The Regular Viewpoint on PA-Processes

Denis Lugiez, Philippe Schnoebelen

► To cite this version:

Denis Lugiez, Philippe Schnoebelen. The Regular Viewpoint on PA-Processes. [Research Report] RR-3403, INRIA. 1998, pp.27. inria-00073287

HAL Id: inria-00073287

<https://hal.inria.fr/inria-00073287>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Regular Viewpoint on PA-Processes

Denis Lugiez, Philippe Schnoebelen

N° 3403

Avril 1998

————— THÈME 2 —————

*apport
de recherche*

The Regular Viewpoint on PA-Processes

Denis Lugiez, Philippe Schnoebelen

Thème 2 — Génie logiciel
et calcul symbolique
Projet protheo

Rapport de recherche n3403 — Avril 1998 — 27 pages

Abstract: PA is the process algebra allowing non-determinism, sequential and parallel compositions, and recursion. We suggest a view of PA-processes as *tree languages*.

Our main result is that the set of iterated predecessors of a regular set of PA-processes is a regular tree language, and similarly for iterated successors. Furthermore, the corresponding tree-automata can be built effectively in polynomial-time. This has many immediate applications to verification problems for PA-processes, among which a simple and general model-checking algorithm.

Key-words: Process Algebra, Infinite State Systems, Verification, Model-Checking, Reachability, Tree Automata

(Résumé : *tsvp*)

Le point de vue rationnel sur les processus PA

Résumé : PA est l'algèbre de processus autorisant le non-déterminisme, les compositions séquentielle et parallèle et la récursion. Nous proposons de voir les processus PA comme des *langages d'arbres*.

Notre résultat principal montre que l'ensemble des prédécesseurs itérés d'un ensemble régulier de processus PA est un langage régulier, et de même pour les successeurs itérés. De plus les automates correspondants sont construits en temps polynomial. Ces résultats ont des applications immédiates pour les problèmes de vérifications liés à l'algèbre Pa, notamment un algorithme de model-checking simple et général.

Mots-clé : Algèbre de Processus, Système à infinité d'états, vérification, model-checking, atteignabilité, automates d'arbres.

Introduction

Verification of Infinite State Processes is a very active field of research today in the concurrency-theory community. Of course, there has always been an active Petri-nets community, but researchers involved in process algebra and model-checking really became interested into infinite state processes after the proof that bisimulation was decidable for normed BPA-processes [BBK87]. This prompted several researchers to investigate decidability issues for BPP and BPA (with or without the normedness hypothesis) (see [CHM94, Mol96, BE97] for a partial survey).

From BPA and BPP to PA: BPA is the “non-determinism + sequential composition + recursion” fragment of process algebra. BPP is the “non-determinism + parallel composition + recursion” fragment. PA (from [BEH95]) combines both and is much less tractable. A few years ago, while more and more decidability results for BPP and BPA were presented, PA was still beyond the reach of the current techniques. Then Mayr showed the decidability of reachability for PA processes [May97c], and extended this into decidability of model-checking for PA w.r.t. the EF fragment of CTL [May97b]. This was an important breakthrough, allowing Mayr to successfully attack more powerful process algebras [May97a] while other decidability results for PA were presented by him and other researchers (e.g. [Kuč96, Kuč97, JKM98]).

A field asking for new insights: The decidability proofs from [May97b] (and the following papers) are certainly not trivial. The constructions are quite complex and hard to check. It is not easy to see in which directions the results and/or the proofs could be adapted or generalized without too much trouble. Probably, this complexity cannot be avoided with the techniques currently available in the field. We believe we are at a point where it is more important to look for new insights, concepts and techniques that will simplify the field, rather than trying to further extend already existing results.

Our contribution: In this paper, we show how **tree-automata techniques** greatly help dealing with PA. Our main results are two **Regularity Theorems**, stating that $Post^*(L)$ (resp. $Pre^*(L)$) the set of configurations reachable from (resp. allowing to reach) a configuration in L is a regular tree language when L is, and giving simple polynomial-time constructions for the associated automata. Many important consequences follow directly, including a simple algorithm for model-checking PA-processes.

Why does it work ? The regularity of $Post^*(L)$ and $Pre^*(L)$ could only be obtained after we had the combination of two main insights:

1. the tree-automata techniques that have been proved very powerful in several fields (see [CKSV97]) are useful for the process-algebraic community as well. After all, PA is just a simple term-rewrite system with a special context-sensitive rewriting strategy, not unlike head-rewriting, in presence of the sequential composition operator.

2. the syntactic congruences used to simplify notations in simple process algebras help one get closer to the intended semantics of processes, but they break the regularity of the behavior. The decidability results are much simpler when one only introduces syntactic congruences at a later stage. (Besides, this is a more general approach.)

Plan of the paper: We start by recalling the basics of tree-automata theory (§ 1) before we introduce our definition for the PA process algebra (§ 2). After we explain how sets of PA processes can be seen as tree languages (§ 3) we give a simple proof showing how $Post^*(t)$ and $Pre^*(t)$ are regular tree languages and start listing applications to verification problems. We then move on to $Post^*(L)$ and $Pre^*(L)$ for L a regular language (§ 5). These are our main technical results and we devote § 6 to the important applications in model-checking. We end up with an extension to reachability and model-checking *under constraints* (§ 7) and some simple but important techniques allowing to deal with PA processes modulo structural equivalence (§ 8).

Related work: Several recent works in the field use tree-automata to describe the *behaviour* of systems. We use them to describe set of configurations.

The set of all reachable configurations of a pushdown automaton form a regular (word) language. This was proven in [Büc64] and extended in [Cau92]. Applications to the model-checking of pushdown automata have been proposed in [FWW97, BEM97].

The decidability of the first-order theory of the rewrite relation induced by a ground term rewrite system relies on *ground tree transducers* [DT90] (note that PA is defined by a *conditional* ground rewrite system).

Among the applications we develop for our regularity theorems, several have been suggested by Mayr's work on PA [May97c, May97b] and/or our earlier work on RPPS [KS97a, KS97b].

1 Regular tree languages and tree automata

We recall some basic facts on tree automata and regular tree languages. For more details, the reader is referred to any classical source (e.g. [CDG⁺97, GS97]).

A *ranked alphabet* is a finite set of symbols $\mathcal{F}$ together with an arity function $\eta : \mathcal{F} \rightarrow \mathbb{N}$. This partitions $\mathcal{F}$ according to arities: $\mathcal{F} = \mathcal{F}_0 \cup \mathcal{F}_1 \cup \mathcal{F}_2 \cup \dots$. We write $\mathcal{T}(\mathcal{F})$ the set of terms over $\mathcal{F}$ and call them *finite trees* or just *trees*.

A *tree language* over $\mathcal{F}$ is any subset of $\mathcal{T}(\mathcal{F})$.

A (finite, bottom-up) *tree automaton* $\mathcal{A}$ is a tuple $\langle \mathcal{F}, Q, F, R \rangle$ where $\mathcal{F}$ is a ranked alphabet, $Q = \{q_1, \dots\}$ is a finite set of *states*, $F \subseteq Q$ is the subset of *final states*, and R is a finite set of *transition rules* of the form $f(q_1, \dots, q_n) \mapsto q$ where $n \geq 0$ is the arity $\eta(f)$ of symbol $f \in \mathcal{F}$. *Tree automata with ε -rules* also allow some transition rules of the form $q \mapsto q'$.

The transition rules define a rewrite relation on terms built on $\mathcal{F} \cup Q$ (seeing states from Q as nullary symbols). This works bottom-up. At first the nullary symbols at the leaves

are replaced by states from Q , and then the quasi-leaf symbols immediately on top of leaves from Q are replaced by states from Q . We write $t \xrightarrow{\mathcal{A}} q$ when $t \in \mathcal{T}(\mathcal{F})$ can be rewritten (in some number of steps) to $q \in Q$ and say t is accepted by $\mathcal{A}$ if it can be rewritten into a final state of $\mathcal{A}$. We write $L(\mathcal{A})$ for the set of all terms accepted by $\mathcal{A}$. Any tree language which coincide with $L(\mathcal{A})$ for some $\mathcal{A}$ is a *regular tree language*. Regular tree languages are closed under complementation, union, etc.

An example: Let $\mathcal{F}$ be given by $\mathcal{F}_0 = \{a, b\}$, $\mathcal{F}_1 = \{g\}$ and $\mathcal{F}_2 = \{f\}$. There is an automaton accepting the set of all $t \in \mathcal{T}(\mathcal{F})$ where g occurs an even number of times in t . $\mathcal{A}$ is given by $Q \stackrel{\text{def}}{=} \{q_0, q_1\}$, $R \stackrel{\text{def}}{=} \{a \mapsto q_0, b \mapsto q_0, g(q_0) \mapsto q_1, g(q_1) \mapsto q_0, f(q_0, q_0) \mapsto q_0, f(q_0, q_1) \mapsto q_1, f(q_1, q_0) \mapsto q_1, f(q_1, q_1) \mapsto q_0\}$ and $F \stackrel{\text{def}}{=} \{q_0\}$.

Let t be $g(f(g(a), b))$. $\mathcal{A}$ rewrites t as follows: $g(f(g(a), b)) \mapsto g(f(g(q_0), q_0)) \mapsto g(f(q_1, q_0)) \mapsto g(q_1) \mapsto q_0$. Hence $t \mapsto q_0 \in F$ so that $t \in L(\mathcal{A})$.

If we replace R by $R' \stackrel{\text{def}}{=} \{a \mapsto q_0, b \mapsto q_0, g(q_0) \mapsto q_1, g(q_1) \mapsto q_0, f(q_0, q_0) \mapsto q_0, f(q_1, q_1) \mapsto q_1\}$ we have an automaton accepting the set of all t where there is an even number of g 's along every path from the root to a leaf.

The *size* of a tree automaton, denoted by $|\mathcal{A}|$, is the number of states of $\mathcal{A}$ augmented by the size of the rules of $\mathcal{A}$ where a rule $f(q_1, \dots, q_n) \mapsto q$ has size $n + 2$. In this paper, we shall never be more precise than counting $|Q|$, the number of states of our automata. Notice that, for a fixed $\mathcal{F}$ where the largest arity is $m \geq 2$, $|\mathcal{A}|$ is in $O(|Q|^m)$.

A tree automaton is *deterministic* if all transition rules have distinct left-hand sides (and there are no ε -rule). Otherwise it is non-deterministic. Given a non-deterministic tree automaton, one can use the classical “subset construction” and build a deterministic tree automaton accepting the same language, but this construction involves a potential exponential blow-up in size. Telling whether $L(\mathcal{A})$ is empty for $\mathcal{A}$ a (non-necessarily deterministic) tree automaton can be done in time $O(|\mathcal{A}|)$. Telling whether a given tree t is accepted by a (non-necessarily deterministic) $\mathcal{A}$ can be done in time polynomial in $|\mathcal{A}| + |t|$.

A tree automaton is *completely specified* (also *complete*) if for each $f \in \mathcal{F}_n$ and $q_1, \dots, q_n \in Q$, there is a rule $f(q_1, \dots, q_n) \mapsto q$. By adding a sink state and the obvious rules, any $\mathcal{A}$ can be extended into a complete one accepting the same language.

