Reachability Analysis of Pushdown Automata: Application to Model-Checking

Ahmed Bouajjani [*‡]   Javier Esparza[†‡]   Oded Maler [*‡]

Abstract

We apply the *symbolic* analysis principle to pushdown systems. We represent (possibly infinite) sets of configurations of such systems by means of finite-state automata. In order to reason in a uniform way about analysis problems involving both existential and universal path quantification (like model-checking for branching-time logics), we consider the more general class of *alternating* pushdown systems and use *alternating* finite-state automata as a representation structure for their sets of configurations. We give a simple and natural procedure to compute sets of predecessors for this representation structure. We apply this procedure and the automata-theoretic approach to model-checking to define new model-checking algorithms for pushdown systems and both linear and branching-time properties. From these results we derive upper bounds for several model-checking problems, and we also provide matching lower bounds, using reductions based on some techniques introduced by Walukiewicz.

1 Introduction

Systems are commonly modeled by various types of transition systems, including finite automata, pushdown automata, Petri nets, timed or hybrid automata, etc. In this framework, most of the system analysis problems (model-checking, synthesis) reduce to (various kinds of) "reachability problems" on these models. It is therefore fundamental for system analysis to develop algorithms that compute the set of *all predecessors* of a given set of states S, i.e., the set of states from which it is possible to reach S.

Let $pre(S)$ denote the set of immediate predecessors of the set S, and let $pre^*(S)$ denote the set of all its predecessors. Clearly, $pre^*(S)$ is the limit of the *infinite* increasing sequence $\{X_i\}_{i\geq 0}$ given by $X_0 = S$ and $X_{i+1} = X_i \cup pre(X_i)$ for every $i \geq 0$.

In the case of finite-state systems, the sets X_i are all finite, and the sequence $\{X_i\}_{i\geq 0}$ is guaranteed to reach a fixpoint, which immediately provides an (usually inefficient, but terminating) algorithm to compute $pre^*(S)$. Unfortunately, these properties no longer hold for any interesting class of infinite-state systems. Our first task is then to find a class of *finite* structures that can represent the infinite sets of states we are interested in. Since boolean combinations of interesting infinite sets of states are usually interesting,

[*]VERIMAG, Centre Equation, 2 avenue de Vignate, 38610 Gieres, France.
[†]Institut für Informatik, Technische Universität München, Arcisstr. 21, 81539 München, Germany
[‡]Ahmed.Bouajjani@imag.fr, esparza@informatik.tu-muenchen.de, Oded.Maler@imag.fr

the class should be closed under boolean operations. Since we wish to check if a given state (for instance the initial state) belongs to an infinite set, the membership problem of the class should be decidable. Once the class has been found, it remains to show that it is (effectively) closed under the pre^* function.

Several instances of systems and corresponding representation structures have been considered in the litterature. For example, in the case of timed automata, special kinds of polyhedra (regions) are used to represent infinite sets of states (vectors of reals corresponding to clock valuations) [AD94]. Polyhedra are also used for linear hybrid systems. However, in this case, there is no algorithm for computing a finite representation of the *exact* set of predecessors (the reachability problem is undecidable), but upper approximations of this set can be calculated [ACH+95]. In [BG96], representation structures called QDD's are introduced for fifo-channel systems. These structures are finite-state automata representing sets of queue contents. As in the case of linear hybrid systems, the procedure for calculating the set of predecessors for these structures is not guaranteed to terminate. Finally, notice that the idea of using representations of sets of states is also used in the finite-state case to overcome the state-explosion problem [McM93]. The usual representation structures in this case are BDD's (binary decision diagrams) [Bry92].

In this paper we consider pushdown systems, and also the more general class of *alternating pushdown systems*, i.e., pushdown systems with both existential and universal nondeterminism (see [Var95] for a survey on alternating automata). This general setting allows to reason in a uniform way about *analysis* problems where existential and universal path quantification must be considered, like model-checking for branching-time temporal logics (see Section 5) and also about *synthesis* problems, like finding winning strategies for 2-player games (see [AMP95]).

A state (we use rather the word "configuration") of a pushdown system is a pair $\langle p, w\rangle$ where p is control location and w is a sequence of stack symbols (w is the stack contents). We propose as a representation structure for sets of configurations the *alternating multi-automaton* (AMA), an alternating finite-state automaton with one initial state for each control location. The automaton recognizes the configuration $\langle p, w\rangle$ if it accepts the word w from the initial state corresponding to p. It is important to remember that an AMA is just a tool to represent a set of configurations, and not to confuse its behaviour with the behaviour of the pushdown system.

It is easy to show that AMA's are closed under boolean operations, and its membership problem is trivially decidable. Our main result is a simple and natural algorithm to compute the pre^* function. As an application, we use it to construct elegant model-checking algorithms for pushdown systems and both linear and branching-time temporal logics. More precisely, we show how to construct AMA's accepting the set of *all* configurations satisfying ω-regular properties of linear-time temporal logics (including all properties expressible in LTL [Pnu77] or the linear-time μ-calculus [Var88]), or properties expressed as formulas of the alternation-free modal μ-calculus.

Our approach also allows us to obtain a number of complexity results: we show that the model-checking problems mentioned above are in DEXPTIME, and that the model-checking problem for pushdown systems and a subset of CTL [CES83] can be solved in PSPACE. Using a technique due to Walukiewicz [Wal96], we complement these results with matching lower bounds, i.e., we show that all these problems are complete for their corresponding complexity classes.

The paper is structured as follows. In Section 2, we give an algorithm which computes the pre^* function for pushdown systems. In this case, the used representation structure is a simple (OR-)multi-automaton (i.e., alternation is not needed). We apply this algorithm in Section 3 to the model-checking problem for linear-time logics. In Section 4, we generalize the algorithm given in Section 2 to alternating pushdown systems. In that case, we use the alternating multi-automaton representation structure. In Section 5, we apply the new algorithm to the model-checking problem for branching-time logics. The appendix contains all the proofs.

2 Reachability in pushdown systems

A pushdown system (PDS for short) is a triplet $\mathcal{P} = (P, \Gamma, \Delta)$ where P is a finite set of *control locations*, Γ is a finite *stack alphabet*, and $\Delta \subseteq (P \times \Gamma) \times (P \times \Gamma^*)$ is a set of *transition rules.* If $((q, \gamma), (q', w)) \in \Delta$ then we write $(q, \gamma) \hookrightarrow (q', w)$ (we reserve $\rightarrow$ to denote the transition relations of finite automata).

Notice that PDS's have no input alphabet. We do not use them as language acceptors but are rather interested in the behaviours they generate.

A *configuration* of $\mathcal{P}$ is a pair $\langle p, w\rangle$ where $p \in P$ is a control location and $w \in \Gamma^*$ is a *stack content.*

If $(q, \gamma) \hookrightarrow (q', w)$, then for every $w' \in \Gamma^*$ the configuration $\langle q, \gamma w'\rangle$ is an *immediate predecessor* of $\langle q', ww'\rangle$, and $\langle q', ww'\rangle$ an *immediate successor* of $\langle q, \gamma w'\rangle$. The *reachability relation* $\Rightarrow$ is the reflexive and transitive closure of the immediate successor relation. A *run* of $\mathcal{P}$ is a maximal sequence of configurations such that for each two consecutive configurations $c_i c_{i+1}$, c_{i+1} is an immediate successor of c_i. The set of all runs of $\mathcal{P}$ is denoted by $Runs_{\mathcal{P}}$.

The predecessor function $pre_{\mathcal{P}} : 2^{P\times\Gamma^*} \rightarrow 2^{P\times\Gamma^*}$ is defined as follows: c belongs to $pre_{\mathcal{P}}(C)$ if some immediate successor of c belongs to C. The reflexive and transitive closure of $pre_{\mathcal{P}}$ is denoted by $pre^*_{\mathcal{P}}$. Clearly, $pre^*_{\mathcal{P}}(C) = \{c \in P \times \Gamma^* \mid \exists c' \in C.\ c \Rightarrow c'\}$. We denote by $pre^+_{\mathcal{P}}$ the function $pre_{\mathcal{P}} \circ pre^*_{\mathcal{P}}$. We will omit the subscript $\mathcal{P}$ and write simply pre, pre^*, and pre^+ when it is clear from the context which system is under consideration.

2.1 Multi-automata

Let $\mathcal{P} = (P, \Gamma, \Delta)$ be a pushdown system where $P = \{p^1, \ldots p^m\}$. A $\mathcal{P}$*-multi-automaton* ($\mathcal{P}$-MA for short, or just MA when $\mathcal{P}$ is clear from the context) is a tuple $\mathcal{A} = (\Gamma, Q, \delta, I, F)$ where Q is a finite set of *states*, $\delta \subseteq Q \times \Gamma \times Q$ is a set of *transitions*, $I = \{s^1, \ldots s^m\} \subseteq Q$ is a set of *initial states* and $F \subseteq Q$ is a set of *final states.*

We define the *transition relation* $\longrightarrow \subseteq Q \times \Gamma^* \times Q$ as the smallest relation satisfying:

- if $(q, \gamma, q') \in \delta$ then $q \xrightarrow{\gamma} q'$,
- $q \xrightarrow{\varepsilon} q$ for every $q \in Q$, and
- if $q \xrightarrow{w} q''$ and $q'' \xrightarrow{\gamma} q'$ then $q \xrightarrow{w\gamma} q'$.

$\mathcal{A}$ *accepts* or *recognizes* a configuration $\langle p^i, w\rangle$ if $s^i \xrightarrow{w} q$ for some $q \in F$. The set of configurations recognized by $\mathcal{A}$ is denoted by $Conf(\mathcal{A})$. A set of configurations is *regular* if it is recognized by some MA.

An w*-run* of $\mathcal{A}$, where $w = \gamma_1 \ldots \gamma_n \in \Gamma^*$, is a sequence $s^i \xrightarrow{\gamma_1} q_1 \ldots \xrightarrow{\gamma_n} q_n$.

2.2 Calculating pre^*

Fix a pushdown system $\mathcal{P} = (P, \Gamma, \Delta)$ where $P = \{p^1, \ldots, p^m\}$. We show in this section that given a regular set of configurations C of $\mathcal{P}$ recognized by a MA $\mathcal{A}$, we can construct another MA $\mathcal{A}_{pre^*}$ recognizing $pre^*(C)$.

By definition, $pre^*(C) = \bigcup_{i \geq 0} X_i$ with $X_0 = C$ and $X_{i+1} = X_i \cup pre(X_i)$ for every $i \geq 0$. Therefore, one may try to calculate $pre^*(C)$ by iteratively constructing the increasing sequence $X_0, X_1, \ldots$. If $X_{i+1} = X_i$ holds for some $i \geq 0$, then it is clear that $X_i = pre^*(C)$.

However, the existence of such a fixed point is not guaranteed in general, and we may never reach the limit of the X_i sequence. Consider for instance the PDS with one state p, one stack symbol γ, and one transition rule $(p, \gamma) \hookrightarrow (p, \varepsilon)$, and take $C = \{\langle p, \varepsilon \rangle\}$. Clearly, we have $X_i = \{\langle p, \varepsilon \rangle, \langle p, \gamma \rangle, \ldots, \langle p, \gamma^i \rangle\}$ and so $X_{i+1} \neq X_i$ for every $i \geq 0$.

To overcome this problem, we calculate $pre^*(C)$ differently, as the limit of another increasing sequence of sets of configurations $Y_0, Y_1, \ldots$ for which we can prove the following properties:

P1. $\exists i \geq 0.\ Y_{i+1} = Y_i$,

P2. $\forall i \geq 0.\ X_i \subseteq Y_i$,

P3. $\forall i \geq 0.\ Y_i \subseteq \bigcup_{j \geq 0} X_j$.

Property (P1) ensures termination of the procedure that computes the sequence of Y_i's. Property (P2) ensures that, by calculating the limit of the Y_i's, we capture (at least) the whole set $pre^*(C)$, and property (P3) ensures that only elements of $pre^*(C)$ are captured.

The Y_i's are formally defined as the sets of configurations recognized by a sequence $\mathcal{A}_0, \mathcal{A}_1 \ldots$ of MA's satisfying for every $i \geq 0$ the following property: $\mathcal{A}_{i+1}$ has *the same states* as $\mathcal{A}_i$, and the same or more transitions. Since an MA with n states and m input symbols can have at most $n^2 \cdot m$ transitions, the Y_i's must converge to a fixpoint.[1]

We start with a MA $\mathcal{A}$ recognizing the regular set of configurations C. We assume without loss of generality that $\mathcal{A}$ has no transition leading to an initial state (every MA can be converted to one having this property). We take $\mathcal{A}_0 = \mathcal{A}$. We denote by $\rightarrow_i$ the transition relation of $\mathcal{A}_i$. For every $i \geq 0$, $\mathcal{A}_{i+1}$ is obtained from $\mathcal{A}_i$ by conserving the same states and transitions, and adding for every transition rule $(p^j, \gamma) \hookrightarrow (p^k, w)$ and every state q such that $s^k \xrightarrow{w}_i q$ a new transition $s^j \xrightarrow{\gamma}_{i+1} q$. Then, for every $i \geq 0$ we define $Y_i = Conf(\mathcal{A}_i)$.

