Regular Symbolic Analysis of Dynamic Networks of Pushdown Systems

Ahmed Bouajjani[1], Markus Müller-Olm[2], and Tayssir Touili[1]

[1] LIAFA, University of Paris 7, 2 place Jussieu, 75251 Paris cedex 5, France
[2] Universität Dortmund, FB 4, LS 5, Baroper Str. 301, 44221 Dortmund, Germany

Abstract. We introduce two abstract models for multithreaded programs based on dynamic networks of pushdown systems. We address the problem of symbolic reachability analysis for these models. More precisely, we consider the problem of computing effective representations of their reachability sets using finite-state automata. We show that, while forward reachability sets are not regular in general, backward reachability sets starting from regular sets of configurations are always regular. We provide algorithms for computing backward reachability sets using word/tree automata, and show how these algorithms can be applied for flow analysis of multithreaded programs.

1 Introduction

Multithreaded programs are an important class of programs, in which parallelism is used routinely in practice. Parallel programming in general is known to be difficult and error prone, and multithreaded programs are no exception. Therefore, the design of methods and techniques for automatic analysis of such programs is an important and a quite challenging issue. For that, we need to define formal models which are adequate for modelling multithreaded programs, and for which it is possible to construct automatic analysis algorithms.

In recent related work, complete analysis algorithms for abstract classes of parallel programs have been studied by several researchers. Mayr [14] establishes a number of decidability and undecidability results for process classes in the so-called PRS (process rewrite system) hierarchy. PRS are able to model sequential as well as parallel phenomena. In fact, they can be seen as combinations of pushdown systems and Petri nets (defined in a term rewriting setting using prefix and multiset rewrite rules). Following the automata-based approach for the symbolic verification of pushdown systems [2, 11], Lugiez and Schnoebelen [13] show how to use tree automata for reachability analysis of PA processes [1], a particularly well-known class in the PRS hierarchy. Their paper has inspired further work that applies tree automata techniques to analysis of more expressive models [6, 7, 3, 4, 21]. Another line of research generalizes fixpoint-based techniques as common in flow analysis to analysis of similar models of parallel programs [20, 15, 16]. Both approaches can be used to solve bitvector problems, a certain type of simple but important data-flow-analysis problems, for flow graph systems with parallel calls of procedures, or, equivalently, parbegin/parend-blocks interprocedurally [9, 10, 20]. While [9, 10] reduce the problem to reachability analysis of PA-processes, [20] uses fixpoint-based techniques.

Unfortunately, these results do *not* cover interprocedural analysis of multithreaded programs because commands that start new threads cannot adequately be modelled by

parallel calls. In a multithreaded program such a command typically returns immediately (see, e.g., the JAVA or POSIX thread API). Therefore the father of a new thread can pursue its execution concurrently to its son and can even terminate or return to its caller while the son is still alive. In contrast, a parallel call returns only when and if all its component processes have terminated, which is a fundamentally different behavior. Indeed we show in Sect. 2 that in presence of procedures, multithreaded programs can have trace languages different from that of any program with parallel calls.

The goal of this paper is to adapt the automata-based approach mentioned above to interprocedural (reachability) analysis of multithreaded programs. For this purpose we propose two models of multithreaded programs, show how to perform reachability analysis for them with automata-theoretic constructions, and discuss their utility for modelling and analysing multithreaded and other classes of parallel programs.

In Sect. 2 we introduce *Dynamic Pushdown Networks (DPNs)* as a basic model of multithreaded programs. Intuitively, a DPN is a network of pushdown processes that run independently in parallel. Each process can create new members of the network as a side effect of a pushdown transition. DPNs thus model a network of threads each of which can perform basic actions, call (recursively) procedures, and *spawn* new processes. We show that while forward reachability of DPNs does not preserve regularity of configuration sets in general, it still preserves context-freeness (Sect. 4). Backward reachability in contrast preserves regularity and we show how to compute the backward reachability set of a regular set of configurations by means of a saturation algorithm in polynomial time (Sect. 4). We also show that DPN allow us to solve bitvector problems interprocedurally for multithreaded programs (Sect. 3), contrary to previously used models in the literature such as PA processes (Sect. 2).

We extend DPNs to *Constrained DPNs (CDPN)* in Sect. 5, a model that combines (indeed even extends) the modelling power of both DPNs and PA (and even the so-called PAD [14]). The new idea is that enabledness of a transition for a process can be made dependent on a *constraint* which is a regular pattern among the sequence of control states of its sons. We require constraints to be *stable* in the sense that further evolution of the sons cannot invalidate a constraint. We show that otherwise we lose the property that backward reachability preserves regularity. Transition rules with stable constraints increase the expressive power considerably over DPNs. In particular they allow us to model, in addition to thread creation and procedure calls, also parallel calls and various types of join commands among other things. It also allows us to return information back from procedures called in parallel to their caller which cannot be handled in PA and not even in PAD. Constrained DPNs inherit from DPNs that forward reachability does not preserve regularity. Therefore, we consider here backward reachability only. We show that the set of configurations that can reach a given regular set of configurations of a CDPN can again be computed by a saturation algorithm. As configurations of CDPNs are given by unbounded width trees rather than by words as in the DPN case—the tree structure captures the father-son relationship—we resort to hedge automata here [8]. The construction is nontrivial and its justification uses in a subtle manner the assumption about the stability of the constraints in the system definition. While the overall complexity of this procedure is exponential—we indeed prove a PSPACE lower bound—it is exponential only in the number of different constraints used

in the rules of the given CDPN, and just polynomial in the other problem parameters. Therefore, if the number of different constraints is bounded, we obtain a polynomial-time analysis algorithm. This in particular holds if we just model (in addition to spawn operations), parallel calls, a fixed selection of join commands, or a combination of these. Due to lack of space, proofs are omitted. They can be found in [5].

2 Dynamic Pushdown Networks

A *Dynamic Pushdown Network* (DPN) is a tuple $M = (Act, P, \Gamma, \Delta)$, where Act is a finite set of visible *actions*, P is a finite set of *control states*, Γ is a finite set of *stack symbols* disjoint from P, and Δ is a finite set of transition rules of the following forms: either (a) $p\gamma \overset{a}{\hookrightarrow} p_1 w_1$, or (b) $p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2$, where $p, p_1, p_2 \in P$, $a \in Act$, $\gamma \in \Gamma$, and $w_1, w_2 \in \Gamma^*$. A DPN can be seen as a collection of identical sequential processes running in parallel, each of them being able to (1) perform pushdown operations and to (2) create processes in the network. Synchronization is not allowed between processes.

A configuration of a DPN M (also called M-configuration) is a word over the alphabet $\Sigma = P \cup \Gamma$ starting with a symbol in P. An M-configuration can be seen as a sequence of (sub)words in $P\Gamma^*$ each of them corresponding to the configuration of one of the processes running in parallel in the network. Let $Conf_M$ be the set of all M-configurations.

For every $a \in Act$, we define $\xrightarrow{a}_M$ to be the smallest relation in $Conf_M \times Conf_M$ s.t. $\forall u, v \in Conf_M$, $u \xrightarrow{a}_M v$ iff (1) there is a rule $p\gamma \overset{a}{\hookrightarrow} p_1 w_1$ in Δ s.t. $u = u_1 p \gamma u_2$ and $v = u_1 p_1 w_1 u_2$, or (2) there is a rule $p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2$ in Δ s.t. $u = u_1 p \gamma u_2$ and $v = u_1 p_2 w_2 p_1 w_1 u_2$. We write $u \rightarrow_M v$ if there exists $a \in Act$ s.t. $u \xrightarrow{a}_M v$.

The semantics above says that rules of the form (a) correspond precisely to pushdown operations (manipulation of the top of the stack) which can be applied anywhere in the configuration (i.e., by any of the processes in the network): if a process is at control state p and has γ as topmost stack symbol, then it can move to control state p_1 and replace γ by w_1 at the top of its stack. Rules of the form (b) allow in addition the creation of new processes: a process with control state p and topmost stack symbol γ can (1) move to state p_1 and modify its stack by replacing γ with w_1, and moreover, (2) create (to its left) a process which starts its execution at the initial configuration $p_2 w_2$.

Given a configuration c, the set of immediate predecessors (resp. successors) of c is $\mathsf{pre}_M(c) = \{c' \in \mathcal{C} \,:\, c' \rightarrow_M c\}$ (resp. $\mathsf{post}_M(c) = \{c' \in \mathcal{C} \,:\, c \rightarrow_M c'\}$). These notations can be generalized straightforwardly to sets of configurations. Let pre^*_M (resp. post^*_M) denote the reflexive-transitive closure of pre_M (resp. post_M). We omit the subscript M when it is understood from the context. Given $\Delta' \subseteq \Delta$, we use $\mathsf{pre}_{\Delta'}$ (resp. $\mathsf{post}_{\Delta'}$) to denote immediate predecessors (resp. successors) using a rule in Δ'. Then, $\mathsf{pre}^*_{\Delta'}$ and $\mathsf{post}^*_{\Delta'}$ denote the corresponding reflexive-transitive closures. Furthermore, $\mathsf{Traces}_M(c) = \{w \in Act^* \,:\, \exists c'.\, c \xrightarrow{w}_M c'\}$ is the set of traces generated by c.

DPN vs. PA Processes: DPNs allow to model multithreaded programs where creation of threads is done using spawn commands (see Sect. 3). This is not the case for other formalisms used in the literature for modelling parallel programs like PA [1]:[3]

[3] PA corresponds to processes definable by a set of rewrite rules of the form $A \rightarrow t$ where A is a process variable, and t is a term built from process variables, sequential composition, and asynchronous parallel composition.

Theorem 1. *Let* $\mathcal{L} = \bigcup\{a^n\left(b^{n'} \otimes (c^m d^{m'})\right) : n \geq n' \geq 0,\ m \geq m' \geq 0\}$*, where* $\otimes$ *denotes the shuffle (or interleaving) operator defined as usual. Then:*

a) There is a DPN M and an M-configuration c such that $\mathsf{Traces}_M(c) = \mathcal{L}$.
b) There is no PA system Δ *and no process variable A such that* $\mathsf{Traces}_\Delta(A) = \mathcal{L}$.

Hence, PA processes are inadequate for capturing the behavior of multithreaded programs with spawn-like creation of threads. It also follows from the proof that trace sets of DPNs cannot be captured by the type of constraint systems used as semantic reference point in the constraint-based approach [20, 15, 16]. Therefore, the methods of [9, 10, 20, 16, 15] for interprocedural analysis of flow graphs with parallel calls do not carry over immediately to multithreaded programs. These inadequacy results are rather strong because any interesting process equivalence would imply equality of traces.

3 Program Analysis Based on DPN

We show hereafter how DPNs can be used to model multithreaded programs and how our results on symbolic reachability analysis can be used in flow analysis of these programs. This is inspired by Esparza et. al. [9, 10].

Flow Graph Systems: As common in program analysis we assume that the program is given by a flow graph system. Let **Proc** be a finite set of procedure names containing **Main**. We assume that the program operates on a set $\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_k\}$ of global variables. We consider the following types of basic statements: assignment statements, $\mathbf{x}_i := e$, where $\mathbf{x}_i \in \mathbf{X}$ and e is some expression; call of a single procedure, $\mathsf{call}(\pi)$, where $\pi \in$ **Proc**; and spawn of a new thread, $\mathsf{spawn}(\pi)$, where $\pi \in$ **Proc**. The intuitive meaning of assignment statements and calls is obvious. The spawn command $\mathsf{spawn}(\pi)$ models creation of a new independent thread. Like the call $\mathsf{call}(\pi)$, $\mathsf{spawn}(\pi)$ starts an instance of procedure π. In contrast to a call, however, the spawn command returns immediately such that the newly created instance of π runs as a new thread concurrently to the statements that are executed after the spawn. Let **Stmt** be the set of basic statements.

The control flow of each procedure $\pi \in$ **Proc** is described by a control flow graph $G_\pi = (N_\pi, E_\pi, e_\pi, x_\pi)$, where N_π is a finite set of program points of procedure π; $E_\pi \subseteq N_\pi \times$ **Stmt** $\times N_\pi$ is a finite set of edges annotated by basic statements; $e_\pi \in N_\pi$ is the entry point of π; and $x_\pi \in N_\pi$ is the exit point of π. We assume that the sets of program points of different procedures are disjoint, $N_\pi \cap N_{\pi'} = \emptyset$ if $\pi, \pi' \in$ **Proc**, $\pi \neq \pi'$, and agree that $N = \bigcup_{\pi \in \mathbf{Proc}} N_\pi$ and $E = \bigcup_{\pi \in \mathbf{Proc}} E_\pi$.

