, we need a finite way of describing infinite sets of configurations. A standard approach is to consider *regular* sets. Below we adapt this approach to the multi-stack scenario we deal with. Namely, instead of languages of words we consider languages of tuples of words, one word per stack.

Consider the configurations of a stateless MPDA, $\mathsf{S} = \mathsf{S}_1^* \times \dots \times \mathsf{S}_k^*$. There is a natural monoid structure in S, with pointwise identity $\langle \varepsilon, \dots, \varepsilon \rangle$ and multiplication

$$\langle \alpha_1, \dots, \alpha_k \rangle \cdot \langle \beta_1, \dots, \beta_k \rangle \;=\; \langle \alpha_1 \beta_1, \dots, \alpha_k \beta_k \rangle.$$

Call a subset $L \subseteq \mathsf{S}$ *regular* if there is a finite monoid M and a monoid morphism $\gamma : \mathsf{S} \to M$ that *recognizes* L, which means that $L = \gamma^{-1}(N)$ for some subset $N \subseteq M$. It is well-known (see for instance [4] and references therein) that a language L is regular iff it is a finite union of languages of the form

$$L_1 \times \dots \times L_k,$$

where L_i is a regular word language over S_i, for $i = 1 \dots k$. Thus one may assume, without loss of generality, that the monoid M is a product of finite monoids

$$M = M_1 \times \dots \times M_k,$$

and that

$$\gamma = \gamma_1 \times \ldots \times \gamma_k \quad \text{where} \quad \gamma_i : \mathsf{S}_i^* \to M_i \text{ for } i = 1 \ldots k.$$

Thus we may use an equivalent but more compact representation of regular sets, based on automata: a regular set L is given by a tuple of (nondeterministic) finite automata $\mathcal{B}_1 \ldots \mathcal{B}_k$ over alphabets $\mathsf{S}_1 \ldots \mathsf{S}_k$, respectively, together with a set

$$F \subseteq Q_1 \times \ldots \times Q_k$$

of accepting tuples of states, where Q_i denotes the state space of automaton $\mathcal{B}_i$.

Unless stated otherwise, in the sequel we always use such representations of regular sets of configurations. If there is more than one state, we assume a representation for every state. In particular, when saying "polynomial wrt. L", for a regular language L, we mean polynomial wrt. the sum of sizes of automata representing L.

Remark 3.1. Clearly, the cardinality of the set F of accepting tuples may be exponential wrt. the cardinalities of state spaces of automata $\mathcal{B}_i$. However, complexities we derive in the sequel will never depend on cardinality of F.

Example 3.2. Assume that there are two stacks. An example of properties we can define is: „odd number of elements on the first stack and symbol A on the top of the second stack, or an even number of the elements on the first stack and the odd number of elements on the second stack". On the other hand, "all stacks have equal size" is not a regular property according to our definition.

Remark 3.3. We have deliberately chosen a notion of regularity of languages of *tuples* of words. Another possible approach could be to consider regular languages of words, over the product alphabet $(\mathsf{S}_1 \cup \perp) \times \ldots \times (\mathsf{S}_k \cup \perp)$, where the additional symbol $\perp$ is necessary for padding. This would yield a larger class, for instance the last language from Example 3.2 would be regular. The price to pay would be however undecidability of the reachability problems. The undecidability will be established in Theorem 4.9.

Remark 3.4. The notion of regularity for languages of tuples of words we work with is known in the literature also under the name *recognizable* languages, see for instance [4, 5]. We prefer to stick to name *regular*, in order to place emphasis on the fact that the sets of configurations are assumed to be represented by tuples of finite automata.

Reachability. In this paper we consider the following reachability problem:

> INPUT: an MPDA $\mathcal{A}$ and two regular sets of configurations $L, K \subseteq \mathsf{S}$.
> QUESTION: is there a path in the configuration graph from L to K?

We will write $L \leadsto_{\mathcal{A}} K$ if a path from L to K exists in the configuration graph of an automaton $\mathcal{A}$. The sets L and K we call *source* and *target* sets, respectively. We will distinguish special cases, when either L or K or both the sets are singletons, thus obtaining four different variants of reachability altogether. For brevity we will use symbol '*1*' for a singleton, and symbol 'REG' for a regular set, and speak of *1*$\leadsto$REG *reachability* (when L is a singleton), REG$\leadsto$REG *reachability* (the unrestricted case), and likewise for REG$\leadsto$*1* and *1*$\leadsto$*1*.

Before stating the results, we note that all the problems we consider here are NP-hard:

Lemma 3.5. *The $1\rightsquigarrow1$ reachability is NP-hard for strongly normed stateless MPDAs, even if all stack alphabets are singletons.*

The above fact follows immediately from NP-completeness of the reachability problem for communication-free Petri nets, see [19] for details.

4. Summary of results

In this section we merely state all our results without any proofs. The proofs occupy all the following sections. Our results come in three groups, each dealing with one of the following questions:

- What is the complexity of the reachability problem in various variants mentioned above? Where the decidability border lies?
- Is the set of forward/backward reachable configurations regular and computable?
- What changes if the relaxed notion of regularity is assumed?

The three groups are discussed in details in Sections 4.1, 4.2 and 4.3, respectively.

Most of our results apply only to weak MPDAs, or even only to stateless MPDAs. There are however few exceptions: decidability of reachability (in Theorem 4.1) and regularity of backward reachability set (in Theorem 4.7), where we consider unrestricted state spaces. We want to emphasise that the results for unrestricted state spaces are rather accidental, and are not the most interesting; the core part of the paper is devoted to weak state spaces.

4.1. **Complexity of reachability.** In presence of states, the $1\rightsquigarrow1$ reachability problem is obviously undecidable, because the model is Turing powerful. Undecidability holds even for normed MPDAs. The problem becomes decidable under the assumption of strong normedness:

Theorem 4.1. *The* REG$\rightsquigarrow$REG *reachability is decidable for strongly normed MPDAs.*

The idea behind this result is that the model is similar to lossy FIFO systems, and hence is susceptible to the same proof methods. We exploit a well ordering of configurations, along the lines of [20].

All our complexity results are outlined in Table 1. The table distinguishes variants of the reachability problem for strongly normed/normed/unnormed MPDAs, and restrictions on control states: stateless/weak/unrestricted. For clarity, we do not distinguish stateless strongly normed case from stateless normed one, as these two cases obviously coincide.

[REG$\rightsquigarrow$1] REG$\rightsquigarrow$REG	strongly normed	normed	unnormed
stateless	NP-compl. (Thm. 4.4)		[NP-compl. (Thm. 4.5)] undecidable (Thm. 4.3)
weak	NP-compl. (Thm. 4.4)	[decidable] undecidable (Thm. 4.3)	[decidable (Thm. 4.6)] undecidable
unrestrictred	decidable (Thm. 4.1)	undecidable	undecidable

Table 1: Summary of complexity results

The remainder of Section 4.1 is devoted to weak state spaces only. We start by observing that out of four combinations of the reachability problem, it is sufficient to consider only two, namely the REG$\rightsquigarrow$1 and REG$\rightsquigarrow$REG cases. Indeed, as far as complexity is concerned, we observe the following collapse:

$$1\rightsquigarrow 1 \; = \; \text{REG}\rightsquigarrow 1 \qquad\qquad 1\rightsquigarrow \text{REG} \; = \; \text{REG}\rightsquigarrow \text{REG} \qquad\qquad (4.1)$$

independently of a restriction on automata. The first equality follows from our first result, that says that for a given target configuration, there is a source configuration, witnessing reachability, of polynomial size:

Lemma 4.2 (Polynomial source configuration). *Suppose $\mathcal{A}$ is a weak unnormed MPDA. Let L be a regular set of configurations of $\mathcal{A}$ and let t be a configuration of $\mathcal{A}$. Then*

$$L \rightsquigarrow_{\mathcal{A}} t \implies s \rightsquigarrow_{\mathcal{A}} t \text{ for some } s \in L \text{ of size polynomial wrt. } \mathcal{A}, L \text{ and } t.$$

Indeed, the reduction from REG$\rightsquigarrow$1 to 1$\rightsquigarrow$1 is by nondeterministic guessing of a source configuration of polynomial size. The second equality (4.1) will follow from our results listed below.

The collapse (4.1) significantly simplifies the summary of results in Table 1. Each entry of the table contains the complexity of REG$\rightsquigarrow$REG reachability problem. Additionally, the complexity of REG$\rightsquigarrow$1 reachability problem is given in cases it is different from the complexity of REG$\rightsquigarrow$REG reachability.

Now we discuss the results in detail. We first observe an apparent decidability frontier witnessed by stateless unnormed MPDAs and weak normed MPDAs:

Theorem 4.3. *The 1$\rightsquigarrow$REG reachability is undecidable for stateless unnormed MPDAs and for weak normed MPDAs.*

The proof is by reduction of the nonemptiness of intersection of context-free languages and uses three stacks. The case of two stacks remains open.