Notice that the new transitions added to $\mathcal{A}_i$ in order to construct $\mathcal{A}_{i+1}$ start at initial states.

To understand the idea behind this construction, observe that $\langle p^k, \gamma w' \rangle$ is an immediate predecessor of $\langle p^j, ww' \rangle$ by the rule $(p^j, \gamma) \hookrightarrow (p^k, w)$. So, if the word ww' is accepted starting from s^k by $\mathcal{A}_i$ ($s^k \xrightarrow{w}_i q \xrightarrow{w'}_i q' \in F$), then the new transition in $\mathcal{A}_{i+1}$ allows to accept $\gamma w'$ starting from s^j ($s^j \xrightarrow{\gamma}_{i+1} q \xrightarrow{w'}_i q' \in F$). Let us illustrate the construction by means of an example.

Let $\mathcal{P}$ be the PDS such that $P = \{p^1, p^2\}$, $\Gamma = \{\gamma_1, \ldots, \gamma_6\}$, and Δ contains the rules

$$(p^2, \gamma_4) \hookrightarrow (p^2, \gamma_1\gamma_2) \qquad (p^1, \gamma_5) \hookrightarrow (p^2, \gamma_4\gamma_3) \qquad (p^1, \gamma_6) \hookrightarrow (p^1, \varepsilon)$$

[1] The idea is inspired by the construction given in [BO93], pages 91-93, of a finite-state automaton recognizing the closure of a regular language under the rewriting relation induced by a *monadic string-rewriting system*.

Consider the set of configurations $C = \{\langle p^2, \gamma_1\gamma_2\gamma_3\rangle\}$. It can be represented by a MA $\mathcal{A}$ such that $Q = \{s^1, s^2, q_1, q_2, q_3\}$, $I = \{s^1, s^2\}$, $F = \{q_3\}$, and δ contains the transitions $s^2 \xrightarrow{\gamma_1} q_1$, $q_1 \xrightarrow{\gamma_2} q_2$, and $q_2 \xrightarrow{\gamma_3} q_3$.

The picture below shows the automaton $\mathcal{A}_{pre^*}$ obtained at the end of the construction.

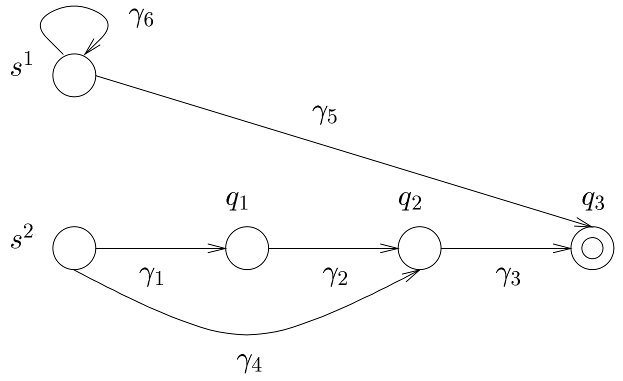

In the first step (from $\mathcal{A}_0$ to $\mathcal{A}_1$) we have $s^2 \xrightarrow{\gamma_1\gamma_2}_0 q_2$, and so we add the transition $s^2 \xrightarrow{\gamma_4}_1 q_2$. Notice that the new automaton now accepts all immediate predecessors of $\langle p^2, \gamma_1\gamma_2\gamma_3\rangle$, namely the configuration $\langle p^2, \gamma_4\gamma_3\rangle$. No other transitions are added.

In the second step, we add the transition $s^1 \xrightarrow{\gamma_5}_2 q_3$, corresponding to the second transition rule of $\mathcal{P}$.

In the third step, we add the transition $s^1 \xrightarrow{\gamma_6}_3 s^1$. At this point the construction stops since no further transition must be added. So, we have $\mathcal{A}_{pre^*} = \mathcal{A}_3$, and

$$pre^*(C) = (\{p^1\} \times \gamma_6^*\gamma_5) \cup (\{p^2\} \times \{\gamma_1\gamma_2\gamma_3, \gamma_4\gamma_3\})$$

Observe that in this example we have $X_1 = Y_1$ and $X_2 = Y_2$, but $X_3 \subset Y_3$. Indeed, in the third step of the construction, after adding $s^1 \xrightarrow{\gamma_6}_3 s^1$, $\mathcal{A}_3$ accepts all the configurations of the form $\langle p^1, \gamma_6^k\gamma_5\rangle$ for every $k \geq 1$, whereas only $\langle p^1, \gamma_6\gamma_5\rangle$ belongs to X_3. However, despite the fact that these configurations are not immediate predecessors of X_2 configurations, they are all in $pre^*(C)$ because $\langle p^1, \gamma_6^k\gamma_5\rangle \in X_{k+2}$ for every $k \geq 1$.

The proofs of the properties (P1), (P2), and (P3) are given in the appendix. We deduce from these properties the following theorem.

Theorem 2.1 *Given a PDS $\mathcal{P}$ and a regular set of configurations recognized by a $\mathcal{P}$-MA $\mathcal{A}$, we can construct a $\mathcal{P}$-MA $\mathcal{A}_{pre^*}$ recognizing $pre^*(Conf(\mathcal{A}))$.*

We finish with a remark on complexity. Define the size of a PDS as the length of a list containing its states, its stack symbols and its production rules. Define the size of an MA as the length of a list containing its states, its alphabet, its production rules, its initial state and its final states. It follows easily from the facts below that the complexity of the algorithm for the construction of $\mathcal{A}_{pre^*}$ is polynomial on the sizes of $\mathcal{P}$ and $\mathcal{A}$:

- $\mathcal{A}_{pre^*}$ has the same states as $\mathcal{A}$,
- a $\mathcal{P}$-MA with k states has $O(n_{\mathcal{P}} \cdot k^2)$ transitions, where $n_{\mathcal{P}}$ is the size of $\mathcal{P}$, and
- during the construction of the sequence $\mathcal{A}_0, \mathcal{A}_i$, polynomial time suffices to decide if a new transition can be added to the current automaton.

3 Model-Checking Linear-Time Temporal Logics

Let $Prop$ be a finite set of atomic propositions, and let $\Sigma = 2^{Prop}$. It is well known that the semantics of properties expressed in linear time temporal logics like LTL or the linear-time μ-calculus are ω-regular sets over the alphabet Σ. Moreover, there exist algorithms which construct Büchi automata to recognize them [VW86, Var95]. This is all we need to know about these logics in this paper in order to give model checking algorithms for PDS's.

Let $\mathcal{P} = (P, \Gamma, \Delta)$ be a PDS, and let $\Lambda \colon P \to \Sigma$ be a labelling function, which intuitively associates to a control location p the set of propositions that are true of it. Given a formula φ of such an ω-regular logic we wish to solve the following problem:

> compute the set of *all* configurations c of $\mathcal{P}$ such that every run starting from c satisfies φ (via the labelling function Λ).

Then, the model checking problem consists in checking whether a given initial configuration belongs to this set of configurations.

We start by constructing a Büchi automaton $\mathcal{B}$ corresponding to the negation of φ. The product of the PDS $\mathcal{P}$ and this Büchi automaton yields a Büchi PDS $\mathcal{BP}$ with a set of *repeating control locations* G.

Then, the original problem reduces to the following *accepting run problem*:

> compute the set $\mathcal{C}$ of configurations c of $\mathcal{BP}$ such that $\mathcal{BP}$ has an accepting run starting from c (i.e., a run which visits infinitely often configurations with control locations in G).

Notice that the emptiness problem of Büchi PDS's (whether the initial configuration has an accepting run) reduces to the accepting states generation problem via the membership problem of MA (finite state automata).

The following proposition shows that the accepting run problem of Büchi PDS's can be reduced to a reachability problem:

Proposition 3.1 *Let c be a configuration of a Büchi PDS $\mathcal{BP}$. $\mathcal{BP}$ has an accepting run starting from c if and only if there exist configurations $\langle p, \gamma\rangle$, $\langle g, u\rangle$, and $\langle p, \gamma v\rangle$, not all three equal, such that $g \in G$ and:*

(1) $c \Rightarrow \langle p, \gamma w\rangle$ *for some* $w \in \Gamma^*$, *and*

(2) $\langle p, \gamma\rangle \Rightarrow \langle g, u\rangle \Rightarrow \langle p, \gamma v\rangle$.

The proof of Proposition 3.1 is given in the appendix.

We can reformulate conditions (1) and (2) of Proposition 3.1 as follows:

(1′) $c \in pre^*(\{p\} \times \gamma\Gamma^*\}$, and

(2′) $\langle p, \gamma\rangle \in pre^+((G \times \Gamma^*) \cap pre^*(\{p\} \times \gamma\Gamma^*))$.

Since $G \times \Gamma^*$ and $\{p\} \times \gamma\Gamma^*$ are regular sets, we can use Theorem 2.1 to construct MA's recognizing the sets $pre^*(\{p\} \times \gamma\Gamma^*)$ and $pre^+((G \times \Gamma^*) \cap pre^*(\{p\} \times \gamma\Gamma^*))$ (for pre^+ we need to define for an MA $\mathcal{A}$ another MA recognizing $pre(Conf(\mathcal{A}))$, which is a simple exercise). Therefore, by Proposition 3.1, we can construct an MA which recognizes the set of all configurations having an accepting run: First, we determine all the configurations $\langle p, \gamma \rangle$ (there are finitely many of them) for which $(2')$ holds, and then we construct an MA recognizing the union of the sets $pre^*(\{p\} \times \gamma\Gamma^*)$ for all such pairs.

The MA's for the sets $G \times \Gamma^*$ and $\{p\} \times \gamma\Gamma^*$ have polynomial size in the size of the Büchi PDS. Then, since the computation of $pre^*_{\mathcal{P}}(Conf(\mathcal{A}))$ for an MA $\mathcal{A}$ takes polynomial time in the size of $\mathcal{P}$ and the number of states of $\mathcal{A}$, we deduce the following result:

Theorem 3.1 *The accepting run problem of Büchi PDS's can be solved in polynomial time.*

Since the membership problem of MA's (finite state automata) can be solved in linear time, a consequence of Theorem 3.1 is that the emptiness problem of Büchi PDS's can also be solved in polynomial time.

We are now able to state our result about the model checking problem for PDS's and linear-time logics.

Theorem 3.2 *The model checking problem for LTL or the linear-time μ-calculus and PDS's is DEXPTIME-complete. The model checking problem for a fixed formula is polynomial in the size of the PDS.*

Proof: Let us first prove membership in DEXPTIME. Let $\mathcal{P}$ be a PDS of size $n_{\mathcal{P}}$ and φ a formula of length n_φ. It is well known that it is possible to construct a Büchi automaton $\mathcal{B}$ for the negation of φ having exponential size in n_φ, and this construction can be done in exponential time [VW86, Var88]. Hence, the product of $\mathcal{P}$ and $\mathcal{B}$ has polynomial size in $n_{\mathcal{P}}$ and exponential size on n_φ. Applying Theorem 3.1 we obtain an exponential time bound. If the formula φ is fixed, then we have a polynomial algorithm in $n_{\mathcal{P}}$.

To prove hardness, we use a reduction from the problem of deciding whether a given linearly bounded alternating Turing machine accepts a given input or not. The details of the reduction are given in the appendix. $\square$

The model-checking problem for LTL or the linear-time μ-calculus and finite-state systems is known to be PSPACE-complete, but polynomial in the size of the system. Since the properties of systems one wishes to check can be usually encoded into short formulas, model-checkers based on linear-time logics, like SPIN [Hol94], have proved to be useful in practice. Theorem 3.2 shows that the complexity of model-checking for PDS's is worse than the complexity for finite-state systems, but not much worse: it remains polynomial in the size of the system.

4 Reachability in Alternating Pushdown Systems

We consider now the problem of computing the set of predecessors of a regular set of configurations of an *alternating* pushdown system. We show that this set is also regular, and we give a procedure constructing a representation of it by means of *alternating* finite-state multi-automata. For that, we generalize the technique described in the Section 2. The construction we give is used in the model checking algorithms for branching-time logics given in the next section.