From Flow Graph Systems to DPN: From a given flow graph system as above we construct a DPN $M = (Act, P, \Gamma, \Delta)$ that captures its operational semantics:

- The actions are given by the assignments that appear in the flow graph system; a special symbol τ is used to signify steps in which no assignment is executed: $Act = \{\mathbf{x} := e \mid \exists u, v : (u, \mathbf{x} := e, v) \in E\} \cup \{\tau\}$;
- we have just one artificial control state #: $P = \{\#\}$;
- we work with a stack of program points; the topmost stack symbol is the current program point of the current procedure, the other stack symbols are the return points of its callers: $\Gamma = N$;

– the transition rules in Δ describe computation steps of the flow graph system:
  1. for every assignment edge $(u, x := e, v) \in E$ we put the rule $\#u \overset{x:=e}{\hookrightarrow} \#v$ to Δ;
  2. for every call edge $(u, \mathsf{call}(\pi), v) \in E$ we put the rule $\#u \overset{\tau}{\hookrightarrow} \#e_\pi v$ to Δ;
  3. for every spawn-edge $(u, \mathsf{spawn}(\pi), v) \in E$ we put the rule $\#u \overset{\tau}{\hookrightarrow} \#v \rhd \#e_\pi$ to Δ,
  4. for each procedure $\pi \in \mathbf{Proc}$, we put the rule $\#x_\pi \overset{\tau}{\hookrightarrow} \#$ to Δ. This rule describes the return from procedure π.

Note that it is possible to extend the semantics above in order to handle local procedure variables and return values from procedure calls. For that, we assume as usual that data values are mapped into a finite abstract domain using standard techniques such as predicate abstraction. Then, abstract values of local variables can be encoded in the stack alphabet and abstract return values can be encoded in the control states.

Solving Bitvector Problems: The operational semantics given above can be used for solving bitvector problems. In order to ease comparison with [10] we discuss detection of live (global) variables. Other bitvector problems can be solved in a similar fashion. Informally, a variable $\mathbf{x} \in \mathbf{X}$ is *live* at a program point $u \in N$ if there is an execution from u in which $\mathbf{x}$ is used before it is over-written. We restrict attention to *reachable* configurations and use a similar definition and notation as Esparza and Podelski [10]. Thus, we define: program variable $\mathbf{x}$ is *live* at a program point $u \in N$ if there is a transition sequence $\#e_{\mathbf{Main}} \xrightarrow{\sigma_1} c_1 \xrightarrow{\sigma_2} c_2 \xrightarrow{\mathbf{y}:=e} c_3$ such that: (1) u is *active* in configuration c_1, i.e., appears as the topmost stack symbol of one of the parallel pushdown processes in the network described by c_1; (2) σ_2 is a sequence of statements that do not modify $\mathbf{x}$ (i.e., do not write to $\mathbf{x}$); and (3) e is an expression in which $\mathbf{x}$ is used.

We denote the set of configurations c in which u is active by At_u, the set of assignments in the given program that modify $\mathbf{x}$ by $\mathsf{Mod}_\mathbf{x} \subseteq Act$, and the set of assignments in the program in which $\mathbf{x}$ is used by $\mathsf{Use}_\mathbf{x} \subseteq Act$. Moreover, we write Δ_A for the set of rules of Δ with an action in a subset $A \subseteq Act$: $\Delta_A = \{(p\gamma \overset{a}{\hookrightarrow} w) \in \Delta \mid a \in A\}$. Using this notation it is not hard to see that $\mathbf{x}$ is live at u if and only if

$$\#e_{\mathbf{Main}} \ \in \ \mathsf{pre}^*(\mathsf{At}_u \cap \mathsf{pre}^*_{\Delta_{Act} \setminus \mathsf{Mod}_\mathbf{x}}(\mathsf{pre}_{\Delta_{\mathsf{Use}_\mathbf{x}}}(Conf_M)))$$

Then, our results concerning backward reachability analysis of DPN given in the next section (see Theorem 3 and Note 1) can be used to decide this property.

4 Reachability Analysis for DPN

We consider the problem of computing representations of the post^* and pre^* images of given sets of configurations. We are interested in the case that sets of configurations are effectively given using automata-based representations.

Computing post^* Images: We show first that post^* does not preserve regularity in general. Consider indeed the DPN $M = (\{a\}, \{p\}, \{\gamma_1, \gamma_2\}, \{p\gamma_1 \overset{a}{\hookrightarrow} p\gamma_1\gamma_1 \rhd p\gamma_2\})$. It is easy to see that $\mathsf{post}^*_M(\{p\gamma_1\}) = \{(p\gamma_2)^n p\gamma_1^{n+1} \ : \ n \geq 0\}$, which is clearly nonregular.

Proposition 1. *There is a DPN M, and a configuration c of M, such that* $\mathsf{post}^*(c)$ *is not a regular set of configurations.*

We prove, however, that post^* preserves context-freeness:

Theorem 2. *For every DPN M and any context-free set C of M-configurations, the set* $\mathsf{post}^*(C)$ *is context-free and effectively constructible in polynomial time.*

Computing pre^* Images: We show now that pre^* preserves regularity. Let M be a DPN and $\mathcal{A}$ be an automaton recognizing a set of M-configurations. We define a polynomial-time algorithm allowing to construct an automaton $\mathcal{A}_{\mathsf{pre}^*}$ s.t. $L(\mathcal{A}_{\mathsf{pre}^*}) = \mathsf{pre}^*_M(L(\mathcal{A}))$. For technical reasons, we require that $\mathcal{A}$ is in a special form we define below.

M-Automata: Let $M = (Act, P, \Gamma, \Delta)$ be a DPN. A finite automaton $\mathcal{A} = (S, \Sigma, \delta, s^0, F)$ is an M-automaton if the following conditions hold:

1. $\Sigma = P \cup \Gamma$ is the finite alphabet,
2. the set of states is partitioned into two sets, $S = S_c \cup S_s$, $S_c \cap S_s = \emptyset$,
3. for every $s \in S_c$ and every $p \in P$, there is a (unique and distinguished) state $s_p \in S_s$,
4. there is a relation $\delta' \subseteq S_s \times \Gamma \times (S_s \setminus \{s_p : s \in S_c, p \in P\}) \cup S_s \times \{\varepsilon\} \times S_c$ such that $\delta = \delta' \cup \{(s, p, s_p) : s \in S_c, p \in P\}$,
5. the initial state $s^0 \in S_c$, and
6. $F \subseteq S$ is the set of final states.

For $\sigma \in \Sigma \cup \{\varepsilon\}$ and $s, s' \in S$, we write $s \xrightarrow{\sigma}_\delta s'$ in lieu of $(s, \sigma, s') \in \delta$. We extend this notation in the obvious manner to sequences of symbols: (1) $\forall s \in S.\ s \xrightarrow{\varepsilon}_\delta s$, and (2) $\forall s, s' \in S.\ \forall \sigma \in \Sigma \cup \{\varepsilon\}.\ \forall w \in \Sigma^*.\ s \xrightarrow{\sigma w}_\delta s'$ iff $\exists s'' \in S.\ s \xrightarrow{\sigma}_\delta s''$ and $s'' \xrightarrow{w}_\delta s'$.

Note that requirement (4) codes a number of conditions on δ: (1) each $s \in S_c$ has s_p as its unique p-successor and has no Γ-transitions, (2) s is the only predecessor of s_p, (3) only ε-moves from states in S_s lead to states $s \in S_c$, (4) states $s \in S_s$ do not have p-successors, for any $p \in P$. So, every path in an M-automaton (starting from the initial state) is the concatenation of paths of the form $s \xrightarrow{p}_\delta s_p \xrightarrow{w}_\delta t \xrightarrow{\varepsilon}_\delta s'$ where $s, s' \in S_c$, $p \in P$, $w \in \Gamma^*$, and all states in the path $s_p \xrightarrow{w}_\delta t$ are in S_s. Note that for every finite automaton $\mathcal{A}$ over the alphabet $P \cup \Gamma$ such that $L(\mathcal{A}) \subseteq \mathit{Conf}_M$, it is possible to construct an M-automaton recognizing the same language.

Constructing the Automaton $\mathcal{A}_{\mathsf{pre}^*}$: Let M be a DPN and $\mathcal{A} = (S, \Sigma, \delta, s^0, F)$ be an M-automaton. The construction of $\mathcal{A}_{\mathsf{pre}^*}$ is in the same spirit as the ones for single pushdown systems (see [2]). It consists in adding iteratively new transitions to the automaton $\mathcal{A}$ according to *saturation* rules (reflecting the backward application of the transition rules in the system), while the set of states remains unchanged. Therefore, we define $\mathcal{A}_{\mathsf{pre}^*}$ to be the finite-state automaton $(S, \Sigma, \delta', s^0, F)$, where δ' is the smallest relation which contains δ (i.e., $\delta \subseteq \delta'$) and satisfies the following conditions:

$R1$: If $(p\gamma \overset{a}{\hookrightarrow} p_1 w_1) \in \Delta$ and $s \xrightarrow{p_1 w_1}_{\delta'} s'$, for $s, s' \in S$, then $(s_p, \gamma, s') \in \delta'$.
$R2$: If $(p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2) \in \Delta$ and $s \xrightarrow{p_2 w_2 p_1 w_1}_{\delta'} s'$, for $s, s' \in S$, then $(s_p, \gamma, s') \in \delta'$.

The relation δ' can be computed as the limit of an increasing sequence of relations obtained by adding transitions to δ that are required by one of the implications above. This procedure terminates after a polynomial number of steps since only a polynomial number of transitions can potentially be added.

Let us explain intuitively the role of the saturation rule (R_1). Consider a path in the automaton of the form $s \xrightarrow{p_1 w_1} s'$. This means, by definition of M-automata, that s is necessarily in S_c and that we have $s \xrightarrow{p_1} s_{p_1} \xrightarrow{w_1} s'$. Then, the rule consists in adding to the automaton the transition $s_p \xrightarrow{\gamma} s'$. Since by definition of M-automata we have $s \xrightarrow{p} s_p$, we obtain a path $s \xrightarrow{p\gamma} s'$ in the automaton. Therefore, if a configuration $u_1 p_1 w_1 u_2$ is recognized by a run $s^0 \xrightarrow{u_1} s \xrightarrow{p_1 w_1} s' \xrightarrow{u_2} s_F$, then its predecessor $u_1 p\gamma u_2$ is also recognized due to the new transition by the run $s^0 \xrightarrow{u_1} s \xrightarrow{p\gamma} s' \xrightarrow{u_2} s_F$. The role of (R_2) is similar.

Theorem 3. $L(\mathcal{A}_{\mathsf{pre}^*}) = \mathsf{pre}^*_M\big(L(\mathcal{A})\big)$.

Note 1. For the sake of completeness, we mention that for every DPN M, and every M-automaton $\mathcal{A}$, the sets $\mathsf{pre}_M(\mathcal{A})$ and $\mathsf{post}_M(\mathcal{A})$ are regular and effectively constructible. The constructions are quite straightforward. For pre_M we take two copies of $\mathcal{A}$. The first copy provides the initial state and the second copy the final states. We then apply the saturation rules to the first copy of the automaton, but let all new transitions lead from states of the first copy to states of the second copy. The post_M construction is similar (it needs adding a finite number of intermediary states).

5 Constrained DPN

We consider in this section an extension of the DPN model introduced in Section 2. In addition to the ability of performing spawn operation as previously, processes are now allowed to observe the control states of their children (processes they have created in the past). This is relevant in particular for handling return values and some kinds of *join* statements between parallel processes. To achieve that, we define a model where the application of a transition rule by some process is conditioned by a (regular language) constraint on the sequence of control states of its children. We need however to impose a *stability* condition (defined below) on the constraints in order to have a model which can be analysed by means of finite-state automata representations. We show later that we lose regularity of the reachability sets if we relax the stability condition.

Stable Regular Languages: Let Σ be a finite alphabet and let $\rho \subseteq \Sigma \times \Sigma$ be a binary relation over Σ. Then, a set of symbols $S \subseteq \Sigma$ is ρ-*stable* iff $\forall s \in S.\ \forall t \in \Sigma.\ (s,t) \in \rho \Rightarrow t \in S$. A ρ-stable regular language over Σ is a subset of Σ^* which is definable by a regular expression of the form:

$$e ::= S, \text{ a } \rho\text{-stable set} \mid e + e \mid e \cdot e \mid e^*$$

We can prove straightforwardly by induction on the structure of regular expressions:

Lemma 1. *Let $\phi \subseteq \Sigma^*$ be a ρ-stable regular language, let $u, v \in \Sigma^*$, and let $a \in \Sigma$ such that $uav \in \phi$. Then, for every $b \in \Sigma$, $(a,b) \in \rho$ implies that $ubv \in \phi$.*

Definition of the Models: A *Constrained Dynamic Pushdown Network* (CDPN) is a tuple $M = (Act, P, \Gamma, \Delta)$, where Act is a finite set of visible *actions*, P is a finite set of *control states*, Γ is a finite set of *stack symbols* disjoint from P, and Δ is a finite set of transition rules of the following forms: either (a) $\phi : p\gamma \stackrel{a}{\hookrightarrow} p_1 w_1$, or (b) $\phi : p\gamma \stackrel{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2$, where $p, p_1, p_2 \in P$, $a \in Act$, $\gamma \in \Gamma$, $w_1, w_2 \in \Gamma^*$, and ϕ is a ρ_Δ-stable regular language over P, with $\rho_\Delta = \{(p,p') \in P \times P :$ there is a rule $\psi : p\delta \stackrel{a}{\hookrightarrow} p'u$ or $\psi : p\delta \stackrel{a}{\hookrightarrow} p'u \rhd p''v$ in $\Delta\}$.

