

Recurrent Reachability Analysis in Regular Model Checking

Anthony Widjaja To and Leonid Libkin

LFCS, School of Informatics, University of Edinburgh
`anthony.w.to,libkin@ed.ac.uk`

Abstract. We consider the problem of recurrent reachability over infinite systems given by regular relations on words and trees, i.e. whether a given regular set of states can be reached infinitely often from a given initial state in the given transition system. Under the condition that the transitive closure of the transition relation is regular, we show that the problem is decidable, and the set of all initial states satisfying the property is regular. Moreover, our algorithm constructs an automaton for this set in polynomial time, assuming that a transducer of the transitive closure can be computed in poly-time. We then demonstrate that transition systems generated by pushdown systems, regular ground tree rewrite systems, and the well-known process algebra PA satisfy our condition and transducers for their transitive closures can be computed in poly-time. Our result also implies that model checking EF-logic extended by recurrent reachability predicate (EGF) over such systems is decidable.

1 Introduction

Infinite-state systems play an important role in verification as they capture many scenarios that cannot be adequately described by standard finite-state models. For example, the behavior of parameterized systems needs to be checked regardless of the number of processes, and this is often most suitably represented by an infinite-state system.

The most common verification problems for such systems can be abstracted as reachability and recurrent reachability [3, 8, 2, 9]. Reachability asks if a given state, or a state in a given set, can be reached from an initial state. Checking these is essential for verifying safety of infinite-state systems, as we want to find counterexamples to specifications saying that bad states cannot be reached. If we have slightly weaker specifications saying that undesirable states can only appear in some initial portion of each execution path, then counterexamples to those are formalized as *recurrent reachability*, i.e. the existence of a witnessing path that infinitely often goes through a given set of states. In the CTL* notation, recurrent reachability for a set L is **EGFL**. Observe that although for finite systems recurrent reachability is reducible to reachability, this is not the case for infinite systems in general, e.g. lossy channel systems (see [1]).

To make the questions of model checking meaningful for infinite systems, they need to have an effective finite representation. Often, the state space is described

by regular word or tree languages, and transitions are given by regular word or tree transducers: this is the general framework of regular model-checking [3, 2]. Without restrictions, this does not guarantee decidability even for the simplest reachability properties. Hence, one normally restricts the class of transducers so that their iterations would remain regular [7, 17]. Then such infinite-state systems have an effective word-automatic or tree-automatic presentation [6, 5]. Some of the most well-known and most studied classes of such systems include pushdown systems [15, 14, 22, 27], prefix-recognizable graphs [11, 24], ground tree rewrite systems [13, 19], and the process algebra PA [4, 20, 22].

For such systems, reachability has been extensively studied [3, 17, 9, 23, 15, 16, 14, 27]. Much less is known about recurrent reachability. Unlike reachability, it is not immediately seen to be decidable even under the assumption that the transitive closure of the transition relation is representable by a regular transducer. One recent result [18, 19] showed that recurrent reachability is decidable for infinite-state transition systems generated by ground tree rewrite systems.

Our main contributions are as follows. We look at arbitrary infinite-state transition systems that have an automatic representation (either word-automatic or tree-automatic) and that further satisfy the condition that the transitive closure of its transition relation is regular. We then show the following:

1. For every regular language L , the set of all states that satisfy a recurrent reachability property **EGFL** is also regular. This observation gives rise to two flavors of the model-checking problem: the global problem is to construct an automaton accepting the set of states satisfying **EGFL**, and the local problem is to verify whether a given word/tree satisfies the property.
2. We give a *generic* poly-time algorithm that solves both model-checking problems for **EGFL** given the following as inputs: word/tree regular transducers defining a transition relation and its transitive closure, and a nondeterministic word/tree automaton defining L . For positive answers, our algorithm also constructs some witnessing infinite paths using Büchi word/tree automata as finite representations. In particular, if the transducer defining the transitive closure can itself be computed in poly-time, we obtain a poly-time algorithm for checking recurrent reachability properties. One can also combine our algorithm with the semi-algorithms for computing iterating transducers developed in regular model checking (e.g. [2, 3, 7, 17]).
3. We then look at some particular examples of transition systems in which the transitive closure of the transition relation is regular for which an iterating transducer is poly-time computable. As corollaries, we obtain poly-time algorithms for recurrent reachability over pushdown systems, ground tree rewrite systems, and PA-processes. These also imply that the extension of the **EF**-logic [19, 23, 8] with the **EGF** operator remains decidable for all those examples. For the first two examples, our results follow from known results [15, 18, 19] proven using specialized methods for pushdown systems and ground tree rewrite systems, although their methods do not show how to compute witnessing paths, which are also of interests in verification. Our results for PA-processes are new.