An *alternating pushdown system* (APDS for short) is a triplet $\mathcal{P} = (P, \Gamma, \Delta)$, where P and Γ are as for PDSs, and Δ is a function that assigns to each element of $P \times \Gamma$ a positive (i.e, negation-free) boolean formula over elements of $(P \times \Gamma^*)$. We assume that boolean formulae are always in disjunctive normal form, which allows us to equivalently define Δ as a subset of the set $(P \times \Gamma) \times 2^{(P \times \Gamma^*)}$ of *transition rules*: for example, instead of writing

$$\Delta(p, \gamma) = ((p_1, w_1) \vee (p_2, w_2)) \wedge (p_3, w_3)$$

we write

$$\{ \ ((p, \gamma), \{(p_1, w_1), (p_3, w_3)\}) \ , \ \ ((p, \gamma), \{(p_2, w_2), (p_3, w_3)\}) \ \} \subseteq \Delta$$

or just

$$(p, \gamma) \hookrightarrow \{(p_1, w_1), (p_3, w_3)\} \ , \ \ (p, \gamma) \hookrightarrow \{(p_2, w_2), (p_3, w_3)\}$$

If $(p, \gamma) \hookrightarrow \{(p_1, w_1), \ldots, (p_n, w_n)\}$, then for every $w \in \Gamma^*$ the configuration $\langle p, \gamma w\rangle$ is an *immediate predecessor* of the set $\{\langle p_1, w_1 w\rangle, \ldots, \langle p_n, w_n w\rangle\}$, and this set is an *immediate successor* of $\langle p, \gamma w\rangle$. Intuitively, at the configuration $\langle p, \gamma w\rangle$ the APDS selects nondeterministically a transition rule of the form $(p, \gamma) \hookrightarrow \{(p_1, w_1), \ldots, (p_n, w_n)\}$, and forks into n copies in the configurations $\langle p_1, w_1 w\rangle, \ldots, \langle p_n, w_n w\rangle$.

A *run* of $\mathcal{P}$ for an initial configuration c is a tree of configurations with root c, such that the children of a node c' are the configurations that belong to one of its immediate successors (nodes of the form $\langle p, \varepsilon\rangle$ have no successors).

We define the *reachability relation* $\Rightarrow \subseteq (P \times \Gamma^*) \times 2^{P \times \Gamma^*}$ between configurations and sets of configurations. Informally, $c \Rightarrow C$ if and only if C is a finite frontier (finite maximal set of incomparable nodes) of a run of $\mathcal{P}$ starting from c. Formally, $\Rightarrow$ is the smallest subset of $(P \times \Gamma^*) \times 2^{P \times \Gamma^*}$ such that:

1. $c \Rightarrow \{c\}$ for every $c \in P \times \Gamma^*$,
2. if c is an immediate predecessor of C, then $c \Rightarrow C$,
3. if $c \Rightarrow \{c_1, \ldots, c_n\}$ and $c_i \Rightarrow C_i$ for each $1 \leq i \leq n$, then $c \Rightarrow (C_1 \cup \ldots \cup C_n)$.

The function $pre_{\mathcal{P}} \colon 2^{P \times \Gamma^*} \to 2^{P \times \Gamma^*}$ is now defined as follows: c belongs to $pre_{\mathcal{P}}(C)$ if some immediate successor of c is contained in C (observe that the immediate successor of c is now a set!). We denote by $pre^*_{\mathcal{P}}$ the transitive closure of $\lambda C.\ (C \cup pre_{\mathcal{P}}(C))$, i.e., given a set of configurations C, $pre^*_{\mathcal{P}}(C) = \bigcup_{i \geq 0} X_i$, where $X_0 = C$ and $X_{i+1} = X_i \cup pre_{\mathcal{P}}(X_i)$, for every $i \geq 0$. As in the case of PDS's, $pre^*_{\mathcal{P}}(C) = \{c \in P \times \Gamma^* \mid \exists C' \subseteq C.\ c \Rightarrow C'\}$.

4.1 Alternating multi-automata

Fix an APDS $\mathcal{P} = (P, \Gamma, \Delta)$. An *alternating $\mathcal{P}$-multi-automaton* ($\mathcal{P}$-AMA for short, or just AMA when $\mathcal{P}$ is clear from the context) is a tuple $\mathcal{A} = (\Gamma, Q, \delta, I, F)$ which differs from an MA only in the nature of δ. δ is now a function that assigns to every pair of $Q \times \Gamma$ a positive boolean formula with Q as set of variables. As in the case of APDSs, we can equivalently represent δ as a set of transitions, which are elements of $(Q \times \Gamma) \times 2^Q$.

The *transition relation* $\rightarrow \subseteq Q \times \Gamma^* \times 2^Q$ is the smallest relation satisfying

- if $(q, \gamma, Q') \in \delta$ then $q \xrightarrow{\gamma} Q'$,

- $q \xrightarrow{\varepsilon} \{q\}$ for every $q \in Q$,
- if $q \xrightarrow{w} \{q_1, \ldots, q_n\}$ and $q_i \xrightarrow{\gamma} Q_i$ for each $1 \leq i \leq n$, then $q \xrightarrow{w\gamma} (Q_1 \cup \ldots \cup Q_n)$.

A configuration $\langle p^i, w \rangle$ is *recognized* by $\mathcal{A}$ if $s^i \xrightarrow{w} Q'$ for some $Q' \subseteq F$. Given a finite sequence $w \in \Gamma^*$ and a state $q \in Q$, a *run* of $\mathcal{A}$ over w starting from q is a finite tree whose nodes are labelled by states in Q and whose edges are labelled by symbols in Γ, such that the root is labelled by q, and the labelling of the other nodes is consistent with δ. Notice that in such a tree each sequence of edges going from the root to the leaves is labelled by w, and hence, all the edges starting at the same level of the tree have the same label, and all the leaves of the tree are at the same height.

It is immediate to show that AMA's are closed under boolean operations. We mention also that the membership problem for AMA's can be solved in polynomial time.

4.2 Calculating pre^*

Let $\mathcal{P} = (P, \Gamma, \Delta)$ be an alternating pushdown system. We show in this section that given a regular set of configurations C of $\mathcal{P}$, recognized by an *alternating*-multi-automaton $\mathcal{A}$, we can construct another AMA $\mathcal{A}_{pre^*}$ such that $Conf(\mathcal{A}_{pre^*}) = pre^*(C)$.

The construction is very similar to that of the non-alternating case. We assume without loss of generality that no transition of $\mathcal{A}$ leads to a set of states containing an initial state. We define a sequence of AMA's $\mathcal{A}_0, \mathcal{A}_1, \ldots$ such that $\mathcal{A}_0 = \mathcal{A}$. For every $i \geq 0$, $\mathcal{A}_{i+1}$ is obtained from $\mathcal{A}_i$ by conserving the same states and transitions, and adding for every transition rule

$$\langle p^j, \gamma \rangle \hookrightarrow \{\langle p^{k_1}, w_1 \rangle, \ldots, \langle p^{k_m}, w_m \rangle\}$$

and every set

$$s^{k_1} \xrightarrow{w_1}_i P_1 \ , \ \ldots \ , \ s^{k_m} \xrightarrow{w_m}_i P_m$$

a new transition

$$s^j \xrightarrow{\gamma}_{i+1} (P_1 \cup \ldots \cup P_m)$$

Then, define $Y_i = Conf(\mathcal{A}_i)$ for every $i \geq 0$.

The intuitive justification of the construction is that we add the configuration $\langle p^j, \gamma w \rangle$ to the set of predecessors of C whenever all the configurations $\langle p^{k_1}, w_1 w \rangle, \ldots, \langle p^{k_m}, w_m w \rangle$ are already in this set. So, if for every $\ell \in \{1, \ldots, m\}$, the word $w_i w$ is accepted by $\mathcal{A}_i$ starting from s^{k_ℓ}, which means that $s^{k_\ell} \xrightarrow{w_\ell}_i P_\ell$ and $\forall p \in P_\ell.\ p \xrightarrow{w}_i Q_i \subseteq F$, then, due to the new transition, the word γw is accepted by $\mathcal{A}_{i+1}$ starting from s^j. Notice that the new transition imposes that only words w that are accepted starting from all the states in the P_ℓ's can be considered (w is in the intersection of the languages of all these states). The use of alternating automata allows to represent this intersection without modification of the number of states of the original automaton $\mathcal{A}$. This is crucial for the termination argument of the construction.

The following theorem, which shows the correcteness of the construction of $\mathcal{A}_{pre^*}$, is proved in the appendix:

Theorem 4.1 *Given an APDS $\mathcal{P}$ and a regular set of configurations recognized by a $\mathcal{P}$-AMA $\mathcal{A}$, we can construct a $\mathcal{P}$-AMA $\mathcal{A}_{pre^*}$ recognizing $pre^*(Conf(\mathcal{A}))$.*

It follows easily from the facts below that the algorithm is polynomial on the size of $\mathcal{P}$ and (singly) exponential in the size of $\mathcal{A}$:

- $\mathcal{A}_{pre^*}$ has the same states as $\mathcal{A}$,
- a $\mathcal{P}$-AMA with k states has $O(n_{\mathcal{P}} \cdot k \cdot 2^k)$ transitions, where $n_{\mathcal{P}}$ is the size of $\mathcal{P}$, and
- during the construction of the sequence $\mathcal{A}_0, \mathcal{A}_1, \ldots$, polynomial time suffices to decide if a new transition can be added to the current automaton.

5 Model-Checking Branching-Time Temporal Logics

5.1 The alternation-free (propositional) μ-calculus

Let $Prop$ be a set of atomic propositions and $\mathcal{X}$ a finite set of variables. The set of formulas of the (propositional) μ-calculus is defined by the following grammar:

$$\varphi ::= \pi \in Prop \mid X \in \mathcal{X} \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists\bigcirc\varphi \mid \mu X.\, \varphi$$

where in formulas of the form $\mu X.\, \varphi$, the variable X must occur in φ under an even number of negations. In addition, we consider the usual abbreviations: the boolean connectives $\wedge$ and $\Rightarrow$, $\forall\bigcirc\varphi = \neg\exists\bigcirc\neg\varphi$, and $\nu X.\, \varphi(X) = \neg\mu X.\, \neg\varphi(\neg X)$. We write $\sigma X.\, \varphi(X)$ for either $\mu X.\, \varphi(X)$ or $\nu X.\, \varphi(X)$.

The notion of free occurrence of a variable in a formula is defined as usual by considering μ and ν as quantifiers. We suppose without loss of generality that in every formula each variable is bound at most once. We write $\varphi(X)$ to precise that X occurs free in φ. A formula φ is *closed* if no variable occurs free in it, otherwise it is *open*.

We interpret formulas on the set of configurations of a PDS $\mathcal{P} = (P, \Gamma, \Delta)$. We need a labelling function $\Lambda : P \to 2^{Prop}$, which intuitively assigns to each control location the set of propositions that hold at it, and a valuation $\mathcal{V}$ which assigns to each variable a set of configurations. The set of configurations of $\mathcal{P}$ that satisfy a formula φ is denoted by $[\![\varphi]\!]_{\mathcal{P}}(\mathcal{V})$ and defined by the following rules:

$$\begin{array}{rcl}
[\![\pi]\!]_{\mathcal{P}}(\mathcal{V}) & = & \Lambda^{-1}(\pi) \times \Gamma^* \\
[\![X]\!]_{\mathcal{P}}(\mathcal{V}) & = & \mathcal{V}(X) \\
[\![\neg\phi]\!]_{\mathcal{P}}(\mathcal{V}) & = & (P \times \Gamma^*) \setminus [\![\phi]\!]_{\mathcal{P}}(\mathcal{V}) \\
[\![\phi_1 \vee \phi_2]\!]_{\mathcal{P}}(\mathcal{V}) & = & [\![\phi_1]\!]_{\mathcal{P}}(\mathcal{V}) \cup [\![\phi_2]\!]_{\mathcal{P}}(\mathcal{V}) \\
[\![\exists\bigcirc\varphi]\!]_{\mathcal{P}}(\mathcal{V}) & = & pre([\![\varphi]\!]_{\mathcal{P}}(\mathcal{V})) \\
[\![\nu X.\phi]\!]_{\mathcal{P}}(\mathcal{V}) & = & \bigcup\{\mathcal{C} \subseteq P \times \Gamma^* \mid \mathcal{C} \subseteq [\![\phi]\!]_{\mathcal{P}}(\mathcal{V}[\mathcal{C}/X])\}
\end{array}$$

where $\mathcal{V}[\mathcal{C}/X]$ is the valuation which coincides with $\mathcal{V}$ for all variables but X, where it takes the value $\mathcal{C}$.