A CDPN consists of a collection of identical sequential processes running in parallel, each of them being modeled as a pushdown system which is able to (1) manipulate its own stack using pushdown rules of the form (a), (2) create a new process (which becomes its youngest son) using rules of the form (b), and (3) observe, under some conditions, the states of its children (processes it created in the past): each transition rule is constrained by the fact that the sequence of control states of the children (given in the decreasing order of their age) must belong to the specified language ϕ.

Since we need to refer to the children of each process, a configuration of a CDPN can be naturally seen as a tree where each vertex is annotated with the configuration of some sequential process (pushdown system), and where the structure corresponds to the relation father-son. Notice that such a tree may have an arbitrary width. We define hereafter a class of terms describing such configurations and we define a transition relation between such terms.

M-Terms: Let $X = \{x_1, \ldots, x_n\}$ be a set of variables. We define the set $\mathcal{T}[X]$ of M-terms over $P \cup \Gamma \cup X$ inductively as follows:

- $X \subseteq \mathcal{T}[X]$,
- If $t \in \mathcal{T}[X]$ and $\gamma \in \Gamma$, then $\gamma(t) \in \mathcal{T}[X]$,
- If $t_1, \ldots, t_n \in \mathcal{T}[X]$ and $p \in P$, then $p(t_1, \ldots, t_n) \in \mathcal{T}[X]$, for $n \geq 0$.

Note that in the last item of this definition, n can be 0 (i.e., p is on a leaf). In that case, we write $p()$ or simply p to represent the corresponding term.

Terms in $\mathcal{T}[\emptyset]$ are called *ground terms*, and will also be denoted by $\mathcal{T}$. A term in $\mathcal{T}[X]$ is linear if each variable occurs at most once. A *context* C is a linear term. Let $t_1, \ldots, t_n$ be n ground terms. Then $C[t_1, \ldots, t_n]$ is the ground term obtained by substituting in C the occurrence of the variable x_i with the term t_i, for $1 \leq i \leq n$.

A term in $\mathcal{T}[X]$ can be seen as a rooted labeled tree of arbitrary width, where (1) an internal node is either of arity 1 (has one successor) if it is labeled with a stack symbol $\gamma \in \Gamma$, or it has an arbitrary arity if it is labeled with a state $p \in P$, and (2) where the leaves are labeled with either variables $x \in X$, or with states $p \in P$.

M-Configurations: We define M-configurations to be the ground M-terms (terms in $\mathcal{T}[X]$ without variables). Given n ground terms $t_1, \ldots, t_n$, the term $\gamma_m \cdots \gamma_1 p(t_1, \ldots, t_n)$ represents a configuration where (1) the common ancestor to all processes is at local control state p and has $\gamma_1 \cdots \gamma_m$ as stack content, where γ_1 is the topmost stack symbol, and (2) this process has n children, the i^{th} of which is described, together with all of its descendants, by the term t_i, for $i = 1, \ldots, n$. A ground term of the form $\gamma_m \cdots \gamma_1 p$ corresponds to the case of one single process without children.

Transition Relation: Given a CDPN M, we define a transition relation $\rightarrow_M$ between M-configurations. We introduce first a notation. Given a configuration t of one of the forms $\gamma_m \cdots \gamma_1 p(t_1,\ldots,t_n)$ or $\gamma_m \cdots \gamma_1 p$, we define $\mathcal{S}(t)$ to be the control state p, i.e., $\mathcal{S}(t)$ is the local control state of the topmost process represented in t. Then, $\rightarrow_M$ is the smallest relation between M-configurations such that:

- If $(\phi : p\gamma \overset{a}{\hookrightarrow} p_1 w_1) \in \Delta$ and $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n) \in \phi$, then
$$C\big[\gamma p(t_1,\ldots,t_n)\big] \rightarrow_M C\big[w_1^R p_1(t_1,\ldots,t_n)\big]$$
- If $(\phi : p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2) \in \Delta$ and $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n) \in \phi$, then
$$C\big[\gamma p(t_1,\ldots,t_n)\big] \rightarrow_M C\big[w_1^R p_1(t_1,\ldots,t_n, w_2^R p_2)\big]$$

where w^R denotes the reverse word (mirror image) of w. The notions of post, pre, post^*, and pre^* are defined as usual.

Modelling Power: Since CDPN generalize DPN, the modelling of programs with spawn operations given in Section 3 is still valid for CDPN. Moreover, stable constraints as preconditions of transition rules increase tremendously the modelling power of our formalism. We discuss some applications in this section.

Parallel Calls: In the data-flow analysis scenario, we can use constraints, e.g., in order to accommodate parallel call commands as another basic primitive for creation of parallelism in addition to spawn commands. A parallel call, $\mathsf{pcall}(\pi,\pi')$ with $\pi,\pi' \in \mathbf{Proc}$ starts an instance of procedure π and an instance of π' and runs them in parallel. It terminates if and when both these instances terminate.

Assume that we extend the flow-graph model of Section 3 by allowing parallel calls as another type of basic statement. In the CDPN model we capture the operational semantics of an edge $(u,\mathsf{pcall}(\pi,\pi'),v)$ as follows: we start two new threads for π and π' and ensure by a transition rule with an appropriate constraint that we can move to v only after both these threads have terminated. For that, both threads indicate termination by moving to a special new "terminated" control state $\natural$ when they see a special new stack symbol \$ that we put at the bottom of their stack upon thread creation. Thus, we have the following rules for modelling $(u,\mathsf{pcall}(\pi,\pi'),v)$:

$$P^* : \#u \overset{\tau}{\hookrightarrow} \#\gamma_1 \rhd \#e_\pi\$ \qquad P^* : \#\gamma_1 \overset{\tau}{\hookrightarrow} \#\gamma_2 \rhd \#e_{\pi'}\$ \qquad P^*\natural^2 : \#\gamma_2 \overset{\tau}{\hookrightarrow} \#v$$

where γ_1,γ_2 are two auxiliary stack symbols chosen fresh for each parallel call. Moreover, the rule $P^* : \#\$ \overset{\tau}{\hookrightarrow} \natural$ allows a thread to move to the state $\natural$ once it has terminated.

Join Statements: Besides parallel calls we can also model different types of join-commands. We use the same technique as above for making termination visible to the father of threads: we now use the rule $\#u \overset{\tau}{\hookrightarrow} \#v \rhd \#e_p\$$ to describe the behavior of a spawn edge $(u,\mathsf{spawn}(p),v) \in E$. Thus, we mark the bottom of the stack with the special symbol \$. We also use the rule $P^* : \#\$ \overset{\tau}{\hookrightarrow} \natural$ from above to make termination visible

in the control state. This allows us to describe the operational semantics of different types of join-command such as for instance (1) $\mathsf{join}_\forall$: proceed if all threads directly created by the current thread have terminated, and (2) $\mathsf{join}_{\exists k}$: proceed if at least k among the threads directly created by the current thread have terminated.

The behavior of an edge (u, j, v) where j is one of the join commands from above is modelled by the rule $\phi : \#u \overset{\tau}{\hookrightarrow} \#v$ where $\phi = \natural^*$ for $j = \mathsf{join}_\forall$, and $\phi = (P^*\natural)^k P^*$ for $j = \mathsf{join}_{\exists k}$. Obviously, these constraints are stable.

Return Values: We can distinguish between different termination conditions by using more than one terminated control state and use regular patterns of such control states in constraints in the father process. This allows us, for instance, to return information back to the caller from procedures called in parallel. Therefore, the modelling power of CDPNs exceeds that of PA and even that of PAD [4] [14]: While in a PAD process (like in a DPN process) we can use control states to return information back to a caller in a normal procedure call, there is no such mechanism for parallel calls. The modelling power for calls and parallel calls is thus more symmetric for CDPNs than for PAD.

Observing Execution Phases: Finally, as we allow *stable* constraints, a creator of a thread can react on situations in which the created thread has achieved some progress already but is not necessarily terminated yet. As an example, let us assume that a process F (the father) creates a number of worker threads that sequentially go through a number of phases, say phases $1, \ldots, n$, before termination. For modelling the worker threads we use new control states from a hierarchy $P_0 \supset P_1 \supset \ldots \supset P_n = \emptyset$ of control states such that a worker thread is in phase i if and only if its control state is in $P_{i-1} \setminus P_i$. This means a worker thread has finished phase i if and only if its control state belongs to P_i. Then, the sets P_i are stable and can be used as building blocks for constraints in transitions of F. Hence, process F can react on situations like "all worker threads have finished phase i" by using the constraint P_i^*, "there is a worker thread that has finished phase i and all other worker threads have finished phase j" by the constraint $P_j^* P_i P_j^*$, etc.

6 Backward reachability analysis of CDPN

Symbolic Representations: We use hedge automata (unbounded width tree automata) [8] to represent infinite sets of CDPN configurations. Let $M = (Act, P, \Gamma, \Delta)$ be a CDPN. An *M-tree automaton* is a tuple $\mathcal{A} = (Q, \delta, F)$, where Q is a set of states, F is the set of final states, and δ is a set of rules of either the form (1) $\gamma(q) \to q'$, where $\gamma \in \Gamma$, and $q, q' \in Q$, or (2) $p(L) \to q$, where L is a regular language over Q, $p \in P$, and $q \in Q$.

In order to define the language recognized by $\mathcal{A}$, we define a *move relation* $\to_\delta$ between terms over $P \cup \Gamma \cup Q$: for every two terms t and t', we have $t \to_\delta t'$ iff there exist a context C and a rule $r \in \delta$ such that $t = C[s]$, $t' = C[s']$, and (1) either $r = \gamma(q) \to q'$, $s = \gamma(q)$, and $s' = q'$, or (2) $r = p(L) \to q$, $s = p(q_1, \ldots, q_n)$, $q_1 \cdots q_n \in L$, and $s' = q$.

Let $\overset{*}{\to}_\delta$ denote the reflexive-transitive closure of $\to_\delta$. A term $t \in \mathcal{T}$ is accepted by $q \in Q$ if $t \overset{*}{\to}_\delta q$. Let $L_q^\delta = \{t \in \mathcal{T} : t \overset{*}{\to}_\delta q\}$. A term t is accepted by $\mathcal{A}$ if there exists a state $q \in F$ such that $t \overset{*}{\to}_\delta q$. Let $L(\mathcal{A})$ be the set of all terms accepted by $\mathcal{A}$.

[4] PAD extends PA by allowing rewrite rules of the form $A \cdot B \to t$

A straightforward adaptation of the proofs in [8] allows to show that:

Theorem 4. *The class of M-tree automata is closed under boolean operations. Moreover, the emptiness problem of M-tree automata is decidable.*

Computing pre^* Images: Let $M = (Act, P, \Gamma, \Delta)$ be a CDPN and let $\mathcal{A} = (Q, \delta, F)$ be an M-tree automaton. We present hereafter an algorithm that allows us to construct an M-tree automaton $\mathcal{A}_{\mathsf{pre}^*}$ recognizing the pre^*-image of $L(\mathcal{A})$. The construction proceeds (similarly to Section 4) by adding new transitions to the original automaton $\mathcal{A}$ corresponding to the backward application of transition rules. In order to deal with the constraints in the transition rules, we need to extend the original automaton.

Propagating Control States: Remember that, by definition of CDPN terms, the configuration of each process is encoded bottom-up in the tree (reading first the control state, and then the stack contents starting from its topmost symbol). Since constraints in CDPN transition rules refer to control states of the children processes, and since hedge automata can check only constraints on immediate successors in trees (which correspond in our case to the bottom symbols in the stacks of the children processes), we need to propagate upward the informations about the control states through the stacks. Therefore, the first step of our construction consists in defining a new automaton $\mathcal{A}_P = (Q_P, \delta_P, F_P)$ such that $L(\mathcal{A}_P) = L(\mathcal{A})$, and where states of Q are labelled by control states $p \in P$. This automaton is given by: $Q_P = Q \times P$, $F_P = F \times P$, and δ_P is the smallest set of rules such that:

- if $p(L) \to s \in \delta$, then $p(L') \to (s,p) \in \delta_P$, where L' is obtained by substituting in the words of L every occurrence of a state $s \in Q$ by $\{(s,p) \mid p \in P\}$;
- if $\gamma(s) \to s' \in \delta$, then for every $p \in P$, $\gamma\big((s,p)\big) \to (s',p) \in \delta_P$.

Lemma 2. *$L(\mathcal{A}_P) = L(\mathcal{A})$, and for every $t \in \mathcal{T}$, $t \xrightarrow{*}_{\delta_P} (s,p)$ iff $t \xrightarrow{*}_{\delta} s$ and $\mathcal{S}(t) = p$.*

Note 2. To avoid confusion, we use in the sequel $p, p', p_1, p_2, \ldots$ to denote elements of P, $s, s', s_1, s_2, \ldots$, to denote states of $\mathcal{A}$, and $q, q', q_1, q_2, \ldots$ to denote states of $\mathcal{A}_P$.