Thus lack of strong normedness combined with a regular target set yields undecidability in case of weak automata. Surprisingly, restricting additionally:

- either the automaton to be strongly normed,
- or the target set to a singleton,

makes a dramatical difference for complexity of the problem, as summarized in Theorems 4.4, 4.5 and 4.6 below. In the first theorem we only assume strong normedness:

Theorem 4.4. *The REG$\rightsquigarrow$REG reachability is NP-complete for strongly normed weak MP-DAs.*

Theorem 4.4 is the main result of this paper. It is proved by showing that reachability is always witnessed by a polynomial witness, obtained by careful elimination of 'irrelevant' transitions.

In the following two theorems we do not assume strong normedness, thus according to Theorem 4.3 we have to restrict target set to singleton. Under such a restriction, we are able to prove NP-completeness only in the class of stateless MPDAs, while for all weak MPDAs we merely state decidability:

Theorem 4.5. *The REG$\rightsquigarrow$1 reachability is NP-complete for stateless unnormed MPDAs.*

Theorem 4.6. *The REG$\rightsquigarrow$1 reachability is decidable for weak unnormed MPDAs.*

Theorem 4.5 is shown similarly to Theorem 4.4, while the proof of Theorem 4.6 is based on a well order, the point-wise extension of a variant of Higman ordering.

Open questions. Except for two few entries in the summarizing table above, where we merely prove decidability, we know the exact complexity of the reachability problem. The remaining open question are summarized in Section 9.

4.2. **Reachability set.** Now we consider the problem of computing the whole reachability set. For a given automaton $\mathcal{A}$, and a set L of configurations, we consider forward and backward reachability sets of L, defined as:

$$\{s : L \leadsto_{\mathcal{A}} s\} \qquad \text{and} \qquad \{s : s \leadsto_{\mathcal{A}} L\},$$

respectively. It turns out that the backward reachability set is regular, whenever L is regular, under the strong normedness assumption, even for unrestricted state spaces.

Theorem 4.7. *For strongly normed MPDAs, the backward reachability set of a regular set is regular.*

Roughly speaking, we show that the backward reachability set is upward closed with respect to the point-wise extension of a suitable variant of Higman ordering. Note that we do not claim however that the reachability set may be effectively computed.

On the other hand, the forward reachability set needs not be regular, even in the case of strongly normed stateless automata, as shown in the following example.

Example 4.8. Consider a strongly normed stateless automaton with two stacks, over alphabets $\{A, X\}$ and $\{B\}$, and the following transition rules:

$$X \to XA, B \qquad X \to \varepsilon, \varepsilon \qquad A \to \varepsilon, \varepsilon \qquad B \to \varepsilon, \varepsilon.$$

The set of configurations reachable from the configuration (X, ε) is not regular:

$$\{(A^i, B^j) : i, j \in \mathbb{N}\} \quad \cup \quad \{(XA^k, B^l) : k \geq l\}.$$

4.3. **Relaxed regularity.** The relaxed definition of regularity, as discusses in Remark 3.3, makes the reachability problem intractable in all cases. The following theorem is shown by reduction from the Post Correspondence Problem:

Theorem 4.9. *The $1\leadsto\mathrm{REG}$ reachability is undecidable for stateless strongly normed MP-DAs, under the relaxed notion of regularity.*

Furthermore, the backward reachability set of a relaxed regular set is not necessarily regular, even in stateless strongly normed MPDAs, as illustrated by the following example.

Example 4.10. The automaton uses two stacks, with alphabets $\{A, X, B\}$ and $\{C\}$. Every symbol has a disappearing rule: $A \to \varepsilon, \varepsilon$, and likewise for X, B and C. Additionally there is a transition rule $B \to C$. Consider the relaxed regular language

$$L = \{(XA^n, C^n) : n \geq 0\}$$

and its backward reachability set. Denote by K the subset of the backward reachability set consisting of configurations with the second stack empty. We claim that the projection of K on the alphabet of the first stack is a word language which is not regular. Indeed, as the only non-disappearing rule is $B \to C$, configurations from K have on the first stack a word of the form wXA^n, with at least n occurrences of B in w.

The remaining sections contain proofs of all the results announced Section 4. For Reader's convenience, every result is restated before its proof.

5. Proof of the polynomial source configuration property

Lemma 4.2. Suppose $\mathcal{A}$ is a weak unnormed MPDA. Let L be a regular set of configurations of $\mathcal{A}$ and let t be a configuration of $\mathcal{A}$. Then

$$L \rightsquigarrow_{\mathcal{A}} t \implies s \rightsquigarrow_{\mathcal{A}} t \text{ for some } s \in L \text{ of size polynomial wrt. } \mathcal{A}, L \text{ and } t.$$

Proof. Consider an MPDA $\mathcal{A}$ and a regular set L of configurations of $\mathcal{A}$. Let $s \in L$ be source configuration and let t be an arbitrary target configuration. Suppose $s \rightsquigarrow_{\mathcal{A}} t$. We will show that the size of s may be reduced, while preserving membership in L. The crucial but simple idea of the proof will rely on an analysis of *relevance* of symbol occurrences, to be defined below.

Symbol occurrences. Suppose that there is a path π from s to t, consisting of consecutive transitions $s \rightarrow s_1 \rightarrow s_2 \ldots \rightarrow s_n = t$. We will consider all individual occurrences of symbols that appear in the configurations. For instance, in the following exemplary sequence of two-stack configurations

$$\langle q, \mathrm{AA}, \mathrm{C} \rangle \rightarrow \langle q, \mathrm{BBA}, \mathrm{DC} \rangle \rightarrow \langle q, \mathrm{ABBA}, \mathrm{DC} \rangle \tag{5.1}$$

there are altogether 14 *symbol occurrences*: 3 in the first configuration, 5 in the second one and 6 in the third one.

Recall that every transition $s_i \rightarrow s_{i+1}$ is induced by some transition rule $q_1, \mathrm{X} \rightarrow q_2, \alpha$ of the automaton. Then there is a distinguished occurrence of symbol X in s_i that is involved in the transition. In the sequel we use the term *symbol occurrence involved in a transition*.

Precisely one occurrence of symbol in s_i is involved in the transition $s_i \rightarrow s_{i+1}$; for every other occurrence of a symbol in s_i there is a *corresponding* occurrence of the same symbol in s_{i+1}. (Note that we always make a difference between corresponding symbol occurrences from different configurations.) All remaining occurrences of symbols in s_{i+1} are created by the transition; we call these occurrences *fresh*.

We define the *descendant* relation as follows. All fresh symbol occurrences in s_{i+1} are descendants of the symbol occurrence in s_i involved in the transition $s_i \rightarrow s_{i+1}$. Moreover, a symbol occurrence in s_{i+1} corresponding to a symbol occurrence in s_i is its descendant too. We will use term *descendant* for the reflexive-transitive closure of the relation defined above and the term *ancestor* for its inverse relation. In particular, every symbol occurrence in t is descendant of a unique symbol occurrence in s. The descendant relation is a forest, i.e., a disjoint union of trees.

Example 5.1. As an example, consider again the sequence of transitions (5.1), with symbol occurrences identified by subscripts $1 \ldots 14$:

$$\langle q, \mathrm{A}_1\mathrm{A}_2, \mathrm{C}_3 \rangle \rightarrow \langle q, \mathrm{B}_4\mathrm{B}_5\mathrm{A}_6, \mathrm{D}_7\mathrm{C}_8 \rangle \rightarrow \langle q, \mathrm{A}_9\mathrm{B}_{10}\mathrm{B}_{11}\mathrm{A}_{12}, \mathrm{D}_{13}\mathrm{C}_{14} \rangle \tag{5.2}$$

Say the transitions are induced by the following two transition rules:

$$q, \mathrm{A} \rightarrow q, \mathrm{BB}, \mathrm{D} \qquad\qquad q, \mathrm{D} \rightarrow q, \mathrm{A}, \mathrm{D}$$

The descendant relation can be presented as the following forest:

$$D_7 \longrightarrow D_{13} \qquad C_3 \longrightarrow C_8 \longrightarrow C_{14}$$

$$A_1 \longrightarrow B_5 \longrightarrow B_{11} \longrightarrow A_9$$

$$B_4 \longrightarrow B_{10} \qquad A_2 \longrightarrow A_6 \longrightarrow A_{12}$$
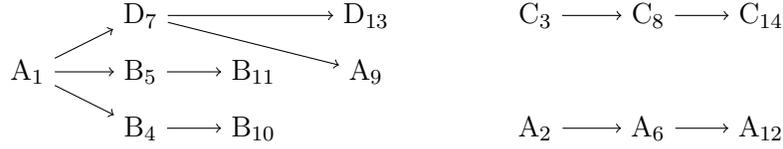
The symbol occurrences involved in the two transitions (5.2) are A_1 in the first configuration and D_7 in the second one. The fresh symbol occurrences are B_4, B_5 and D_7 in the second configuration, and A_9 and D_{13} in the third one.

Relevant symbol occurrences. As the automaton $\mathcal{A}$ is weak, the number of transitions in π that change state is bounded by the number of states of $\mathcal{A}$. All remaining transitions in π do not change state.