2 The PA process algebra

For our presentation of PA, we explicitly refrain from writing terms modulo some simplification laws (e.g. the neutral laws for 0). Hence our use of the *IsNil* predicate (see below), inspired by [Chr93].

This viewpoint is in agreement with the earliest works on (general) process algebras like CCS, ACP, etc. It is a key condition for the results of the next section, and it clearly does not prevent considering terms modulo some structural congruence at a later stage, as we demonstrate in section 8.

2.1 Syntax

$Act = \{a, b, c, \dots\}$ is a set of *action names*.

$Var = \{X, Y, Z, \dots\}$ is a set of *process variables*.

$E_{PA} = \{t, u, \dots\}$ is the set of PA-terms, given by the following abstract syntax

$$t, u ::= 0 \mid X \mid t.u \mid t \parallel u$$

Given $t \in E_{PA}$, we write $Var(t)$ the set of process variables occurring in t and $Subterms(t)$ the set of all subterms of t (includes t).

A guarded PA *declaration* is a finite set $\Delta = \{X_i \xrightarrow{a_i} t_i \mid i = 1, \dots, n\}$ of *process rewrite rules*. Note that the X_i 's need not be distinct.

We write $Var(\Delta)$ for the set of process variables occurring in Δ , and $Subterms(\Delta)$ the union of all $Subterms(t)$ for t a right- or a left-hand side of a rule in Δ .

$\Delta_a(X)$ denotes $\{t \mid \text{there is a rule } "X \xrightarrow{a} t" \text{ in } \Delta\}$ and $\Delta(X)$ is $\bigcup_{a \in Act} \Delta_a(X)$. $Var_\emptyset \stackrel{\text{def}}{=} \{X \in Var \mid \Delta(X) = \emptyset\}$ is the set of variables for which Δ provides no rewrite.

In the following, we assume a fixed Var and Δ .

2.2 Semantics

A PA declaration Δ defines a labeled transition relation $\rightarrow_\Delta \subseteq E_{PA} \times Act \times E_{PA}$. We always omit the Δ subscript when no confusion is possible, and use the standard notations and abbreviations: $t \xrightarrow{w} t'$ with $w \in Act^*$, $t \xrightarrow{k} t'$ with $k \in \mathbb{N}$, $t \xrightarrow{*} t'$, $t \rightarrow, \dots \rightarrow_\Delta$ is inductively defined via the following SOS rules:

$$\begin{array}{c} \frac{t_1 \xrightarrow{a} t'_1}{t_1 \parallel t_2 \xrightarrow{a} t'_1 \parallel t_2} \qquad \frac{t_1 \xrightarrow{a} t'_1}{t_1.t_2 \xrightarrow{a} t'_1.t_2} \qquad \frac{}{X \xrightarrow{a} t} (X \xrightarrow{a} t) \in \Delta \\[10pt] \frac{t_2 \xrightarrow{a} t'_2}{t_1 \parallel t_2 \xrightarrow{a} t_1 \parallel t'_2} \qquad \frac{t_2 \xrightarrow{a} t'_2}{t_1.t_2 \xrightarrow{a} t_1.t'_2} IsNil(t_1) \end{array}$$

where the $IsNil(\dots)$ predicate is inductively defined by

$$\begin{aligned} IsNil(t_1 \parallel t_2) &\stackrel{\text{def}}{=} IsNil(t_1) \wedge IsNil(t_2), & IsNil(0) &\stackrel{\text{def}}{=} true, \\ IsNil(t_1.t_2) &\stackrel{\text{def}}{=} IsNil(t_1) \wedge IsNil(t_2), & IsNil(X) &\stackrel{\text{def}}{=} \begin{cases} true & \text{if } \Delta(X) = \emptyset, \\ false & \text{otherwise.} \end{cases} \end{aligned}$$

The $IsNil$ predicate is a syntactic test for termination, and indeed

Lemma 2.1. *The following three properties are equivalent:*

1. $IsNil(t) = true$,
2. $t \not\rightarrow$ (i.e. t is terminated),
3. $Var(t) \subseteq Var_\emptyset$.

Proof. (3 $\Rightarrow$ 2) Assume $t \rightarrow t'$. This derivation used has some $X_i \xrightarrow{a_i} t_i$ with $X_i \in \text{Var}(t)$.
 (2 $\Rightarrow$ 1) Use induction over t to prove that $\text{IsNil}(t) = \text{false}$ implies that $t \rightarrow t'$ for some t' .
 (1 $\Rightarrow$ 3) is obvious from the definition. $\square$

3 E_{PA} as a tree language

We shall use tree automata to recognize sets of terms from E_{PA} . This is possible because E_{PA} is just a $\mathcal{T}(\mathcal{F})$ for $\mathcal{F}$ given by $\mathcal{F}_0 = \{0, X, Y, \dots\}$ ($= \{0\} \cup \text{Var}$) and $\mathcal{F}_2 = \{., ||\}$. Of course, we shall keep using the usual infix notation for terms built with “.” or “||”.

We begin with one of the simplest languages in E_{PA} :

Proposition 3.1. *For any t , the singleton tree language $\{t\}$ is regular, and an automaton for $\{t\}$ needs only have $|t|$ states.*

Similarly, an immediate consequence of Lemma 2.1 is

Corollary 3.2. *L° , the set of terminated processes, is a regular tree language, and an automaton for L° needs only have one state.*

4 Regularity of the reachability set

For $t \in E_{\text{PA}}$, we let $\text{Pre}^*(t) \stackrel{\text{def}}{=} \{t' \mid t' \xrightarrow{*} t\}$ (resp. $\text{Post}^*(t) \stackrel{\text{def}}{=} \{t' \mid t \xrightarrow{*} t'\}$) denote the set of iterated predecessors (resp. the set of iterated successors, also called the *reachability set*) of t .

These notions do not take into account the sequences $w \in \text{Act}^*$ of action names allowing to move from some t to some t' in $\text{Post}^*(t)$. Indeed, we will forget about action names until section 7 which is devoted to $\text{Pre}^*[C](t)$ and $\text{Post}^*[C](t)$ for $C \subseteq \text{Act}^*$.

Given two tree languages $L, L' \subseteq E_{\text{PA}}$, we let

$$L.L' \stackrel{\text{def}}{=} \{t.t' \mid t \in L, t' \in L'\}, \quad L \parallel L' \stackrel{\text{def}}{=} \{t \parallel t' \mid t \in L, t' \in L'\}.$$

4.1 Regularity of $\text{Pre}^*(t)$

We define $(L_t)_{t \in E_{\text{PA}}}$, an infinite family of tree languages, as the least solution of the following set of recursive equations. The intuition is that these are *quasi-regular* equations satisfied

by $Pre^*(t)$.

$$\begin{aligned}
L_0 &= \{0\} \cup \bigcup_{Y \xrightarrow{a} 0 \in \Delta} L_Y, & L_{t||t'} &= L_t || L_{t'} \cup \bigcup_{Y \xrightarrow{a} t||t' \in \Delta} L_Y, \\
L_X &= \{X\} \cup \bigcup_{Y \xrightarrow{a} X \in \Delta} L_Y, & L_{t.t'} &= \begin{cases} L_t.L_{t'} \cup \bigcup_{Y \xrightarrow{a} t.t' \in \Delta} L_Y, & \text{if } t \in L^\emptyset, \\ L_t.\{t'\} \cup \bigcup_{Y \xrightarrow{a} t.t' \in \Delta} L_Y, & \text{otherwise.} \end{cases} \quad (1)
\end{aligned}$$

Observe that all equations define L_t as containing all L_Y 's for Y a process variable allowing a one step transition $Y \xrightarrow{a} t$ into t .

Lemma 4.1. *For any $t \in E_{PA}$, $L_t = Pre^*(t)$.*

Proof. (Sketch) The proof that $u \xrightarrow{*} t$ implies $u \in L_t$ is an induction over the length of the transition sequence from u to t , then a case analysis of which SOS rule gave the last transition, and then an induction over the structure of t .

The proof that $u \in L_t$ implies $u \xrightarrow{*} t$ is a fixpoint induction, followed by a case analysis over which summand of which equation is used, and relies on simple lemmas about reachability, such as “ $t_1 \xrightarrow{*} u_1$ implies $t_1 || t_2 \xrightarrow{*} u_1 || t_2$ ”. $\square$

The equations from (1) can easily be transformed into *regular equations*, just by introducing new variables for sets $\{t\}$ in the definitions for the $L_{t.t'}$'s. Now, because any given L_t only depends on a finite number of L_u 's and $\{u\}$'s, namely only for u 's in $Subterms(t) \cup Subterms(\Delta)$, we have ¹

Corollary 4.2. *For any $t \in E_{PA}$, the set L_t is a regular tree language.*

and the corresponding tree automaton has $O(|\Delta| + |t|)$ states. This entails

Theorem 4.3. *For any $t \in E_{PA}$, $Pre^*(t)$, $Pre(t)$ and $Pre^+(t)$ are regular tree languages.*

4.2 Regularity of $Post^*(t)$

We define $(L'_t, L''_t)_{t \in E_{PA}}$, two infinite families of tree languages, as the least solution of the following set of recursive equations. Our aim is that L'_t (resp. L''_t) should coincide with

¹In section 5.1, we shall see that Corollary 4.2 holds even when Δ is infinite (but $Var(\Delta)$ must be finite).

$Post^*(t)$ (resp. $Post^*(t) \cap L^\emptyset$).

$$\begin{aligned}
L'_0 &= \{0\}, & L''_0 &= \{0\}, \\
L'_X &= \{X\} \cup \bigcup_{X \xrightarrow{a} t \in \Delta} L'_t, & L''_X &= \begin{cases} \{X\} & \text{if } X \in Var_\emptyset, \\ \bigcup_{X \xrightarrow{a} t \in \Delta} L''_t, & \text{otherwise,} \end{cases} \\
L'_{t||t'} &= L'_t \parallel L'_{t'}, & L''_{t||t'} &= L''_t \parallel L''_{t'}, \\
L'_{t.t'} &= L'_t \cdot \{t'\} \cup L''_t \cdot L'_{t'}, & L''_{t.t'} &= L''_t \cdot L''_{t'}.
\end{aligned} \tag{2}$$

Again, these can easily be turned into regular equations. Again, any given L'_t or L''_t only depends on a finite number of L'_u 's, L''_u 's and $\{u\}$'s.

Corollary 4.4. *For any $t \in E_{PA}$, the sets L'_t and L''_t are regular tree languages.*

and the corresponding tree automata have $O(|\Delta| + |t|)$ states.

As with $Pre^*(t)$, we can easily show

Lemma 4.5. *For any $t \in E_{PA}$, $L'_t = Post^*(t)$ and $L''_t = Post^*(t) \cap L^\emptyset$.*

hence the corollary

Theorem 4.6. *For any $t \in E_{PA}$, $Post^*(t)$, $Post(t)$ and $Post^+(t)$ are regular tree languages that can be constructed effectively.*

Theorems 4.3 and 4.6 will be generalized in sections 5 and 7. However, we found it enlightening to give simple proofs of the simplest variants of our regularity results.