The set of formulas in *positive normal form* is defined by the following syntax:

$$\varphi ::= \pi \mid \neg\pi \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \exists\bigcirc\varphi \mid \forall\bigcirc\varphi \mid \mu X.\, \varphi \mid \nu X.\, \varphi$$

It is easy to show that every formula is equivalent to a formula in positive normal form (push negations inside).

A σ-subformula of a formula $\sigma X.\, \phi(X)$ is *proper* if it does not contain any occurrence of X. The alternation-free μ-calculus is the set of formulas φ in positive normal form such that for every σ-subformula ϕ of φ the following holds:

- if ϕ is a μ-formula, then all its ν-subformulas are proper, and
- if ϕ is a ν-formula, then all its μ-subformulas are proper.

Given a formula φ, we define its *closure* $cl(\varphi)$ as the smallest set of formulas containing φ and such that

- if $\phi_1 \vee \phi_2 \in cl(\varphi)$ or $\phi_1 \wedge \phi_2 \in cl(\varphi)$ then $\phi_1 \in cl(\varphi)$ and $\phi_2 \in cl(\varphi)$,
- if $\exists\bigcirc\phi \in cl(\varphi)$ or $\forall\bigcirc\phi \in cl(\varphi)$ then $\phi \in cl(\varphi)$,
- if $\sigma X.\ \phi(X) \in cl(\varphi)$ then $\phi(\sigma X.\ \phi(X)) \in cl(\varphi)$.

It is easy to see that the closure of a formula is always a finite set, and that its cardinality is bounded by the length of the formula.

5.1.1 The Model-Checker

Consider a PDS $\mathcal{P} = (P, \Gamma, \Delta)$ and a labelling function $\Lambda : P \to 2^{Prop}$. Let φ be a formula of the alternation-free μ-calculus, and let $\mathcal{V}$ be a valuation of the free variables in φ.

We show how to construct an AMA $\mathcal{A}_\varphi$ recognizing $[\![\varphi]\!]_\mathcal{P}(\mathcal{V})$. From now on we drop the indices and write just $[\![\varphi]\!]$.

We start by considering the case where all the σ-subformulas of φ are μ-formulas. We construct an APDS $\mathcal{AP}$ which is, roughly speaking, the product of $\mathcal{P}$ and the alternating automaton corresponding to φ [Eme96]; we then reduce the problem of computing $[\![\varphi]\!]$ to computing the value of $pre^*_{\mathcal{AP}}$ for a certain regular set of configurations. Intuitively, a configuration $\langle [p, \phi], w\rangle$ belongs to this set if ϕ is a basic formula of the form π, $\neg\pi$, or X, for X free in ϕ, and the configuration $\langle p, w\rangle$ of $\mathcal{P}$ satisfies ϕ. Observe that whether $\langle p, w\rangle$ satisfies ϕ or not can be decided by direct inspection of the labelling function Λ and the valuation $\mathcal{V}$. The AND-branching in the transition rules of $\mathcal{AP}$ is due to conjunctions and universal path quantifications (in $\forall\bigcirc$ operators) occurring in the formula φ.

Formally, we define the APDS $\mathcal{AP} = (P^\varphi_\mathcal{P}, \Gamma, \Delta^\varphi_\mathcal{P})$ where

- $P^\varphi_\mathcal{P} = P \times cl(\varphi)$,
- $\Delta^\varphi_\mathcal{P}$ is the smallest set of transition rules satisfying the following conditions for every control location $[p, \phi]$ and every stack symbol γ:
  - if $\phi = \phi_1 \vee \phi_2$, then $([p, \phi], \gamma) \hookrightarrow ([p, \phi_1], \gamma)$ and $([p, \phi], \gamma) \hookrightarrow ([p, \phi_2], \gamma)$,
  - if $\phi = \phi_1 \wedge \phi_2$, then $([p, \phi], \gamma) \hookrightarrow \{\ ([p, \phi_1], \gamma), ([p, \phi_2], \gamma)\ \}$,
  - if $\phi = \mu X.\ \psi(X)$, then $([p, \phi], \gamma) \hookrightarrow ([p, \psi(\phi)], \gamma)$,
  - if $\phi = \exists\bigcirc\psi$ and $(p, \gamma) \hookrightarrow (q, w)$ is a transition rule of $\mathcal{P}$, then $([p, \phi], \gamma) \hookrightarrow ([q, \psi], w)$,
  - if $\phi = \forall\bigcirc\psi$ then $([p, \phi], \gamma) \hookrightarrow \{([q, \psi], w) \mid (p, \gamma) \hookrightarrow (q, w)\}$.

Let $\mathcal{C}_t$ (where the index t stands for true) be the subset of configurations of $\mathcal{AP}$ containing all configurations of the form

- $\langle [p, \pi], w\rangle$, where $\pi \in \Lambda(p)$,

- $\langle [p, \neg\pi], w\rangle$, where $\pi \notin \Lambda(p)$,
- $\langle [p, X], w\rangle$, where X is free in φ and $\langle p, w\rangle \in \mathcal{V}(X)$.

Clearly, if $\mathcal{V}(X)$ is a regular set of configurations for every variable X free in φ, then $\mathcal{C}_t$ is also a regular set of configurations.

The following result can be easily proved using standard techniques based on the notion of signature [Bra92]:

Proposition 5.1 *Let $\mathcal{AP}$ be the APDS obtained from $\mathcal{P}$ and φ using the construction above. A configuration $\langle p, w\rangle$ of $\mathcal{P}$ belongs to $[\![\varphi]\!]$ iff the configuration $\langle [p, \varphi], w\rangle$ of $\mathcal{AP}$ belongs to $pre^*_{\mathcal{AP}}(\mathcal{C}_t)$.*

Applying Theorem 2.1 we obtain a procedure to compute an AMA $\mathcal{A}_\varphi$ which accepts exactly the configurations of $\mathcal{P}$ that satisfy φ.

The case in which all the σ-subformulas of φ are ν-subformulas is now easy to solve: the negation of φ is equivalent to a formula φ' in positive normal form whose σ-subformulas are all μ-subformulas. Applying Theorem 2.1 we construct an AMA which accepts the configurations of $\mathcal{P}$ that satisfy φ'. We then just use the fact that AMA's are closed under complementation.

Let us now consider the general case of in which φ is an arbitrary formula of the alternation-free μ-calculus. We can assume without loss of generality that φ is a σ-formula (otherwise a "dummy" fixpoint can be added). The following property (which does *not* hold for the full μ-calculus) follows easily from the definitions, and allows us to construct the AMA $\mathcal{A}_\varphi$. We use the following notation: given a family $\Phi = \{\phi_i\}_{i=1}^n$ of subformulae of φ, pairwise incomparable with respect to the subformula relation, and a family $U = \{U_i\}_{i=1}^n$ of fresh variables, $\varphi[U/\Phi]$ denotes the result of simultaneously substituting U_i for ϕ_i in φ.

Proposition 5.2 *Let φ be a μ-formula (ν-formula) of the alternation-free μ-calculus, and let $\Phi = \{\phi_i\}_{i=1}^n$ be the family of maximal ν-subformulas (μ-subformulas) of ϕ with respect to the subformula relation. Then*

$$[\![\varphi]\!] = [\![\varphi[U/\Phi]]\!](\mathcal{V}')$$

where $U = \{U_i\}_{i=1}^n$ is a suitable family of fresh variables, and $\mathcal{V}'$ is the valuation which extends $\mathcal{V}$ by assigning to each U_i the set $[\![\phi_i]\!]$.

Observe that if φ is a μ-formula (ν-formula), then all the σ-subformulas of $\varphi[U/\Phi]$ are also μ-formulas (ν-formulas). Together with Proposition 5.1, this leads immediately to a recursive algorithm for computing $\mathcal{A}_\varphi$: for every $\phi \in \Phi$, compute recursively AMA's $\mathcal{A}_\phi$ recognizing $[\![\phi]\!]$, and then use them and Proposition 5.2 to compute $\mathcal{A}_\varphi$. So we have the following theorem:

Theorem 5.1 *Let $\mathcal{P}$ be a PDS, let φ a formula of the alternation-free μ-calculus, and let $\mathcal{V}$ be a valuation of the free variables of φ. We can construct an AMA $\mathcal{A}_\varphi$ such that $Conf(\mathcal{A}_\varphi) = [\![\varphi]\!]_\mathcal{P}(\mathcal{V})$.*

5.1.2 Complexity

Walukiewicz has shown in [Wal96] that there exists a formula of the alternation-free μ-calculus such that the model checking problem for PDS's and this formula is DEXPTIME-complete. This implies that all model-checking algorithms must have exponential complexity in the size of the system. We show that the algorithm we have obtained (which is very different from Walukiewicz's) has this complexity.

Let $n_{\mathcal{P}}$ be the size of $\mathcal{P}$ and let n_φ be the length of φ. We define a tree of σ-subformulas of φ: the root of the tree is φ; the children of a μ-subformula (ν-subformula) ϕ are the maximal ν-subformulas (μ-subformulas) of ϕ. Clearly, the number of nodes of the tree does not exceed n_φ.

Let ϕ be a leaf of the tree. The AMA $\mathcal{A}_\phi$ recognizing $[\![\phi]\!]$ is obtained by applying the pre^* construction to the AMA recognizing the set $\mathcal{C}_t$. Since the latter has $O(n_{\mathcal{P}} \cdot n_\varphi)$ states, $\mathcal{A}_\phi$ has also $O(n_{\mathcal{P}} \cdot n_\varphi)$ states.

Now, let ϕ be an internal node of the tree with children $\phi_1, \ldots, \phi_k$. If the AMA recognizing $[\![\phi_i]\!]$ has n_i states, then the AMA recognizing $[\![\phi]\!]$ has $O(\Sigma_{i=1}^{n} n_i + n_{\mathcal{P}} \cdot n_\varphi)$.

Since the number of nodes of the tree does not exceed n_φ, the AMA $\mathcal{A}_\varphi$ recognizing $[\![\varphi]\!]$ has $O(n_{\mathcal{P}} \cdot n_\varphi^2)$ states.

Since each AMA can be constructed in exponential time in the number of states, the algorithm is singly exponential in $n_{\mathcal{P}}$ and n_φ.

5.2 The logic EF

The alternation-free μ-calculus is a rather powerful logic. Proper sublogics, like CTL, are considered to be sufficiently expressive for many applications. This raises the question whether the model-checking problem for PDS's and some interesting fragment of the alternation-free μ-calculus may lie in some complexity class below DEXPTIME. In this section we show that this is the case: we prove that the model-checking problem for the logic EF (propositional logic plus the temporal operator EF) is in PSPACE.[2] However, the problem turns out to be PSPACE-complete, even PSPACE-complete in the size of the system. Therefore, the complexity gap between the alternation-free μ-calculus and its sublogics is rather small.

Given a set $Prop$ of atomic propositions, the set of formulas of EF is defined by the following grammar:

$$\varphi ::= \pi \in Prop \mid \neg\varphi \mid \varphi \vee \varphi \mid \exists\bigcirc\varphi \mid EF\varphi$$

The semantics of formulas of the form π, $\neg\varphi$, $\varphi_1 \vee \varphi_2$, and $\exists\bigcirc\varphi$ is defined as for the alternation-free μ-calculus. A configuration c satisfies a formula $EF\varphi$ if there exists a configuration c' reachable from c that satisfies φ.

We consider the proof of membership in PSPACE first, since this is the part that makes use of our reachability analysis. The hardness part is a standard reduction from the acceptance problem for linearly bounded Turing machines.

Fix a PDS $\mathcal{P}(P, \Gamma, \Delta)$ and a configuration c of $\mathcal{P}$. Denote by $R(c)$ the set of words $pw \in P\Gamma^*$ such that $\langle p, w\rangle$ is reachable from c. We have the following result (see the proof in the appendix):

[2] We assume $PSPACE \neq DEXPTIME$

Theorem 5.2 *The set $R(c)$ is regular. Moreover, $R(c)$ is recognized by a finite multi automaton having polynomially many states in the sum of the sizes of c and $\mathcal{P}$.*

Given a formula φ of the alternation-free μ-calculus, denote by $R(c, \varphi)$ the set of words $pw \in P\Gamma^*$ such that $\langle p, w \rangle$ is reachable from c *and* satisfies φ. By Theorem 5.2 and Theorem 5.1, we have:

Corollary 5.1 *Let φ be a formula of the alternation-free μ-calculus. The set $R(c, \varphi)$ is regular, and is recognized by an alternating finite automaton having polynomially many states in the sum of the sizes of c, $\mathcal{P}$ and φ.*

We can now obtain our first result:

Theorem 5.3 *The model checking problem for the logic EF and pushdown automata is in PSPACE.*

Proof: Let $\mathcal{P}$ be a PDS and let φ be a formula of EF. We show by induction on the structure of φ that the problem of deciding if a given configuration c of $\mathcal{P}$ satisfies ϕ can be solved in nondeterministic polynomial space in the size of c, $\mathcal{P}$ and φ.