Consider all the occurrences of all symbols in all configurations along the path π, including configurations s and t themselves. A symbol occurrence is called *relevant* if some of its descendants:

• belongs to the target configuration t; or
• is involved in some transition in π that changes state.

Otherwise, a symbol occurrence is *irrelevant*. In particular, all symbol occurrences in t are relevant. Referring back to our example, all symbol occurrences appearing in (5.2) are relevant.

Note that if t is not the empty configuration then every configuration in π contains at least one relevant symbol occurrence. On the other side, in every configuration, the number of relevant occurrences is always bounded by the sum of the size of t and the number of states of $\mathcal{A}$.

Small source configuration. So prepared, we are ready to prove that there is a configuration $s' \in L$ of polynomial size with $s' \leadsto_{\mathcal{A}} t$. We will rely on the following lemma:

Lemma 5.2. *For any configuration s' obtained from s by removing some irrelevant symbol occurrences, it holds $s' \leadsto_{\mathcal{A}} t$.*

The lemma follows from the following two observations: (1) all the transitions in π involving symbol occurrences remaining in s' and their descendants may be re-done; (2) the resulting configuration will be exactly t, as only irrelevant symbol occurrences have been removed from s.

Recall that the language L is represented by a tuple $\mathcal{B}_1 \ldots \mathcal{B}_k$ of deterministic finite automata, one automaton per stack. Consider the content of a fixed ith stack in s, say $w \in A_i^*$. Let n be the number of states of $\mathcal{B}_i$. The run of the automaton $\mathcal{B}_i$ over w labels every position of w by some state. We will use a standard pumping argument to argue that every block of consecutive irrelevant symbol occurrences in s may be reduced in length to at most n. Indeed, upon every repetition of a state of $\mathcal{B}_i$, the word w may be shortened by removing the induced infix, while preserving membership in L. By repeating the pumping argument for all blocks of consecutive irrelevant symbol occurrences in all stacks in s, one obtains a configuration s', still belonging to L, of quadratic size. By Lemma 5.2 we know that $s' \leadsto t$, as required. $\qquad \square$

6. NP-completeness

In this section we prove the following two theorems of Section 4:

Theorem 4.4. The REG$\rightsquigarrow$REG reachability is NP-complete for strongly normed weak MP-DAs.

Theorem 4.5. The REG$\rightsquigarrow$1 reachability is NP-complete for stateless unnormed MPDAs.

NP-hardness in both cases follows from Lemma 3.5. Concerning Theorem 4.4, the proof of membership in NP relies on the following two core lemmas:

Lemma 6.1. *The 1$\rightsquigarrow$1 reachability problem is in NP for strongly normed weak MPDAs.*

Lemma 6.2. *Let $\mathcal{A}$ be a strongly normed weak MPDA and let L, K be regular sets of configurations. If $L \rightsquigarrow K$ then $s \rightsquigarrow t$ for some $s \in L$ and $t \in K$ of size polynomial wrt. the sizes of $\mathcal{A}$, L and K.*

Indeed, the two lemmas easily yield a decision procedure for REG$\rightsquigarrow$REG reachability: guess configurations $s \in L$ and $t \in K$ of size bounded by a polynomial deduced from the proof of Lemma 6.2, and then apply the procedure of Lemma 6.1 to check if $s \rightsquigarrow t$.

Concerning Theorem 4.5, membership in NP follows easily by Lemma 4.2 together with the following:

Lemma 6.3. *The 1$\rightsquigarrow$1 reachability problem is in NP for stateless unnormed MPDAs.*

The proof of Lemma 6.3 is similar to the proof of Lemma 6.1, as the irrelevant symbol occurrences must necessarily be normed, because they do not contribute to the target configuration. We thus skip this proof, and devote the rest of this section to the proofs of Lemmas 6.1 and 6.2.

Proof of Lemma 6.1. Consider an MPDA $\mathcal{A}$ and two configurations s and t. We will define a nondeterministic polynomial-time decision procedure for $s \rightsquigarrow_{\mathcal{A}} t$.

Stateless assumption. For simplicity, we assume that both s and t have the same control state. Thus we can treat transitions that lead from s to t as stateless transitions. At the very end of the proof, we will discuss how to generalize it to the general case of strongly normed weak MPDAs.

Polynomial witness. Our aim is to show that if there is a path from s to t then there is a path of polynomial length. So stated, the above claim may not be verbally true, even in the case of context-free graphs, as witnessed by the following simple example.

$$X_1 \rightarrow X_2 X_2 \qquad X_2 \rightarrow X_3 X_3 \qquad \ldots \qquad X_{n-1} \rightarrow X_n X_n \qquad X_n \rightarrow \varepsilon \qquad (6.1)$$

The example scales with respect to n, and thus the shortest path from the configuration X_1 to X_n is of exponential length. As a conclusion, one must use some subtle analysis in order to be able to reduce the length of a witness of existence of the path as required. Note that X_1 is relevant and thus can not be simply omitted.

Proof idea. As a first step towards a polynomial bound on the witness of the path from s to t, we will modify the notion of transition. Intuitively speaking, our aim is to consider exclusively relevant symbol occurrences.

By a *subword* we mean any subsequence of a given word. For instance, *aaccbc* is a subword of *aacabbcbcbc*. Further, by a *subtransition* of $X \rightarrow \alpha_1 \ldots \alpha_k$ we mean any $X \rightarrow \beta_1 \ldots \beta_k$ such that the following conditions hold:

- *subword*: β_i is a subword of α_i, for all $i \in \{1 \ldots k\}$; and
- *nonemptiness*: $\beta_1 \ldots \beta_k \neq \varepsilon$, i.e., at least one of words β_i is nonempty.

Note that relying on the notion of relevance one easily deduces that whenever there is a sequence of transitions from s to t, then there is also sequence of subtransitions. Indeed, it is sufficient to remove irrelevant symbol occurrences in all transitions along the path from s to t.

Clearly, the converse implication is not true in general. For instance, if we add symbols X_0, A and transitions $X_0 \rightarrow X_1A$, $A \rightarrow \varepsilon$ to the example (6.1), there is a sequence of subtransitions from the configuration X_0 to X_n, but no sequence of transitions from X_0 to X_n. Our aim now it to modify the notion of subtransition in such a way that the converse implication does hold as well, i.e., that existence of a sequence of subtransitions implies existence of a sequence of transitions. This requires certain amount of laborious book-keeping, as defined in detail below.

Marked subtransitions. We will need an additional copy of every stack alphabet A_i, denoted by $\bar{A}_i$, for $i = 1 \ldots k$. Thus for every $a \in A_i$ there is a corresponding marked symbol $\bar{a} \in \bar{A}_i$. Formally, let the ith stack alphabet be $A_i \cup \bar{A}_i$.

A *marked subword* of a word $w \in A_i^*$ is any word in $(A_i \cup \bar{A}_i)^*$ that may be obtained from w by the following *marking procedure*:

- color arbitrary occurrences in w (the idea is to color irrelevant symbol occurrences);
- choose a prefix of w and mark every occurrence in this prefix; the prefix should contain all occurences that are followed by some colored occurrence;
- and finally remove colored occurrences.

For instance, the following four words are all the marked subwords of AACABBCBCBC, with respect to the coloring A**AC**A**BB**C**B**CBC:

$$\bar{A}\bar{A}\bar{C}CBC \qquad \bar{A}\bar{A}\bar{C}\bar{C}BC \qquad \bar{A}\bar{A}\bar{C}\bar{C}\bar{B}C \qquad \bar{A}\bar{A}\bar{C}\bar{C}\bar{B}\bar{C}. \tag{6.2}$$

Recall that a word $w \in A_i^*$ represents a content of the ith stack, with the left-most symbol being the top-most. Intuitively, the idea behind the notion of marked subword is to keep track of removed occurrences that are covered by other symbols on the stack. For technical reasons in the second point above we allow additionally marking some further symbols that are not followed by a colored occurrence.

A notion of *marked subtransition* is a natural adaptation of the notion of subtransition. Compared to subtransitions, there are two differences: 'subword' is replaced with 'marked subword'; and whenever the left-side symbol is marked, then it may only put marked symbols on its stack. Formally, a marked subtransition of $X \rightarrow \alpha_1 \ldots \alpha_k$ is any $X \rightarrow \beta_1 \ldots \beta_k$ such that the following conditions hold:

- *subword*: β_i is a marked subword of α_i, for all $i \in \{1 \ldots k\}$;
- *nonemptiness*: $\beta_1 \ldots \beta_k \neq \varepsilon$, i.e., at least one of words β_i is nonempty; and
- *marking inheritance*: if X is marked, say $X \in \bar{A}_i$, then all symbols in β_i are marked.