Outline of the paper In Section 2, we recall some basic definitions. In Section 3 we prove our results for transition systems that have word-automatic presentations, and provide applications to pushdown systems. In Section 4 we prove results for tree-automatic presentations, and provide applications to ground tree rewrite systems and PA-processes. We conclude in Section 5 with future work.

2 Preliminaries

Transition systems Let $\text{AP} = \{P_1, \dots, P_n\}$ be a finite set of *atomic propositions*. A *transition system* over AP is

$$\mathcal{S} = \langle S, \rightarrow, \lambda \rangle,$$

where S is a set of *states*, $\rightarrow \subseteq S \times S$ is a *transition relation*, and $\lambda : \text{AP} \rightarrow 2^S$ is a function defining which states satisfy any particular atomic proposition. The set S is not required to be finite.

We write \rightarrow^+ (resp. \rightarrow^*) to denote the transitive (resp. transitive-reflexive) closure of \rightarrow . If $S' \subseteq S$, then $\text{pre}^*(S')$ (resp. $\text{post}^*(S')$) denotes the set of states s that can reach (resp. be reached from) some state in S' .

Recurrent reachability Given a transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ and a set $S' \subseteq S$, we write $s \rightarrow^\omega S'$ iff there exists an infinite sequence $\{s_i\}_{i \in \mathbb{N}}$ such that $s_0 = s$ and $s_i \in S'$ and $s_{i-1} \rightarrow^+ s_i$ for all $i > 0$. By transitivity of \rightarrow^+ , every infinite subsequence of such a sequence $\{s_i\}_{i \in \mathbb{N}}$ that starts with s_0 is also a witness for $s \rightarrow^\omega S'$. We write $\text{Rec}(S')$ to denote the set of states s such that $s \rightarrow^\omega S'$. We will also write $\text{Rec}(S', \rightarrow^+)$ to emphasize the transitive binary relation in use.

Words, Trees, and Automata We assume basic familiarity with automata on finite and infinite words and trees (see [25, 12]). Fix a finite alphabet Σ . For each finite word $w = w_1 \dots w_n \in \Sigma^*$, we write $w[i, j]$, where $1 \leq i \leq j \leq n$, to denote the segment $w_i \dots w_j$. Given an automaton $\mathcal{A} = (\Sigma, Q, \delta, q_0, F)$, a run of \mathcal{A} on w is a function $\rho : \{0, \dots, n\} \rightarrow Q$ with $\rho(0) = q_0$ that obeys δ . In this case, the length $\|\rho\|$ of ρ is n . The last state $\rho(n)$ appearing in ρ is denoted by $\mathbf{last}(\rho)$; the first state $\rho(0)$ is denoted by $\mathbf{first}(\rho)$. We also define *run segments* to be runs that do not necessarily start from q_0 . Given a run segment $\rho' : \{0, \dots, m\} \rightarrow Q$ such that $\mathbf{first}(\rho') = \mathbf{last}(\rho)$, we may concatenate ρ and ρ' to obtain a new run $\rho \odot \rho' : \{0, \dots, n + m\} \rightarrow Q$ defined in the obvious way. We also use the notation $\rho[i, j]$ to denote the segment $\rho(i) \dots \rho(j)$. A run on an ω -word $w \in \Sigma^\omega$ is a function $\rho : \mathbb{N} \rightarrow Q$ with $\rho(0) = q_0$ that obeys δ . We use abbreviations NWA and NBWA for nondeterministic (Büchi) word automata.

Given a finite direction alphabet Υ , a *tree domain* is a non-empty prefix-closed set $D \subseteq \Upsilon^*$. The empty word (denoted by ε) is referred to as the root. Words $u \in D$ so that no ui is in D are called *leaves*. A *tree* T is a pair (D, τ) , where D is a tree domain and τ is a node-labeling function mapping D to Σ .