The cases $\varphi = \pi, \varphi_1 \vee \varphi_2, \neg\varphi_1$ are trivial. So let $\varphi = EF\varphi_1$, and assume that we can decide if a configuration c satisfies φ_1 using nondeterministic polynomial space in the size of c, $\mathcal{P}$ and φ_1.

By the definition of the semantics of EF, $\mathcal{P}$ satisfies φ iff there exists a configuration c_1 reachable from c that satisfies φ_1.

If an AMA with n states recognizes a nonempty set, then it recognizes some word of length at most n. Therefore, by Corollary 5.1, we can assume that c_1 has polynomial size in $\mathcal{P}$ and φ_1. The following nondeterministic algorithm decides in polynomial space if $\mathcal{P}$ satisfies φ:

- Guess a configuration c_1 of polynomial size in $\mathcal{P}$ and φ;
- check in polynomial time in the size of c_1 and $\mathcal{P}$ that c_1 is reachable from c;
- check in polynomial space in the size of c_1, $\mathcal{P}$ and φ_1 that c_1 satisfies φ_1.

The membership of the model checking in PSPACE follows now from NPSPACE = PSPACE. $\square$

We give in the appendix the proof of the following hardness result.

Theorem 5.4 *The model checking problem for the logic EF and pushdown automata is PSPACE-hard.*

Finally, a simple look at the proof of Theorem 5.4 shows that the following stronger result also holds:

Corollary 5.2 *There is a formula φ of EF such that the problem of deciding if a PDS satisfies φ is PSPACE-complete.*

6 Conclusion

We have applied the "*symbolic*" analysis principle to a class of infinite state systems, namely pushdown systems. We have represented (possibly infinite) sets of configurations using finite-state automata, and have proposed a simple procedure to compute sets of predecessors. Using this procedure and the automata-theoretic approach to model-checking, we have obtained model-checking algorithms for both linear and branching-time properties. From these results we have derived upper bounds for several model-checking problems. We have also provided matching lower bounds by means of some reductions based on Walukiewicz's ideas [Wal96].

The model-checking problem for pushdown systems and the modal mu-calculus (or its alternation-free fragment) has been studied in several papers [BS92, BS95, Wal96]. The main advantage of our approach (apart from an homogeneous treatment of both branching-time and linear-time logics) is the simplicity of our algorithms: only well known concepts from automata theory are needed to understand them. They are also smooth generalizations of global model-checking algorithms for branching-time logics and finite-state systems.

We do not know whether our approach can be extended to the full modal mu-calculus. The exact complexity of the model-checking problem for pushdown systems and CTL is also open: it lies somewhere between PSPACE and DEXPTIME. These questions, and the extension of the symbolic analysis principle to other classes of systems with infinite state spaces, are left for future investigations.

References

[ACH+95] R. Alur, C. Courcoubetis, N. Halbwachs, T. Henzinger, P. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The Algorithmic Analysis of Hybrid Systems. *TCS*, 138, 1995.

[AD94] R. Alur and D. Dill. A Theory of Timed Automata. *TCS*, 126, 1994.

[AMP95] E. Asarin, O. Maler, and A. Pnueli. Symbolic Controller Synthesis for Discrete and Timed Systems. In *Hybrid Systems II*. LNCS 999, 1995.

[BG96] B. Boigelot and P. Godefroid. Symbolic Verification of Communication Protocols with Infinite State Spaces using QDDs. In *CAV'96*. to appear, 1996.

[BO93] R.V. Book and F. Otto. *String-Rewriting Systems*. Springer-Verlag, 1993.

[Bra92] J.C. Bradfield. *Verifying Temporal Properties of Systems*. Birkhauser, 1992.

[Bry92] R. Bryant. Symbolic Boolean Manipulation with Ordered Binary-Decision Diagrams. *ACM Computing Surveys*, 24, 1992.

[BS92] O. Burkart and B. Steffen. Model Checking for Context-Free Processes. In *CONCUR'92*, 1992. LNCS 630.

[BS95] O. Burkart and B. Steffen. Composition, Decomposition and Model-Checking of Pushdown Processes. *Nordic Journal of Computing*, 2, 1995.

[CES83] E.M. Clarke, E.A. Emerson, and E. Sistla. Automatic Verification of Finite State Concurrent Systems using Temporal Logic Specifications: A Practical Approach. In *POPL'83*. ACM, 1983.

[Eme96] E.A. Emerson. Automated Temporal Reasoning about Reactive Systems. In *Logics for Concurrency*. LNCS 1043, 1996.

[Hol94] Holzmann. Basic SPIN manual. Technical report, Bell Laboratories, 1994.

[McM93] K.L. McMillan. *Symbolic Model-Checking: an Approach to the State-Explosion Problem*. Kluwer, 1993.

[Pnu77] A. Pnueli. The Temporal Logic of Programs. In *FOCS'77*. IEEE, 1977.

[Var88] M.Y. Vardi. A Temporal Fixpoint Calculus. In *POPL'88*. ACM, 1988.

[Var95] M.Y. Vardi. Alternating Automata and Program Verification. In *Computer Science Today*. LNCS 1000, 1995.

[VW86] M.Y. Vardi and P. Wolper. An Automata-Theoretic Approach to Automatic Program Verification. In *LICS'86*. IEEE, 1986.

[Wal96] I. Walukiewicz. Pushdown Processes: Games and Model Checking. In *CAV'96*. LNCS 1102, 1996.

Appendix

6.1 Calculating pre^* for Pushdown Systems

We prove hereafter the correcteness of the construction given in Section 2.2.

The following invariants hold for all $i \geq 0$:

(1) The states, initial states, and final states of $\mathcal{A}_i$ and $\mathcal{A}_0$ coincide.

(2) $\rightarrow_i \subseteq \rightarrow_{i+1}$.

(3) $Y_i \subseteq Y_{i+1}$ (follows immediately from (1) and (2)).

(4) Every transition of $\mathcal{A}_{i+1}$ that does not belong to $\mathcal{A}_i$ starts at an initial state.

(5) No transition of $\mathcal{A}_i$ leads from a non-initial state to an initial state.

We start by proving that the sequence of the Y_i's eventually reaches its limit.

Lemma 6.1 $\exists i \geq 0.\ Y_{i+1} = Y_i$.

Proof: Follows immediately from the fact that $\mathcal{A}_{i+1}$ is constructed from $\mathcal{A}_i$ by adding new transitions, but no new states. Since the number of possible transitions for a fixed set of states and input alphabet is finite, there is an $i \geq 0$ such that $\mathcal{A}_i = \mathcal{A}_{i+n}$ for every $n \geq 0$. $\square$

Now, let us show the inclusion of the X_i's in the Y_i's.

Lemma 6.2 $\forall i \geq 0.\ X_i \subseteq Y_i$.

Proof: By induction on i. The case $i = 0$ is trivial. Assume $\langle p^j, \gamma w\rangle \in X_{i+1}$. Then, there exists a transition rule $(p^j, \gamma) \hookrightarrow (p^k, w')$ such that $\langle p^k, w'w\rangle \in X_i$. By induction hypothesis, $s^k \xrightarrow{w'}_i q' \xrightarrow{w}_i q$ for some state q and some final state q. By the definition of $\mathcal{A}_{i+1}$, we have $s^j \xrightarrow{\gamma}_{i+1} q' \xrightarrow{w}_{i+1} q$. So $\langle p^j, \gamma w\rangle \in Y_{i+1}$. $\square$

It remains to prove that $\forall i \geq 0$, $Y_i \subseteq pre^*(C)$. For that, we need the following key lemma:

Lemma 6.3 *If $s^j \xrightarrow{w}_i q$, then $\langle p^j, w\rangle \Rightarrow \langle p^k, v\rangle$ for some p^k and v such that $s^k \xrightarrow{v}_0 q$.*

Proof:

We prove a stronger result: if q is an initial state s^l, then not only the lemma holds, but also $\langle p^j, w\rangle \Rightarrow \langle p^l, \varepsilon\rangle$, i.e, we may take $p^k = p^l$ and $v = \varepsilon$.

The proof is by induction on i.

Basis. $i = 0$. Take $k = j$ and $v = w$. If q is an initial state, then, since the automaton $\mathcal{A}_0 = \mathcal{A}$ has no transitions leading to an initial state, we must have $w = \varepsilon$ and $q = s^j$. This implies $\langle p^j, w\rangle \Rightarrow \langle p, \varepsilon\rangle$.

Step. $i \geq 1$. Fix an w-run ρ leading from s^j to q. The proof is by induction on the number of times that ρ uses the transitions added to $\mathcal{A}_{i-1}$ to yield $\mathcal{A}_i$. If this number is 0, then we have $s^j \xrightarrow{w}_{i-1} q$, and the result follows from the induction hypothesis. So assume that $w = w_1\gamma w_2$ and ρ satisfies

$$s^j \xrightarrow{w_1}_i s^{k_0} \xrightarrow{\gamma}_i q' \xrightarrow{w_2}_i q \qquad (1)$$

where the transition $s^{k_0} \xrightarrow{\gamma} q'$ does *not* belong to $\mathcal{A}_{i-1}$. Notice that we can safely write $s^{k_0} \xrightarrow{\gamma}_i q'$ because all new transitions start at an initial state.

The application of the induction hypothesis to $s^j \xrightarrow{w_1}_i s^{k_0}$ yields

$$\langle p^j, w_1 \rangle \Rightarrow \langle p^{k_0}, \varepsilon \rangle \tag{2}$$

From $s^{k_0} \xrightarrow{\gamma}_i q'$ and the definition of $\mathcal{A}_i$ we have

$$\langle p^{k_0}, \gamma \rangle \Rightarrow \langle p^{k_1}, v_1 \rangle \text{ for some } p^{k_1} \text{ and } v_1 \text{ such that } s^{k_1} \xrightarrow{v_1}_{i-1} q' \tag{3}$$

We now apply the induction hypothesis (induction on i) to $s^{k_1} \xrightarrow{v_1}_{i-1} q'$, and obtain:

$$\langle p^{k_1}, v_1 \rangle \Rightarrow \langle p^{k_2}, v_2 \rangle \text{ for some } p^{k_2} \text{ and } v_2 \text{ such that } s^{k_2} \xrightarrow{v_2}_0 q' \tag{4}$$

By 1 and 4 there is a run of $\mathcal{A}_i$ satisfying $s^{k_2} \xrightarrow{v_2}_0 q' \xrightarrow{w_2}_i q$. Since this run contains less applications of new transitions than $s^j \xrightarrow{w}_i q$, we apply the induction hypothesis and obtain:

$$\langle p^{k_2}, v_2 w_2 \rangle \Rightarrow \langle p^k, v \rangle \text{ for some } p^k \text{ and } v \text{ such that } s^k \xrightarrow{v}_0 q \tag{5}$$

Putting 2 to 5 together, we get

$$\langle p^j, w_1 \gamma w_2 \rangle \Rightarrow \langle p^k, v \rangle \text{ for some } p^k \text{ and } v \text{ such that } s^k \xrightarrow{v}_0 q$$

This is the result we wished to prove.

Assume now that q is an initial state s^l. If ρ uses no new transitions, then we have $s^j \xrightarrow{w}_{i-1} q$, and the result follows from the induction hypothesis. So assume that $w = w_1 \gamma w_2$ and that ρ satisfies

$$s^j \xrightarrow{w_1}_i s^{k_0} \xrightarrow{\gamma}_i q' \xrightarrow{w_2}_i s^l$$

Since $\mathcal{A}_i$ contains no transitions from non-initial to initial states, q' must also be an initial state s^{k_1}. We then have (a) $\langle p^j, w_1 \rangle \Rightarrow \langle p^{k_0}, \varepsilon \rangle$ by induction hypothesis, (b) $\langle p^{k_0}, \gamma \rangle \Rightarrow \langle p^{k_1}, \varepsilon \rangle$ because s^{k_1} is an initial state, and (3) $\langle s^m, v \rangle \Rightarrow \langle s^l, \varepsilon \rangle$ by induction hypothesis. So $\langle p^j, w_1 \gamma w_2 \rangle \Rightarrow \langle p^l, \varepsilon \rangle$. □

Lemma 6.4 *$\forall i \geq 0$, $Y_i \subseteq pre^*(C)$.*

Proof: The proof is by induction on i. The case $i = 0$ is trivial since $Y_0 = C$. Let $\langle p^j, w \rangle \in Y_i$, $i \geq 1$. By the definition of $\mathcal{A}_i$, we have $s^j \xrightarrow{w}_i q$ for some $q \in F$. By Lemma 6.3, $\langle p^j, w \rangle \Rightarrow \langle p^k, v \rangle$ for some p^k and v such that $s^k \xrightarrow{v}_0 q$. Since $Y_0 = C$, $\langle p^k, v \rangle \in C$. So $\langle p^j, w \rangle \in pre^*(C)$. □

6.2 Proof of Proposition 3.1

($\Rightarrow$): Let $\rho = c_0 c_1 c_2 \ldots$ be an accepting run of $\mathcal{BP}$ starting from $c_0 = c$. For every $i \geq 0$, let $\rho^{(i)}$ be the suffix of ρ starting at c_i, and let $m^{(i)}$ be the minimal length of the configurations of $\rho(i)$, where the length of a configuration is defined as the length of its stack content.