Observe that in the *marking inheritance* condition we exploit possibility to mark additional occurrences (in fact all occurrences, cf. (6.2)). As an example, consider a transition:

$$A \ \rightarrow \ AACABBCBCBC, \ DDED$$

for $\{A, B, C\}$ and $\{D, E\}$ the alphabets of the two stacks. Then according to the colorings

$$A\mathbf{AC}ABB\mathbf{C}BCBC \qquad \text{and} \qquad D\mathbf{DE}D,$$

there are two marked subtransitions with $\bar{A}$ the left-side symbol:

$$\bar{A} \ \rightarrow \ \bar{A}\bar{A}\bar{C}\bar{C}\bar{B}\bar{C}, \ \bar{D}D \qquad \text{and} \qquad \bar{A} \ \rightarrow \ \bar{A}\bar{A}\bar{C}\bar{C}\bar{B}\bar{C}, \ \bar{D}\bar{D}.$$

Note that there are exponentially many different marked subtransitions, but each one is of polynomial size. Finally, note that every subtransition is obtained from some transition by the marking procedure as above, applied to every stack separately.

By the nonemptiness assumption on marked subtransitions we obtain a simple but crucial observation:

Lemma 6.4. *Along a sequence of marked subtransitions, the size of configuration can not decrease.*

A *marked subconfiguration* of a configuration $\langle \alpha_1, \ldots, \alpha_k \rangle$ is any tuple $\langle \beta_1, \ldots, \beta_k \rangle$ such that β_i is a marked subword of α_i for all $i \in \{1 \ldots k\}$.

Lemma 6.5. *For two configurations s and t, the following conditions are equivalent:*
(1) *there is a sequence of transitions from s to t,*
(2) *there is a sequence of marked subtransitions from u to t, for some marked subconfiguration u of s.*

Proof. The implication from (1) to (2) follows immediately. The sequence of marked subtransitions is obtained by application of the marking procedure to all transitions. For every transition, color in the marking procedure precisely those symbol occurrences that are irrelevant.

Now we show the implication from (2) to (1). The proof uses strong normedness.

Assume a sequence π of marked subtransitions from u to t, for some marked subconfiguration u of s. Recall that each subtransition in π has its original transition of $\mathcal{A}$. We claim that there is a sequence of transitions from s to t, that contains the original transitions of all the marked subtransitions appearing in π, and *canceling sequences*

$$q \ X \ \rightarrow \ \ldots \ \rightarrow \ \langle q, \varepsilon, \ldots, \varepsilon \rangle \tag{6.3}$$

for some stack symbols X, existing due to strong normedness assumption.

The sequence of transitions from s to t is constructed by reversing the marking procedure. For the ease of presentation, beside letters from A_i, we will also use colored letters.

Start with the configuration s, and choose any coloring of symbol occurrences in s that induces u as the outcome of the marking procedure. Then consecutively apply the following rule:

- If the top-most symbol X on some stack is colored, apply a canceling sequence for X.
- Otherwise, apply the original transition of the next subtransition from π, using again some coloring that could have been used in the marking procedure.

For correctness, we need to show that all colored occurrences of symbols are eventually canceled out, as this guarantees that the final configuration is precisely t.

Let's inspect π. As no symbol in t is marked, every marked symbol occurrence eventually disappears as a result of firing of some subtransition. Recall that marking of a symbol $\bar{\text{X}}$ disappears only if the subtransition pushes nothing on the stack of $\bar{X}$. As a consequence, every colored symbol occurrence will eventually appear on the top of its stack. Thus the canceling sequence for X will eventually be applied.   □

Lemma 6.6. *For two configurations u and v, if there is a sequence of marked subtransitions from u to v, then there is such a sequence of polynomial length wrt. the sizes of u, v and $\mathcal{A}$.*

Proof. From now on, we will write 'subtransitions' instead of 'marked subtransitions'. As we will primarily work with subtransitions, we will use the stack alphabets $\mathsf{A}_i \cup \bar{\mathsf{A}}_i$ for $i \in \{1 \ldots k\}$.

The number of subtransitions that change state is bounded by the number of states of $\mathcal{A}$, as $\mathcal{A}$ is assumed to be weak. Thus it is sufficient to prove the lemma under the assumption that the subtransitions do not change state. In other words, wlog. we may assume $\mathcal{A}$ to be stateless.

The size of the right-hand side of a marked subtransition is at least 1. Distinguish subtransitions with the size of the right-hand side equal 1, and call them *singleton subtransitions*. Clearly, the number of non-singleton subtransitions appearing in the sequence in the above claim is at most equal to the size of v, thus it is sufficient to concentrate on the following claim:

Claim 6.7. If there is a sequence of *singleton* subtransitions from a configuration u to v then there is such a sequence of polynomial length.

Note that the sizes of u and v in the above claim are necessarily the same.

Now we analyze in more detail the singleton subtransitions. Note that they have the form

$$\text{X} \to \text{Y} \tag{6.4}$$

as the right-hand sides contain precisely one occurrence of a symbol. Consider the strongly connected components in the induced graph, with symbols as vertices, and singleton subtransitions (6.4) as edges.

Distinguish those singleton subtransitions (6.4) that stay inside one strongly connected component (in other words, such that there is a sequence of subtransitions from Y back to X) and call them *inner singleton subtransitions*. Note that the number of non-inner subtransitions that appear in the sequence of the last claim is polynomial (at most quadratic), thus the last claim is equivalent to the following one:

Claim 6.8. If there is a sequence of *inner singleton* subtransitions from a configuration u to v then there is such a sequence of polynomial length.

The rest of the proof is devoted to showing the last claim.

We start by doing a sequence of simplifying assumption without losing generality.

First, wlog. we may assume that the relation (6.4) is transitive, as we only care about the length of the sequence of subtransitions up to a polynomial. Thus every strongly connected component is a directed clique.

By the *type* of a clique we mean the set of stacks that are represented in the clique, i.e., the stacks that have at least one symbol in the alphabet that belongs to the clique. We may assume that there is no clique of singleton type. Indeed, otherwise the stack is essentially

inactive along the path, except for the top-most symbol, and thus may be ignored in our analysis.

Further, wlog. we may also assume that every clique has *at most* one symbol belonging to every stack alphabet. Indeed, two different symbols from the same clique and the same stack alphabet can easily mutate from one into the other, when being the top-most symbol of the stack. And every symbol X may be easily made top-most by popping all symbols above X to other stacks (this can be done due to the assumption that type of cliques are not singletons).

The simplifications lead us to the following reformulation of the last claim. Let $k \geq 1$ be an integer. Assume a finite set of symbols A, each symbol $X \in \mathsf{A}$ coming with its type $\mathrm{type}(X) \subseteq \{1 \ldots k\}$ of cardinality at least 2. Consider the set of k-tuples of stacks $(\mathsf{A}^*)^k$ satisfying the following consistency condition: if X appears in the ith sequence then $i \in \mathrm{type}(X)$. Consider the following transition rules: the top-most letter of some stack may be moved to the top of some other stack, as long as the consistency is preserved.

Claim 6.9. If there is a sequence of transitions from some configuration $u \in (\mathsf{A}^*)^k$ to some configuration $v \in (\mathsf{A}^*)^k$, then there is such a sequence of polynomial length.

So formulated, the claim is not so hard.

We will show a polynomial sequence of transitions that starts in u and ends in a configuration u' that has the same bottom-most symbol as v on some stack. This is sufficient, as the same thing may be done for all other occurrences of symbols in v (formally, an induction over stack depth is involved here).

Note that we do not assume that different symbols have different types. Two symbols we call *siblings* if they have the same type and this type has two elements (thus the symbols may be placed only on two stacks).

Choose an arbitrary stack that is nonempty in v, say the ith stack, with the bottom-most symbol X. We may assume wlog. that X does not appear in u on the ith stack (otherwise, i.e., if some occurrences of X in u are on the ith stack, move all the occurrences of X, together with all other symbols above them, to arbitrary other stacks).

Let the jth stack in u contain an occurrence of symbol X, for some $j \neq i$.

The sequence of steps from u to u' is the following:

(1) Move all symbols above the chosen occurrence of X from the jth stack to other stacks.
(2) Move all symbols from the ith stack to other stacks such that X is still on the top of the jth stack.
(3) Move the chosen occurrence of symbol X to the destination ith stack.

Clearly step 1. may be always done. We will show that step 2. may be always done as well. We distinguish two cases.

If the symbol X is not a sibling, every other symbol may be moved, from the ith stack, to a stack different than the ith one, in such a way that after this operation X will be still on the top of the jth stack. Indeed, assume that a symbol Y is on the top of the ith stack. If Y can be moved to a stack different than the jth one, we are done. Otherwise, Y can only occur on the ith and jth stacks. According to the assumption, X and Y are not siblings, thus there is another kth stack to which X can be moved. Then we proceed as follows: X is moved from jth to the kth stack, next Y is moved from the ith stack to the jth stack, and finally X is moved back from the kth stack to the jth stack.

As the second case, assume that X is a sibling. Assume that the top-most occurrence of X on the jth stack has been chosen. As $j \neq i$, and there is a sequence of steps from u to

v, one easily observes that no sibling of X may occur in u either on the ith stack, or above X on the jth stack. Thus step 2. clearly can be done.