Already, Theorems 4.3 and 4.6 and the effective constructibility of the associated automata have many applications.

4.3 Some applications

Theorem 4.7. *The reachability problem “is t reachable from t' ?” is in P .*

Proof. Combine the cost of membership testing for non-deterministic tree automata and the regularity of $Pre^*(t')$ or the regularity of $Post^*(t)$. $\square$

For a different presentation of PA and $\rightarrow_\Delta$, [May97c] shows that the reachability problem is NP-complete. In section 8, we describe how to get his result as a byproduct of our approach.

Many other problems are solved by simple application of Theorems 4.3 and 4.6:

boundedness. Is $Post^*(t)$ infinite ?

covering. (a.k.a. control-state reachability). Can we reach a t' in which $Y_1, \dots, Y_m$ occur (resp. do not occur).

inclusion. Are all states reachable from t_1 also reachable from t_2 ? Same question modulo a regularity preserving operation (e.g. projection).

liveness. where a given $\Delta' \subseteq \Delta$ is live if, in all reachable states, at least one transition from Δ' can be fired.

5 Regularity of $Post^*(L)$ and $Pre^*(L)$ for a regular language L

In this section we prove the regularity of $Pre^*(L)$ and $Post^*(L)$ for a regular language L .

For notational simplicity, given two states q, q' of an automaton $\mathcal{A}$, we denote by $q \parallel q'$ (resp. $q \cdot q'$) any state q'' such that $q \parallel q' \xrightarrow{\mathcal{A}} q''$ (resp. $q \cdot q' \xrightarrow{\mathcal{A}} q''$), possibly using ε -rules.

5.1 Regularity of $Pre^*(L)$

Ingredients for $\mathcal{A}_{Pre^*}$: Assume $\mathcal{A}_L$ is an automaton recognizing $L \subseteq E_{PA}$. $\mathcal{A}_{Pre^*}$ is a new automaton combining several ingredients:

- $\mathcal{A}_\emptyset$ is a *completely specified* automaton accepting terminated processes (see Corollary 3.2).
- $\mathcal{A}_L$ is the automaton accepting L .
- We also use a boolean to record whether some rewriting steps have been done.

States of $\mathcal{A}_{Pre^*}$: A state of $\mathcal{A}_{Pre^*}$ is a 3-tuple $(q_\emptyset \in Q_{\mathcal{A}_\emptyset}, q_L \in Q_{\mathcal{A}_L}, b \in \{true, false\})$ where $Q_{\dots}$ denotes the set of states of the relevant automaton.

Transition rules of $\mathcal{A}_{Pre^*}$: The transition rules of $\mathcal{A}_{Pre^*}$ are defined as follows:

type 0: all rules of the form $0 \mapsto (q_\emptyset, q_L, false)$ s.t. $0 \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $0 \xrightarrow{\mathcal{A}_L} q_L$.

type 1a: all rules of the form $X \mapsto (q_\emptyset, q_L, true)$ s.t. there exists some $u \in Post^+(X)$ with $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $u \xrightarrow{\mathcal{A}_L} q_L$.

type 1b: all rules of the form $X \mapsto (q_\emptyset, q_L, false)$ s.t. $X \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $X \xrightarrow{\mathcal{A}_L} q_L$.

type 2: all rules of the form $(q_\emptyset, q_L, b) \parallel (q'_\emptyset, q'_L, b') \mapsto (q_\emptyset \parallel q'_\emptyset, q_L \parallel q'_L, b \vee b')$.

type 3a: all rules of the form $(q_\emptyset, q_L, b) \cdot (q'_\emptyset, q'_L, b') \mapsto (q_\emptyset \cdot q'_\emptyset, q_L \cdot q'_L, b \vee b')$ s.t. $q_\emptyset$ is a final state of $\mathcal{A}_\emptyset$.

type 3b: all rules of the form $(q_\emptyset, q_L, b).(q'_\emptyset, q'_L, \text{false}) \mapsto (q_\emptyset.q'_\emptyset, q_L.q'_L, b)$.

Lemma 5.1. *For any $t \in E_{\text{PA}}$, $t \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset, q_L, b)$ iff there is some $u \in E_{\text{PA}}$ and some $p \in \mathbb{N}$ such that $t \xrightarrow{p} u$, $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $u \xrightarrow{\mathcal{A}_L} q_L$ and $(b = \text{false} \text{ iff } p = 0)$.*

Proof. By structural induction over t . There are three cases:

1. $t = 0$ or $t = X$: Because $\mathcal{A}_{\text{Pre}}^*$ has no ε -rules, we only have to observe that its rules of type 0, 1a and 1b exactly correspond to what the lemma requires.

2. $t = t_1.t_2$: ($\Rightarrow$): the rewrite $t \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset, q_L, b)$ required that, for $i = 1, 2$, we have $t_i \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset^i, q_L^i, b^i)$ and there is a type 3 rule $(q_\emptyset^1, q_L^1, b^1).(q_\emptyset^2, q_L^2, b^2) \mapsto (q_\emptyset, q_L, b)$.

The induction hypothesis entails there are $t_1 \xrightarrow{p_1} u_1$ and $t_2 \xrightarrow{p_2} u_2$ corresponding to the rewrite of t_1 and t_2 by $\mathcal{A}_{\text{Pre}}^*$. Now if $\mathcal{A}_{\text{Pre}}^*$ used a type 3b rule, then $b_2 = \text{false}$ hence $p_2 = 0$, $u_2 = t_2$, $p_1 = p$ and $t_1.t_2 \xrightarrow{p_1} u_1.t_2 = u_1.u_2$. If we used a type 3a rule, then $q_\emptyset^1$ is a final state, therefore $u_1 \in L_\emptyset$ is a terminated process, hence $t_1.t_2 \xrightarrow{p_1} u_1.t_2 \xrightarrow{p_2} u_1.u_2$ and $(b = b_1 \vee b_2 = \text{false} \text{ iff } p_1 + p_2 = 0)$.

($\Leftarrow$): Conversely, assume $t = t_1.t_2 \xrightarrow{p} u$ with $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $u \xrightarrow{\mathcal{A}_L} q_L$. Then u is some $u_1.u_2$ and either (1) $u_2 = t_2$ and $t_1 \xrightarrow{p} u_1$, or (2) $u_1 \in L_\emptyset$ and $t_1.t_2 \xrightarrow{p_1} u_1.t_2 \xrightarrow{p_2} u_1.u_2$ for $p_1 + p_2 = p$.

In the first case the ind. hyp. entails $t_1 \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset^1, q_L^1, b^1)$ with $u_1 \xrightarrow{\mathcal{A}_L} q_L^1$, and $t_2 = u_2 \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset^2, q_L^2, \text{false})$. Now we can use a type 3b rule to show $t \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset^1.q_\emptyset^2, q_L^1.q_L^2, b_1)$ with $u \xrightarrow{\mathcal{A}_L} q_L^1.q_L^2$.

In the second case, $u_1 \in L_\emptyset$ entails $t_1 \xrightarrow{\mathcal{A}_L} (q_\emptyset^1, q_L^1, b_1)$ with $q_\emptyset^1$ a final state of $\mathcal{A}_\emptyset$. We can use a type 3a rule to show $t \xrightarrow{\mathcal{A}_{\text{Pre}}^*} (q_\emptyset^1.q_\emptyset^2, q_L^1.q_L^2, b_1 \vee b_2)$.

3. $t = t_1 \parallel t_2$: This case is similar to the previous one (actually it is simpler).

□

If we now let the final states of $\mathcal{A}_{\text{Pre}}^*$ be all states $(q_\emptyset, q_L, b)$ s.t. q_L is a final state of $\mathcal{A}_L$, then $t \xrightarrow{*} u$ for some u accepted by $\mathcal{A}_L$ iff $\mathcal{A}_{\text{Pre}}^*$ accepts t (this is where we use the assumption that $\mathcal{A}_\emptyset$ is completely specified.)

Theorem 5.2. (Regularity)

- (1) If L is a regular subset of E_{PA} , then $\text{Pre}^*(L)$ is regular.
- (2) Furthermore, from an automaton $\mathcal{A}_L$ recognizing L , is it possible to construct (in polynomial time) an automaton $\mathcal{A}_{\text{Pre}}^*$ recognizing $\text{Pre}^*(L)$. If $\mathcal{A}_L$ has k states, then $\mathcal{A}_{\text{Pre}}^*$ needs only have $4k$ states.

Proof. (1) is an immediate consequence of Lemma 5.1. Observe that the result does not need the finiteness of Δ (but $\text{Var}(\Delta)$ must be finite).

(2) Building $\mathcal{A}_{Pre^*}$ effectively requires an effective way of listing the type 1a rules. This can be done by computing a product of $\mathcal{A}_X$, an automaton for $\text{Post}^+(X)$, with $\mathcal{A}_\emptyset$ and $\mathcal{A}_L$. Then there exists some $u \in \text{Post}^+(X)$ with $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $u \xrightarrow{\mathcal{A}_L} q_L$ iff the language accepted by the final states $\{(q_X, q_\emptyset, q_L) \mid q_X \text{ a final state of } \mathcal{A}_X\}$ is not-empty. This gives us the pairs $q_\emptyset, q_L$ we need for type 1a rules. Observe that we need the finiteness of Δ to build the $\mathcal{A}_X$'s. $\square$

5.2 Regularity of $\text{Post}^*(L)$

Ingredients for $\mathcal{A}_{Post^*}$: Assume $\mathcal{A}_L$ is an automaton recognizing $L \subseteq E_{PA}$. $\mathcal{A}_{Post^*}$ is a new automaton combining several ingredients:

- Automata $\mathcal{A}_\emptyset$ and $\mathcal{A}_L$ as in the previous construction, but this time *we need to assume each of them is a completely specified automata*.
- $\mathcal{A}_\Delta$ is a completely specified automaton recognizing the subterms of Δ . It has all states q_s for $s \in \text{Subterms}(\Delta)$. We ensure “ $t \xrightarrow{\mathcal{A}_\Delta} q_s$ iff $s = t$ ” by taking as transition rules $0 \mapsto q_0$ if $0 \in \text{Subterms}(\Delta)$, $X \mapsto q_X$ if $X \in \text{Subterms}(\Delta)$, $q_s \parallel q_{s'} \mapsto q_{s \parallel s'}$ (resp. $q_s.q_{s'} \mapsto q_{s.s'}$) if $s \parallel s'$ (resp. $s.s'$) belongs to $\text{Subterms}(\Delta)$. In addition, the automaton has a sink state $q_\perp$ and the obvious transitions so that it is a completely specified automaton.
- Again, we use a boolean b to record whether rewrite steps have occurred.

States of $\mathcal{A}_{Post^*}$: The states of $\mathcal{A}_{Post^*}$ are 4-uples $(q_\emptyset \in Q_{\mathcal{A}_\emptyset}, q_L \in Q_{\mathcal{A}_L}, q_\Delta \in Q_{\mathcal{A}_\Delta}, b \in \{true, false\})$.