The tree T is said to be *finite* if D is finite; otherwise, it is said to be *infinite*. The tree T is said to be *complete* if, whenever $u \in D$, if $ui \in D$ for some $i \in \mathcal{Y}$, then $uj \in D$ for all $j \in \mathcal{Y}$. If T is infinite, it is said to be *full* if $D = \mathcal{Y}^*$. The set of all finite trees over \mathcal{Y} and Σ is denoted by $\text{TREE}_{\mathcal{Y}}(\Sigma)$. If $\mathcal{Y} = \{1, 2\}$, we write $\text{TREE}_2(\Sigma)$ for $\text{TREE}_{\mathcal{Y}}(\Sigma)$.

A (top-down) *tree automaton* $\mathcal{A} = (\Sigma, Q, \delta, q_0, F)$ over finite Σ -labeled trees has a transition function $\delta : Q \times \Sigma \rightarrow 2^{Q^m}$, where $m = |\mathcal{Y}|$. For our constructions, it will be most convenient to define runs on trees with virtual leaf nodes. We define $\mathbf{virt}(T)$ to be the $(\Sigma \cup \perp)$ -labeled \mathcal{Y} -tree (D', τ') such that $D' := D \cup \{vi : v \in D, i \in \mathcal{Y}\}$ and if $u \in D$, then $\tau'(u) := \tau(u)$; if $u \notin D$, then $\tau'(u) = \perp$. Notice that $\mathbf{virt}(T)$ is complete. A *run* of \mathcal{A} on T , i.e. a mapping $\rho : D' \rightarrow Q$, starts in the initial state q_0 and for each node u labeled a with children $u1, \dots, um$, we have $(\rho(u1), \dots, \rho(um)) \in \delta(q, a)$. A run is *accepting* if $\rho(u) \in F$ for each leaf $u \in D'$.

We abbreviate nondeterministic tree automata as NTA, and write NBTA for tree automata over full infinite trees with a Büchi acceptance condition (that will be sufficient for our purposes). For all kinds of automata, $L(\mathcal{A})$ stands for the word or tree languages accepted by \mathcal{A} . Also for all types of automata \mathcal{A} , we write \mathcal{A}^q for \mathcal{A} in which the initial state is set to q .

Transducers These will be given by *letter-to-letter automata* that accept binary (and, more generally, k -ary) relations over words and trees (cf. [6, 5]). We start with words. Given two words $w = w_1 \dots w_n$ and $w' = w'_1 \dots w'_m$ over the alphabet Σ , we define a word $w \otimes w'$ of length $k = \max\{n, m\}$ over alphabet $\Sigma_{\perp} \times \Sigma_{\perp}$, where $\Sigma_{\perp} = \Sigma \cup \{\perp\}$ and $\perp \notin \Sigma$, as follows:

$$w \otimes w' = \begin{bmatrix} a_1 \\ b_1 \end{bmatrix} \dots \begin{bmatrix} a_k \\ b_k \end{bmatrix}, \text{ where } a_i = \begin{cases} w_i & i \leq n \\ \perp & i > n, \end{cases} \text{ and } b_i = \begin{cases} w'_i & i \leq m \\ \perp & i > m. \end{cases}$$

In other words, the shorter word is padded with \perp 's, and the i th letter of $w \otimes w'$ is then the pair of the i th letters of padded w and w' . A letter-to-letter automaton is simply an automaton over $\Sigma_{\perp} \times \Sigma_{\perp}$, and a binary relation R over Σ^* is *regular* if the set $\{w \otimes w' : (w, w') \in R\}$ is accepted by a letter-to-letter automaton \mathcal{R} . We shall refer to such an automaton as a *transducer* over Σ^* , since it can be alternatively viewed as mapping words $w \in \Sigma^*$ nondeterministically into words w' so that $w \otimes w'$ is accepted by \mathcal{R} .