From Constraints over P to Constraints over Q_P: Given a constraint ϕ and n terms $t_1, \ldots, t_n$ such that $t_i \xrightarrow{*}_{\delta_P} q_i$ for $1 \leq i \leq n$, we need also to be able to get the information whether $\mathcal{S}(t_1) \cdots \mathcal{S}(t_n) \in \phi$ from the states $q_1, \ldots, q_n$. For that, we associate with each constraint ϕ over P a constraint $\langle \phi \rangle$ over Q_P such that $\mathcal{S}(t_1) \cdots \mathcal{S}(t_n) \in \phi$ if and only if $q_1 \cdots q_n \in \langle \phi \rangle$. The definition of $\langle \phi \rangle$ is straightforward by induction on the structure of regular expressions for stable languages: (1) $\langle S \rangle = \{(s,p) : s \in Q, p \in S\}$, (2) $\langle \phi_1 \cdot \phi_2 \rangle = \langle \phi_1 \rangle \cdot \langle \phi_2 \rangle$, (3) $\langle \phi_1 + \phi_2 \rangle = \langle \phi_1 \rangle + \langle \phi_2 \rangle$, and (4) $\langle \phi^* \rangle = \langle \phi \rangle^*$.

Closed Set of Constraints: During the construction of the automaton, new transition rules of the form $p(L') \to q$ are added where L' are languages which are built from languages L appearing in the rules of the original automaton $\mathcal{A}$, and constraints ϕ appearing in the transition rules of the CDPN M, using intersection and right-quotient operations. Intersections $L \cap \langle \phi \rangle$ allow us to check that the guarding constraint for the application

of a transition rule is satisfied at the considered position in the tree. Right-quotients $Lq^{-1} = \{w : wq \in L\}$ allow us to get immediate predecessors by a spawn operation of trees where the children of the spawning process are recognized by a sequence of states in L, and the youngest son among these children (i.e., the one created by the spawn operation and which is the right-most one in the list of children) is recognized by the state q. Then, let us define Λ to be the smallest family of languages over Q_P such that:

- If $(p(L) \to q) \in \delta_P$, then $L \in \Lambda$.
- If $L \in \Lambda$, and $(\phi : p\gamma \overset{a}{\hookrightarrow} p_1w_1 \rhd p_2w_2) \in \Delta$, then $L \cap \langle \phi \rangle \in \Lambda$.
- If $L \in \Lambda$ and $q \in Q_P$, then $Lq^{-1} \in \Lambda$.

Lemma 3. *The family Λ is finite. Assuming that all languages and constraints appearing in rules δ_P and Δ are given by backward-deterministic finite-state automata of size at most K, the number of elements of Λ is in $O(K^{n+1})$ where n is the number of different constraints appearing in the rules of Δ.*

Constructing $\mathcal{A}_{\mathsf{pre}^*}$: We define $\mathcal{A}_{\mathsf{pre}^*}$ to be the M-tree automaton (Q', δ', F') such that (1) $Q' = Q_P \cup \{q_p^L : p \in P, L \in \Lambda\}$, (2) $F' = F_P$, and (3) δ' is the smallest set of rules such that $\delta'_0 = \delta_P \cup \{p(L) \to q_p^L : p \in P, L \in \Lambda\} \subseteq \delta'$ and:

R_1: If $(\phi : p\gamma \overset{a}{\hookrightarrow} p'w) \in \Delta$, $p'(L) \to q \in \delta'_0$, and $w^R(q) \overset{*}{\to}_{\delta'} q'$, then $\left(\gamma(q_p^{L \cap \langle \phi \rangle}) \to q'\right) \in \delta'$.

R_2: If $(\phi : p\gamma \overset{a}{\hookrightarrow} p'w_1 \rhd p''w_2) \in \Delta$, $p'(L) \to q'' \in \delta'_0$, $w_1^R(q'') \overset{*}{\to}_{\delta'} q'$, and $w_2^R(p'') \overset{*}{\to}_{\delta'} q$, then $\left(\gamma(q_p^{Lq^{-1} \cap \langle \phi \rangle}) \to q'\right) \in \delta'$.

Note that the states q_p^L, for $p \in P$, and $L \in \Lambda$, are added to the automaton in order to recognize precisely all the terms having p at the root and such that the sequence of children of the root is recognized by a sequence of states in the language L. Note also that all the transitions added by the construction are Γ-transitions, and therefore they do not add P-transitions to the automaton.

The set of rules δ' can be computed iteratively as the limit of an increasing sequence $\delta'_0 \subseteq \delta'_1 \cdots$ such that δ'_{i+1} contains at most one transition more than δ'_i added by applying either (R_1) or (R_2). Note that δ' is necessarily finite since (by Lemma 3) the number of triples (γ, q_p^L, q), for $\gamma \in \Gamma$, $p \in P$, $L \in \Lambda$, and $q \in Q'$ is finite.

Lemma 4. *For every $q \in Q_P$, $L_q^{\delta'} = \mathsf{pre}^*(L_q^{\delta_P})$.*

The lemma above says that the construction ensures that every state recognizes the set of all predecessors of its original language (i.e., in the automaton before saturation). Let us give some intuitive explanations about the role of the saturation rules, and let us consider the rule (R_1) (since the role of (R_2) is similar). Consider a term $w^R p'(t_1, \ldots, t_n)$ such that $t_i \overset{*}{\to}_{\delta'} q_i$, for $i \in \{1, \ldots, n\}$. Assume that $p'(L) \to q$ is a rule of the automaton. This means that after recognizing each of the terms t_i and labelling their roots by the states q_i, the automaton can label the term $p'(t_1, \ldots, t_n)$ by q if the sequence $q_1 \cdots q_n$ is in L. Assume furthermore that $w^R(q) \overset{*}{\to}_{\delta'} q'$. This means that the automaton can proceed by reading upward the word w and label the term $w^R p'(t_1, \ldots, t_n)$ by q'. Therefore,

if $(\phi : p\gamma \overset{a}{\hookrightarrow} p'w)$ is a transition rule of the system, and if the sequence of control states $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n)$ is in ϕ, then we must add the term $\gamma p(t_1,\ldots,t_n)$ (which is the immediate predecessor of $w^R p'(t_1,\ldots,t_n)$ by the transition rule) to the language of q' (to say that this term is a predecessor of some term which was recognized by q' in the original automaton). This is achieved by applying the saturation rule which adds to the automaton the transition $(\gamma(q_p^{L\cap\langle\phi\rangle}) \rightarrow q')$. The justification of this is in fact subtle. First, if $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n) \in \phi$, we must have $q_1\cdots q_n \in \langle\phi\rangle$. Since states recognize predecessors of terms in their original language, each state q_i is a pair (s_i, p'_i) such that $p'_i = \mathcal{S}(t'_i)$ for some t'_i such that $t_i \in \mathsf{pre}^*(t'_i)$. Now, here is the point where the stability property of ϕ plays a crucial role: it ensures that backward transitions cannot make a term satisfy new constraints (or equivalently, that forward transitions cannot falsify a constraint). Therefore, since $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n) \in \phi$, we must have also $\mathcal{S}(t'_1)\cdots\mathcal{S}(t'_n) \in \phi$, which implies that $q_1\cdots q_n \in \langle\phi\rangle$. On the other hand, assume that $\mathcal{S}(t_1)\cdots\mathcal{S}(t_n) \notin \phi$ but $q_1\cdots q_n \in \langle\phi\rangle$ because $\mathcal{S}(t'_1)\cdots\mathcal{S}(t'_n) \in \phi$. We can show that $\gamma p(t_1,\ldots,t_n)$ is actually in the pre^* image of the original language. Indeed, it is possible in this case to start by rewriting each term t_i to its successor t'_i, which makes the transition rule $(\phi : p\gamma \overset{a}{\hookrightarrow} p'w)$ applicable.

Theorem 5. *For every CDPN M, and for every M-tree automaton $\mathcal{A}$, we can construct an M-tree automaton $\mathcal{A}_{\mathsf{pre}^*}$ such that $L(\mathcal{A}_{\mathsf{pre}^*}) = \mathsf{pre}^*\big(L(\mathcal{A})\big)$.*

Note 3. It is easy to show that, given an M-tree automaton $\mathcal{A}$, the set $\mathsf{pre}_M(\mathcal{A})$ (and in fact also the set $\mathsf{post}_M(\mathcal{A})$) is an effectively M-tree automata definable set.

Then, based on the modelling described in Sections 3 and 5, we can apply Theorems 5 and 4 to check reachability properties and solve flow analysis problems (such as bitvector problems) for multithreaded programs.

Complexity Issues: By Lemma 3, we know that the size of the automaton $\mathcal{A}_{\mathsf{pre}^*}$ is at most exponential in the number of constraints appearing in the given CDPN. In fact, we can prove the following PSPACE lower bound by a reduction of the satisfiability problem for quantified Boolean formulas (QBF).

Theorem 6. *It is at least PSPACE-hard to decide for a given CDPN M, a regular set of M-configurations R and an M-configuration c, whether $c \in \mathsf{pre}^*(R)$ or not.*

Despite the hardness result above, in many interesting cases, we only need a *fixed* number of constraints, which leads to polynomial analysis algorithms. For instance, this is the case when only trivial constraints (i.e., of the form P^*) are used, which corresponds to the case of DPN models. Also, to model parallel calls only one additional constraint is needed, namely $P^*\natural^2$, as we have seen in Section 5. Similarly, we only need one additional constraint for each type of join statement such as $\mathsf{join}_\forall$ or $\mathsf{join}_{\exists k}$. Note that the automata for these constraints can easily be defined by backward deterministic automata of very small sizes. Also for typical properties such as bitvector problems (see Section 3), the initial automaton is always the one recognizing the set of all configurations. Therefore, for an important fragment of CDPN which subsumes (in modelling power) existing formalisms such as PA and PAD, and allows us in addition to model spawn operations, our construction leads to a polynomial analysis algorithm.

However, when return values from parallel processes are taken into account, our construction becomes exponential in the number of used abstract data values. This price is unavoidable since dealing with an unfixed domain of return values is precisely the feature which makes our model complex (see the proof of Theorem 6). Such complexity does not appear for weaker models such as PA or PAD (which have polynomial analysis algorithms [13, 10, 6]) since they cannot handle return values from parallel processes.

Relaxing Stability: We end this section by mentioning the fact that relaxing the stability condition on the constraints appearing in the transition rules of CDPN leads to a model for which pre^* images are not regular in general.

Theorem 7. *There exists a CDPN M with nonstable constraints, and a regular set T of M-configurations such that $\mathsf{pre}^*_M(T)$ is not definable by an M-tree automaton.*

Actually, we can define M s.t. all its transition rules are of the form $\phi : p\gamma \hookrightarrow p'\gamma$ (i.e., without stack manipulation and dynamic creation of processes), and where ϕ is of the simple form pP^*, for $p \in P$. This shows that it is hard to relax the stability condition in the definition of CDPN without losing the property that pre^* preserves regularity.

7 Conclusion

We have defined new formalisms (DPN and CDPN), based on word/term rewrite systems, allowing to model adequately spawn-like commands in multithreaded programs. We have shown that (1) they are more suitable for modelling these commands than previously proposed formalisms (such as PA and PAD), and that (2) they subsume in fact in modelling power these models (concerning CDPN), and allow to handle features these models cannot handle such as return values from parallel processes, various join commands, etc.

We have defined automata-based techniques for computing backward reachability sets of our models. In the case of the basic model of DPN, word automata can be used for this purpose and the construction is simple. In the case of CDPN where constraints on the children are used, the problem of reachability analysis becomes much more delicate. The condition of stability we impose in CDPN on the constraints (guards) appearing in the transition rules seems to be necessary in order to have regular backward reachability sets. Concerning complexity, our construction is exponential in the number of different constraints used in the model, but significant classes of parallel programs can be modelled using a fixed number of constraints (often representable using small automata), and therefore they can be analysed in polynomial time.

Future work includes the extension of our models and our approach to handle synchronisation between parallel processes. Of course, the reachability analysis becomes undecidable in general, but reasonable classes of programs with particular synchronisation policies can be considered (see e.g., [18]), and generic frameworks for defining abstractions (and refining them) can be developed based on our models and our techniques, e.g., following the approaches of [3, 4, 15]. We think also that our techniques could be used to handle models which extend those considered in this paper by allowing a bounded number of context switches, in the spirit of the approach of [19].