Transition rules of $\mathcal{A}_{Post^*}$: The transition rules are:

type 0: all rules of the form $0 \mapsto (q_\emptyset, q_L, q_\Delta, false)$ s.t. $0 \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $0 \xrightarrow{\mathcal{A}_L} q_L$ and $0 \xrightarrow{\mathcal{A}_\Delta} q_\Delta$.

type 1: all rules of the form $X \mapsto (q_\emptyset, q_L, q_\Delta, false)$ s.t. $X \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $X \xrightarrow{\mathcal{A}_L} q_L$, and $X \xrightarrow{\mathcal{A}_\Delta} q_\Delta$.

type 2: all ε -rules of the form $(q_\emptyset, q'_L, q_s, b') \mapsto (q_\emptyset, q_L, q_X, true)$ s.t. $X \rightarrow s$ is a rule in Δ with $X \xrightarrow{\mathcal{A}_L} q_L$.

type 3: all rules of the form
 $(q_\emptyset, q_L, q_\Delta, b) \parallel (q'_\emptyset, q'_L, q'_\Delta, b') \mapsto (q_\emptyset \parallel q'_\emptyset, q_L \parallel q'_L, q_\Delta \parallel q'_\Delta, b \vee b')$

type 4a: all rules of the form
 $(q_\emptyset, q_L, q_\Delta, b).(q'_\emptyset, q'_L, q'_\Delta, false) \mapsto (q_\emptyset.q'_\emptyset, q_L.q'_L, q_\Delta.q'_\Delta, b).$

type 4b: all rules of the form

$$(q_{\emptyset}, q_L, q_{\Delta}, b) \cdot (q'_{\emptyset}, q'_L, q'_{\Delta}, b') \mapsto (q_{\emptyset} \cdot q'_{\emptyset}, q_L \cdot q'_L, q_{\Delta} \cdot q'_{\Delta}, b \vee b') \text{ s.t. } q_{\emptyset} \text{ is a final state of } \mathcal{A}_{\emptyset}.$$

Lemma 5.3. *For any $t \in E_{PA}$, $t \xrightarrow{\mathcal{A}_{Pos}^*} (q_{\emptyset}, q_L, q_{\Delta}, b)$ iff there is some $u \in E_{PA}$ and some $p \in \mathbb{N}$ such that $u \xrightarrow{p} t$, $u \xrightarrow{\mathcal{A}_L} q_L$, $u \xrightarrow{\mathcal{A}_{\Delta}} q_{\Delta}$, ($b = \text{false}$ iff $p = 0$) and $t \xrightarrow{\mathcal{A}_{\emptyset}} q_{\emptyset}$.*

Proof. We first prove the ($\Rightarrow$) direction by induction over the length k of the rewrite $t \xrightarrow{\mathcal{A}_{Pos}^*} (q_{\emptyset}, q_L, q_{\Delta}, b)$. We distinguish four cases:

1. $k = 1$: Then $t = 0$ or $t = X$ and we used a type 0 or type 1 rule. Taking $u = t$ and $p = 0$ satisfies the requirements.
2. $k > 1$ **and the last rewrite step used a type 2 ε -rule:** Then the rewrite has the form

$$\overbrace{t \mapsto (q'_{\emptyset}, q'_L, q_s, b')}^{k-1 \text{ steps}} \mapsto (q_{\emptyset}, q_L, q_X, \text{true}).$$
 By ind. hyp., there is a u' and a p' s.t. $u' \xrightarrow{p'} t$. Now $u' \xrightarrow{\mathcal{A}_{\Delta}} q_s$ entails $u' = s$. The existence of the type 2 rule entails $X \rightarrow s \in \Delta$. Hence $X \xrightarrow{p'+1} t$. Taking $u = X$ and $p = p' + 1$ satisfies the requirements.
3. $k > 1$ **and the last rewrite step used a type 4 rule:** Then t is some $t_1.t_2$ and the type 4 rule applied on top of two rewrite sequences $t_i \mapsto (q^i_{\emptyset}, q^i_L, q^i_{\Delta}, b^i)$ for $i = 1, 2$. The ind. hyp. gives us, for $i = 1, 2$, some u_i and p_i s.t. $u_i \xrightarrow{p_i} t_i$.

If the last rule was a type 4a rule, then $b^2 = \text{false}$ so that $p_2 = 0$ and $u_2 = t_2$. Then $u_1.u_2 \xrightarrow{p_1} t_1.u_2 = t$. Taking $u = u_1.u_2$ and $p = p_1$ satisfies the requirements.

Otherwise the last rule was a type 4b rule. Then $q_{\emptyset}^1$ is a final state and $t_1 \xrightarrow{\mathcal{A}_{\emptyset}} q_{\emptyset}^1$ entails that t_1 is a terminated process. Hence $u_1.u_2 \xrightarrow{p_1} t_1.u_2 \xrightarrow{p_2} t_1.t_2 = t$. Again, taking $u = u_1.u_2$ (with $p = p_1 + p_2$) satisfies the requirements.

4. $k > 1$ **and the last rewrite step used a type 3 rule:** This case is similar (actually simpler) to the previous one.

For the ($\Leftarrow$) direction, we assume $u \xrightarrow{p} t$ with the accompanying conditions (a.c.), and proceed by induction over the length of the transition sequence (i.e. over p), followed by structural induction over u . There are five cases:

1. $u = 0$: Then $t = u$ and the a.c.'s ensure we can use a type 0 rule to show $t \xrightarrow{\mathcal{A}_{Pos}^*} (q_{\emptyset}, q_L, q_{\Delta}, \text{false})$.
2. $u = X$ **and** $p = 0$: Like the previous case but with a type 1 rule.
3. $u = X$ **and** $p > 0$: Then the sequence has the form $X \xrightarrow{1} u' \xrightarrow{p-1} t$. Here the a.c.'s read $q_{\Delta} = q_X$ and $b = \text{true}$. $X \xrightarrow{1} u' \in \Delta$ entails $u' \in \text{Subterms}(\Delta)$. If we now take a q'_L

s.t. $u' \xrightarrow{\mathcal{A}_L} q'_L$ (one such q'_L must exist) and let b' be *false* iff $p - 1 = 0$, the ind. hyp. gives us $t \xrightarrow{\mathcal{A}_{Post}^*} (q_\emptyset, q'_L, q_{u'}, b')$. Now, there must be a type 2 ε -rule $(q_\emptyset, q'_L, q_{u'}, b') \mapsto (q_\emptyset, q_L, q_X, true)$. We use it to show $t \xrightarrow{\mathcal{A}_{Post}^*} (q_\emptyset, q_L, q_X, true)$.

4. $u = u_1.u_2$: Then t is some $t_1.t_2$ and $u \xrightarrow{p} t$ is a combination of some $u_1 \xrightarrow{p_1} t_1$ and $u_2 \xrightarrow{p_2} t_2$ with $p = p_1 + p_2$. Additionally, if $p_2 > 0$ then $t_1 \in L_\emptyset$.

For $i = 1, 2$, The rewrites $t \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $u \xrightarrow{\mathcal{A}_L} q_L$ and $u \xrightarrow{\mathcal{A}_\Delta} q_\Delta$ used some $t_i \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset^i$, $u_i \xrightarrow{\mathcal{A}_L} q_L^i$ and $u_i \xrightarrow{\mathcal{A}_\Delta} q_\Delta^i$ (for $i = 1, 2$). If we now define b^i according to p_i , the ind. hyp. entails that, for $i = 1, 2$, $t_i \xrightarrow{\mathcal{A}_{Post}^*} (q_\emptyset^i, q_L^i, q_\Delta^i, b^i)$.

There are two cases. If $t_1 \in L_\emptyset$ then $q_\emptyset^1$ is a final state of $\mathcal{A}_\emptyset$ and $\mathcal{A}_{Post}^*$ has a type 4b rule $(q_\emptyset^1, q_L^1, q_\Delta^1, b^1).(q_\emptyset^2, q_L^2, q_\Delta^2, b^2) \mapsto (q_\emptyset, q_L, q_\Delta, b)$ that we can use. If $t_1 \notin L_\emptyset$, then $p_2 = 0$ and $b^2 = false$. There is a type 4a rule that we can use.

5. $u = u_1 \parallel u_2$: Similar to the previous case (actually it is simpler).

□

If we now let the final states of $\mathcal{A}_{Post}^*$ be all states $(q_\emptyset, q_L, q_\Delta, b)$ s.t. q_L is a final state of $\mathcal{A}_L$, then $\mathcal{A}_{Post}^*$ accepts a term t iff $u \xrightarrow{*} t$ for a u accepted by $\mathcal{A}_L$ iff t belongs to $Post^*(L)$.

Theorem 5.4. (Regularity)

- (1) If L is a regular subset of E_{PA} , then $Post^*(L)$ is regular.
 (2) Furthermore, from an automaton $\mathcal{A}_L$ recognizing L , is it possible to construct (in polynomial time) an automaton $\mathcal{A}_{Post}^*$ recognizing $Post^*(L)$. If $\mathcal{A}_L$ has k states, then $\mathcal{A}_{Pre}^*$ needs only have $O(k \cdot |\Delta|)$ states.

Proof. Obvious from the previous construction. □

Our results relate t and $Pre^*(t)$ (resp. $Post^*(t)$). A natural question is to ask if the relation “ $\xrightarrow{*}$ ” (i.e. $\{(t, u) \mid t \xrightarrow{*} u\}$) is recognizable in some sense. The most relevant notion of recognizability related to our problem is linked to *ground tree transducers*, GTT’s for short, see [CDG⁺97] for details. Since it can be shown that the $\xrightarrow{*}$ relation induced by a ground rewrite system is recognizable by a GTT, we tried to extend this result to our PA case where the rules are ground rewrite rules with simple left hand sides, but where there is a notion of prefix rewriting. Unfortunately, this prefix rewriting entails that our $\xrightarrow{*}$ is not stable under contexts and the natural extensions of GTT that could handle such conditional rules are immediately able to recognize any recursively enumerable relation.

6 Model-checking PA processes

In this section we show a simple approach to the model-checking problem solved in [May97b]. We see this as one more immediate application of our main regularity theorems.

We consider a set $Prop = \{P_1, P_2, \dots\}$ of *atomic propositions*. For $P \in Prop$, Let $Mod(P)$ denotes the set of PA processes for which P holds. We only consider propositions P such that $Mod(P)$ is a regular tree-language. Thus P could be “ t can make an a -labeled step right now”, “there is at least two occurrences of X inside t ”, “there is exactly one occurrence of X in a non-frozen position”, ...