Given two trees $T_1 = (D_1, \tau_1)$ and $T_2 = (D_2, \tau_2)$, we define $T = T_1 \otimes T_2$ as a tree over the labeling alphabet Σ_{\perp}^2 similarly to the definition of $w \otimes w'$. That is, the domain of T is $D_1 \cup D_2$, and the labeling $\tau : D_1 \cup D_2 \rightarrow \Sigma_{\perp}^2$ is defined as $\tau(u) = (a_1, a_2)$ so that $a_i = \tau_i(u)$ if $u \in D_i$ and \perp otherwise, for $i = 1, 2$.

As for words, a binary relation R over $\text{TREE}_{\mathcal{Y}}(\Sigma)$ is regular if there is a tree automaton \mathcal{R} over $\text{TREE}_{\mathcal{Y}}(\Sigma_{\perp}^2)$ accepting the set $\{T_1 \otimes T_2 \mid (T_1, T_2) \in R\}$. We also view it as a transducer that nondeterministically assigns to a tree T_1 any tree T_2 so that $(T_1, T_2) \in R$. If the binary relation R defined by a transducer \mathcal{R} is transitive, we shall refer to \mathcal{R} itself as being transitive.

Automatic transition systems In this paper, we deal with infinite transition systems that can be finitely represented by word or tree automata. We say that a transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ over AP is *word-automatic* if, for some finite alphabet Σ , we have $S = \Sigma^*$, the relation \rightarrow is a regular relation on S , and each $\lambda(P_i)$ is a regular subset of S . Likewise, a transition system \mathcal{S} over AP is *tree-automatic* if, for some Υ and Σ , we have $S = \text{TREE}_\Upsilon(\Sigma)$, and all of \rightarrow and $\lambda(P_i)$'s are regular tree relations/languages over $\text{TREE}_\Upsilon(\Sigma)$.

We measure the size of such a word- or tree-automatic transition system \mathcal{S} as the total size of the transducer for \rightarrow , and the automata for S and $\lambda(P_i)$, for $P_i \in \text{AP}$. We shall assume that these are nondeterministic.

As mentioned already, pushdown systems and prefix-recognizable graphs are examples of word-automatic infinite transition systems, while PA-processes and graphs generated by ground tree rewrite systems are examples of tree-automatic transition systems.

3 Recurrent reachability: the word case

We call a word-automatic transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ *transitive* if the relation \rightarrow^+ is regular. As we shall see shortly, if \mathcal{S} is transitive, then the set $\text{Rec}(L)$ is regular too, for an arbitrary regular language L . This gives rise to two variants of the model-checking problem for recurrent reachability: in the global model-checking problem, we are given \mathcal{S} and a language L represented by an NWA \mathcal{A} , and we want to construct an NWA accepting $\text{Rec}(L(\mathcal{A}))$. In the local version, we also have a word w , and we must check whether $w \in \text{Rec}(L(\mathcal{A}))$. That is,

GLOBAL MODEL-CHECKING:	INPUT: 1) A transitive word-automatic \mathcal{S} 2) An NWA \mathcal{A} OUTPUT: A description of $\text{Rec}(L(\mathcal{A}))$
LOCAL MODEL-CHECKING:	INPUT: 1) A transitive word-automatic \mathcal{S} 2) An NWA \mathcal{A} 3) a word w OUTPUT: <i>yes</i> , if $w \rightarrow^\omega L(\mathcal{A})$ <i>no</i> , otherwise

Throughout this section, we assume that the transition relation \rightarrow of transitive \mathcal{S} is given by a transducer \mathcal{R} , and that \mathcal{R}^+ is the transducer for \rightarrow^+ (which exists by the transitivity assumption). We shall also use the transducer for \rightarrow^* , denoted by \mathcal{R}^* . It can be obtained from \mathcal{R}^+ by letting it accept pairs $w \otimes w$.

Theorem 1. *Given a transitive word-automatic transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ and an NWA \mathcal{A} , the set $\text{Rec}(L(\mathcal{A}))$ of states w such that $w \rightarrow^\omega L(\mathcal{A})$ is regular.*

Moreover, if the transducer \mathcal{R}^+ for \rightarrow^+ is computable in time $t(|\mathcal{R}|)$, then one can compute an NWA recognizing $\text{Rec}(L(\mathcal{A}))$ of size $O(|\mathcal{R}^+|^2 \times |\mathcal{A}|)$ in time $t(|\mathcal{R}|) + O(|\mathcal{R}^+|^3 \times |\mathcal{A}|^2)$.