References

1. J. Baeten and W. Weijland. Process Algebra. In *Cambridge Tracts in Theoretical Computer Science*, volume 18, 1990.
2. A. Bouajjani, J. Esparza, and O. Maler. Reachability Analysis of Pushdown Automata: Application to Model Checking. In *CONCUR'97*. LNCS 1243, 1997.
3. A. Bouajjani, J. Esparza, and T. Touili. A Generic Approach to the Static Analysis of Concurrent Programs with Procedures. In *POPL'03*. ACM, 2003.
4. A. Bouajjani, J. Esparza, and T. Touili. Reachability Analysis of Synchronised PA systems. In *INFINITY'04*. to appear in ENTCS, 2004.
5. A. Bouajjani, M. Müller-Olm, and T. Touili. Regular Symbolic Analysis of Dynamic Networks of Pushdown Processes. Technical report, LIAFA lab No 2005-05, and University of Dortmund No 798, June 2005.
6. A. Bouajjani and T. Touili. Reachability Analysis of Process Rewrite Systems. In *FSTTCS'03*. LNCS 2914, 2003.
7. A. Bouajjani and T. Touili. On Computing Reachability Sets of Process Rewrite Systems. In *RTA'05*. LNCS, 2005.
8. A. Bruggemann-Klein, M. Murata, and D. Wood. Regular Tree and Regular Hedge Languages over Unranked Alphabets. Research report, 2001.
9. J. Esparza and J. Knoop. An Automata-Theoretic Approach to Interprocedural Data-Flow Analysis. In *FoSSaCS'99*, volume 1578 of *LNCS*, 1999.
10. J. Esparza and A. Podelski. Efficient Algorithms for pre* and post* on Interprocedural Parallel Flow Graphs. In *POPL'00*. ACM, 2000.
11. A. Finkel, B. Willems, and P. Wolper. A Direct Symbolic Approach to Model Checking Pushdown Systems. In *Infinity'97, ENTCS 9*. Elsevier Sci. Pub., 1997.
12. J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
13. D. Lugiez and P. Schnoebelen. The Regular Viewpoint on PA-Processes. *Theoretical Computer Science*, 274(1-2):89–115, 2002.
14. R. Mayr. Decidability and Complexity of Model Checking Problems for Infinite-State Systems. Phd. thesis, Technical University Munich, 1998.
15. M. Müller-Olm. Variations on Constants. Habilitationsschrift, Fachbereich Informatik, Universität Dortmund, 2002.
16. M. Müller-Olm. Precise Interprocedural Dependence Analysis of Parallel Programs. *Theoretical Computer Science*, 311:325–388, 2004.
17. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
18. S. Qadeer, S. Rajamani, and J. Rehof. Procedure Summaries for Model Checking Multithreaded Software. In *POPL'04*, 2004.
19. S. Qadeer and J. Rehof. Context-Bounded Model-Checking of Concurrent Software. In *TACAS'05*. LNCS 3440, 2005.
20. H. Seidl and B. Steffen. Constraint-based Interprocedural Analysis of Parallel Programs. In *ESOP'2000*. LNCS 1782, 2000.
21. T. Touili. Dealing with Communication for Dynamic Multithreaded Recursive Programs. In *1st VISSAS workshop*, March 2005. Invited Paper.

Appendix

Without further ado, we present here the proofs that were omitted due to lack of space. We assume the notation and definitions from the main body of the paper.

A Proof of Theorem 1,a

Theorem 1,a. There is a DPN M with a single control state, and an M-configuration c such that $\mathsf{Traces}_M(c) = \mathcal{L}$.

Proof: Consider the DPN $M = (Act, P, \Gamma, \Delta)$ with

- $Act = \{a, b, c, d\}$,
- $P = \{p\}$,
- $\Gamma = \{A, B, C, D\}$,
- $\Delta = \{pA \overset{a}{\hookrightarrow} pAB, pA \overset{a}{\hookrightarrow} pB \rhd pC, pB \overset{b}{\hookrightarrow} p, pC \overset{c}{\hookrightarrow} pCD, pC \overset{c}{\hookrightarrow} pD, pD \overset{d}{\hookrightarrow} p\}$.

Ignoring the spawn in the second rule, the first three rules describe a push down automaton that from control state p and initial stack contents A accepts the language $a^n b^n$, $n > 0$ upon empty stack. This push down automaton generates the traces $a^n b^{n'}$ with $n \geq n' \geq 0$. When moving from the mode in which it generates a's to the mode in which it generates b's, i.e., when executing the second rule, it spawns a second push down automaton that accepts $c^m d^m$, $m > 0$, upon empty stack and which has the traces $c^m d^{m'}$ for $m \geq m' \geq 0$. These traces are arbitrarily interleaved with the sequence of b's generated by the first push down automaton. Hence we have: $\mathsf{Traces}_M(pA) = \mathcal{L}$. □

B Proof of Theorem 1,b

Theorem 1,b. There is no PA system Δ and no variable A such that $\mathsf{Traces}_\Delta(A) = \mathcal{L}$.

In order to prove Theorem 1,b, we need to introduce some definitions first. Let $r = a_1 \cdots a_l \in Act^*$ be a trace and $I = \{i_1, \ldots, i_k\}$ be a subset of positions in r such that $1 \leq i_1 < i_2 < \cdots < i_k \leq l$. Then, $r|I$ denotes the trace $a_{i_1} \cdots a_{i_k}$. We write $|r|$ for the length of r.

Given two traces $r, r' \in Act^*$, the shuffle language $r \otimes r' \subseteq Act^*$ is given by: $r \otimes r' = \{s \in Act^* : \exists I \subseteq \{1, \ldots, |s|\}.\ s|I = r \text{ and } s|(\{1, \ldots, |s|\} \setminus I) = r'\}$. This definition is lifted to sets of traces in the obvious way.

A *shuffle-constraint system* S over a finite set of variables Y—the variables in Y range over subsets of Act^*—is a finite set of subset constraints of the form $A \supseteq t_i$, where A is a variable from Y and t_i is a term formed with the operators "$\cdot$" (concatenation) and "$\otimes$" (shuffle) from the variables in Y and the languages $\{a\}$, for $a \in Act$, and $\{\varepsilon\}$. Both concatenation and shuffle operations distribute over arbitrary unions and are thus monotonic and even continuous. Thus, each shuffle-constraint system S has a smallest solution $\mu S : Y \to 2^{Act^*}$ by the Knaster-Tarski fixpoint theorem. It can be seen that shuffle-constraint systems generalize context-free grammars to a parallel setting. It can

be shown [20, 16, 15, 4] that these systems generate precisely the set of traces that can be generated by PA processes: processes definable by a set of rewrite rules of the form $A \to t$ where A is a process variable, and t is a term built from process variables, sequential composition, and asynchronous parallel composition [1].

Then, part b of Theorem 1 is an immediate consequence of the following result:

Theorem 8. *There is no shuffle-constraint system S over some set of variables Y such that $(\mu S)(A) = \mathcal{L}$ for some $A \in Y$.*

Proof: Assume there is a shuffle-constraint system S over Y such that $(\mu S)(A) = \mathcal{L}$ for some $A \in Y$.

Our goal is to exhibit a pumping argument in order to deduce a contradiction. We first argue that for each word in $z \in (\mu S)(A)$ we can find a *justifying tree* similar to a derivation tree for words of a context-free language. In a derivation tree in a context-free language, we just need to keep the non-terminals and the terminals, because only the concatenation operator is applied. Here we need to distinguish applications of concatenation from applications of shuffle. Therefore, we put these operators at inner nodes of the tree. We also keep the variables/non-terminals in order to get a handle for the pumping argument. Thus, justifying trees are obtained in the following way:

1. we start from a tree that consists just of a root annotated A.
2. then we iterate the following step until all leafs of the tree are annotated by an action: we replace a nonterminal B at a leaf by a subtree with root B that has a single successor that is the root of a tree corresponding to t for some constraint $B \supseteq t$ of S.

We can assign a set of action sequences $L(T)$ to each justifying tree T in a natural inductive way: a leaf annotated with $a \in Act$ is assigned the language $\{a\}$; to an inner node annotated with a concatenation or shuffle operator, we assign the language obtained as the concatenation or shuffle, respectively of the languages associated with the subtrees; finally for an inner node anotated with a variable, we associate the language of its subtree. Clearly, the language of each justifying tree is contained in $(\mu S)(A)$ because it is the language of a finite unfolding of the constraints. Conversely, we can find for each word $z \in (\mu S)(A)$ a justifying tree T: this follows from the well-known fixpoint theorem of Kleene. All the operators used in a shuffle-constraint system are continuous. Therefore, by Kleene's fixpoint theorem, for each action sequence $w \in (\mu S)(A)$ there is $k \geq 0$ such that z is contained in the k-fold unfolding of the constraint system. This k-fold unfolding gives rise to a justifying tree.

Due to the presence of shuffle operators in the tree we do not necessarily find z at the frontier of a justifying tree for z but the frontier is always some reordering of z. Note, however, that the word at the frontier of a justifying tree T always belongs to $L(T)$ because the concatenation of two languages is contained in their shuffle.

Consider now a justifying tree T for the word $z \stackrel{\text{def}}{=} a^p c^p b^p d^p$ for a $p > 0$. If we choose p big enough, we can find analogously to the well-known Ogden lemma for contextfree languages [12] in the tree the situation pictured in Fig. 1,a such that vx contains the letter b. Then we can "pump" the part beween the two occurences of A. In particular, $w_n \stackrel{\text{def}}{=} uv^n wx^n y \in L(T) \subseteq \mathcal{L}$ for all $n \geq 0$. As $L(T) \subseteq \mathcal{L}$, vx must contain at

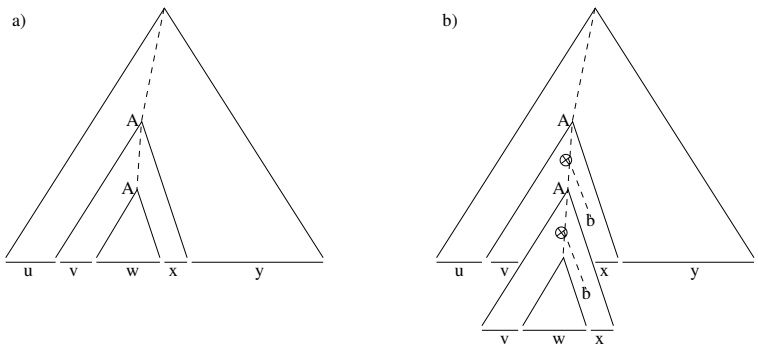

Fig. 1. Repetition of variable A on the path to a b symbol.

least as many a's than b's. By looking at the leftmost v in the word $w_2 = uvvwxxy \in \mathcal{L}$ we infer that v cannot contain any of the letters b, c, or d because $\mathcal{L}$ contains no words in which one of the letters b, c or d is left of an a (and a is found in either v or x).

If x would contain a d, vx must contain a c for counting reasons as $w_n \in \mathcal{L}$ for all n. As v cannot contain a c (see above), such a c must appear in x. But then in w_2 a c would occur to the right of a d which is not allowed by $\mathcal{L}$. Hence x does not contain a d.

Of course there is also no d in u because otherwise we would have a d left of an a already in w_1 which is forbidden by $\mathcal{L}$.

There can also be no d in w: assume there would be a d in w. In w all the d symbols appear right of all b symbols. Therefore, there must be a shuffle operator between the two occurence of A in order to generate w with T. But then we can generate with the tree for w_2 also a word in which an a appear right of a b in contradiction to $\mathcal{L}$ (see Fig. 1,b).

We have seen so far that all the d symbols are contained in y. By a second application of the analog of the Ogden lemma we can find another repetition in the tree as pictured in Fig. 2,a such that $v'x'$ contains the letter d. By pumping, we then have that $w'_n \stackrel{\text{def}}{=} u'v'^n w'x'^n y \in L(T) \subseteq \mathcal{L}$ for all $n \geq 0$. Again by a counting argument, $v'x'$ must contain at least as many c symbols as d symbols and by looking at the word $w'_2 \in \mathcal{L}$ we see that v' contains a c symbol but no d symbol and x' a d symbol but no c symbol.

We now distinguish three cases how the two nodes annotated B in Fig. 2,a are situated in the tree relative to the nodes annotated A in Fig. 1,a.

- Both B-nodes lie on the path to the upper A-node (Fig. 2,b). Then vwx is contained in w' and w' contains an a. But as v' contains a c, this implies that $w_1 \notin \mathcal{L}$. Contradiction!
- The upper B-node lies on the path to the upper A-node but the lower B-node does not (Fig. 2,c). As all d symbols are in y the other B-symbol must generate a subtree with a frontier contained in y. Thus, vwx is contained in v'. But this implies that in w'_2 there is a b left of an a which contradict $w'_2 \in \mathcal{L}$.
- None of the B-nodes lies on the path to the upper A node (Fig. 2,d). As all d symbols are in y, this implies that $v'w'x'$ is contained in y. As $v'w'x'$ contains a c and all the c symbols are between the a and the b symbols in the word $a^p c^p b^p d^p \in L(T)$ there

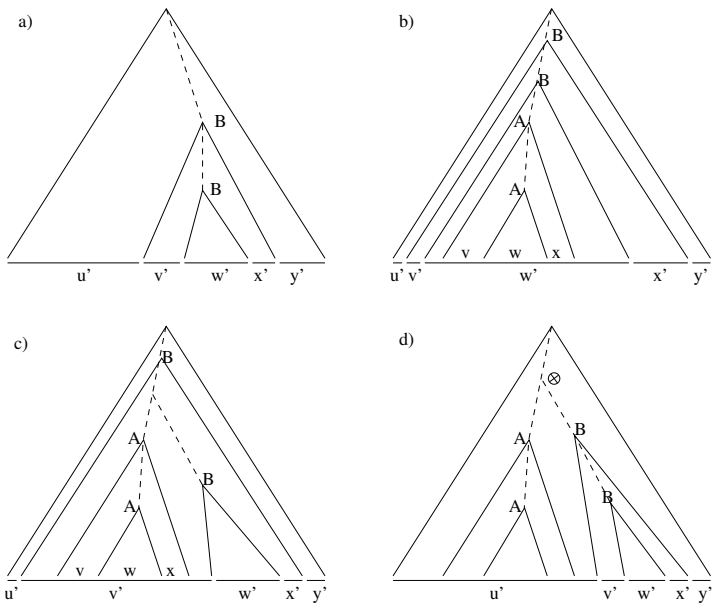

Fig. 2. Repetition of variable B on the path to a d symbol.

must be a shuffle operator at the join point of the two pathes from the upper A and upper B to the root of the tree. But then $L(T)$ also contains a word in which a c is left of an a which contradicts $L(T) \subseteq \mathcal{L}$.