The logic EF has the following syntax:

$$\varphi ::= P \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \text{EX}\varphi \mid \text{EF}\varphi$$

and semantics

$$\begin{aligned} t \models P &\stackrel{\text{def}}{\iff} t \in Mod(P), & t \models \text{EX}\varphi &\stackrel{\text{def}}{\iff} t' \models \varphi \text{ for some } t \rightarrow t', \\ t \models \neg\varphi &\stackrel{\text{def}}{\iff} t \not\models \varphi, & t \models \text{EF}\varphi &\stackrel{\text{def}}{\iff} t' \models \varphi \text{ for some } t \xrightarrow{*} t'. \\ t \models \varphi \wedge \varphi' &\stackrel{\text{def}}{\iff} t \models \varphi \text{ and } t \models \varphi', \end{aligned}$$

Thus $\text{EX}\varphi$ reads “it is possible to reach in one step a state s.t. φ ” and $\text{EF}\varphi$ reads “it is possible to reach (via some sequence of steps) a state s.t. φ ”.

Definition 6.1. *The model-checking problem for EF over PA has as inputs: a given Δ , a given t in E_{PA} , a given φ in EF. The answer is yes iff $t \models \varphi$.*

If we now define $Mod(\varphi) \stackrel{\text{def}}{=} \{t \in E_{PA} \mid t \models \varphi\}$, we have

$$\begin{aligned} Mod(\neg\varphi) &= E_{PA} - Mod(\varphi) & Mod(\text{EX}\varphi) &= Pre(Mod(\varphi)) \\ Mod(\varphi \wedge \varphi') &= Mod(\varphi) \cap Mod(\varphi') & Mod(\text{EF}\varphi) &= Pre^*(Mod(\varphi)) \end{aligned} \quad (3)$$

Theorem 6.2. (1) *For any EF formula φ , $Mod(\varphi)$ is a regular tree language.*
 (2) *If we are given tree-automata $\mathcal{A}_P$'s recognizing the regular sets $Mod(P)$, then a tree-automaton $\mathcal{A}_\varphi$ recognizing $Mod(\varphi)$ can be built effectively.*

Proof. A corollary of (3) and the regularity theorems. $\square$

This gives us a decision procedure for the model-checking problem: build an automaton for $Mod(\varphi)$ and check whether it accepts t . We can estimate the complexity of this approach in term of $|\varphi|$ and $n_{alt}(\varphi)$.

We define $n_{alt}(\varphi)$ the number of alternation of negations and temporal connectives in φ as

$$\begin{aligned} n_{alt}(P) &= 0 & n_{alt}(\neg P) &= 1 \\ n_{alt}(\varphi \wedge \psi) &= \max(n_{alt}(\varphi), n_{alt}(\psi)) & n_{alt}(\neg(\varphi \wedge \psi)) &= \max(n_{alt}(\neg\varphi), n_{alt}(\neg\psi)) \\ n_{alt}(\text{EF}\varphi) &= n_{alt}(\varphi) & n_{alt}(\neg\text{EF}\varphi) &= 1 + n_{alt}(\varphi) \\ n_{alt}(\text{EX}\varphi) &= n_{alt}(\varphi) & n_{alt}(\neg\text{EX}\varphi) &= 1 + n_{alt}(\varphi) \\ & & n_{alt}(\neg\neg\varphi) &= n_{alt}(\varphi) \end{aligned}$$

Theorem 6.3. (Model-checking) *An automaton for $Mod(\varphi)$ can be computed in time*

$$2^{|\varphi||\Delta|2^{O(|\varphi||\Delta|)}} \left. \vphantom{2^{|\varphi||\Delta|2^{O(|\varphi||\Delta|)}}} \right\} n_{alt}(\varphi).$$

Proof. We assume all automata for the $Mod(P)$'s have size bounded by M (a constant). We construct an automaton for $Mod(\varphi)$ by applying the usual automata-theoretic constructions for intersection, union, complementation of regular tree languages, and by invoking our regularity theorems for Pre and Pre^* . All constructions are polynomial except for complementation. With only polynomial constructions, we would have a $2^{O(|\varphi|)}$ size for the resulting automaton. The negations involving complementation are the cause of the non-elementary blowup.

Negations can be pushed inward except that they cannot cross the temporal connectives EF and EX. Here we have one exponential blowup for determinization at each level of alternation. This is repeated $n_{alt}(\varphi)$ times, yielding the given bound on the number of states hence the overall complexity. $\square$

The procedure described in [May97b] is non-elementary and today the known lower bound is PSPACE-hard. Observe that computing a representation of $Mod(\varphi)$ is more general than just telling whether a given t belongs to it. Observe also that our results allow model-checking approaches based on combinations of forward and backward methods (while Theorem 6.2 only relied on the standard backward approach.)

7 Reachability under constraints

In this section, we consider *reachability under constraints*. Let $C \subseteq Act^*$ be a (word) language over action names. We write $t \xrightarrow{C} t'$ when $t \xrightarrow{w} t'$ for some $w \in C$, and we say that t' can be reached from t under the constraint C . We extend our notations and write $Pre^*[C](L)$, $Post^*[C](L)$, $\dots$ with the obvious meaning.

Observe that, even if we assume C is regular, the problem of telling whether $t \xrightarrow{C}$, i.e. whether $Post^*[C](t)$ is not empty, is undecidable for the PA algebra. This can be proved by a reduction from the intersection problem for context-free languages as follows: Let Σ be an alphabet and $\#$ some distinguished symbol. We use two copies $\underline{a}, \overline{a}$ of every letter a in $\Sigma \cup \{\#\}$. Context-free languages can be defined in BPA (PA without $\|$), that is, for any context-free language L_1 (resp. L_2) on Σ , we can define PA rules such that $X_1 \xrightarrow{\underline{w}, \#} w$ iff $w \in L_1$ (resp. $X_2 \xrightarrow{\overline{w}, \#} w$ iff $w \in L_2$). These rules don't overlap. We now introduce the regular constraint $C \stackrel{\text{def}}{=} (\underline{a}_1.\overline{a}_1 + \dots + \underline{a}_n.\overline{a}_n)^*\underline{\#}.\overline{\#}$. Then $(X_1 \parallel X_2) \xrightarrow{C}$ holds iff $L_1 \cap L_2 \neq \emptyset$, which is undecidable.

In this section we give sufficient conditions over C so that the problem becomes decidable (and so that we can compute the C -constrained Pre^* of a regular tree language).

Recall that the *shuffle* $w \parallel w'$ of two finite words is the set of all words one can obtain by interleaving w and w' in an arbitrary way.

Definition 7.1. • $\{(C_1, C'_1), \dots, (C_m, C'_m)\}$ is a (finite) seq-decomposition of C iff for all $w, w' \in Act^*$ we have

$$w.w' \in C \quad \text{iff} \quad (w \in C_i, w' \in C'_i \text{ for some } 1 \leq i \leq m).$$

• $\{(C_1, C'_2), \dots, (C_m, C'_m)\}$ is a (finite) paral-decomposition of C iff for all $w, w' \in Act^*$ we have

$$C \cap (w \parallel w') \neq \emptyset \quad \text{iff} \quad (w \in C_i, w' \in C'_i \text{ for some } 1 \leq i \leq m).$$

The crucial point of the definition is that a seq-decomposition of C must apply to all possible ways of splitting any word in C . It even applies to a decomposition $w.w'$ with $w = \varepsilon$ (or $w' = \varepsilon$) so that one of the C_i 's (and one of the C'_i 's) contains ε . Observe that the formal difference between seq-decomposition and paral-decomposition comes from the fact that $w \parallel w'$, the set of all shuffles of w and w' usually contains several elements.

Definition 7.2. A family $\mathbb{C} = \{C_1, \dots, C_n\}$ of languages over Act is a finite decomposition system iff every $C \in \mathbb{C}$ admits a seq-decomposition and a paral-decomposition only using C_i 's from $\mathbb{C}$.

Not all $C \subseteq Act^*$ admit finite decompositions, even in the regular case. Consider $C = (ab)^*$ and assume $\{(C_1, C'_1), \dots\}$ is a finite paral-decomposition. Then for every k , there is a shuffle of a^k and b^k in C . Hence there must be a i_k s.t. $a^k \in C_{i_k}$ and $b^k \in C'_{i_k}$. Now if $i_k = i_{k'}$ then there must exist a shuffle w'' of a^k and $b^{k'}$ with $w'' \in C$. This is only possible if $k = k'$. Hence all i_k 's are distinct, contradicting finiteness.

A simple example of a finite decomposition system is $\mathbb{C} = \{\{w\} \mid |w| \leq k\}$, i.e. the set of all singleton languages with words shorter than k . Here the paral-decomposition of $\{w\}$ is $\{(\{w_1\}, \{w'_1\}), \dots, (\{w_m\}, \{w'_m\})\}$ where the w_i 's are all subwords² of w (and w'_i is the corresponding remainder). This example shows that decomposability is not composability: not all pairs from $\mathbb{C}$ appears in the decomposition of some member of C .

More generally, for any linear weight function θ of the form $\theta(w) \stackrel{\text{def}}{=} \sum_i n_i |w|_{a_i}$ with $n_i \in \mathbb{N}$, for any $k, k' \in \mathbb{N}$, the sets $C_{(\theta=k)} \stackrel{\text{def}}{=} \{w \mid \theta(w) = k\}$, $C_{(\theta < k)} \stackrel{\text{def}}{=} \dots$, $C_{(\theta > k)}$, and $C_{(\theta \equiv k \pmod{k'})}$ belong to finite decomposition system.

Assume $\mathbb{C}$ is a finite decomposition system. We shall show

Theorem 7.3. (Regularity)

For any regular $L \subseteq E_{PA}$ and any $C \in \mathbb{C}$, $Pre^*[C](L)$ and $Post^*[C](L)$ are regular tree languages.

²A subword of w is any w' obtained by erasing letters from w at any position.

Ingredients for $\mathcal{A}_{Post^*[\mathbb{C}]}$: We build $\mathcal{A}_{Post^*[\mathbb{C}]}$ in the same way as $\mathcal{A}_{Post^*}$ but states contain a new $C \in \mathbb{C}$ component.

States of $\mathcal{A}_{Post^*[\mathbb{C}]}$: The states of $\mathcal{A}_{Post^*[\mathbb{C}]}$ are 5-uples $(q_\emptyset \in Q_{\mathcal{A}_\emptyset}, q_L \in Q_{\mathcal{A}_L}, q_\Delta \in Q_{\mathcal{A}_\Delta}, b \in \{true, false\}, C \in \mathbb{C})$.