Corollary 2. *Given a transitive word-automatic transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ and an NWA \mathcal{A} , such that the transducer \mathcal{R}^+ is poly-time computable, both global and local model-checking for recurrent reachability are solvable in poly-time.*

As another corollary, consider the **EF**, **EX**-fragment of CTL, known as the EF-logic [23, 27]. Its formulae over $\text{AP} = \{P_1, \dots, P_n\}$ are given by

$$\varphi, \varphi' := \top \mid P_i, i \leq n \mid \varphi \vee \varphi' \mid \neg\varphi \mid \mathbf{EX}\varphi \mid \mathbf{EF}\varphi.$$

Each formula, evaluated over a transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$, defines a set $\llbracket \varphi \rrbracket_{\mathcal{S}} \subseteq S$ as follows:

- (1) $\llbracket \top \rrbracket_{\mathcal{S}} = S$;
- (2) $\llbracket P_i \rrbracket_{\mathcal{S}} = \lambda(P_i)$;
- (3) $\llbracket \varphi \vee \varphi' \rrbracket_{\mathcal{S}} = \llbracket \varphi \rrbracket_{\mathcal{S}} \cup \llbracket \varphi' \rrbracket_{\mathcal{S}}$;
- (4) $\llbracket \neg\varphi \rrbracket_{\mathcal{S}} = S - \llbracket \varphi \rrbracket_{\mathcal{S}}$;
- (5) $\llbracket \mathbf{EX}\varphi \rrbracket_{\mathcal{S}} = \{s \mid \exists s' : s \rightarrow s' \text{ and } s' \in \llbracket \varphi \rrbracket_{\mathcal{S}}\}$;
- (6) $\llbracket \mathbf{EF}\varphi \rrbracket_{\mathcal{S}} = \{s \mid \exists s' : s \rightarrow^* s' \text{ and } s' \in \llbracket \varphi \rrbracket_{\mathcal{S}}\}$.

If \rightarrow^* is given by a regular transducer, then $\llbracket \varphi \rrbracket_{\mathcal{S}}$ is clearly effectively regular [6], and so the model-checking problem for EF-logic is decidable. We now extend this to the (EF+EGF)-logic, defined as the extension of EF-logic with the formulae **EGF** φ with the semantics

$$\llbracket \mathbf{EGF}\varphi \rrbracket_{\mathcal{S}} = \text{Rec}(\llbracket \varphi \rrbracket_{\mathcal{S}}, \rightarrow^+) = \{s \mid s \rightarrow^\omega \llbracket \varphi \rrbracket_{\mathcal{S}}\}.$$

Theorem 1 extends decidability to (EF+EGF)-logic:

Corollary 3. *If $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ is a transitive word-automatic transition system such that the transducer \mathcal{R}^+ is computable, then for each formula φ of (EF+EGF)-logic, the set $\llbracket \varphi \rrbracket_{\mathcal{S}}$ is regular, and an NWA defining $\llbracket \varphi \rrbracket_{\mathcal{S}}$ can be effectively constructed.*

We now prove Theorem 1. Throughout the proof, we let \mathcal{M} stand for \mathcal{R}^+ and use unambiguous abbreviations such as $\text{Rec}(\mathcal{A}, \mathcal{M})$ for $\text{Rec}(L(\mathcal{A}), L(\mathcal{M}))$. By definition, we have $w \in \text{Rec}(\mathcal{A}, \mathcal{M})$ iff there exists a sequence $\{s_i\}_{i \in \mathbb{N}}$ of words with $s_0 = w$ such that $s_{i-1} \otimes s_i \in L(\mathcal{M})$ and $s_i \in L(\mathcal{A})$ for all $i > 0$. We now divide $\text{Rec}(\mathcal{A}, \mathcal{M})$ into two sets $\text{Rec}_1(\mathcal{A}, \mathcal{M})$ and $\text{Rec}_2(\mathcal{A}, \mathcal{M})$, where $\text{Rec}_1(\mathcal{A}, \mathcal{M})$ contains words with a witnessing infinite sequence $\{s_i\}_{i \in \mathbb{N}}$ that satisfies $s_j = s_k$ for some $j < k$, and $\text{Rec}_2(\mathcal{A}, \mathcal{M})$ contains words with a witnessing infinite sequence $\{s_i\}_{i \in \mathbb{N}}$ that satisfies $s_j \neq s_k$ for all distinct $j, k \in \mathbb{N}$. We shall write Rec_1 and Rec_2 when the intended automata \mathcal{A} and \mathcal{M} are clear from the context. Now notice that $\text{Rec}(L(\mathcal{A})) = \text{Rec}_1 \cup \text{Rec}_2$. It is easy to construct an NWA that recognizes Rec_1 . Observe that, for all word w , we have $w \in \text{Rec}_1$ iff there exists a word w' such that $w \rightarrow^* w'$, $w' \rightarrow^+ w'$, and $w' \in L(\mathcal{A})$. By taking a product and then applying projection (e.g. see [6]), we can compute an NWA \mathcal{A}_1 that recognizes Rec_1 in time $O(|\mathcal{R}^+|^2 \times |\mathcal{A}|)$ with $|\mathcal{A}| = O(|\mathcal{R}^+|^2 \times |\mathcal{A}|)$.