□

C Proof of Theorem 2

Theorem 2. For every DPN M and any context-free set C of M-configurations, the set $\mathsf{post}^*(C)$ is context-free and effectively constructible in polynomial time.

Proof: Let $M = (Act, P, \Gamma, \Delta)$ be a DPN and C be a regular set of configurations. First, we show that, for every pair $(p,\gamma) \in P \times \Gamma$, the set $\mathsf{post}^*(\{p\gamma\})$ is effectively definable by means of a context-free grammar.

We define a set of nonterminal symbol V_N as the smallest set such that:

- If $p, p' \in P$ and $\gamma \in \Gamma$, then $\langle p,\gamma\rangle \in V_N$ and $\langle p,\gamma,p'\rangle \in V_N$,
- If $p\gamma \overset{a}{\hookrightarrow} p_1 w_1[\triangleright p_2 w_2] \in \Delta$ and $p' \in P$, then $\langle p_1, w_1\rangle \in V_N$, $[\langle p_2, w_2\rangle \in V_N,]$ and $\forall p' \in P$, $\langle p_1, w_1, p'\rangle \in V_N$.

The set of productions is the smallest set such that:

- If $p \in P$ and $\gamma \in \Gamma$, then we have the production

$$\langle p,\gamma\rangle \to p\gamma$$

– For every rule
$$p\gamma \xhookrightarrow{a} p_1\alpha_1\cdots\alpha_n[\triangleright p_2\beta_1\cdots\beta_m]$$
we have the productions:
$$\begin{aligned}\langle p,\gamma\rangle &\to [\langle p_2,\beta_1\cdots\beta_m\rangle]\langle p_1,\alpha_1\cdots\alpha_n\rangle\\ \langle p,\gamma,p'\rangle &\to [\langle p_2,\beta_1\cdots\beta_m\rangle]\langle p_1,\alpha_1\cdots\alpha_n,p'\rangle \text{ where } p'\in P\end{aligned}$$
– $\forall p,p'\in P$, $\forall\gamma_1\cdots\gamma_n\in\Gamma$, we have the productions
$$\begin{aligned}\langle p,\varepsilon\rangle &\to p\\ \langle p,\gamma_1\cdots\gamma_n\rangle &\to \langle p,\gamma_1\rangle\gamma_2\cdots\gamma_n \text{ where } n\geq 1\\ \langle p,\gamma_1\cdots\gamma_n\rangle &\to \langle p,\gamma_1,q_1\rangle\langle q_1,\gamma_2,q_2\rangle\cdots\langle q_{i-1},\gamma_i,q_i\rangle\langle q_i,\gamma_{i+1}\rangle\gamma_{i+2}\cdots\gamma_n\\ &\qquad \text{where } n\geq 2, i\in\{1,\ldots,n-1\}, \text{ and } q_1,\ldots,q_i\in P\end{aligned}$$
and the productions
$$\begin{aligned}\langle p,\varepsilon,p\rangle &\to \varepsilon\\ \langle p,\gamma_1\cdots\gamma_n,p'\rangle &\to \langle p,\gamma_1,q_1\rangle\langle q_1,\gamma_2,q_2\rangle\cdots\langle q_{n-1},\gamma_n,p'\rangle\\ &\qquad \text{where } n\geq 1, \text{ and } q_1,\ldots,q_{n-1}\in P.\end{aligned}$$

Then, it can be checked that teh following holds.

Lemma 5. *$\forall p,p'\in P$, $\forall w\in\Gamma^*$, we have*

– $L(\langle p,w\rangle)=\mathsf{post}^*(\{pw\})$.
– $L(\langle p,w,p'\rangle)=\left(\mathsf{post}^*(\{pw\})\cap\Sigma^*p'\right)(p')^{-1}$. □

Now, we define a transducer τ which associates with every configuration $c\in\Sigma^*$ the set $\mathsf{post}^*(c)$. The transducer τ has a finite set of states and transitions labeled by pairs of the form (w,L) where w is an input word, and L is a context-free set of output words. The set of states of τ is $\{q_0,q_{copy}\}\cup\{\widehat{p}\ :\ p\in P\}$, the state q_0 is the unique initial and accepting state, and the set of transitions is as follows:
$$\begin{aligned}q_0 &\xrightarrow{(p\gamma,\langle p,\gamma\rangle)} q_{copy}\\ q_0 &\xrightarrow{(p\gamma,\langle p,\gamma,p'\rangle)} \widehat{p'}\\ \widehat{p} &\xrightarrow{(\gamma,\langle p,\gamma\rangle)} q_{copy}\\ \widehat{p} &\xrightarrow{(\gamma,\langle p,\gamma,p'\rangle)} \widehat{p'}\\ q_{copy} &\xrightarrow{(p,p)} q_{copy}\\ q_{copy} &\xrightarrow{(\gamma,\gamma)} q_{copy}\\ q_{copy} &\xrightarrow{(\varepsilon,\varepsilon)} q_0\end{aligned}$$

The result follows immediatly from the fact that context-free languages are closed under context-free transductions, i.e., given a context-free set of configurations C (effectively defined by, e.g., a context-free grammar), the set $\tau(C)$ is context-free and effectively constructible. □

D Proof of Theorem 3

Let M be a DPN, $\mathcal{A}$ be an M-automaton and $\mathcal{A}_{\mathsf{pre}^*}$ be the automaton obtained by the saturation procedure described in Section 4.

Theorem 3. $L(\mathcal{A}_{\mathsf{pre}^*}) = \mathsf{pre}^*_M\big(L(\mathcal{A})\big)$.

Let us first consider the easier inclusion.

Lemma 6. $\mathsf{pre}^*_M(L(\mathcal{A})) \subseteq L(\mathcal{A}_{\mathsf{pre}^*})$.

Proof: Clearly, $\mathsf{pre}^*_M(L(\mathcal{A})) = \bigcup_{k\geq 0} \mathsf{pre}^k_M(L(\mathcal{A}))$. We show by induction on k that $\mathsf{pre}^k_M(L(\mathcal{A})) \subseteq L(\mathcal{A}_{\mathsf{pre}^*})$ for all $k \geq 0$.

The induction base, $k = 0$, is obvious as the saturation procedure just adds transitions such that $pre^0(L(\mathcal{A})) = L(\mathcal{A}) \subseteq L(\mathcal{A}_{\mathsf{pre}^*})$.

Now, suppose $k \geq 0$ is given and assume that $\mathsf{pre}^k_M(L(\mathcal{A})) \subseteq L(\mathcal{A}_{\mathsf{pre}^*})$ (induction hypothesis). Consider an arbitrary configuration $c \in \mathsf{pre}^{k+1}_M(L(\mathcal{A}))$. Then there is a configuration $d \in \mathsf{pre}^k_M(L(\mathcal{A}))$ and an action $a \in Act$ such that $c \xrightarrow{a} d$, i.e., there is a rule $p\gamma \overset{a}{\hookrightarrow} p_1w_1$ or $p\gamma \overset{a}{\hookrightarrow} p_1w_1 \rhd p_2w_2$ in Δ as well as $u, v \in (P \cup \Gamma)^*$ such that $c = up\gamma v$ and $d = urv$ for $r = p_1w_1$ or $r = p_2w_2p_1w_1$, respectively. By the induction hypothesis, d is accepted by $\mathcal{A}_{\mathsf{pre}^*}$, i.e., there are states s, s', s'' such that:

$$s^0 \xrightarrow{u}_{\delta'} s \xrightarrow{r}_{\delta'} s' \xrightarrow{v}_{\delta'} s'' \in F\,.$$

In particular we have $s \xrightarrow{r}_{\delta'} s'$, which implies $(s_p, \gamma, s') \in \delta'$ because the two implications (R1) and (R2) are valid upon termination of the saturation algorithm. Thus, we have

$$s^0 \xrightarrow{u}_{\delta'} s \xrightarrow{p} s_p \xrightarrow{\gamma}_{\delta'} s' \xrightarrow{v}_{\delta'} s'' \in F\,.$$

This shows that $c = up\gamma v$ is accepted by $\mathcal{A}_{\mathsf{pre}^*}$. □

The crucial lemma for the remaining inclusion is this.

Lemma 7. *Suppose $w \in Conf_M$, $t \in S_c$, $p \in P$. The following is true for all transition relations $\bar{\delta}$ that appear as intermediate values in the saturation algorithm:*

If $s^0 \xrightarrow{w}_{\bar{\delta}} t_p$ then there is a w^* with $s^0 \xrightarrow{w^*}_{\bar{\delta}} t$ and $w \to_\Delta w^*p$.

Proof: The transition relation δ' of $\mathcal{A}_{\mathsf{pre}^*}$ is obtained by successively adding transitions to δ. We show that the property claimed in the lemma is valid for δ and remains true under each single addition of a transition.

Firstly, we persuade ourselves that it is true for δ: as the p-transition from t is the only possible transition to t_p, w must be of the from $w = w^*p$ and we must have $s^0 \xrightarrow{w^*}_{\delta} t$ as required; $w \to_\Delta w^*p$ holds trivially if $w = w^*p$.

Secondly, suppose the property is true for a relation $\bar{\delta}$ and assume that $\bar{\delta}'$ is obtained from $\bar{\delta}$ by a single saturation step. Assume this saturation step considers the rule $p_0\gamma \overset{a}{\hookrightarrow} p_1w_1$ or $p_0\gamma \overset{a}{\hookrightarrow} p_1w_1 \rhd p_2w_2$ and the states s, s' with $s \xrightarrow{r}_{\bar{\delta}} s'$ for $r = p_1w_1$ or $r = p_2w_2p_1w_1$, respectively, such that $\bar{\delta}' = \bar{\delta} \cup \{(s_{p_0}, \gamma, s')\}$. We show the property

claimed in the lemma for $\bar{\delta}'$ by induction over the number n of applications of the new transition (s_{p_0}, γ, s') in transition sequences $s^0 \xrightarrow{w}_{\bar{\delta}'} t_p$.

If the new transition is not used in a transition sequence $s^0 \xrightarrow{w}_{\bar{\delta}'} t_p$ ($n = 0$, Base Case) we also have $s^0 \xrightarrow{w}_{\bar{\delta}} t_p$ and we are done by the assumption that $\bar{\delta}$ satisfies the property.

So assume for some $n > 0$ that we are given a transition sequence $s^0 \xrightarrow{w}_{\bar{\delta}'} t_p$ that uses the new transition (s_{p_0}, γ, s') n times. Assume that a word w^* with the properties claimed in the lemma exists for all transition sequences from s^0 to t_p that use the new transition less than n times (Induction Hypothesis). By considering the first time the new transition is used in the transition sequence, we can write w as $w = w_1 x w_2$ such that

$$s^0 \xrightarrow{w_1}_{\bar{\delta}} s_{p_0} \xrightarrow{\gamma}_{\bar{\delta}'} s' \xrightarrow{w_2} t_p\,.$$

By the assumption that $\bar{\delta}$ satisfies the property, we can find w_1^* with $s^0 \xrightarrow{w_1^*}_{\delta} s$ and $w \rightarrow_\Delta w_1^* p_0$. From the transitions we have seen up to now, we can construct the transition sequence

$$s^0 \xrightarrow{w_1^*}_{\delta} s \xrightarrow{r}_{\bar{\delta}} s' \xrightarrow{w_2}_{\bar{\delta}'} t_p$$

for the word $w_1^* r w_2$ that uses the new transition (s_p, x, s') only $n-1$ times. (As $\delta \subseteq \bar{\delta} \subseteq \bar{\delta}'$ the transition sequence really consists of $\bar{\delta}'$ transitions.) From the induction hypothesis we can now infer that there is w^* with $s^0 \xrightarrow{w^*}_{\delta} t$ and $w_1^* r w_2 \rightarrow_\Delta w^* p$. Combining the Δ-transitions we have seen so far and applying the rule $p_0\gamma \overset{a}{\hookrightarrow} p_1 w_1$ or $p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2$, respectively, we get

$$w = w_1 \gamma w_2 \rightarrow_\Delta w_1^* p_0 \gamma w_2 \rightarrow_\Delta w_1^* r w_2 \rightarrow_\Delta w^* p$$

such that w^* has all the required properties. □

We are now well prepared for the proof of the remaining inclusion.

Lemma 8. $L(\mathcal{A}_{\mathsf{pre}^*}) \subseteq \mathsf{pre}_M^*(L(\mathcal{A}))$.

Proof: Let δ_i be the transition relation obtained after i transitions have been added to δ in the saturation procedure and let $\mathcal{A}_i$ be the automaton $\mathcal{A}_i = (Q, \Sigma, \delta_i, s^0, F)$. We show by induction over i that $L(\mathcal{A}_i) \subseteq \mathsf{pre}_M^*(L(\mathcal{A}))$.

For $\mathcal{A}_0 = \mathcal{A}$ this is trivially true, as $L(\mathcal{A}) \subseteq \mathsf{pre}_M^*(L(\mathcal{A}))$.