Transition rules of $\mathcal{A}_{Post^*[\mathbb{C}]}$: The transition rules are:

- type 0:** all rules of the form $0 \mapsto (q_\emptyset, q_L, q_\Delta, false, C)$ s.t. $0 \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $0 \xrightarrow{\mathcal{A}_L} q_L$, $0 \xrightarrow{\mathcal{A}_\Delta} q_\Delta$ and $\varepsilon \in C$.
- type 1:** all rules of the form $X \mapsto (q_\emptyset, q_L, q_\Delta, false, C)$ s.t. $X \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $X \xrightarrow{\mathcal{A}_L} q_L$, $X \xrightarrow{\mathcal{A}_\Delta} q_\Delta$ and $\varepsilon \in C$.
- type 2:** all ε -rules of the form $(q_\emptyset, q'_L, q_s, b', C'') \mapsto (q_\emptyset, q_L, q_X, true, C)$ s.t. $X \xrightarrow{a} s$ is a rule in Δ with $X \xrightarrow{\mathcal{A}_L} q_L$, and $a \in C'$ for some C' s.t. (C', C'') appears in the seq-decomposition of C .
- type 3:** all rules of the form $(q_\emptyset, q_L, q_\Delta, b, C) \parallel (q'_\emptyset, q'_L, q'_\Delta, b', C') \mapsto (q_\emptyset \parallel q'_\emptyset, q_L \parallel q'_L, q_\Delta \parallel q'_\Delta, b \vee b', C'')$ s.t. (C, C') appears in the paral-decomposition of C'' .
- type 4a:** all rules of the form $(q_\emptyset, q_L, q_\Delta, b, C) \cdot (q'_\emptyset, q'_L, q'_\Delta, false, C') \mapsto (q_\emptyset \cdot q'_\emptyset, q_L \cdot q'_L, q_\Delta \cdot q'_\Delta, b, C)$.
- type 4b:** all rules of the form $(q_\emptyset, q_L, q_\Delta, b, C) \cdot (q'_\emptyset, q'_L, q'_\Delta, b', C') \mapsto (q_\emptyset \cdot q'_\emptyset, q_L \cdot q'_L, q_\Delta \cdot q'_\Delta, b \vee b', C'')$ s.t. $q_\emptyset$ is a final state of $\mathcal{A}_\emptyset$ s.t. (C, C') appears in the seq-decomposition of C'' .

Lemma 7.4. For any $t \in E_{PA}$, $t \xrightarrow{\mathcal{A}_{Post^*[\mathbb{C}]}} (q_\emptyset, q_L, q_\Delta, b, C)$ iff there is some $u \in E_{PA}$ and some $w \in C$ such that $u \xrightarrow{w} t$, $u \xrightarrow{\mathcal{A}_L} q_L$, $u \xrightarrow{\mathcal{A}_\Delta} q_\Delta$, $(b = false \text{ iff } |w| = 0)$ and $t \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$.

Proof. $\mathcal{A}_{Post^*[\mathbb{C}]}$ is $\mathcal{A}_{Post^*}$ equipped with a new component and the proof follows exactly the lines of the proof of Lemma 5.3. We refer to this earlier proof and only explain how we deal with the new C components.

The $(\Rightarrow)$ direction is as in lemma 5.3. The new observations in the 4 cases are:

1. $k = 1$: The type 0 and type 1 rules entail $\varepsilon \in C$, so that we can take $w = \varepsilon$.
2. $k > 1$ and the last rewrite step used a type 2 ε -rule: Use the fact that $w' \in C''$ entail $a.w' \in C$.
3. $k > 1$ and the last rewrite step used a type 4 rule: Use the fact that $C.C' \subseteq C''$.
4. $k > 1$ and the last rewrite step used a type 3 rule: Use the fact that $w_1 \in C$ and $w_2 \in C'$ entail that there exists at least one shuffling w of w_1 and w_2 s.t. $w \in C''$.

The ($\Leftarrow$) direction is as in lemma 5.3. The new observations in the 5 cases are:

1. $u = 0$: The type 0 rules allow all C 's containing ε .
2. $u = X$ and $p = 0$: *Idem*.
3. $u = X$ and $p > 0$: Then the sequence has the form $X \xrightarrow{a} u' \xrightarrow{w'} t$. Now if $w = a.w' \in C$, there must be a (C', C'') in the seq-decomposition of C s.t. $a \in C'$ and $w' \in C''$. So that there is a type 2 rule $(\dots, q_{u'}, b', C'') \mapsto (\dots, q_X, \text{true}, C)$ we can use.
4. $u = u_1.u_2$: Here $u_1 \xrightarrow{w_1} t_1$, $u_2 \xrightarrow{w_2} t_2$ and $w_1.w_2.w \in C$. If $t_1 \notin L_\emptyset$ then $w_2 = \varepsilon$, $w_1 \in C$ and we have the type 4a rule we need. Otherwise there is a pair (C_1, C_2) in the seq-decomposition of C s.t. $w_i \in C_i$ ($i = 1, 2$). This pair gives us the type 4b rule we need.
5. $u = u_1 \parallel u_2$: Here $u_1 \xrightarrow{w_1} t_1$, $u_2 \xrightarrow{w_2} t_2$ and $w \in C$ is some shuffle of w_1 and w_2 . Therefore there is a (C_1, C_2) in the paral-decomposition of C s.t. $w_i \in C_i$ ($i = 1, 2$). This pair gives us the type 3 rule we need.

□

If we now let the final states of $\mathcal{A}_{Post^*[\mathbb{C}]}$ be all $(q_\emptyset, q_L, q_\Delta, b, C)$ s.t. q_L is a final state of $\mathcal{A}_L$, then $\mathcal{A}_{Post^*[\mathbb{C}]}$ accepts a term t iff $t \in Post^*[C](t)$. (The set of final states can easily be adapted so that we recognize $Post^+[C](L)$.)

Ingredients for $\mathcal{A}_{Pre^*[\mathbb{C}]}$: Same as in the construction of $\mathcal{A}_{Pre^*}$, with an additional $C \in \mathbb{C}$ component.

States of $\mathcal{A}_{Pre^*}$: A state of $\mathcal{A}_{Pre^*}$ is a 4-tuple $(q_\emptyset \in Q_{\mathcal{A}_\emptyset}, q_L \in Q_{\mathcal{A}_L}, b \in \{\text{true}, \text{false}\}, C \in \mathbb{C})$.

The final states are all $(q_\emptyset, q_L, b, C_i)$ s.t. q_L is a final state of $\mathcal{A}_L$ and C_i the constraint to satisfy.

Transition rules of $\mathcal{A}_{Pre^*}$: The transition rules of $\mathcal{A}_{Pre^*}$ are defined as follows:

type 0: all rules of the form $0 \mapsto (q_\emptyset, q_L, \text{false}, C)$ s.t. $0 \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $0 \xrightarrow{\mathcal{A}_L} q_L$ and $\varepsilon \in C$.

type 1a: all rules of the form $X \mapsto (q_\emptyset, q_L, \text{true}, C)$ s.t. there exists some $u \in Post^+[C](X)$ with $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$ and $u \xrightarrow{\mathcal{A}_L} q_L$.

type 1b: all rules of the form $X \mapsto (q_\emptyset, q_L, \text{false}, C)$ s.t. $X \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $X \xrightarrow{\mathcal{A}_L} q_L$ and $\varepsilon \in C$.

type 2: all rules of the form $(q_\emptyset, q_L, b, C) \parallel (q'_\emptyset, q'_L, b', C') \mapsto (q_\emptyset \parallel q'_\emptyset, q_L \parallel q'_L, b \vee b', C'')$ s.t. (C, C') appears in the paral-decomposition of C'' .

type 3a: all rules of the form

$$(q_\emptyset, q_L, b, C).(q'_\emptyset, q'_L, b', C') \mapsto (q_\emptyset.q'_\emptyset, q_L.q'_L, b \vee b', C'') \text{ s.t. } q_\emptyset \text{ is a final state of } \mathcal{A}_\emptyset \text{ and } (C, C') \text{ appears in the seq-decomposition of } C''$$

type 3b: all rules of the form

$$(q_\emptyset, q_L, b, C).(q'_\emptyset, q'_L, \text{false}, C') \mapsto (q_\emptyset.q'_\emptyset, q_L.q'_L, b, C).$$

Lemma 7.5. *For any $t \in E_{\text{PA}}$, $t \xrightarrow{\mathcal{A}_{\text{Pre}^*[\mathbb{C}]}} (q_\emptyset, q_L, b, C)$ iff there is some $u \in E_{\text{PA}}$ and some $w \in C$ such that $t \xrightarrow{w} u$, $u \xrightarrow{\mathcal{A}_\emptyset} q_\emptyset$, $u \xrightarrow{\mathcal{A}_L} q_L$ and $(b = \text{false} \text{ iff } |w| = 0)$.*

Proof. $\mathcal{A}_{\text{Pre}^*[\mathbb{C}]}$ is $\mathcal{A}_{\text{Pre}^*}$ equipped with a new component and the proof follows exactly the lines of the proof of Lemma 5.3. We refer to this earlier proof and only explain how we deal with the new C components.

1. $t = 0$ or $t = X$: The conditions on the C component for the existence of rules of type 0, 1a and 1b agree with the statement of the lemma.
2. $t = t_1.t_2$: ($\Rightarrow$): Now, for $i = 1, 2$, we have $t_i \xrightarrow{\mathcal{A}_{\text{Pre}^*}} (q_\emptyset^i, q_L^i, b^i, C^i)$ and there is a type 3 rule $(\dots, C^1).(\dots, C^2) \mapsto (\dots, C)$. Also, the ind. hyp. gives $t_i \xrightarrow{w_i} u_i$ ($i = 1, 2$) with $w_1 \in C^1$. In the type 3b case, $w_1 \in C$. In the type 3a case, we use $C^1.C^2 \subseteq C$.

($\Leftarrow$): Here we have either (1) $u_2 = t_2$ and $t_1 \xrightarrow{w} u_1$, or (2) $u_1 \in L_\emptyset$ and $t_1.t_2 \xrightarrow{w_1} u_1.t_2 \xrightarrow{w_2} u_1.u_2$ with $w = w_1.w_2$.

In the first case we apply the induction hypothesis with C itself on t_1 and some C' containing ε on t_2 , then we can use a type 3b rule. In the second case, there must be a pair (C^1, C^2) in the seq-decomposition of C , with $w_i \in C^i$ and we just have to use the ind. hyp. and a type 3a rule.

3. $t = t_1 \parallel t_2$: This case is similar to the previous one. The ($\Leftarrow$) direction uses the pair accounting for w_1, w_2 in the paral-decomposition of C . The ($\Rightarrow$) direction uses the crucial fact that whenever $t_i \xrightarrow{w_i} u_i$ for $i = 1, 2$, we have $t_1 \parallel t_2 \xrightarrow{w} u_1 \parallel u_2$ for any shuffling w of w_1 and w_2 , in particular for the w that C must contain.

□

7.1 Applications to model-checking

The above results let us apply the model-checking method from section 6 to an extended EF logic where we now allow all $\langle C \rangle \varphi$ formulas for decomposable C . The semantics is given by $\text{Mod}(\langle C \rangle \varphi) \stackrel{\text{def}}{=} \text{Pre}^*[C](\text{Mod}(\varphi))$.

Decomposability of C is a quite general condition. It excludes the undecidable situations that would exist in the general regular case and immediately includes the extensions proposed in [May97b].

Observe that it is possible to combine decomposable constraints already in the model-checking algorithm: when $C \in \mathbb{C}$ and $C' \in \mathbb{C}'$ are decomposable, we can deal with $\langle C \cap C' \rangle \varphi$ directly (i.e. without constructing a finite decomposition system containing C and C') because it is obvious how to extend the construction for $\mathcal{A}_{Pre^*[\mathbb{C}]}$ to some $\mathcal{A}_{Pre^*[\mathbb{C}, \mathbb{C}']}$ where several C components are dealt with simultaneously.