This completes the proof of Lemma 6.6 and thus also the proof of Lemma 6.1, under the stateless assumption. □

Decision procedure. Now we drop the stateless assumption. Note that the notion of marked subconfiguration and marked subtransition may be easily adapted to transitions that change state. We do not impose however the nonemptiness condition on transitions that change state, which is in accordance with the intuition that irrelevant symbol occurrences are removed in the marking procedure. Using Lemmas 6.4, 6.5 and 6.6 we will define the nondeterministic decision procedure for strongly normed weak MPDAs.

Let the two given configurations s and t have control states q and p, respectively. In the first step, the algorithm guesses a number of marked subconfigurations $t_1 \ldots t_{n-1}$, where n is not greater than the number of states of $\mathcal{A}$, and marked subtransitions that change state:

$$t_1 \rightarrow s_1 \qquad\qquad t_2 \rightarrow s_2 \qquad\qquad \ldots \qquad\qquad t_{n-1} \rightarrow s_{n-1}$$

such that s_i and t_{i+1} have the same control states for $i \in \{0 \ldots n-1\}$. Note that size of all configurations s_i and t_i is necessarily bounded by $\text{size}(t) + |\mathcal{A}|$, as the size of a configuration along the run may be decreased only during the change of state and even then only by one. For convenience, we write s_0 instead of s and t_n instead of t. In particular, we assume that the control state of t_1 is q, and the control state of s_{n-1} is p. Relying on Lemma 6.4, it is sufficient to consider configurations of sizes satisfying the following inequalities:

$$\text{size}(s_i) \leq \text{size}(t_{i+1}) \qquad \text{for } i \in \{1 \ldots n-1\}. \tag{6.5}$$

In the second phase, the algorithm guesses, for $i \in \{0 \ldots n-1\}$, a sequence of subtransitions from s_i to t_{i+1} of length bounded by polynomial derived from the proof of Lemma 6.6; and checks that the respective sequences of subtransitions lead from s_i to t_{i+1}, as required by Lemma 6.5. □

Proof of Lemma 6.2. Suppose $\mathcal{A}$ is a strongly normed weak MPDA. Let L, K be regular sets of configurations of $\mathcal{A}$ and let π be a sequence of transitions from some configuration $s \in L$ to some configuration $t \in K$. We will demonstrate existence of configurations $s' \in L$ and $t' \in K$ such that t' is of polynomial size and $s' \rightsquigarrow t'$. Importantly, we do not have to provide any bound on the size of s', as the polynomial bound follows by Lemma 4.2.

Recall the coloring discipline used in the proof of Lemma 6.1. There we used just one color; here we will use an unbounded number of different colors, as described below.

The coloring discipline will apply to all configurations appearing in π. We start by coloring all symbol occurrences in the first configuration s with different colors. Then, for every transition $s_1 \rightarrow s_2$ in π, assumed that s_1 has already been colored, we stipulate the following coloring rule for s_2 (recall that symbol occurrences in s_2 are divided into those corresponding to symbol occurrences in s_1, and fresh ones):

- If a symbol occurrence corresponds to a symbol occurrence in s_1, its color is the same as the color of corresponding symbol occurrence.
- Let c be the color of the unique occurrence of symbol in s_1, say symbol X, that is involved in the transition. All fresh symbol occurrences in s_2 that appear on the same stack as X are colored with c; we say that they *inherit* the color from X. All other fresh symbol occurrences in s_2 are colored by new fresh colors: one for each other stack. Thus at most

$k - 1$ new fresh colors is needed for coloring fresh occurrences on other stacks, where k is the number of stacks.

For any color used, and for any fixed configuration, the set of all symbol occurrences colored with that color we call *line*. Note that a line is always a subset of symbol occurrences on a single stack. However, there are typically many lines sharing their color, but this happens only to lines from different configurations. Further, note that the cardinality of a line is not bounded in principle, due to the inheritance of color.

We now aim at reducing the size of the destination configuration t'. Roughly speaking, we will prove that $s' \rightsquigarrow t'$, for some $s' \in L$ and $t' \in K$ such that both the number of different lines in t', and the cardinality of all lines in t', are polynomially bounded.

For convenience, we split colors into two disjoint subsets. A color c is called *active* if some symbol occurrence labeled by c:

- either is involved in some transition in π,
- or is a fresh symbol occurrence created by some transition in π.

Otherwise, a color is called *inactive*, i.e., occurrences of this color are present in s and are not involved in any transition during the run. Likewise, a line is also called active or inactive, according to its color. Note that inactive colors label suffixes of stacks in every configuration in π, and these suffixes do not change along π. Inactive lines are clearly singletons.

Bounding the number of active lines. Consider content of some stack, say the ith stack, in the destination configuration $t \in K$. Denote by $w \in \mathsf{A}_i^*$ its prefix colored by active colors. Every active line on the ith stack corresponds to an infix of w, and thus the coloring induces a factorization

$$w = w_1 \cdot w_2 \cdot \ldots \cdot w_m$$

determined by some $m - 1$ positions in w.

Recall that the language K is represented by a tuple $\mathcal{B}_1 \ldots \mathcal{B}_k$ of finite automata, one automaton per stack. A run of the automaton $\mathcal{B}_i$ over the ith stack of t labels each of the $m - 1$ distinguished positions in w by a state. By a standard pumping argument, there is a subword w' of w, obtained by removing a number of lines from w, that contains at most as many lines as the number of states of $\mathcal{B}_i$, and such that $\mathcal{B}_i$ reaches the same state after reading w and w'. By repeating the pumping argument for all stacks, one obtains a configuration t' still belonging to K, that contains only a polynomial (in fact, at most quadratic) number of active lines, as required.

We only need to show that $s \rightsquigarrow t'$. In this part of the proof we will use the canceling sequences (6.3), available due to strong normedness. Observe that every active line that appears in π appears as the top-most one on its stack at some configuration in π. We apply the canceling sequence for all symbol occurrences in every active line not appearing in t'. In order to keep the reachability $s \rightsquigarrow t'$, we apply the canceling sequence in the last configuration in π where this line is the top-most one. Thus the disappearance of a line has no effect for the remaining lines.

Bounding the number of inactive lines. We repeat a pumping argument similar to the above one. Let L and K be represented by $\mathcal{A}_1 \ldots \mathcal{A}_k$ and $\mathcal{B}_1 \ldots \mathcal{B}_k$, respectively. Consider some $i \leq k$ and runs of automata $\mathcal{A}_i$ and $\mathcal{B}_i$ over the inactive suffix of the ith stack of t' (or t). The runs label every position by a pair of states of $\mathcal{A}_i$ and $\mathcal{B}_i$, respectively. Upon a repetition of the same pair of states, a standard pumping applies. Thus the length of every inactive suffix in v may be reduced to at most quadratic.

Bounding the cardinalities of active lines. Consider the configuration t' obtained by now, and a single active line on some ith stack in this configuration. Let $v \in \mathsf{A}_i^*$ be the word representing the line. Thus the ith stack in t' is of the form:

$$w_1\, v\, w_2$$

for some words w_1, w_2. Similarly as before, we aim at applying pumping inside v, to reduce its length.

Let's focus on symbol occurrences in v in configuration t' and on the corresponding symbol occurrences in other configurations in π. Observe that all symbol occurrences in v satisfy the following condition:

> the corresponding symbol occurrence in some previous configuration was freshly
> created in some transition.

Some of the above-mentioned transitions have created new lines, and some not. Among symbols in v, distinguish a subset containing only those occurrences that satisfy the following strengthened condition:

> the corresponding symbol occurrence in some previous configuration was freshly
> created in some transition that created a new line that is represented in t'.

The distinguished symbol occurrences call *non-local*, the others call *local*.

The overall number of lines in t' is polynomially bounded, thus the same bound applies to the number of non-local symbol occurrences in v. We thus obtain:

Claim 6.10. There is only polynomially many non-local symbol occurrences in v.

Thus, it is sufficient to reduce the length of any block of local occurrences in v. From now on we focus on a single maximal infix u of v that contains only local symbol occurrences.

Those transitions in π that involve a symbol occurrence corresponding to a symbol occurrence in u use only the ith stack. Thus this set of transitions is essentially a stateless pushdown automaton. We will use a well-known fact (see for instance [7]):

Claim 6.11. For a pushdown automaton one can construct a finite automaton of polynomial size that recognizes the language of all reachable configurations of the pushdown automaton.

We will now use a standard pumping to reduce the length of u. As said above, a run of $\mathcal{B}_i$ over the ith stack of v labels each position of u with a state. Likewise for the automaton of the above claim. Upon a repetition of a pair of states, a standard pumping applies, as usual. This completes the proof of Lemma 6.2. $\qquad\square$

7. Decidability

In this section we provide proofs of Theorems 4.6, 4.7 and 4.1.

Theorem 4.6. The REG$\leadsto$1 reachability is decidable for weak unnormed MPDAs.

Proof. By virtue of Lemma 4.2 we may focus on the 1$\leadsto$1 reachability only. Fix an MPDA $\mathcal{A}$ and two configurations s and t. We will describe an algorithm to decide whether $s \leadsto_{\mathcal{A}} t$. Roughly speaking, our approach is to define a suitable well order compatible with transitions, and then apply a standard algorithm for reachability of a downward-closure of a

configuration[3]. However, to apply the standard framework we need to introduce some additional structure in configurations. This additional structure will be intuitively described as coloring of symbols, similarly as marking in the proof of Lemma 6.1.