Thus, it remains to construct the automaton \mathcal{A}_2 for Rec_2 . We shall first compute a Büchi automaton \mathcal{B} that recognizes an ω -word which represents the witnessing infinite sequence for membership in Rec_2 . Once \mathcal{B} is constructed, it is easy to obtain \mathcal{A}_2 as we shall see later. The most obvious representation of the

infinite sequence $\{s_i\}_{i \in \mathbb{N}}$ is $s_0 \otimes s_1 \otimes \dots$. The problem with this representation is that it requires an infinite alphabet, and possibly infinitely many copies of the automata \mathcal{A} and \mathcal{M} to check whether $s_i \in L(\mathcal{A})$ and $s_{i-1} \rightarrow^+ s_i$ for all $i > 0$. Therefore, the first step towards solving the problem is to analyze the infinite witnessing paths and to show that it is sufficient to consider only infinite sequences of a special form. For the rest of this section, we let $\mathcal{A} = (Q_1, \delta_1, q_0^1, F_1)$ and $\mathcal{M} = \mathcal{R}^+ = (Q_2, \delta_2, q_0^2, F_2)$.

Lemma 4. *For every word $w \in \Sigma^*$, it is the case that $w \in \text{Rec}_2(\mathcal{A})$ iff there exist two infinite sequences $\{\alpha_i\}_{i \in \mathbb{N}}$ and $\{\beta_i\}_{i \in \mathbb{N}}$ of words over Σ such that*

1. $\alpha_0 = w$ and $|\alpha_i| > 0$ for all $i > 0$,
2. $|\alpha_i| = |\beta_i|$ for all $i \in \mathbb{N}$,
3. there exists an infinite run r of \mathcal{A} on $\beta_0\beta_1\dots$ such that, for all $i \in \mathbb{N}$, the automaton \mathcal{A}^q accepts α_{i+1} , where $q = r(|\beta_0\dots\beta_i|)$,
4. there exists an infinite run r' of \mathcal{M} on $(\beta_0 \otimes \beta_0)(\beta_1 \otimes \beta_1)\dots$ such that, for all $i \in \mathbb{N}$, \mathcal{M}^q accepts $\alpha_i \otimes \beta_i\alpha_{i+1}$ where $q = r'(|\beta_0\dots\beta_{i-1}|)$.

One direction of the lemma is easy: if 1)–4) hold, then from the infinite sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ we can form a new sequence $\{s_i\}_{i \geq 0}$ with $s_i := \beta_0\dots\beta_{i-1}\alpha_i$. Condition (3) ensures that $s_i \in L(\mathcal{A})$ for all $i > 0$, and condition (4) implies that $s_i \rightarrow^+ s_{i+1}$ for all $i \geq 0$. This implies that $w \in \text{Rec}_2(\mathcal{A})$ and thus proving sufficiency in Lemma 4.