We can also deal with $\langle C \cup C' \rangle \varphi$ and $\langle C.C' \rangle \varphi$ directly since $Pre^*[C \cup C'](L)$ and $Pre^*[C.C'](L)$ are $Pre^*[C](L) \cup Pre^*[C'](L)$ and $Pre^*[C](Pre^*[C'](L))$ for any C, C' and L .

8 Structural equivalence of PA terms

In this section we investigate the congruence $\equiv$ induced on PA terms by the following equations:

$$\begin{array}{llll}
 (C_{\parallel}) & t \parallel t' & \equiv & t' \parallel t \\
 (A_{\parallel}) & (t \parallel t') \parallel t'' & \equiv & t \parallel (t' \parallel t'') \\
 (A_{\cdot}) & (t.t').t'' & \equiv & t.(t'.t'') \\
 (N_1) & t.0 & \equiv & t \\
 (N_2) & 0.t & \equiv & t \\
 (N_3) & t \parallel 0 & \equiv & t \\
 (N_4) & 0 \parallel t & \equiv & t
 \end{array}$$

This choice of equations is motivated by the fact that several recent works on PA (and extensions) only consider processes up-to this same congruence. Our techniques could deal with variants.

It is useful to explain how our definition of PA compares with the definition used in [May97c, May97b]. We consider a transition system between terms from E_{PA} . The terms Mayr considers for his transition system can be seen as equivalence classes, modulo $\equiv$, of our E_{PA} terms. Write $[t]_{\equiv}$ for the set $\{t' \mid t \equiv t'\}$. The transition relation used by Mayr coincides with a transition relation defined by

$$[t]_{\equiv} \xrightarrow{a} [u]_{\equiv} \iff \exists t' \in [t]_{\equiv}, u' \in [u]_{\equiv} \text{ s.t. } t' \xrightarrow{a} u'. \quad (4)$$

In the following, we speak of “PA $_{\equiv}$ ” when we mean the transition system one obtains with $\equiv$ -classes of terms as states, and transitions given by (4).

Our approach is more general in the sense that we can define the other approach in our framework. By contrast, if one reasons modulo $\equiv$ right from the start, one loses the information required to revert to the other approach.

For example, the reachability problem “do we have $t \xrightarrow{*} u$?” from Theorem 4.7 asks for a very precise form for u . The reachability problem solved in [May97c] asks for u modulo $\equiv$. In our framework, this can be stated as “given t and u , do we have $t' \xrightarrow{*} u'$ for some $t' \equiv t$ ”

and $u' \equiv u$?" (see below). In the other framework, it is impossible to state our problem. (But of course, the first motivation for our framework is that it allows the two regularity theorems.)

The rest of this section is devoted to some applications of our tree-automata approach to problems for $\text{PA}_{\equiv}$. The aim is not exhaustivity. Rather, we simply want to show that our framework allows solving (not just stating) problems from the other framework and its variants.

8.1 Structural equivalence and regularity

$(A.)$, $(C_{||})$ and $(A_{||})$ are the associativity-commutativity axioms satisfied by $.$ and $||$. We call them the *permutative axioms* and write $t =_P u$ when t and u are permutatively equivalent.

(N_1) to (N_4) are the axioms defining 0 as the neutral element of $.$ and $||$. We call them the *simplification axioms* and write $t \searrow u$ when u is a simplification of t , i.e. u can be obtained by applying the simplification axioms *from left to right* at some positions in t . Note that $\searrow$ is a (well-founded) partial ordering. We write $\swarrow$ for $(\searrow)^{-1}$. The *simplification normal form* of t , written $t\downarrow$, is the unique u one obtains by simplifying t as much as possible (no permutation allowed).

Such axioms are classical in rewriting and have been extensively studied [BN98]. $\equiv$ coincide with $(=_P \cup \searrow \cup \swarrow)^*$. Now, because the permutative axioms commute with the simplification axioms, we have

$$t \equiv t' \quad \text{iff} \quad t \searrow u =_P u' \swarrow t' \text{ for some } u, u' \quad \text{iff} \quad t\downarrow =_P t'\downarrow. \quad (5)$$

This lets us decompose questions about $\equiv$ into questions about $=_P$ and questions about $\searrow$. We start with $=_P$.

Lemma 8.1. *For any t , the set $[t]_{=_P} \stackrel{\text{def}}{=} \{u \mid t =_P u\}$ is a regular tree language, and an automaton for $[t]_{=_P}$ needs only have $m.(m/2)!$ states if $|t| = m$.*

Proof. (Sketch) This is because $[t]_{=_P}$ is a finite set with at most $(m/2)!$ elements. (The exponential blowup cannot be avoided.) $\square$

The simplification axioms do not have the nice property that they only allow finitely many combinations, but they behave better w.r.t. regularity. Write $[L]_{\searrow}$ for $\{u \mid t \searrow u \text{ for some } t \in L\}$, $[L]_{\swarrow}$ for $\{u \mid u \searrow t \text{ for some } t \in L\}$, and $[L]_{\downarrow}$ for $\{t\downarrow \mid t \in L\}$.

Lemma 8.2. *For any regular L , the sets $[L]_{\swarrow}$, $[L]_{\searrow}$, and $[L]_{\downarrow}$ are regular tree languages. From an automaton $\mathcal{A}_L$ recognizing L , we can build automata of size $O(|\mathcal{A}|)$ for these three languages in polynomial time.*

Proof. 1. $[L]_{\swarrow}$: u is in $[L]_{\swarrow}$ iff u is some $t \in L$ with additional 0 's that can be simplified out. Hence an automaton accepting $[L]_{\swarrow}$ is obtained from $\mathcal{A}_L$ by adding a new state q_0 for the subterms that will be simplified. We also add rules $0 \mapsto q_0$, $q_0 \parallel q_0 \mapsto q_0$, and

$q_0.q_0 \mapsto q_0$ for accepting these subterms, and, for any q in $\mathcal{A}_L$, rules $q.q_0 \mapsto q$, $q_0.q \mapsto q$, $q \parallel q_0 \mapsto q$ and $q_0 \parallel q \mapsto q$ for simulating simplification.

2. $[L]_{\searrow}$: u is in $[L]_{\searrow}$ iff u is some $t \in L$ where some 0's have been simplified. A simple way to obtain an automaton for $[L]_{\searrow}$ is to synchronize the automaton $\mathcal{A}_L$ accepting L with the complete automaton $\mathcal{A}_0$ recognizing terms built with 0, . and $\parallel$ only. $\mathcal{A}_0$ has only two states: q^0 and $q^{\neq 0}$.

Once the two automata are synchronized, we have $t \mapsto (q, q')$ iff $t \xrightarrow{\mathcal{A}_L} q$ and $t \xrightarrow{\mathcal{A}_0} q'$. We simulate simplification of nullable terms with additional ε -rules. Namely, whenever there is a rule $(q_1, q'_1) \parallel (q_2, q'_2) \mapsto (q_3, q'_3)$ with $q'_2 = q^0$, we add an ε -rule $(q_1, q'_1) \mapsto (q_3, q'_3)$. We add a symmetric rule if $q'_1 = q^0$ and do the same for . instead of $\parallel$.

Now a routine induction on the length of derivations shows that $s \mapsto (q, q')$ iff $\exists t \in L$ s.t. $t \searrow s$ and $t \xrightarrow{\mathcal{A}_L} q$.

3. $[L]_{\downarrow}$: The simplest way to see regularity is to note that $[L]_{\downarrow}$ is $[L]_{\searrow} \cap [E_{PA}]_{\downarrow}$. □

Note that for a regular L , $[L]_{=P}$ and $[L]_{\equiv}$ are not necessarily regular [GD89]. However we have

Proposition 8.3. *For any t , the set $[t]_{\equiv}$ is a regular tree language, and an automaton for $[t]_{\equiv}$ needs only have $m.(m/2)!$ states if $|t| = m$.*

Proof. Combine (5) with lemmas 8.1 and 8.2. □

8.2 Structural equivalence and behaviour

Seeing terms modulo $\equiv$ does not modify the observable behaviour because of the following standard result:

Proposition 8.4. *$\equiv$ is a bisimulation relation, i.e. for all $t \equiv t'$ and $t \xrightarrow{a} u$ there is a $t' \xrightarrow{a} u'$ with $u \equiv u'$ (and vice versa).*

The proof is standard but tedious. We shall only give a proof sketch.

Proof. For any single equation $l = r$ in the definition of $\equiv$, we show that the set $R = \{(\iota\sigma, r\sigma)\}$ of all instances of the equation is a bisimulation relation. A complete proof of this for $(A_{\parallel})$ takes the better part of p 95 of the book [Mil89] and the other equations can be dealt with similarly, noting that $IsNil()$ is compatible with $\equiv$. Then there only remains to prove that the generated congruence is a bisimulation. This too is standard: the SOS rules for PA obey a format ensuring that the behaviour of a term depends on the behaviour of its subterms, not their syntax. □

We may now define a new transition relation between terms: $t \xRightarrow{a} t'$ iff $t \equiv u \xrightarrow{a} u' \equiv t'$ for some u, u' . This amounts to the " $[t]_{\equiv} \xrightarrow{a} [u]_{\equiv}$ " from (4) and is the simplest way to translate

problems for $\text{PA}_\equiv$ into problems for our set of terms.

We adopt the usual abbreviations $\xrightarrow{*}, \xrightarrow{w}, \xrightarrow{k}$ for $w \in \text{Act}^*, k \in \mathbb{N}$, etc.

Proposition 8.5. *For any $w \in \text{Act}^*, t \xrightarrow{w} u$ iff $t \xrightarrow{w} u'$ for some $u' \equiv u$.*

Proof. By induction on the length of w , and using Proposition 8.4. $\square$

8.3 Reachability modulo $\equiv$

Now it is easy to prove decidability of the reachability problem modulo $\equiv$: $t \xrightarrow{*} u$ iff $\text{Post}^*(t) \cap [u]_\equiv \neq \emptyset$. Recall that $[u]_\equiv$ and $\text{Post}^*(t)$ are regular tree-languages one can build effectively. Hence it is decidable whether they have a non-empty intersection.