Construct a subsequence $c_{i_0}c_{i_1}\dots$ of ρ as follows: c_{i_0} is the first configuration of ρ of length $m^{(0)}$; for every $j > 1$, c_{i_j} is the first configuration of $\rho^{(i_j+1)}$ of length $m(i_j+1)$.

Since the number of control locations is finite, there exists a subsequence $c_{j_1}c_{j_2}\dots$ of $c_{i_1}c_{1_2}\dots$ whose elements have all the same control location p, and the same symbol γ on top of the stack.

Since ρ is an accepting run, there is a number $k \geq 0$ and a configuration c_g with control location $g \in G$ such that

$$c_0 \Rightarrow c_{j_1} \Rightarrow c_g \Rightarrow c_{j_k}$$

Let $c_{j_1} = \langle p, \gamma w\rangle$. Then, $c \Rightarrow \langle p, \gamma w\rangle$, and so (1) holds. Due to the definition of the subsequences $c_{i_1}c_{1_2}\dots$ and $c_{j_1}c_{j_2}\dots$, all the configurations of ρ between c_{j_1} and c_{j_k} have stack contents of the form $w'w$. In particular, c_g is of the form $\langle g, uw\rangle$ and c_{j_k} is of the form $\langle p, \gamma v w\rangle$. This implies

$$\langle p, \gamma\rangle \Rightarrow \langle g, u\rangle \Rightarrow \langle p, \gamma v\rangle$$

and so (2) holds, which concludes the proof.

($\Leftarrow$): By (1) and (2), we have $c \Rightarrow \langle p, \gamma w\rangle$ and

$$\langle p, \gamma v^i w\rangle \Rightarrow \langle g, uv^i w\rangle \Rightarrow \langle p, \gamma v^{i+1} w\rangle$$

for every $i \geq 0$, which allows to construct an accepting run.

6.3 Proof of Theorem 3.2

We give here the hardness proof. For that, we use a reduction from the problem of deciding whether a given linearly bounded alternating Turing machine accepts a given input or not. We suppose that the reader is familiar with the temporal logic LTL [Pnu77].

An alternating Turing machine is a tuple $(Q, \Sigma, \delta, q_0, F, \lambda)$, where λ is a function that labels states as existential, universal, accepting or rejecting. Without loss of generality, we only consider machines where $|\delta(q, a)| = 2$ for every universal state q and symbol a. We choose an arbitrary order on the two elements of $\delta(q, a)$, so that later we can speak of the first and the second successor configurations of a universal configuration.

Let $M = (Q, \Sigma, \delta, q_0, F, \lambda)$ be an alternating Turing Machine satisfying the following property: for every word in the language of M of length n, M has an accepting computation which uses at most $k \cdot n$ space, where k is a constant independent of the input.

Let w be an input of M of length n. We assume that the tape of M has length $k \cdot n$, and initially M's head is scanning the first cell of the tape. We construct in polynomial time a PDS $\mathcal{A}$ and a formula ϕ of such that M accepts w iff $\mathcal{A}$ satisfies ϕ.

Define a configuration of M on the input w as a word of $\Sigma^*(Q \times \Sigma)\Sigma^*$ of length $k \cdot n$. In a configuration $\alpha(q, a)\beta$, α is the word representing the contents of the tape to the left of the head (including blanks for the cells that have not been visited yet), q is the current state, a is the symbol scanned by the head, and β represents the contents of the tape to the left of the head. A configuration $\alpha q \beta$ is accepting if q is an accepting state.

The automaton $\mathcal{A}$ guesses nondeterministically a finite or infinite tree of configurations, in which every node has at most two successors (if the tree is infinite, then the automaton just never stops guessing). Since the PDS generates words, and not trees, it is necessary to encode a tree of configurations as a string. We use the following coding: the tree is traversed in infix order; when we enter a node c coming from its parent node (or when we initially enter the root), we write $\#fc\#$; when we enter a node c coming from one of its children, we write $\#bc^r\#$, where f and b are special markers and c^r is the reverse of c.

The root of the trees guessed by $\mathcal{A}$ is always the initial configuration of M for input w, the internal nodes are non-accepting configurations, and the leaves are accepting configurations. In order to guess a tree, $\mathcal{A}$ keeps in its stack the path from the root to the actual configuration by pushing the new guessed configurations and popping when backtracking along the subtree guessed so far. $\mathcal{A}$ accepts by empty stack.

It is easy to see that M accepts w iff $\mathcal{A}$ satisfies the following property: $\mathcal{A}$ accepts a string $\#d_0\#d_1\#\ldots\#d_n$ such that for every $0 \leq i \leq n-1$

- if $d_i = fc$ and c is an existential configuration, then $d_{i+1} = fc'$ and c' is a successor configuration of c;
- if $d_i = fc$ and c is a universal configuration, then $d_{i+1} = fc'$ and c' is the first successor configuration of c;
- if $d_i = bc^r$ and $d_{i+1} = fc'$, then c is a universal configuration, and c' is the second successor configuration of c.

The formula ϕ is just an encoding of this property in LTL. For each symbol $a \in \Sigma \cup (Q \times \Sigma) \cup \{f, b, \#\}$ we introduce a proposition p_a. We use the abbreviation $\bigcirc^i\psi$ for $\underbrace{\bigcirc \ldots \bigcirc \phi}_{i}$. We set

$$\phi = \Box((p_\# \wedge \bigcirc^{n+2} p_\#) \Rightarrow (\phi_1 \wedge \phi_2 \wedge \phi_3))$$

where $p_\# \wedge O^{n+2}p_\#$ expresses that the current position is a $\#$-position different from the last, and ϕ_1, ϕ_2, ϕ_3 encode parts (1), (2) and (3) of the property above.

We only construct ϕ_1, since ϕ_2 and ϕ_3 are similar.

- "$d_i = fc$" is encoded as $\bigcirc p_f$;
- "c is an existential configuration" is encoded as

  $$\bigvee_{j=1}^{n} (\bigvee_{q \in Q_e, a \in \Sigma} p_{q,a})$$

  where Q_e is the set of existential states;
- "$d_{i+1} = fc'$" is encoded as $\bigcirc^{n+2} p_f$;
- "c' is a successor configuration of c" is encoded as a conjunction of formulae, one for each possible successor configuration of c. These formulae are in turn a conjunction of formulae of the form

  $$(\bigcirc^j p_{a_0} \wedge \bigcirc^{j+1} p_{a_1} \wedge \bigcirc^{j+2} p_{a_2}) \Rightarrow \bigcirc^{n+2=j} p_a$$

  for every $2 \leq j \leq n-1$ and every string $a_0a_1a_2$ of actions of $\mathcal{A}$. The symbol a is determined by the transition relation δ of the Turing machine. We just give an example for the case $a_0a_1a_2 = 0(q,1)1$ and $(q', 0, L) \in \delta(q, 1)$. The corresponding formula is

  $$(\bigcirc^j p_0 \wedge \bigcirc^{j+1} p_{(q,1)} \wedge \bigcirc^{j+2} p_1) \Rightarrow \bigcirc^{n+2=j} p_0$$

6.4 Calculating pre^* for Alternating Pushdown Systems

We prove the correctness of the algorithm given in Section 4.2 for the computation of the pre^* function for APDSs.

Let us start by introducing some notations. Given an AMA $\mathcal{A} = (\Gamma, Q, \delta, I, F)$, we denote by $Run(q, w)$ the set of runs starting from q over the word w. Given a run $\rho \in Run(q, w)$, we denote by $Leaves(\rho)$ the set of leaves of ρ. It is easy to see that $q \xrightarrow{w} Q'$ iff there is a run ρ of $\mathcal{A}$ over w starting from q such that $Leaves(\rho) = Q'$.

Given a run $\rho \in Run(q, w)$, we write $\rho : q' \xrightarrow{u} Q'$ if there exists a subrun ρ' of ρ over u which starts from q' and such that $Leaves(\rho') = Q'$. Notice that u must be a factor of w, i.e., $w = u_1 u u_2$ for some words u_1 and u_2.

Let us prove now the correcteness of the construction of the automaton $\mathcal{A}_{pre^*}$.

The invariants (1) to (4) of the non-alternating case are still valid. Invariant (5) has to be slightly reformulated:

(5) No transition of $\mathcal{A}_i$ leads from a non-initial state to a set containing an initial s tate.

Lemma 6.1 and 6.2 can be easily generalized to the alternating case. We omit their proofs.

It remains to prove that $\forall i \geq 0$, $Y_i \subseteq pre^*(C)$. For that, we need the following key lemma, which generalizes Lemma 6.3:

Lemma 6.5 *If $s^j \xrightarrow{w}_i U$, then there exist two sets of configurations $C_1, C_2 \subseteq \mathcal{C}$ such that $\langle p^j, w\rangle \Rightarrow C_1 \cup C_2$, and*

- *$C_1 = \{\langle p^k, \varepsilon\rangle \ : \ s^k \in U\}$, and*
- *for every $\langle p^k, v\rangle \in C_2$ there exists $V_c \subseteq U$ such that $s^k \xrightarrow{v}_0 V_c$; moreover, $\bigcup_{c \in C_2} V_c = U \setminus I$.*

Proof: The proof is by induction on i. Let us start by the case $i = 0$. If $w = \varepsilon$ then $U = \{s^j\}$ (thus $U \setminus I = \emptyset$), and in this case we can take $C_1 = \{\langle p^j, \varepsilon\rangle\}$ and $C_2 = \emptyset$ (recall that $\langle p^j, \varepsilon\rangle \Rightarrow \{\langle p^j, \varepsilon\rangle\}$). If $w \neq \varepsilon$ then, since we have supposed that $\mathcal{A}_0 = \mathcal{A}$ has no transitions leading to I, $s^j \xrightarrow{w}_i U$ implies that $U \setminus I = U$, which means that we must take $C_1 = \emptyset$, and then it suffices to take $C_2 = \{\langle p^j, w\rangle\}$.

Now, for $i \geq 1$, we have to prove that

> $\forall r \in \mathcal{R}$, $\forall j \in \{1, \ldots, m\}$, $\forall w \in \Gamma^*$, $\forall U \subseteq Q$, if $s^j \xrightarrow{w}{}_i^r U$, then $\exists C_1, C_2 \subseteq \mathcal{C}$ such that $\langle p^j, w\rangle \Rightarrow C_1 \cup C_2$, $C_1 = \{\langle p^k, \varepsilon\rangle \ : \ k \in \{1, \ldots, m\}$ and $s^k \in U\}$, and
> $\forall c = \langle p^k, v\rangle \in C_2$ with $k \in \{1, \ldots, m\}$, $\exists V_c \subseteq U$. $s^k \xrightarrow{v}_0 V_c$ such that $\bigcup_{c \in C_2} V_c = U \setminus I$.

Let $r = (\langle p^\ell, \gamma\rangle, \{\langle p^{i_1}, w_1\rangle, \ldots, \langle p^{i_n}, w_n\rangle\}) \in \mathcal{R}$. By definition, the transition table of Y_i^r is obtained by adding to δ_{i-1} new γ-transitions (each of them is starting from s^ℓ and going to all the states appearing in a same conjunction in the DNF of the expression $\bigwedge_{k=1}^{n} \delta_{i-1}^*(s^{i_k}, w_k)$). Then, the proof of the statement above is carried out by induction on the number of times runs starting from initial states (corresponding to transitions of the form $s^j \xrightarrow{w}{}_i^r U$) use one of these new γ-transitions.

Let us suppose that $s^j \xrightarrow{w}{}_i^r U$ and fix a run ρ of $\mathcal{A}_i^r$ over w starting from s^j and such that $Leaves(\rho) = U$. If ρ does not use new γ-transitions, this means that $s^j \xrightarrow{w}_{i-1} U$. Hence, the fact holds by induction hypothesis.