Recall the notion of relevant symbol occurrences, introduced in Section 5. The idea of the proof will be based on the observation that removing some irrelevant symbol occurrences has no impact on reachability of a fixed target configuration (cf. Lemma 5.2 from Section 5).

Fix the target configuration t. We will define *colored configurations* and modified transitions between colored configurations. The basic intuition is that irrelevant symbol occurrences will be colored. Note however that we don't know in advance which symbol occurrences in a given configuration s are relevant and which are not, as we do not even know if $s \rightsquigarrow t$. Thus a coloring will have to be guessed.

Let n be the number of states of $\mathcal{A}$ and let m be the size of t. By a *colored configuration* we mean a configuration with some symbol occurrences colored, such that the number of uncolored symbol occurrences is smaller than $n + m$. Formally, coloring is implemented by extending the alphabet of every stack with its colored copy. We define an ordering on colored configurations: $r' \preceq r$ if r' is obtained from r by removing some colored symbol occurrences. (In particular, if $r' \preceq r$ then both configurations are identical, when restricted to uncolored symbols. Thus an uncolored configuration corresponds to a downward-closed set.) As the number of uncolored occurrences is bounded, the number of blocks of colored occurrences is bounded likewise. The ordering $\preceq$ is like a Higman ordering on words (i.e., a word v is smaller than a word w if v can be obtained from w by removing some letters), extended in the point-wise manner to blocks of colored occurrences. Thus one easily shows, using Higman's lemma:

Claim 7.1. The ordering $\preceq$ is a well order on colored configurations, i.e. for any infinite sequence of configurations $s_1 s_2 \ldots$ there are indices $i < j$ such that $s_i \preceq s_j$.

Now we define the transition rules for colored configurations. Consider any original transition rule δ of $\mathcal{A}$. This transition rule will give rise to a number of new transition rules that will be applicable to colored configurations. One new transition is obtained by coloring all symbols in δ, i.e., both the left-hand side symbol and all the right-hand side symbol occurrences. In all other new transitions arising from δ, the left-hand side symbol is kept uncolored. On the other hand, an arbitrary subset of the right-hand side symbol occurrences may be colored, under the following restriction:

> if the transition δ does not change state then at least one of right-hand side symbols must be kept uncolored.

This corresponds to the intuition that uncolored symbol occurrences correspond to relevant ones.

We have thus now two transition systems: the original transition system and the colored one. The relationship between reachability in these two systems is stated in the following claim (recall that the configuration t is fixed and contains no colored symbols):

Claim 7.2. For any configuration s, $s \rightsquigarrow t$ if and only if there is some coloring s' of s such that $s' \rightsquigarrow t$.

---

[3]The algorithm works actually in any well-structured transition system under suitable assumptions, see [20]. Theorem 4.6 is in fact a special case of Theorem 5.5 in [20].

Indeed, the only if direction is obtained by coloring precisely irrelevant symbol occurrences in s. The if direction also follows immediately, by replacing the colored transitions with their uncolored original transitions.

Basing on the above claim, the algorithm for $s \rightsquigarrow t$ simply guesses a coloring s' of s and then checks if $s' \rightsquigarrow t$ in the colored transition system. It thus only remains to show that the latter problem is decidable. For this we will need a compatibility property of the colored transitions with respect to the well order:

Claim 7.3. For every colored configurations r', r and u, if $r' \preceq r \rightarrow u$ then
- either there is a colored configuration u' with $r' \rightarrow u' \preceq u$,
- or $r' \preceq u$.

In other words, $\preceq$ is a variant of backward simulation with respect to $\rightarrow$. Indeed, if the symbol occurrence involved in $r \rightarrow u$ is uncolored, the transition may be also fired from r'. Otherwise, suppose that the symbol occurrence involved in $r \rightarrow u$ is colored (recall that in this case all fresh symbol occurrences are colored). If this occurrence appears also in r', it may be fired similarly as above. On the other hand, if this occurrence does not appear in r', we have $r' \preceq u$, as required.

Using the last claim we easily show decidability. The algorithm explores exhaustively a portion of the tree of colored configurations reachable from s', with the following termination condition. As the ordering $\preceq$ is a well order, we know that on every path eventually two colored configurations appear, say u' and u, such that u' precedes u and $u' \preceq u$. Such a pair we call *domination pair*. Whenever a domination pair is found on some path, the algorithm stops lengthening this path. The well order guarantees thus that the algorithm terminates, after computing a finite tree of colored configurations (finiteness of the tree follows by König's lemma). The algorithm answers 'yes' if the configuration t appears in the tree.

Now we prove correctness of the algorithm. Towards contradiction, suppose t is reachable from s' but t is not found in the tree. Consider the shortest path π from s' to t, and the domination pair $u' \preceq u$ on that path. Thus $u \rightsquigarrow t$. Using the compatibility condition, we deduce that $u \rightsquigarrow t$ implies $u' \rightsquigarrow t'$ for some $t' \preceq t$, along a path not longer that the path from u to t. By the definition of $\preceq$ we obtain $t' = t$. Thus the fragment of path π from s' to u', composed with the path from u' to t yields a path strictly shorter than π, a contradiction.   □

Theorem 4.7. For strongly normed MPDAs, the backward reachability set of a regular set is regular.

Proof. Consider a strongly normed MPDA, and a regular set L of configurations. Denote by B the backward reachability set of L. We aim at proving that B is regular. We do not claim however that a representation of this regular set may be effectively computed in general.

The proof is in two steps. In the first one we introduce and justify a simplification of the problem. As the second step, we prove regularity of the backward reachability in the simplified case.

Restriction to fully active paths. Let L_q (resp. B_q) denote restrictions of L (resp. B) to configurations with control state q. It is enough to prove regularity of the set

$$B_{p,q} = \{s \in B_p : s \leadsto L_q\}$$

for every pair of states p, q, due to the following equalities:

$$B_p = \bigcup_q B_{p,q} \qquad B = \bigcup_p B_p$$

and closure of regular languages under finite unions. Fix two control states p and q from now on. Denote by S_q the set of all configurations with the control state q.

Recall that the set L_q is recognized by some monoid homomorphism, i.e., there is $h : \mathsf{S}_q \to M$ for some finite monoid M and for every q, $L_q = h^{-1}(N_q)$ for some subset $N_q \subseteq M$. Note that, formally speaking, the state q is a part of every configuration of L_q, hence the domain of the homomorphism are all configurations with state q. Thus formally speaking the homomorphism respects the multiplication of configurations defined as:

$$\langle q, \alpha_1, \ldots, \alpha_k \rangle \cdot \langle q, \beta_1, \ldots, \beta_k \rangle = \langle q, \alpha_1 \beta_1, \ldots, \alpha_k \beta_k \rangle,$$

for k the number of stacks. As L_q is recognized by h, we have

$$L_q = \bigcup_{m,n} h^{-1}(m)\, h^{-1}(n), \tag{7.1}$$

where m and n range over pairs of elements of M satisfying $m \cdot n \in N_q$.

A path from a configuration s to some target configuration is *fully active* if some descendant of every symbol occurrence in s is involved in some transition. For $m \in M$, denote by F_m the set of all configurations s with control state p such that there is a fully active path from s to some configuration t in $h^{-1}(m)$ (the control state of t is necessarily q, by the definition of h). We claim that the set $B_{p,q}$ is equal to the following union:

$$B_{p,q} = \bigcup_{m,n} F_m\, h^{-1}(n),$$

with m, n ranging over all pairs $m, n \in M$ such that $m \cdot n \in N_q$, similarly as in (7.1).

Indeed, if a configuration s belongs to $F_m\, h^{-1}(n)$ for some $m \cdot n \in N_q$, then it can reach a configuration in $h^{-1}(m)\, h^{-1}(n) \subseteq L_q$. On the other side, if some configuration s with control state p can reach a configuration $t \in L_q$, we can split s and t into two parts

$$s = s_1 \cdot u \qquad t = t_1 \cdot u$$

such that

- considering the transitions on the path $s \leadsto t$, they form a fully active path from s_1 to t_1,
- the symbol occurrences appearing in u do not make any transition.

Therefore, for $m = h(t_1)$ we obtain $s_1 \in F_m$. In consequence $s \in F_m\, h^{-1}(n)$ for some $m, n \in M$ such that $m \cdot n \in N_q$, as required.

We have thus demonstrated that we can safely restrict to fully active paths.

The proof for fully active paths. Under the restriction to fully active paths, the proof is fairly easy. Let $\mathcal{A}$ be an MPDA and let L be a regular set of configurations.

The *bottom-fixed Higman ordering* over words relates w' and w if the words have the same last letter (recall that the last letter corresponds to the bottom element of a stack), and the remaining prefixes of w' and w are related by the ordinary Higman ordering. Order the configurations as follows: the order, denoted by $\preceq$, relates two configurations if

- the control states are the same, and
- the stack contents are related by the point-wise extension of the bottom-fixed Higman ordering.