This gives us a simple algorithm using exponential time (because of the size of $[u]_\equiv$). Actually we can have a better result ³:

Theorem 8.6. *The reachability problem in $\text{PA}_\equiv$, “given t and u , do we have $t \xrightarrow{*} u$?”, is in NP.*

Proof. NP-easiness is straightforward in the automata framework: We have $t \xrightarrow{*} u$ iff $t \xrightarrow{*} u'$ for some u' s.t. $u' \downarrow =_P u \downarrow$. Write u'' for $u' \downarrow$ and note that $|u''| \leq |u|$. A simple algorithm is to compute $u \downarrow$, then *guess non-deterministically* a permutation u'' , then build automata $\mathcal{A}_1$ for $[u'']_{\setminus \downarrow}$ and $\mathcal{A}_2$ for $\text{Post}^*(t)$. These automata have polynomial-size. There remains to check whether $\mathcal{A}_1$ and $\mathcal{A}_2$ have a non-empty intersection to know whether the required u' exists. $\square$

Corollary 8.7. *The reachability problem in $\text{PA}_\equiv$ is NP-complete.*

Proof. NP-hardness of reachability for BPP's is proved in [Esp97] and the proof idea can be reused in our framework. We reduce 3SAT to reachability in $\text{PA}_\equiv$. Consider an instance P of 3SAT. P has m variables and n clauses, so that it is some $\bigwedge_{i=1}^n \bigvee_{j=1}^3 \varepsilon_{i,j} x_{r_{i,j}}$ where, for every i, j , $1 \leq r_{i,j} \leq m$ and $\varepsilon_{i,j} \in \{+, -\}$. We define the following Δ_P :

$$\Delta_P \stackrel{\text{def}}{=} \begin{cases} \text{(R1)} & X_r \rightarrow X_r^\varepsilon & \text{for } 1 \leq r \leq m \text{ and } \varepsilon \in \{+, -\}, \\ \text{(R2)} & X_r^\varepsilon \rightarrow 0 & \text{for } 1 \leq r \leq m \text{ and } \varepsilon \in \{+, -\}, \\ \text{(R3)} & X_{r_{i,j}}^{\varepsilon_{i,j}} \rightarrow C_j \parallel X_{r_{i,j}}^{\varepsilon_{i,j}} & \text{for } 1 \leq i \leq n \text{ and } 1 \leq j \leq 3. \end{cases}$$

(Note that $|\Delta_P| = O(|P|)$.) The (R1) rules pick a valuation v for the X_r 's, the (R3) rules use v to list satisfied clauses, the (R2) rules discard unnecessary elements. Finally

$$(X_1 \parallel (X_2 \parallel (\dots \parallel X_m) \dots)) \xrightarrow{*} (C_1 \parallel (C_2 \parallel (\dots \parallel C_n) \dots)) \text{ iff } P \text{ is satisfiable.}$$

$\square$

Other applications are possible, e.g.:

³First proved in [May97c]

Proposition 8.8. *The boundedness problem in $\text{PA}_\equiv$ is decidable in polynomial-time.*

Proof. $[t]_\equiv$ can only reach a finite number of states in $\text{PA}_\equiv$ iff t can only reach a finite number of non- $\equiv$ terms in PA. Now because the permutative axioms only allow finitely many variants of any given term, $\text{Post}^*(L)$ contains a finite number of non- $\equiv$ processes iff $[\text{Post}^*(L)]_\downarrow$ is finite. $\square$

8.4 Model-checking modulo $\equiv$

The model-checking problem solved in [May97b] considers the EF logic over $\text{PA}_\equiv$. Translated into our framework, this amounts to interpret the temporal connectives in terms of $\Rightarrow$ instead of $\rightarrow$: if we write $\text{Mod}_\equiv(\varphi)$ for the interpretation modulo $\equiv$, we have

$$\text{Mod}_\equiv(\langle C \rangle \varphi) \stackrel{\text{def}}{=} \{t \mid t \xrightarrow{w} u \text{ for some } u \in \text{Mod}_\equiv(\varphi) \text{ and some } w \in C\}.$$

Additionally, we only consider atomic propositions P compatible with $\equiv$, i.e. where $t \models P$ and $t \equiv u$ imply $u \models P$.

Model-checking in $\text{PA}_\equiv$ is as simple as model-checking in PA:

Lemma 8.9. *For any EF-formula φ we have $\text{Mod}_\equiv(\varphi) = \text{Mod}(\varphi) = [\text{Mod}(\varphi)]_\equiv$.*

Proof. By structural induction over φ , using Prop. 8.5 and closure w.r.t. $\equiv$ for the $\langle C \rangle \varphi$ case. $\square$

The immediate corollary is that we can use exactly the same approach for model-checking in PA with or without $\equiv$.

Conclusion

In this paper we showed how tree-automata techniques are a powerful tool for the analysis of the PA process algebra. Our main results are two general Regularity Theorems with numerous immediate applications, including model-checking of PA with an extended EF logic.

The tree-automata viewpoint has many advantages. It gives simpler and more general proofs. It helps understand why some problems can be solved in P-time, some others in NP-time, etc. It is quite versatile and many variants of PA can be attacked with the same approach.

Acknowledgments We thank H. Comon and R. Mayr for their numerous suggestions, remarks and questions about this work.

References

- [BBK87] J. C. M. Baeten, J. A. Bergstra, and J. W. Klop. Decidability of bisimulation equivalence for processes generating context-free languages. In *Proc. Parallel Architectures and Languages Europe (PARLE'87)*, Eindhoven, NL, June 1987, vol. II: *Parallel Languages*, volume 259 of *Lecture Notes in Computer Science*, pages 94–111. Springer-Verlag, 1987.
- [BE97] O. Burkart and J. Esparza. More infinite results. In *Proc. 1st Int. Workshop on Verification of Infinite State Systems (INFINITY'96)*, Pisa, Italy, Aug. 30–31, 1996, volume 5 of *Electronic Notes in Theor. Comp. Sci.* Elsevier, 1997.
- [BEH95] A. Bouajjani, R. Echahed, and P. Habermehl. Verifying infinite state processes with sequential and parallel composition. In *Proc. 22nd ACM Symp. Principles of Programming Languages (POPL'95)*, San Francisco, CA, USA, Jan. 1995, pages 95–106, 1995.
- [BEM97] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *Proc. 8th Int. Conf. Concurrency Theory (CONCUR'97)*, Warsaw, Poland, Jul. 1997, volume 1243 of *Lecture Notes in Computer Science*, pages 135–150. Springer-Verlag, 1997.
- [BN98] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- [Büc64] J. R. Büchi. Regular canonical systems. *Arch. Math. Logik Grundlag.*, 6:91–111, 1964.
- [Cau92] D. Caucal. On the regular structure of prefix rewriting. *Theoretical Computer Science*, 106(1):61–86, 1992.
- [CDG⁺97] H. Comon, M. Dauchet, R. Gilleron, D. Lugiez, S. Tison, and M. Tommasi. Tree automata and their application, 1997. A preliminary version of this (yet unpublished) book is available at <http://13ux02.univ-lille3.fr/tata>.
- [CHM94] S. Christensen, Y. Hirshfeld, and F. Møller. Decidable subsets of CCS. *The Computer Journal*, 37(4):233–242, 1994.
- [Chr93] S. Christensen. Decidability and decomposition in process algebras. PhD thesis CST-105-93, Dept. of Computer Science, University of Edinburgh, UK, 1993.
- [CKSV97] H. Comon, D. Kozen, H. Seidl, and M. Y. Vardi, editors. *Applications of Tree Automata in Rewriting, Logic and Programming*, Dagstuhl-Seminar-Report number 193. Schloß Dagstuhl, Germany, 1997.
- [DT90] M. Dauchet and S. Tison. The theory of ground rewrite systems is decidable. In *Proc. 5th IEEE Symp. Logic in Computer Science (LICS'90)*, Philadelphia, PA, USA, June 1990, pages 242–248, 1990.
- [Esp97] J. Esparza. Petri nets, commutative context-free grammars, and basic parallel processes. *Fundamenta Informaticae*, 31(1):13–25, 1997.
- [FWW97] A. Finkel, B. Willems, and P. Wolper. A direct symbolic approach to model checking pushdown systems (extended abstract). In *Proc. 2nd Int. Workshop on Verification of Infinite State Systems (INFINITY'97)*, Bologna, Italy, July 11–12, 1997, volume 9 of *Electronic Notes in Theor. Comp. Sci.* Elsevier, 1997.
- [GD89] R. Gilleron and A. Deruyver. The reachability problem for ground TRS and some extensions. In *Proc. Int. Joint Conf. Theory and Practice of Software Development*

- (TAPSOFT'89), *Barcelona, Spain, March 1989, Volume 1*, pages 227–243. Springer-Verlag, 1989.
- [GS97] F. Gécseg and M. Steinby. Tree languages. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, chapter 1, pages 1–68. Springer-Verlag, 1997.
 - [JKM98] P. Jančar, A. Kučera, and R. Mayr. Deciding bisimulation-like equivalences with finite-state processes. Tech. Report TUM-I9805, Institut für Informatik, TUM, Munich, Germany, February 1998.
 - [KS97a] O. Kouchnarenko and Ph. Schnoebelen. A model for recursive-parallel programs. In *Proc. 1st Int. Workshop on Verification of Infinite State Systems (INFINITY'96)*, Pisa, Italy, Aug. 1996, volume 5 of *Electronic Notes in Theor. Comp. Sci.* Elsevier, 1997.
 - [KS97b] O. Kushnarenko and Ph. Schnoebelen. A formal framework for the analysis of recursive-parallel programs. In *Proc. 4th Int. Conf. Parallel Computing Technologies (PaCT'97)*, Yaroslavl, Russia, Sep. 1997, volume 1277 of *Lecture Notes in Computer Science*, pages 45–59. Springer-Verlag, 1997.
 - [Kuč96] A. Kučera. Regularity is decidable for normed PA processes in polynomial time. In *Proc. 16th Conf. Found. of Software Technology and Theor. Comp. Sci. (FST&TCS'96)*, Hyderabad, India, Dec. 1996, volume 1180 of *Lecture Notes in Computer Science*, pages 111–122. Springer-Verlag, 1996.
 - [Kuč97] A. Kučera. How to parallelize sequential processes. In *Proc. 8th Int. Conf. Concurrency Theory (CONCUR'97)*, Warsaw, Poland, Jul. 1997, volume 1243 of *Lecture Notes in Computer Science*, pages 302–316. Springer-Verlag, 1997.
 - [May97a] R. Mayr. Combining Petri nets and PA-processes. In *Proc. 4th Int. Symp. Theoretical Aspects Computer Software (TACS'97)*, Sendai, Japan, Sep. 1997, volume 1281 of *Lecture Notes in Computer Science*, pages 547–561. Springer-Verlag, 1997.
 - [May97b] R. Mayr. Model checking PA-processes. In *Proc. 8th Int. Conf. Concurrency Theory (CONCUR'97)*, Warsaw, Poland, Jul. 1997, volume 1243 of *Lecture Notes in Computer Science*, pages 332–346. Springer-Verlag, 1997.
 - [May97c] R. Mayr. Tableaux methods for PA-processes. In *Proc. Int. Conf. Automated Reasoning with Analytical Tableaux and Related Methods (TABLEAUX'97)*, Pont-à-Mousson, France, May 1997, volume 1227 of *Lecture Notes in Artificial Intelligence*, pages 276–290. Springer-Verlag, 1997.
 - [Mil89] R. Milner. *Communication and Concurrency*. Prentice Hall Int., 1989.
 - [Mol96] F. Moller. Infinite results. In *Proc. 7th Int. Conf. Concurrency Theory (CONCUR'96)*, Pisa, Italy, Aug. 1996, volume 1119 of *Lecture Notes in Computer Science*, pages 195–216. Springer-Verlag, 1996.

Unite de recherche INRIA Lorraine, Technople de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LS NANCY
Unite de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unite de recherche INRIA Rhone-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unite de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unite de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

diteur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399