So suppose we are given $i > 0$ and assume that $L(\mathcal{A}_{i-1}) \subseteq \mathsf{pre}_M^*(L(\mathcal{A}))$. Assume that the i'th saturation step considers the rule $p\gamma \overset{a}{\hookrightarrow} p_1 w_1$ or $p\gamma \overset{a}{\hookrightarrow} p_1 w_1 \rhd p_2 w_2$ and the states s, s' with $s \xrightarrow{r}_{\bar{\delta}} s'$ for $r = p_1 w_1$ or $r = p_2 w_2 p_1 w_1$, respectively. Then $\delta_i = \delta_{i-1} \cup \{(s_p, \gamma, s')\}$. We show that for all accepting runs $s^0 \xrightarrow{w}_{\delta_i} s_f \in F$ we have $w \in \mathsf{pre}_M^*(L(\mathcal{A}))$. We do so by induction over the number n of applications of the new transition (s_p, γ, s') in the accepting run $s^0 \xrightarrow{w}_{\delta_i} s_f \in F$.

If there is no application of the new transition in a given accepting run $s^0 \xrightarrow{w}_{\delta_i} s_f \in F$ (Base Case), this run is also an accepting run of $\mathcal{A}_{i-1}$. Hence $w \in \mathsf{pre}_M^*(L(\mathcal{A}))$ follows from the assumption $L(\mathcal{A}_{i-1}) \subseteq \mathsf{pre}_P^*(L(\mathcal{A}))$.

If there are $n > 0$ applications of the new transition, we can write w as $w = u\gamma v$ and the accepting run as

$$s^0 \xrightarrow{u}_{\delta_{i-1}} s_p \xrightarrow{\gamma} s' \xrightarrow{v}_{\delta_i} s_f \in F$$

by focusing on the first application of the new transition. By Lemma 7 there is u^* with $s^0 \xrightarrow{u^*}_{\delta} s$ and $u \rightarrow_\Delta u^* p$. Consequently, we have

$$s^0 \xrightarrow{u^*}_{\delta} s \xrightarrow{r}_{\delta_{i-1}} s' \xrightarrow{v}_{\delta_i} s_f \in F$$

such that the word $u^* rv$ is accepted by $\mathcal{A}_i$ with less than n applications of the new transition. By the induction hypothesis, we have thus $u^* rv \in \mathsf{pre}^*(L(\mathcal{A}))$. On the other hand, we can show that w can evolve to this word:

$$w = u\gamma v \rightarrow_\Delta u^* p\gamma v \rightarrow_\Delta u^* rv \in \mathsf{pre}^*(L(\mathcal{A}))\,.$$

Consequently, $w \in \mathsf{pre}^*(L(\mathcal{A}))$. □

E Proof of Lemma 3

Lemma 3. The family Λ is finite. Assuming that all languages and constraints appearing in rules δ_P and Δ are given by backward-deterministic finite-state automata of size at most K, the number of elements of Λ is in $O(K^{n+1})$ where n is the number of different constraints appearing in the rules of Δ.

Proof (Sketch):

Let $\phi_1, \ldots, \phi_n$ be all the constraints that appear in the rules of Δ. Let $B_1, \ldots, B_n$ be n word automata that recognize $\langle\phi_1\rangle, \ldots, \langle\phi_n\rangle$, respectively. Let $S_1, \ldots, S_n$ be the sets of states of $B_1, \ldots, B_n$ respectively.

Suppose w.l.o.g. that δ_P contains a unique rule of the form $p(L) \rightarrow q$. Let $D = (S, S_I, S_F, T)$ be a word automaton that recognizes L, where S is the set of states, S_I and S_F are respectively the set of initial and final states, and T is the set of transitions.

Consider first the case where $\phi_1 = \cdots = \phi_n = P^*$. In this case, it is easy to see that the elements of Λ are recognized by word automata of the form $D^i = (S, S_I, S_F^i, T)$, that differ from D only by the final state (it is unique since the automata are backward-deterministic). Indeed, performing the right-quotient corresponds to changing the final states. For example, if $F = \{e_f\}$, and if T has transitions from e_i to e_f labeled with q for $1 \leq i \leq k$, then Lq^{-1} is recognized by the automaton $D' = (S, S_I, S'_F, T)$, where $S'_F = \{e_1, \ldots, e_k\}$. It is then easy to see that in this case, Λ has at most $O(|S|)$ elements (since each automaton has a single final state due to the fact that the automata are backward-deterministic).

Let us consider now the general case. It is easy to see that the languages of Λ can be recognized by automata having at most $S \times S_1 \times \cdots \times S_n$ as states. Indeed, the intersection corresponds to automata products, and the right-quotient corresponds to changing the final state as explained above. Therefore, Λ contains then at most $O(|S||S_1| \cdots |S_n|)$ elements. □

F Proof of Lemma 4

Lemma 4. For every $q \in Q_P$, $L_q^{\delta'} = pre^*(L_q^{\delta_P})$.

Proof:
$\subseteq$: First, we show that

$$t \xrightarrow{*}_{\delta'} q \Rightarrow t \in pre^*(L_q^{\delta_P})$$

For that, we show by induction on i that:

$$t \xrightarrow{*}_{\delta'_i} q \Rightarrow t \in pre^*(L_q^{\delta_P})$$

- The case where $i = 0$ is straightforward, since any possible derivation $t \xrightarrow{*}_{\delta'_0} q$ contains only rules from δ. Therefore, in this case, we have that $t \in L_q$.
- $i > 0$. Let $t \xrightarrow{*}_{\delta'_i} q$, and let n be the number of applications of a rule in $\delta'_i \setminus \delta'_{i-1}$ in this derivation. We write: $t \xrightarrow{*}{}^{n}_{\delta'_i} q$. We proceed by induction on n:
  - If $n = 0$, this means that only the rules of δ'_{i-1} are used, and we get the result by induction on i.
  - Let $n > 0$. There are two cases depending on the rule of $\delta'_i \setminus \delta'_{i-1}$. Suppose that the rule of $\delta'_i \setminus \delta'_{i-1}$ is added by (α_1), the case where it is added by (α_2) is similar.
    The rule of $\delta'_i \setminus \delta'_{i-1}$ is then of the form $\gamma(q_p^{L \cap \langle \phi \rangle}) \to q'$, added to δ' because there exist a rule $(\phi : p\gamma \xhookrightarrow{a} p'w)$ in Δ and a state $q \in Q_P$ such that $p'(L) \to q \in \delta'_0$ and $w^R(q) \xrightarrow{*}_{\delta'_{i-1}} q'$.
    Let then $t_1, \ldots, t_m$ be m terms such that $\mathcal{S}(t_i) = p_i$ $1 \leq i \leq m$, and C be a context such that

$$t = C\big[\gamma p(t_1, \ldots, t_m)\big]$$

and let $q_1, \ldots, q_m$ be states s.t. $t_i \xrightarrow{*}{}^{n-1}_{\delta'_i} q_i$ for $1 \leq i \leq m$, $q_1 \cdots q_m \in L \cap \langle \phi \rangle$, and

$$t = C\big[\gamma p(t_1, \ldots, t_m)\big] \xrightarrow{*}{}^{n-1}_{\delta'_i} C\big[\gamma p(q_1, \ldots, q_m)\big] \to_{\delta'_0} C\big[\gamma(q_p^{L \cap \langle \phi \rangle})\big] \xrightarrow{*}_{\delta'_i}$$

$$C(q') \xrightarrow{*}_{\delta'_{i-1}} q$$

Since for $1 \leq i \leq m$, $t_i \xrightarrow{*}{}^{n-1}_{\delta'_i} q_i$, we get by induction that $t_i \in pre^*(L_{q_i}^{\delta_P})$. There exist then m terms $t'_1, \ldots, t'_m$ such that $t'_i \xrightarrow{*}_{\delta_P} q_i$ and $t_i \in pre^*(t'_i)$. Therefore, q_i is of the form (s_i, p'_i) where $s_i \in Q$ and $p'_i = \mathcal{S}(t'_i)$
Then, since $q_1 \cdots q_m \in L$, we have:

$$t' = C\big[w^R p'(t'_1, \ldots, t'_m)\big] \xrightarrow{*}{}^{n-1}_{\delta'_i} C\big[w^R p'(q_1, \ldots, q_m)\big] \to_{\delta'_0} C(w^R(q)) \xrightarrow{*}_{\delta'_{i-1}}$$

$$C(q') \xrightarrow{*}_{\delta'_{i-1}} q$$

i.e., $t' = C\big[w^R p'(t'_1,\ldots,t'_m)\big] \xrightarrow{*}{}^{n-1}_{\delta'_i} q$. We get then by induction that

$$C\big[w^R p'(t'_1,\ldots,t'_m)\big] \in pre^*(L_q^{\delta_P})$$

and therefore,

$$t \in pre^*(L_q^{\delta_P})$$

since:

$$C\big[\gamma p(t_1,\ldots,t_m)\big] \in pre^*\Big(C\big[\gamma p(t'_1,\ldots,t'_m)\big]\Big)$$

and

$$C\big[\gamma p(t'_1,\ldots,t'_m)\big] \in pre\Big(C\big[w^R p'(t'_1,\ldots,t'_m)\big]\Big)$$

because $\mathcal{S}(t'_1)\cdots\mathcal{S}(t'_m) \in \phi$ (since $\mathcal{S}(t'_1)\cdots\mathcal{S}(t'_m) = p'_1\cdots p'_m$ and $q_1\cdots q_m = (s_1,p'_1)\cdots(s_m,p'_m) \in \langle\phi\rangle$), and therefore, we can apply the rule $(\phi : p\gamma \xhookrightarrow{a} p'w)$ to $C\big[\gamma p(t'_1,\ldots,t'_m)\big]$ and obtain $C\big[w^R p'(t'_1,\ldots,t'_m)\big]$.

$\supseteq$: For the other direction, since $\delta_P \subseteq \delta'$, we show that if $t \xrightarrow{*}_{\delta'} q$, and $t' \in pre(t)$, then $t' \xrightarrow{*}_{\delta'} q$. Let then such t and t'. There are two cases:

1. t is obtained from t' after a rewriting step using a rule $(\phi : p'\gamma \xhookrightarrow{a} pw)$. Let then C be a context, and $t_1,\ldots,t_n$ be n terms such that $\mathcal{S}(t_i) = p_i$, $p_1\cdots p_n \in \phi$,

   $$t = C\big[w^R p(t_1,\ldots,t_n)\big]$$

   and

   $$t' = C\big[\gamma p'(t_1,\ldots,t_n)\big]$$

   Let then the states $q_1,\ldots,q_n,q,q',q''$, and the rule $p(L) \to q'$ of δ'_0 be such that $q_1\cdots q_n \in L$, and:

   $$t = C\big[w^R p(t_1,\ldots,t_n)\big] \xrightarrow{*}_{\delta'} C\big[w^R p(q_1,\ldots,q_n)\big] \xrightarrow{*}_{\delta'} C\big[w^R(q')\big] \xrightarrow{*}_{\delta'} C[q] \xrightarrow{*}_{\delta'} q''$$

   Then, since Δ contains the rule $(\phi : p'\gamma \xhookrightarrow{a} pw)$, $p(L) \to q' \in \delta'_0$, and $w^R(q') \xrightarrow{*}_{\delta'} q$; the rules α_1 infer that δ' contains also the rule $\gamma(q_p^{L\cap\langle\phi\rangle}) \to q$.

   Therefore we have the following:

   $$t' = C\big[\gamma p'(t_1,\ldots,t_n)\big] \xrightarrow{*}_{\delta'} C\big[\gamma p'(q_1,\ldots,q_n)\big] \xrightarrow{*}_{\delta'} C\big[\gamma(q_{p'}^{L\cap\langle\phi\rangle})\big] \xrightarrow{*}_{\delta'} C[q] \xrightarrow{*}_{\delta'} q''$$

   Indeed, the sequence of states $q_1\cdots q_n$ is in $L\cap\langle\phi\rangle$ since we already know that it is in L, and we show in what follows that it is in $\langle\phi\rangle$:
   Let $s_1,\ldots,s_n \in Q$ and $p'_1,\ldots,p'_n \in P$ be such that $q_i = (s_i,p'_i)$ for $1 \le i \le n$. Then since $t_i \xrightarrow{*}_{\delta'} (s_i,p'_i)$, it follows from the previous direction that $t_i \in pre^*(L^{\delta_P}_{(s_i,p'_i)})$.
   Let then n terms $t'_1,\ldots,t'_n$ such that $t'_i \xrightarrow{*}_{\delta_P} (s_i,p'_i)$, $t_i \in pre^*(t'_i)$, and $p'_i = \mathcal{S}(t'_i)$ for $1 \le i \le n$ (Lemma 2). Since $(p_i,p'_i) \in \rho^*_\Delta$ and $p_1\cdots p_n \in \phi$, Lemma 1 infers that $p'_1\cdots p'_n \in \phi$, and therefore that $q_1\cdots q_n = (s_1,p'_1)\cdots(s_n,p'_n) \in \langle\phi\rangle$.