Now, suppose that the new γ-transitions are used at least once in the run ρ. Then, let $w = u\gamma v$, and let ρ_u be a subrun of ρ over u with $W = Leaves(\rho_u)$, such that the new γ-transitions are used at some state in W (notice that this state is necessarily s^ℓ).

Since ρ_u involves less new γ-transitions than the whole run ρ, by induction hypothesis, we have $\langle p^j, u\gamma v\rangle \Rightarrow C_1' \cup C_2'$ such that $C_1' = \{\langle p^\ell, \gamma v\rangle\} \cup \{\langle p^k, \gamma v\rangle \ :\ k \in \{1, \dots, m\}, k \neq \ell$ and $s^k \in W\}$ and $\forall c = \langle p^k, v_k\gamma v\rangle \in C_2'$ with $k \in \{1, \dots, m\}$, $\exists V_c \subseteq W.\ s^k \xrightarrow{v_k}_0 V_c$ such that $\bigcup_{c\in C_2'} V_c = W \setminus I$. Then, we have to show that $\exists C_1, C_2 \subseteq \mathcal{C}$ such that $C_1' \cup C_2' \subseteq pre^*(C_1 \cup C_2)$, $C_1 = \{\langle p^k, \varepsilon\rangle \ :\ k \in \{1, \dots, m\}$ and $s^k \in U\}$, and $\forall c = \langle p^k, v\rangle \in C_2$ with $k \in \{1, \dots, m\}$, $\exists V_c \subseteq U.\ s^k \xrightarrow{v}_0 V_c$ such that $\bigcup_{c\in C_2} V_c = U \setminus I$. For that, we decompose the proof into three parts according to the kind of configurations in $C_1' \cup C_2'$. First, we consider the particular configuration $\langle p^\ell, \gamma v\rangle \in C_1'$, then the C_1' configurations that are different from $\langle p^\ell, \gamma v\rangle$, and finally the C_2' configurations.

Case 1 Consider the configuration $\langle p^\ell, \gamma v\rangle \in C_1'$. Let $V \subseteq Q$ be the set of states such that $\rho : s^\ell \xrightarrow{\gamma}{}_i^r V$. For every $q \in V$, let ρ_q be the subrun of ρ over v starting from q, and let $U_q \subseteq U$ such that $U_q = Leaves(\rho_q)$. By construction of $\mathcal{A}_i^r$, we have necessarily $\forall k \in \{1, \dots, n\}$, $\exists V_k \subseteq V$, $s^{i_k} \xrightarrow{w_k}_{i-1} V_k$, and $V = \bigcup_{k=1}^n V_k$. Let $U_{(1,k)} = \bigcup_{q\in V_k} U_q$ and $U_1 = \bigcup_{k=1}^n U_{(1,k)}$. Thus, $\rho : s^\ell \xrightarrow{\gamma v} U_1$, and $\forall k \in \{1, \dots, n\}$, there exists a run ρ_k over $w_k v$ starting from s^{i_k} and such that $Leaves(\rho_k) = U_{(1,k)}$ ($\rho_k : s^{i_k} \xrightarrow{w_k v}{}_i^r U_{(1,k)}$), with $\rho_k : s^{i_k} \xrightarrow{w_k}{}_i^r V_k$ and $\forall q \in V_k$, the subrun of ρ_k over v starting from q is exactly ρ_q ($\rho_k : q \xrightarrow{v}{}_i^r U_q$). The following picture illustrate this case assuming for simplicity that $n = 2$.

It is easy to see that the number of applications of the new γ-transitions in each run ρ_k is equal to the one in all its subruns ρ_q's, and that this latter is strictly smaller than the one in the whole run ρ. Then, by induction hypothesis, we have $\forall k \in \{1, \dots, n\}$, $\exists C_1^{(1,k)}, C_2^{(1,k)} \subseteq \mathcal{C}$, $\langle p^{i_k}, w_k v\rangle \Rightarrow C_1^{(1,k)} \cup C_2^{(1,k)}$, $C_1^{(1,k)} = \{\langle p^{k'}, \varepsilon\rangle \ :\ k' \in \{1, \dots, m\}$ and $s^{k'} \in U_{(1,k)}\}$, and $\forall c = \langle p^{k'}, v\rangle \in C_2^{(1,k)}$ with $k' \in \{1, \dots, m\}$, $\exists V_c \subseteq U_{(1,k)}.\ s^{k'} \xrightarrow{v}_0 V_c$ such that $\bigcup_{c\in C_2^{(1,k)}} V_c = U_{(1,k)} \setminus I$.

By applying the transition rule r, we have $\langle p^\ell, \gamma v\rangle \Rightarrow \{\langle p^{i_1}, w_1 v\rangle, \dots, \langle p^{i_n}, w_n v\rangle\}$, and thus $\langle p^\ell, \gamma v\rangle \Rightarrow \bigcup_{k=1}^n (C_1^{(1,k)} \cup C_2^{(1,k)})$. Then, let $C_1^1 = \bigcup_{k=1}^n C_1^{(1,k)}$ and $C_2^{(1,k)} = \bigcup_{k=1}^n C_2^{(1,k)}$. Clearly, $C_1^1 = \{\langle p^k, \varepsilon\rangle \ :\ k \in \{1, \dots, m\}$ and $s^k \in U_1\}$, and $\forall c = \langle p^k, v\rangle \in C_2^1$, $\exists V_c \subseteq U_1.\ s^k \xrightarrow{v}_0 V_c$ such that $\bigcup_{c\in C_2^1} V_c = U_1 \setminus I$.

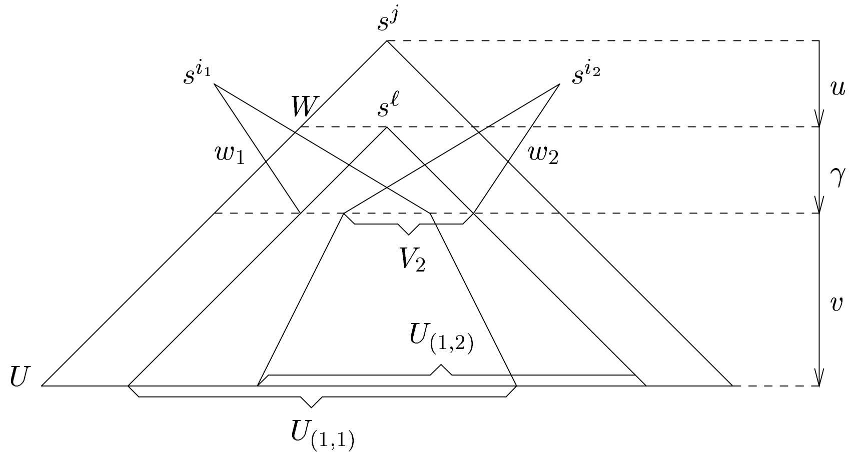

Case 2 Let $\langle p^k, \gamma v\rangle$ be a configuration of C'_1 such that $k \in \{1, \ldots, m\}$, $k \neq \ell$, and $s^k \in W$. Consider the subrun ρ_k of ρ such that $\rho_k : s^k \xrightarrow{\gamma v}{}^r_i U_{(2,k)}$ (for some $U_{(2,k)} \subseteq U$). It is clear that ρ_k contains strictly less new γ-transitions than ρ. Then, by induction hypothesis, we have immediately $\exists C_1^{(2,k)}, C_2^{(2,k)} \subseteq \mathcal{C}$, $\langle p^k, \gamma v\rangle \Rightarrow C_1^{(2,k)} \cup C_2^{(2,k)}$, $C_1^{(2,k)} = \{\langle p^{k'}, \varepsilon\rangle \ : \ k' \in \{1, \ldots, m\}$ and $s^{k'} \in U_{(2,k)}\}$, and $\forall c = \langle p^{k'}, v\rangle \in C_2^{(2,k)}$ with $k' \in \{1, \ldots, m\}$, $\exists V_c \subseteq U_{(2,k)}$. $s^{k'} \xrightarrow{v}_0 V_c$ such that $\bigcup_{c \in C_2^{(2,k)}} V_c = U_{(2,k)} \setminus I$.

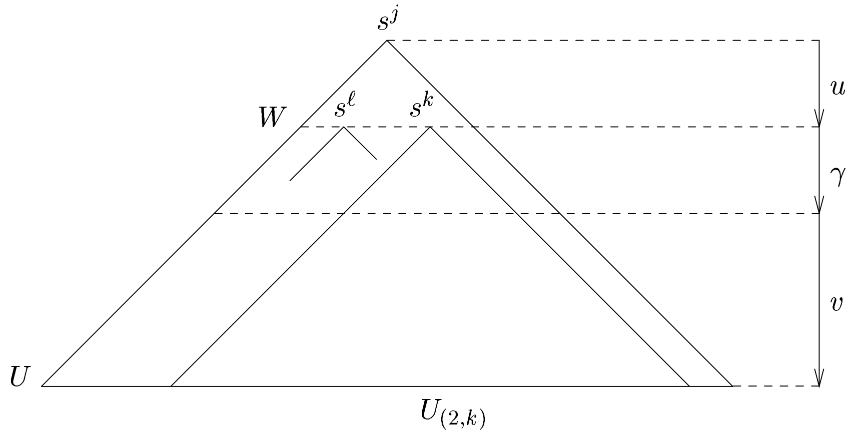

Case 3 Let $\langle p^k, v_k\gamma v\rangle$ be a configuration in C'_2. Consider the set of states $V \subseteq W \setminus I$ such that $s^k \xrightarrow{v_k}_0 V$. Recall that we have supposed that the original automaton $\mathcal{A}_0$ has no transitions leading to I. Then, it is easy to see that, by construction, for every $i \geq 1$, the automaton $\mathcal{A}_i$ has no transitions from $Q \setminus I$ to I (because all the new transitions start from a state in I). Consequently, for every $q \in V$, there exists $U_q \subseteq U$ such that $q \xrightarrow{\gamma v}_0 U_q$. Then, we obtain $s^k \xrightarrow{v_k\gamma v}_0 U_{(3,k)} = \bigcup_{q \in V} U_q$.

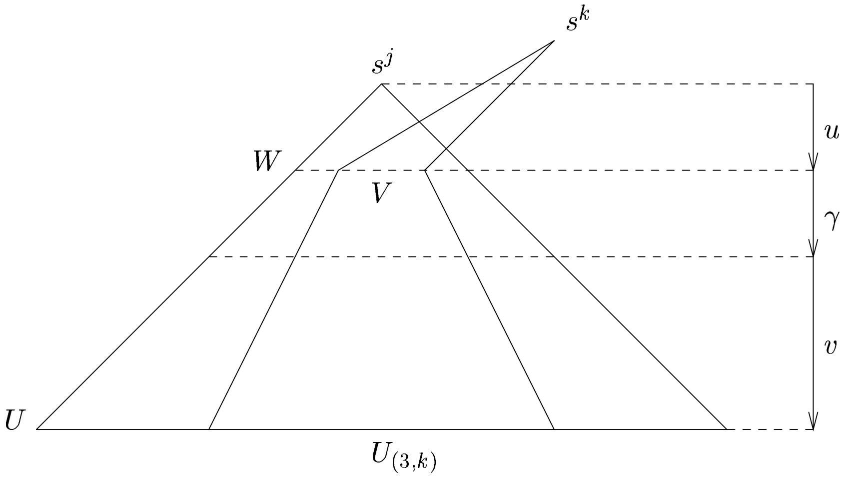

By induction hypothesis (since the property holds for $i = 0$), we have $\exists C_1^{(3,k)}, C_2^{(3,k)} \subseteq \mathcal{C}$, $\langle p^k, v_k\gamma v\rangle \Rightarrow C_1^{(3,k)} \cup C_2^{(3,k)}$, $C_1^{(3,k)} = \{\langle p^{k'}, \varepsilon\rangle \ : \ k' \in \{1,\ldots,m\}$ and $s^{k'} \in U_{(3,k)}\}$, and $\forall c = \langle p^{k'}, v\rangle \in C_2^{(3,k)}$ with $k' \in \{1,\ldots,m\}$, $\exists V_c \subseteq U_{(3,k)}$. $s^{k'} \xrightarrow{v}_0 V_c$ such that $\bigcup_{c\in C_2^{(3,k)}} V_c = U_{(3,k)} \setminus I$.