Observe that this order is a well order.

Claim 7.4. Under the restriction to fully active paths, the backward reachability set of a regular set L is upward closed with respect to $\preceq$.

Indeed, assuming $s' \preceq s$ and $s' \rightsquigarrow t \in L$, one deduces $s \rightsquigarrow t$ by applying the canceling sequences. These sequences are applicable to some descendant of every symbol occurrence in s, as the path is fully active. (It is here where the strong normedness assumption is needed.) Note the reason why we have chosen the bottom-fixed variant of Higman ordering: this prevents s to have some additional symbols below some bottom symbol of s'.

By the above claim, the backward reachability set is determined by the minimal configurations with respect to $\preceq$. As $\preceq$ is a well order there is only finitely many minimal configurations, and thus the backward reachability set is regular. $\square$

Theorem 4.1. The REG$\rightsquigarrow$REG reachability is decidable for strongly normed MPDAs.

Proof. We are going to use Theorem 4.7. Consider two regular sets L, K of configurations. The decision procedure for $L \rightsquigarrow K$ runs two partial procedures. The positive semi-procedure searches for a witness for reachability, by an exhaustive enumeration of all finite paths starting in some configuration of L and ending in some configuration of K. The core of the difficulty lies in the negative semi-procedure, to be described below.

The negative semi-procedure searches for a *separator*. Knowing that the backward reachability set of K is regular, we define separator as any regular set M of configurations satisfying the following conditions:

- $K \subseteq M$,
- $L \cap M = \emptyset$,
- M is backward closed: whenever $s \rightarrow M$ then $s \in M$.

Clearly, if $L \not\rightsquigarrow K$ then the backward reachability set of K satisfies all the three conditions above. By Theorem 4.7 we know that the backward reachability set is regular, and thus

$$L \not\rightsquigarrow K \qquad \text{iff} \qquad \text{there is some separator } M.$$

The negative semi-procedure enumerates all regular sets M, and for every of them checks the three conditions above. The first two are easily shown decidable as boolean operations on regular languages yield regular languages and are effectively computable. Decidability of the third condition follows from an observation that if M is regular than the predecessor set $\{s : s \rightarrow M\}$ is also regular, and may be effectively constructed. $\square$

8. Undecidability

Here we provide proofs of Theorems 4.3 and 4.9.

Theorem 4.3. The $1\rightsquigarrow\text{REG}$ reachability is undecidable for stateless unnormed MPDAs and for weak normed MPDAs.

Proof. We start by considering stateless unnormed MPDAs. We reduce the problem of checking if the intersection of two context-free languages is empty.

Assume two context-free grammars in Greibach Normal Form over an input alphabet A. We will construct an MPDA with three stacks. Two stacks will be used to simulate derivations of the two grammars, and the other third stack will be used for storage of the input word. Formally, the alphabet of the first and second stack are the nonterminals of the two grammars, and the alphabet of the third stack contains two symbols a_1 and a_2 for every terminal symbol a of the grammars. For every production

$$X \rightarrow a\,\alpha \tag{8.1}$$

of the first grammar, there will be a transition

$$X \ \rightarrow \ \alpha,\ \varepsilon,\ a_1$$

that drops α on the first stack and a on the third one. Likewise, for every production (8.1) of the second grammar, there is a transition

$$X \ \rightarrow \ \varepsilon,\ \alpha,\ a_2.$$

The initial configuration is $\langle X_1, X_2, \varepsilon \rangle$, where X_i is the initial symbol of the ith grammar. Finally, the regular language L of target configurations constraints the first two stacks to be empty, and the third one to:

$$\{a_1 a_2 : a \in A\}^*.$$

One easily verifies that the intersection of the two grammars is nonempty if and only if some configuration from L is reachable from the initial configuration.

Now we turn to weak normed MPDAs. It turns out that normedness assumption does not make reachability problem easier, in case of weak automata. Indeed, the case of stateless unnormed MPDAs easily reduces to the case of weak normed MPDAs. It is sufficient to add an additional sink state, and for every symbol X two additional transitions, to enforce normedness. The first one allows X to change state to the sink state. The other one allows X to disappear in the sink state. (This is in fact a reduction of the whole case of weak unnormed MPDAs.) □

Theorem 4.9. The $1\rightsquigarrow\text{REG}$ reachability is undecidable for stateless strongly normed MPDAs, under the relaxed notion of regularity.

Proof. The proof is by reduction from the Post Correspondence Problem (PCP). For a given instance of PCP, consisting of a finite set of pairs (s_i, t_i) of words, $i \in \{1\dots n\}$, we construct a stateless strongly normed MPDA $\mathcal{A}$ and a relaxed-regular set L such that the PCP instance has a solution

$$s_{i_1}\, s_{i_2} \dots s_{i_k} \ \ = \ \ t_{i_1}\, t_{i_2} \dots t_{i_k} \qquad (i_j \in \{1\dots n\} \text{ for } j \in \{1\dots k\})$$

if and only if there exists a path from the initial configuration of $\mathcal{A}$ to L. Roughly speaking, a run of $\mathcal{A}$ will simply guess a PCP solution, and the target language L will be used to check its correctness.

The main difficulty to overcome is the strong normedness requirement, which implies that every symbol may always disappear and not contribute to the target configuration.

Half-solution. We start by restricting to only the left-hand side words s_i of the PCP instance. We will construct an MPDA $\mathcal{A}_1$, and a relaxed-regular language L_1 of configurations, so that the reachable configurations of $\mathcal{A}_1$ belonging to L_1 are essentially of the form (two stacks):

$$(i_1 \, i_2 \ldots i_k, \quad s_{i_1} \, s_{i_2} \ldots s_{i_k}). \tag{8.2}$$

In other words, one of stacks contains the sequence of indexes, and the other one contains the concatenation of the corresponding words s_i.

For technical reasons we will however need four stacks and few auxiliary nonterminal symbols. The nonterminals of $\mathcal{A}_1$ are following (superscripts indicate the stack number of every nonterminal):

- G^1 and G^4, used for 'guarding' symbols on their stacks, as described below;
- i^1 and i^2, for $i \in \{1 \ldots n\}$, representing the ith word s_i;
- a^3 and a^4, for $a \in \Sigma$, representing alphabet letters of the PCP instance.

The initial configuration is $(G^1, \, \varepsilon, \, \varepsilon, \, G^4)$.

For a word $w = a_1 a_2 \ldots a_m \in \Sigma^*$, we write w^3 to mean the word $a_1^3 a_2^3 \ldots a_m^3$. Likewise for w^4. The transition rules of $\mathcal{A}_1$ are the following. For $i \in \{1 \ldots n\}$, there are rules:

$$G^1 \;\rightarrow\; G^1 \, i^1, \; \varepsilon, \; s_i^3, \; \varepsilon \qquad\qquad G^4 \;\rightarrow\; \varepsilon, \; i^2, \; \varepsilon, \; G^4 \, s_i^4.$$

Additionally, to fulfill the strong normedness restriction we add *disappearing transition rules* of the form $X \;\rightarrow\; \varepsilon, \, \varepsilon, \, \varepsilon, \, \varepsilon$ for all nonterminal symbols.

The target set L_1 is defined to contain all configurations of the form

$$(G^1 \, \alpha_1, \; \alpha_2, \; \alpha_3, \; G^4 \, \alpha_4),$$

with α_1 almost equal to α_2 and α_3 almost equal to α_4. By 'almost equal' we mean equality modulo (ignoring) the superscripts. The set L_1 is clearly relaxed-regular.

Let's analyze possible ways of reaching a configuration from L_1. Surely G^1 and G^4 cannot fire the disappearing transitions, because their presence is required by L_1. As G^1 and G^4 are always top-most on their stacks, all other symbols on these stacks are 'guarded' – they can not fire a disappearing transition neither. A key observation is that no symbol from other two stacks could fire a disappearing transition:

Lemma 8.1. *Every path from the initial configuration to L contains no disappearing transitions.*

Proof. The precise proof of this fact needs a certain effort. Let us define the *weight* of a nonterminal. The intuition behind this notion is that it counts for how many letters in words s_i the particular nonterminal is responsible. The definition is the following:

- $\text{weight}(G^1) = \text{weight}(G^4) = 0$
- $\text{weight}(i^1) = \text{weight}(i^2) = \text{length}(s_i)$
- $\text{weight}(a^3) = \text{weight}(a^4) = 1$

Weight of a word is defined as the sum of weights of its letters. Note now that any configuration $\alpha = (G_1\,\alpha_1,\ \alpha_2,\ \alpha_3,\ G_4\,\alpha_4)$ reachable from $(G_1,\ \varepsilon,\ \varepsilon,\ G_4)$ satisfies the following inequalities:

$$\begin{aligned}
\mathrm{SInv}_1(\alpha) &= \mathrm{weight}(\alpha_1) - \mathrm{weight}(\alpha_3) \geq 0 \\
\mathrm{SInv}_2(\alpha) &= \mathrm{weight}(\alpha_4) - \mathrm{weight}(\alpha_2) \geq 0.
\end{aligned}$$