2. t is obtained from t' after a rewriting step using a rule $(\phi : p'\gamma \overset{a}{\hookrightarrow} pw \rhd p''w_1)$. Let then C be a context, and $t_1,\ldots,t_n$ be n terms such that $\mathcal{S}(t_i) = p_i$, $p_1 \cdots p_n \in \phi$,

$$t = C\big[w^R p(t_1,\ldots,t_n,w_1^R p'')\big]$$

and

$$t' = C\big[\gamma p'(t_1,\ldots,t_n)\big]$$

Let then the states $q_1,\ldots,q_n,q_{n+1},q,q',q''$, and the rule $p(L) \to q'$ of δ_0' such that $q_1 \cdots q_n q_{n+1} \in L$, and:

$$t = C\big[w^R p(t_1,\ldots,t_n,w_1^R p'')\big] \xrightarrow{*}_{\delta'} C\big[w^R p(q_1,\ldots,q_n,q_{n+1})\big] \xrightarrow{*}_{\delta'} C\big[w^R(q')\big] \xrightarrow{*}_{\delta'} C[q] \xrightarrow{*}_{\delta'} q''$$

Then, since Δ contains the rule $(\phi : p'\gamma \overset{a}{\hookrightarrow} pw \rhd p''w_1)$, $p(L) \to q' \in \delta_0'$, $w^R(q') \xrightarrow{*}_{\delta'} q$, and $w_1^R p'' \xrightarrow{*}_{\delta'} q_{n+1}$; the rules α_2 infer that δ' contains also the rule $\gamma(q_{p'}^{Lq_{n+1}^{-1}\cap\langle\phi\rangle}) \to q$.

Therefore, we have the following:

$$t' = C\big[\gamma p'(t_1,\ldots,t_n)\big] \xrightarrow{*}_{\delta'} C\big[\gamma p'(q_1,\ldots,q_n)\big] \xrightarrow{*}_{\delta'} C\big[\gamma(q_{p'}^{Lq_{n+1}^{-1}\cap\langle\phi\rangle})\big] \xrightarrow{*}_{\delta'} C[q] \xrightarrow{*}_{\delta'} q''$$

Indeed, $Lq_{n+1}^{-1} \cap \langle\phi\rangle$ is in Λ and $p'(Lq_{n+1}^{-1} \cap \langle\phi\rangle) \to q_{p'}^{Lq_{n+1}^{-1}\cap\langle\phi\rangle}$ is in δ_0'. Moreover, since $p_1 \cdots p_n \in \phi$, we can show as previously using Lemma 1 that $q_1 \cdots q_n \in Lq_{n+1}^{-1} \cap \langle\phi\rangle$.

□

G Proof of Theorem 6

Theorem 6. It is at least PSPACE-hard to decide for a given CDPN M, a regular set of M-configurations R and an M-configuration c, whether $c \in \mathsf{pre}^*(R)$ or not.

Proof: We exhibit a reduction of QBF (quantified Boolean formulas), a well-known PSPACE-complete problem [17]. A QBF-instance I is a Boolean formula of the form

$$\exists \mathbf{x}_1 \forall \mathbf{x}_2 \ldots Q_k \mathbf{x}_k : c_1 \wedge \cdots \wedge c_n\,,$$

where Q_k is the quantifier "$\exists$" is n is odd and "$\forall$" if k is even, $X = \{\mathbf{x}_1,\ldots,\mathbf{x}_k\}$ is a set of k Boolean variables that are quantified alternatingly by "$\exists$" and "$\forall$", and each c_i is a disjunction of *literals*, where each literal is a negated or non-negated variable from X. QBF asks us to decide, whether a given QBF-instance is satisfied or not.

Before we describe how to reduce a QBF-instance to a DPN reachability problem we introduce some notation. Let $\sigma : X \overset{part.}{\longrightarrow} \mathbb{B}$ be a partial truth assignment (where $\mathbb{B} = \{\mathsf{tt},\mathsf{ff}\}$ is the set of truth values) and ψ be a Boolean formula the free variables of which are contained in $\mathsf{dom}(\sigma)$. We write $\sigma \models \psi$ if σ satisfies ψ which is defined as usual. For a closed formula, we write $\models \psi$ if $\sigma \models \psi$ for some (and thus all) truth assignments.

From a given QBF-instance I as above we construct the following DPN $M = (Act, P, \Gamma, \Delta)$:

- *Act* consists of a single default action τ; for clarity we omit τ when defining the rules below.
- P contains control states x_i, t_i, and f_i for all $i \in \{1,\ldots,k\}$ and c_j for all $j \in \{1,\ldots,n+1\}$; all these states are distinct. For convenience we refer to state c_1 also as x_{k+1}.
- Γ contains distinct symbols Y_i for all even $i \in \{1,\ldots,k\}$.
- Finally, Δ consists of the following rules: For each $i \in \{1,\ldots,k\}$ we have the two rules

$$x_i \hookrightarrow x_{i+1} \rhd t_j \quad \text{and } x_i \hookrightarrow x_{i+1} \rhd f_j\,, \qquad \text{if } i \text{ is odd and}$$
$$x_i \hookrightarrow x_{i+1} Y_i \rhd t_j \text{ and } c_{k+1} Y_i \hookrightarrow x_{i+1} \rhd f_j\,, \quad \text{if } i \text{ is even.}$$

  For each $j \in \{1,\ldots,n\}$ and each literal l in clause c_j we have the rule

$$\phi : c_j \hookrightarrow c_{j+1}\,,$$

  where $\phi = P^* t_i (P \setminus \{t_i, f_i\})^*$ if $l = x_i$ and $\phi = P^* f_i (P \setminus \{t_i, f_i\})^*$ if $l = \neg x_i$. It is not hard to see that these constraints are stable.

For $w = (p_1,\ldots,p_l) \in P^*$ and $p \in P$ we write $p(w)$ for the term $p(p_1(),\ldots,p_l())$ by a little abuse of notation. It is easy to see that $R \stackrel{\text{def}}{=} \{c_{n+1}(w) \mid w \in P^*\}$ is a regular set of configurations.[5] We claim:

$$x_1(\varepsilon) \in \mathsf{pre}^*(R) \quad \text{iff} \quad \models \exists \mathbf{x}_1 \forall \mathbf{x}_2 \ldots Q_k \mathbf{x}_k : c_1 \wedge \cdots \wedge c_n\,. \tag{1}$$

Clearly, M as well as a DPN tree automaton for R can be constructed from I in logarithmic space such that (1) proves Theorem 6.

Before we prove (1) we discuss the intuition of the construction. From initial configuration $x_1(\varepsilon)$ the process successively chooses truth values for the variables $\mathbf{x}_1,\ldots,\mathbf{x}_k$. The choice tt (ff) for variable $\mathbf{x}_i$ is recorded by creating a son with control state t_i (f_i). For odd i, i.e. for variables quantified existentially, the choice is non-deterministic by the two transition rules $x_i \hookrightarrow x_{i+1} \rhd t_i$ and $x_i \hookrightarrow x_{i+1} \rhd f_i$. For even i, however, i.e. for variables quantified universally, the process must first choose the value tt as the transition $x_i \hookrightarrow x_{i+1} Y_i \rhd t_i$ is the only transition from state x_i. The transition also records by putting Y_i onto the stack, that is has to choose f_i later. Once validity of the first choice tt has been confirmed the transition $c_{n+1} Y_i \hookrightarrow x_{i+1} \rhd f_i$ is executed that chooses ff as the value for x_i and initiates new choices for the more innermost variables $x_{i'}$, $i < i'$ by going to control state x_{i+1} again. In order to allow overwriting the first choise of a value for x_i by a later choice, the current truth value of x_i is determined by the rightmost son, i.e., the son created last, that has either control state t_i (for tt) or f_i (for ff). In order to prepare for the formal proof, we capture this by defining for a word $w \in P^*$ (representing the control states of the sons) the partial truth assignment $\sigma_w : X \stackrel{part.}{\longrightarrow} \mathbb{B}$:

$$\sigma_w(\mathbf{x}_i) = \begin{cases} \mathsf{tt} & \text{if } w \in P^* t_i (P \setminus \{t_i, f_i\})^* \\ \mathsf{ff} & \text{if } w \in P^* f_i (P \setminus \{t_i, f_i\})^* \\ \text{undefined} & \text{otherwise} \end{cases}$$

[5] R is the language of the M-tree automaton $(\{q_1, q_2\}, \{p(\{\varepsilon\}) \to q_1 \mid p \in P\} \cup \{c_{n+1}(q_1^*) \to q_2\}, \{q_2\})$.

After truth values have been chosen for all the variables, the process is in state $x_{k+1} = c_1$. The transitions from c_j to c_{j+1} are defined in such a way that they are enabled if and only if the clause c_j is satisfied by the current choice of truth values for the variables. Hence, there is a transition sequence bringing the process from c_1 to c_{n+1} if and only if $c_1 \wedge \cdots \wedge c_n$ is satisfied for the current choice of truth values for the variables.

In order to prove claim (1) formally, we first show by induction on j that for all $j \in \{0,\ldots,n\}$ and $w \in P^*$ with $\mathsf{dom}(\sigma_w) = \{\mathbf{x}_1,\ldots,\mathbf{x}_k\}$ the following two properties are valid:

a) $\exists w' \in P^* : c_1(w) \rightarrow^* c_{j+1}(w')$ if and only if $\sigma_w \models c_1 \wedge \cdots \wedge c_j$.
b) $c_1(w) \rightarrow^* c_{j+1}(w')$ implies $w = w'$ for all $w' \in P^*$.

We then use the case $j = n$ as the base case in an inductive proof of the following claim (recall that we write x_{k+1} for c_1): for all $i \in \{1,\ldots,k+1\}$, $w \in P^*$ with $\mathsf{dom}(\sigma_w) \supseteq \{x_1,\ldots,x_{i-1}\}$:

c) $\exists w' \in P^* : x_i(w) \rightarrow^* c_{n+1}(w')$ if and only if $\sigma_w \models Q_i\mathbf{x}_i \ldots Q_k\mathbf{x}_k : c_1 \wedge \cdots \wedge c_n$.
d) $x_i(w) \rightarrow^* c_{n+1}(w')$ implies $\sigma_w(x_l) = \sigma_{w'}(x_l)$ for all $l \in \{1,\ldots,i-1\}$, $w' \in P^*$.

Here we perform the induction downwards, i.e., we start with $i = k+1$ as the base case and argue inductively downwards towards the case $i = 1$. The details of these inductions are left to the reader.

Finally, property c) reads for $i = 1$ and $w = \varepsilon$ as follows:

$$\exists w' \in P^* : x_1(\varepsilon) \rightarrow^* c_{n+1}(w') \quad \text{iff} \quad \sigma_\varepsilon \models Q_1\mathbf{x}_1 \ldots Q_k\mathbf{x}_k : c_1 \wedge \cdots \wedge c_n$$

This equivalence implies the equivalence (1). □

H Proof of Theorem 7

Theorem 7. There exists a CDPN $M = (Act, P, \Gamma, \Delta)$ with nonstable constraints, and a regular set T of M-configurations such that $\mathsf{pre}^*_M(T)$ is not definable by an M-tree automaton.

Proof: Consider the CDPN $M = (Act = \{a\}, P = \{p,q,s,t,s',t'\}, \Gamma = \{\$\}, \Delta)$ where Δ consists of the following transition rules:

(1)	$P^* : p\$ \overset{a}{\hookrightarrow} q\$$	(5)	$P^* : r\$ \overset{a}{\hookrightarrow} p\$$
(2)	$(p+q')P^* : q'\$ \overset{a}{\hookrightarrow} q\$$	(6)	$(r+q)P^* : q\$ \overset{a}{\hookrightarrow} q'\$$
(3)	$(q'+t')P^* : t'\$ \overset{a}{\hookrightarrow} t\$$	(7)	$(q+t)P^* : t\$ \overset{a}{\hookrightarrow} t'\$$
(4)	$(q'+t')P^* : s'\$ \overset{a}{\hookrightarrow} s\$$	(8)	$(q+t)P^* : t\$ \overset{a}{\hookrightarrow} s'\$$

Notice that the transition rules of the model M above do not use (modify) the stacks, and do not create new processes. We consider M-configurations which are in fact 1-ary trees. Therefore, such configurations are sequences of the form $\$^* p_1 \$^* p_2 \cdots \$^* p_n$ (with

the interpretation that the process of index i has the process $i+1$ as unique son). Then, it is easy to check that

$$\mathsf{pre}^*\big((\$s)^*(\$q)^*\big) \cap (\$s)^*(\$t)^*(\$q)^*(\$r)^* = \\ \{(\$s)^i(\$t)^j(\$q)^k(\$r)^\ell \,:\, 0 \leq i,k, \text{ and } 0 \leq j \leq \ell\}$$

which is clearly a nonregular language.

Indeed, starting from a sequence of the form $(\$s)^n(\$q)^m$, the down-most (right-most in the word representation) q in the tree can be rewritten (backward) to p using the rule (1). This information can be transmitted to the top of the tree using constraints: the rules (2) and (3) can be used to propagate upward (to the right according to the word representation) a q' through the q's, and then a t' through the t's (if any), until the first (down-most in the tree, or right-most in the word representation) s is reached and transformed into an s' using the rule (4).

Then, the p which is down in the tree can be rewritten to an r using rule (5). Then, using rules (6) and (7), the states q' are rewritten again to q and states t' to t until s' is reached (which is between the states s and the states t) and transformed to a t.

Of course, several rewriting sequences like the one described above can be running simultaneously since the application of rule (1) can occur at any time and not necessarily at the down-most q in the tree. However, if we restrict our view to reachable configurations of the regular form $(\$s)^*(\$t)^*(\$q)^*(\$r)^*$, then we get precisely the non-regular language given above. □