Let $U_2 = \bigcup\{U_{(2,k)} \ : \ \langle p^k, \gamma v\rangle \in C'_1, k \in \{1,\ldots,m\}, k \neq \ell,$ and $s^k \in W\}$ and $U_3 = \bigcup\{U_{(3,k)} \ : \ \langle p^k, v_k\gamma v\rangle \in C'_2\}$. It is easy to see that $U = \bigcup_{i=1}^3 U_i$. We define also, $i = 1, 2$, $C_i^2 = \bigcup\{C_i^{(2,k)} \ : \ \langle p^k, \gamma v\rangle \in C'_1, k \in \{1,\ldots,m\}, k \neq \ell,$ and $s^k \in W\}$ and $C_i^3 = \bigcup\{C_i^{(3,k)} \ : \ \langle p^k, v_k\gamma v\rangle \in C'_2\}$. Then, for $i = 1, 2$, let $C_i = \bigcup_{j=1}^3 C_i^j$. We have shown by considering the three cases above that $C'_1 \cup C'_2 \subseteq pre^*(C_1 \cup C_2)$, and that $C_1 = \{\langle p^k, \varepsilon\rangle \ : \ k \in \{1,\ldots,m\}$ and $s^k \in U\}$, and $\forall c = \langle p^k, v\rangle \in C_2$ with $k \in \{1,\ldots,m\}$, $\exists V_c \subseteq U$. $s^k \xrightarrow{v}_0 V_c$ such that $\bigcup_{c\in C_2} V_c = U \setminus I$. □

We can now prove:

Lemma 6.6 *$\forall i \geq 0$, $Y_i \subseteq pre^*(C)$.*

Proof: The proof follows the same scheme as the one of Lemma 6.5. It is carried out by induction on i. The case $i = 0$ is trivial since $Y_0 = C$. For $i \geq 1$, we have to show that

> $\forall j \in \{1,\ldots,m\}$, $\forall w \in \Gamma^*$, $\forall r \in \mathcal{R}$, $\forall W \subseteq F$, if $s^j \xrightarrow{w}{}_i^r W$ then $\langle p^j, w\rangle \in pre^*(C)$.

Let $r = (\langle p^\ell, \gamma\rangle, \{\langle p^{i_1}, w_1\rangle, \ldots, \langle p^{i_n}, w_n\rangle\}) \in \mathcal{R}$. As for Lemma 6.5, we prove the fact above by induction on the number of times runs starting from initial states use one of the new γ-transitions introduced (due to the rule r) in the construction of $\mathcal{A}_i^r$ from $\mathcal{A}_{i-1}$.

Let us suppose that $s^j \xrightarrow{w}{}_i^r W$ and fix a run ρ of $\mathcal{A}_i^r$ over w starting from s^j and such that $Leaves(\rho) = W$. If ρ does not use new γ-transitions, we have $s^j \xrightarrow{w}_{i-1} W$, or equivalently $\langle p^j, w\rangle \in Y_{i-1}$. Then, since by induction hypothesis $Y_{i-1} \subseteq pre^*(C)$, we obtain $\langle p^j, w\rangle \in pre^*(C)$.

Consider now the case when the new γ-transitions are used at least once in ρ. Let $w = u\gamma v$, and let ρ_u be a subrun of ρ over u with $U = Leaves(\rho_u)$, such that the new γ-transitions are used for the first time at some state in U (this state is s^ℓ). By Lemma 6.5, $\langle p^j, u\gamma v\rangle \Rightarrow C'_1 \cup C'_2$ such that $C'_1 = \{\langle p^\ell, \gamma v\rangle\} \cup \{\langle p^k, \gamma v\rangle \ : \ k \in \{1,\ldots,m\}, k \neq$

ℓ and $s^k \in U\}$ and $\forall c = \langle p^k, v_k\gamma v\rangle \in C'_2$ with $k \in \{1,\ldots,m\}$, $\exists V_c \subseteq U$. $s^k \xrightarrow{v_k}_0 V_c$ such that $\bigcup_{c\in C_2} V_c = U \setminus I$. So, let us show that $C'_1 \cup C'_2 \subseteq pre^*(C)$. We decompose the proof into three cases:

Case 1 Consider the configuration $\langle p^\ell, \gamma v\rangle \in C'_1$. Let $V \subseteq Q$ be the set of states such that $\rho : s^\ell \xrightarrow{\gamma}{}^r_i V$. For every $q \in V$, let ρ_q be the subrun of ρ over v starting from q, and let $W_q \subseteq W$ such that $W_q = Leaves(\rho_q)$. By construction of $\mathcal{A}^r_i$, we have $\forall k \in \{1,\ldots,n\}$, $\exists V_k \subseteq V$, $s^{i_k} \xrightarrow{w_k}_{i-1} V_k$, and $V = \bigcup_{k=1}^n V_k$. Let $W_k = \bigcup_{q\in V_k} W_q$ and $W' = \bigcup_{k=1}^n W_k$. Thus, we have $\rho : s^\ell \xrightarrow{\gamma v}{}^r_i W'$, and $\forall k \in \{1,\ldots,n\}$, there exists a run ρ_k over $w_k v$ starting from s^{i_k} such that $Leaves(\rho_k) = V_k$ ($\rho_k : s^{i_k} \xrightarrow{w_k v}{}^r_i W_k$), with $\rho_k : s^{i_k} \xrightarrow{w_k}{}^r_i V_k$ and $\forall q \in V$, the subrun of ρ_k over v starting from q is exactly ρ_q ($\rho_k : q \xrightarrow{v}{}^r_i W_q$). The number of applications of the new γ-transitions in ρ_k is equal to the one in all the ρ_q's, which is itself strictly smaller than the one in the whole run ρ. Then, by induction hypothesis, we have $\forall k \in \{1,\ldots,n\}$, $\langle p^{i_k}, w_k v\rangle \in pre^*(C)$. Finally, since by applying the transition rule r we have $\langle p^\ell, \gamma v\rangle \Rightarrow \{\langle p^{i_1}, w_1 v\rangle, \ldots, \langle p^{i_k}, w_k v\rangle\}$, we obtain $\langle p^\ell, \gamma v\rangle \in pre^*(C)$.

Case 2 Let $\langle p^k, \gamma v\rangle$ be a configuration of C'_1 such that $k \in \{1,\ldots,m\}$, $k \neq \ell$, and $s^k \in U$. Consider the subrun ρ_k of ρ such that $\rho_k : s^k \xrightarrow{\gamma v}{}^r_i W'$ (for some $W' \subseteq W$). Clearly, ρ_k contains strictly less applications of new γ-transition rules than ρ. Then, by induction hypothesis, we have immediately $\langle p^k, \gamma v\rangle \in pre^*(C)$.

Case 3 Let $\langle p^k, v_k\gamma v\rangle$ be a configuration in C'_2. Consider the set of states $V \subseteq U \setminus I$ such that $s^k \xrightarrow{v_k}_0 V$. As we have said in the proof of Lemma 6.5 (Case 3), all the automata $\mathcal{A}_i$ have no transitions from $Q \setminus I$ to I. Thus, for every $q \in V$, there exists $W_q \subseteq W$ such that $q \xrightarrow{\gamma v}_0 W_q$. Then, $s^k \xrightarrow{v_k\gamma v}_0 \bigcup_{q\in V} W_q$, and by induction hypothesis, we obtain $\langle p^k, v_k\gamma v\rangle \in pre^*(C)$. □

Theorem 4.1 follows now from the generalization of Lemma 6.1 and 6.2 to the alternating case, and from Lemma 6.4.

6.5 Proof of Theorem 5.2

We need some preliminary definitions and notations.

- Let W be the set of words $w \in \Gamma^*$ such that some transition rule of $\mathcal{P}$ is of the form $(p,\gamma) \hookrightarrow (q, vw)$. Observe that W has $O(n^2_{\mathcal{P}})$ elements, where $n_{\mathcal{P}}$ is the size of $\mathcal{P}$.
- For every $p \in P$ and every $w \in \Gamma^*$ let $N(p,w)$ be the set of control locations p such that $\langle p, w\rangle \Rightarrow \langle p, \varepsilon\rangle$.

We obtain the following recursive equations for the sets $R(\langle p, w\rangle)$ with $w \in W$:

$$\begin{aligned} R(\langle p,\varepsilon\rangle) &= \{p\} \\ R(\langle p,\gamma\rangle) &= \{p\gamma\} \cup \{R(\langle q,w\rangle) \mid (p,\gamma) \hookrightarrow (q,w)\} \\ R(\langle p,w\gamma\rangle) &= R(\langle p,w\rangle)\gamma \cup \{R(\langle q,\gamma\rangle) \mid q \in N(p,w)\} \end{aligned}$$

Clearly, these equations define a left-linear grammar of size $O(n^3_{\mathcal{P}})$. Therefore, the sets $R(\langle p, w\rangle)$ are regular, and are recognized by a finite automaton with $O(n^3_{\mathcal{P}})$ states.

Let $c = \langle p_c, \gamma_1 \dots \gamma_n \rangle$. To compute $R(c)$, add to the equation system above one more equation for each $1 \leq i \leq n-1$:

$$R(\langle p, \gamma_1 \dots \gamma_{i+1} \rangle) = R(\langle p_c, \gamma_1 \dots \gamma_i \rangle)\gamma_{i+1} \cup \{R(\langle p, \gamma_{i+1} \rangle) \mid p \in N(p_c, \gamma_1 \dots \gamma_i)\}$$

The new equation system has size $O(n_{\mathcal{P}}^3 + n_{\mathcal{P}} \cdot n)$.

6.6 Proof of Theorem 5.4

We use a reduction from the problem of deciding whether given a linearly bounded nondeterministic Turing machine accepts a given input or not.

Let $M = (Q, \Sigma, \delta, q_0, F)$ be a nondeterministic Turing Machine with one tape satisfying the following property: for every word in the language of M of length n, M has an accepting computation which uses at most $k \cdot n$ space, where k is a constant independent of the input. We also assume that the machine stops whenever it enters a final state.

Let w be an input of M of length n. We assume that the tape of M has length $k \cdot n$, and initially M's head is scanning the first cell of the tape. We construct in polynomial time a PDS $\mathcal{P}$ and a formula ϕ of EF such that M accepts w iff $\mathcal{P}$ satisfies ϕ.

Define a configuration of M on the input w as a word of $\Sigma^*(Q \times \Sigma)\Sigma^*$ of length $k \cdot n$. In a configuration $\alpha(q, a)\beta$, α is the word representing the contents of the tape to the left of the head (including blanks for the cells that have not been visited yet), q is the current state, a is the symbol scanned by the head, and β represents the contents of the tape to the left of the head. A configuration $\alpha q \beta$ is accepting if q is a final state.

The PDS $\mathcal{P}$ guesses nondeterministically a word $\#c_0\#c_1\# \dots \#c_m$, where c_i is a configuration of M for every $0 \leq i \leq m$, c_0 is the initial configuration of M for input w, and c_m is a final configuration, and pushes it in its stack; then it enters a state *Test*. Notice that the sequence of configurations does *not* have to correspond to the computation of M on w.

For expository reasons, rename the contents of the stack as $a_1 a_2 \dots a_l$, where $l = ((k \cdot n) + 1) \cdot m)$ and a_1 is the top of the stack. From state *Test*, $\mathcal{P}$ starts to pop the contents of the stack. At some point, it nondeterministically enters a state *Store*, after which it pops the next three stack symbols, say a_{i+1}, a_i, a_{i-1}, making sure that $a_i \neq \#$ (if $a_i = \#$, then it stores a_i, a_{i-1}, a_{i-2}; if there are no 3 symbols left then it enters a state *Accept*), and enters a state $Test(a_{i+1} a_i a_{i-1})$.

Now, all the stack symbols in the string $a_{i+1} a_i a_{i-1}$ wich are different from $\#$ belong to the same configuration c_j. $\mathcal{P}$ computes the stack symbol a that would be at position $i - (k \cdot n + 1)$ in the stack if c_j were a successor configuration of c_{j-1} according to the transition relation δ (observe that this requires only a finite amount of information that can be stored in $\mathcal{P}$'s finite control). It pops n symbols from the stack (if there are no n symbols left, then it enters *Accept*), and compares a and $a_{i-(k \cdot n+1)}$. If $a = a_{i-(k \cdot n+1)}$ then it enters *Accept*, otherwise it enters a state *Reject*.

Clearly, $\#c_0\#c_1\# \dots \#c_m$ is the computation of M on input w iff it passes all tests that $\mathcal{P}$ can perform from the state *Test*, and in this case M accepts w because c_m is an accepting configuration. So M accepts w iff $\mathcal{P}$ can perform a sequence of guesses leading to the state *Test* such that all possible computations of $\mathcal{P}$ from this state visit the state *Accept*. This condition can be easily encoded as a formula of EF:

$$EF(Test \wedge AG(EF Accept))$$

where $AG = \neg EF \neg$.