To see this it is enough to observe this both semi-invariants SInv_1 and SInv_2 equal 0 in the initial configuration and that they never decrease due to performing a transition. In particular, every disappearing transition on the second or third stack increases one of the semi-invariants. Finally, every configuration $\alpha \in L$ satisfies the equality:

$$\mathrm{SInv}_1(\alpha) + \mathrm{SInv}_2(\alpha) = 0,$$

as $\mathrm{weight}(\alpha_1) = \mathrm{weight}(\alpha_2)$ and $\mathrm{weight}(\alpha_3) = \mathrm{weight}(\alpha_4)$. Therefore both semi-invariants are necessary equal to 0, and thus there is no possibility for disappearing transitions to be fired. □

As a conclusion we obtain:

Corollary 8.2. *Consider any configuration in L_1 that is reachable from the initial configuration, and suppose that its first and forth stacks have the form:*

$$G^1\,i_1^1 \ldots i_k^1 \qquad\qquad G^4\,a_1^4 \ldots a_m^4.$$

Then it holds:

$$s_{i_1} \ldots s_{i_k} = a_1 \ldots a_m.$$

Complete solution. Similarly as above, one may construct an MPDA $\mathcal{A}_2$ and a language L_2 for the right-hand side words t_i from the PCP instance. Essentially (i.e., ignoring the technical details) the reachable configurations of $\mathcal{A}_2$ intersected with L_2 are (cf. (8.2)):

$$(i_1\,i_2 \ldots i_k,\ \ t_{i_1}\,t_{i_2} \ldots t_{i_k}).$$

Our final solution is to appropriately combine both MPDAs and both languages.

The MPDA $\mathcal{A}$ is obtained by merging $\mathcal{A}_1$ and $\mathcal{A}_2$, but the first stacks are identified. Thus $\mathcal{A}$ will have seven stacks altogether. In particular, symbols i^1 and i^2 represent now the ith pair (s_i, t_i). All transitions are exactly as described above, however with a different numbering of stacks. The language L imposes the requirements of L_1 and L_2, and additionally requires that the fourth stack of $\mathcal{A}_1$ is almost equal to the fourth stack of $\mathcal{A}_2$.

For describing the missing details we have to fix a new numbering of stacks. Let the first four stacks correspond to the stacks of $\mathcal{A}_1$, and the remaining three stacks correspond to the stacks of $\mathcal{A}_2$ different than the first one. The initial configuration of $\mathcal{A}$ is

$$(G^1,\ \varepsilon,\ \varepsilon,\ G^4,\ \varepsilon,\ \varepsilon,\ G^7).$$

Except for the disappearing transitions, $\mathcal{A}$ has the following transition rules:

$$\begin{aligned}
G^1 &\rightarrow G^1\,i^1,\ \varepsilon,\ s_i^3,\ \varepsilon,\ \varepsilon,\ t_i^6,\ \varepsilon \\
G^4 &\rightarrow \varepsilon,\ i^2,\ \varepsilon,\ G^4\,s_i^4,\ \varepsilon,\ \varepsilon,\ \varepsilon \\
G^7 &\rightarrow \varepsilon,\ \varepsilon,\ \varepsilon,\ \varepsilon,\ i^5,\ \varepsilon,\ G^7\,t_i^7.
\end{aligned}$$

The language L contains configurations of the form:

$$(G^1\,\alpha_1,\ \alpha_2,\ \alpha_3,\ G^4\,\alpha_4,\ \alpha_5,\ \alpha_6,\ G^7\,\alpha_7)$$

satisfying the following almost equalities:

$$\alpha_1 \;=\; \alpha_2 = \alpha_5 \qquad \alpha_3 \;=\; \alpha_4 = \alpha_6 \;=\; \alpha_7.$$

One can easily observe that L is reachable from the initial configuration if and only if the PCP instance has a solution, using exactly the same techniques as before. $\square$

9. Conclusions

In this paper we have investigated decidability and complexity of the reachability problem for multi-stack pushdown automata. As the model is undecidable in general, we have mostly focused on the case of weak control states. We have provided an almost complete map of complexity results for the reachability problem, depending on the natural restrictions of source and target sets of configurations, and on normedness assumption. Few remaining open cases are listed below:

- The most interesting open question is the exact complexity of $1 \rightsquigarrow 1$ reachability problem for (normed and unnormed) weak MPDAs. We only know that the problem is NP-hard. Note that the REG$\rightsquigarrow$1 reachability has the same complexity.
- Another interesting question is whether three stacks are necessary for undecidability of $1 \rightsquigarrow$REG reachability problem for normed weak MPDA. In other words, we do not know the status of the problem when there are just two stacks.
- A similar question is still open for the cases when reachability is in NP: is the problem still NP-complete if one restricts the number of stacks to two?
- We do not know exact complexity of reachability for strongly normed MPDA.
- Finally, we do not know whether the construction of backward reachability set of a regular set for strongly normed MPDAs, even in the weak case, could be performed effectively.

Acknowledgements. We are grateful to anonymous reviewers for careful reading and many valuable comments.

References

[1] M. Faouzi Atig. From multi to single stack automata. In *CONCUR*, pages 117–131, 2010.
[2] M. Faouzi Atig and A. Bouajjani. On the reachability problem for dynamic networks of concurrent pushdown systems. In *RP*, pages 1–2, 2009.
[3] M. Faouzi Atig, A. Bouajjani, and S. Qadeer. Context-bounded analysis for concurrent programs with dynamic creation of threads. *Logical Methods in Computer Science*, 7(4), 2011.
[4] Pablo Barceló, Diego Figueira, and Leonid Libkin. Graph logics with rational relations and the generalized intersection problem. In *LICS*, pages 115–124, 2012.
[5] J. Berstel. *Transductions and Context-Free Languages*. Teubner Verlag, 1979.
[6] A. Bouajjani and M. Emmi. Analysis of recursively parallel programs. In *POPL*, pages 203–214, 2012.
[7] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR*, pages 135–150, 1997.
[8] A. Bouajjani, J. Esparza, S. Schwoon, and J. Strejcek. Reachability analysis of multithreaded software with asynchronous communication. In *Software Verification: Infinite-State Model Checking and Static Program Analysis*, 2006.
[9] A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *CONCUR*, pages 473–487, 2005.
[10] O. Burkart, D. Caucal, F Moller, and B. Steffen. Verification of infinite structures. In *Handbook of Process Algebra*, pages 545–623. Elevier, 2001.

[11] R. Chadha, P. Madhusudan, and M. Viswanathan. Reachability under contextual locking. In *TACAS*, pages 437–450, 2012.

[12] S. Christensen. *Decidability and Decomposition in process algebras.* PhD thesis, Dept. of Computer Science, University of Edinburgh, UK, 1993.

[13] W. Czerwiński. *Partially-commutative context-free graphs.* PhD thesis, Dept. of Mathematics, Informatics and Mechanics, University of Warsaw, 2013.

[14] W. Czerwiński, S. Fröschle, and S. Lasota. Partially-commutative context-free processes. In *Proc. CONCUR'09*, volume 5710 of *LNCS*, pages 259–273. Springer-Verlag, 2009.

[15] W. Czerwiński, S. Fröschle, and S. Lasota. Partially-commutative context-free processes: expressibility and tractability. *Information and Computation*, 209:782–798, 2011.

[16] W. Czerwiński, P. Hofman, and S. Lasota. Reachability problem for weak multi-pushdown automata. In *CONCUR*, pages 53–68, 2012.

[17] W. Czerwiński and S. Lasota. Partially-commutative context-free languages. In *DCM*, pages 35–48, 2012.

[18] V. Diekert and G. Rozenberg. *The book of traces.* World Scientific, 1995.

[19] J. Esparza. Petri nets, commutative context-free grammars, and basic parallel processes. *Fundam. Inform.*, 31(1):13–25, 1997.

[20] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001.

[21] Dung T. Huynh. Commutative grammars: The complexity of uniform word problems. *Information and Control*, 57(1):21–39, 1983.

[22] V. Kahlon. Reasoning about threads with bounded lock chains. In *CONCUR*, pages 450–465, 2011.

[23] V. Kahlon, F. Ivancic, and A. Gupta. Reasoning about threads communicating via locks. In *CAV*, pages 505–518, 2005.

[24] M. Kretínský, V. Rehák, and J. Strejcek. Reachability is decidable for weakly extended process rewrite systems. *Inf. Comput.*, 207(6):671–680, 2009.

[25] A. Lal and T. W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. *Formal Methods in System Design*, 35(1):73–97, 2009.

[26] D. Lugiez and Ph. Schnoebelen. The regular viewpoint on PA-processes. *Theor. Comput. Sci.*, 274(1-2):89–115, 2002.

[27] R. Mayr. Process rewrite systems. *Inf. Comput.*, 156(1-2):264–286, 2000.

[28] S. Qadeer and J. Rehof. Context-bounded model checking of concurrent software. In *TACAS*, pages 93–107, 2005.