The idea of the proof of Theorem 1 is that the sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ compactly represent a sequence $\{s_i\}_{i \geq 0}$ witnessing $w \in \text{Rec}_2(\mathcal{A})$. We shall later construct a Büchi automaton that recognizes precisely all ω -words of the form

$$(\alpha_0 \otimes \beta_0) \left[\begin{smallmatrix} \# \\ \# \end{smallmatrix} \right] (\alpha_1 \otimes \beta_1) \left[\begin{smallmatrix} \# \\ \# \end{smallmatrix} \right] (\alpha_2 \otimes \beta_2) \left[\begin{smallmatrix} \# \\ \# \end{smallmatrix} \right] \dots \quad (*)$$

such that the sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ satisfy r.h.s. of Lemma 4. From such an automaton \mathcal{B} it is easy to obtain an automaton recognizing $\alpha_0 = w \in \text{Rec}_2$.

Now we shall prove the other direction in Lemma 4: that the sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ exist under the assumption $w \in \text{Rec}_2(\mathcal{A})$. We will first need to extend the definition of $\text{Rec}_2(\mathcal{N}, \mathcal{T})$ to allow not necessarily transitive transducers \mathcal{T} : $w \in \text{Rec}_2(\mathcal{N}, \mathcal{T})$ iff there exists a sequence $\{s_i\}_{i \geq 0}$ of words such that $s_0 = w$, $s_i \neq s_{i'}$ for all distinct $i, i' \in \mathbb{N}$, $s_i \in L(\mathcal{N})$ for all $i > 0$, and $s_j \otimes s_k \in L(\mathcal{T})$ for all $k > j \geq 0$.

Lemma 5. *Suppose \mathcal{N} and \mathcal{T} are, respectively, an automaton and a transducer over Σ . For every word $w \in \Sigma^*$, if $w \in \text{Rec}_2(\mathcal{N}, \mathcal{T})$, then there exists a word $w'w''$ such that*

1. $|w'| = |w|$ and $|w''| > 0$,
2. $w \otimes w'w'' \in L(\mathcal{T})$,
3. there exist an accepting run r of \mathcal{N} on $w'w''$, and a run r' of \mathcal{T} on $w' \otimes w'$ such that $w'' \in \text{Rec}_2(\mathcal{N}^{q_1}, \mathcal{T}^{q'_1})$, where $q_1 = r(|w|)$ and $q'_1 = r'(|w|)$.

Proof. Suppose that $w \in \text{Rec}_2(\mathcal{N}, \mathcal{T})$. Then, there exists an infinite sequence $\sigma = \{s_i\}_{i \in \mathbb{N}}$ such that $s_0 = w$, $s_i \neq s_{i'}$ for all distinct $i, i' \in \mathbb{N}$, and it is the case that, for all $i > 0$, the word s_i is in $L(\mathcal{N})$ with accepting run η_i , and for all distinct pair of indices $i' > i \geq 0$, we have $s_i \otimes s_{i'} \in L(\mathcal{T})$. As there are only finitely many different words of length $|w|$ but infinitely many different words in σ , we may assume that $|s_i| > |w|$ for all $i \geq 1$; for, otherwise, we may simply omit these words from σ . Now every word s_i , where $i > 0$, can be written as $s_i = u_i v_i$ for some words u_i, v_i such that $|u_i| = |w|$ and $|v_i| > 0$. As there are only finitely many different words of length $|w|$ and finitely many different runs of \mathcal{N} of length $|w|$, by pigeonhole principle there must exist $k > 0$ such that $u_j = u_k$ and $\eta_j[0, |w|] = \eta_k[0, |w|]$ for infinitely many $j > 0$. Let $w' := u_k$ and $\eta := \eta_k[0, |w|]$. Therefore, we may discard all words s_i in σ with $i \geq 1$ such that $u_i \neq w'$ or η is not a prefix of η_i . By renaming indices, call the resulting sequence $\sigma = \{s_i\}_{i \in \mathbb{N}}$ and, for all $i \geq 1$, denote by η_i the accepting run of \mathcal{N} on s_i that has η as a prefix. Notice that σ is still a witness for $w \in \text{Rec}_2(\mathcal{N}, \mathcal{T})$. So, let $\theta_{j,k}$, where $0 \leq j < k$, be the accepting run of \mathcal{T} on $s_j \otimes s_k$. Let \mathcal{C} be the finite set of all runs of \mathcal{T} on $w' \otimes w'$. Notice that it is not necessarily the case that $|\mathcal{C}| = 1$ as \mathcal{T} is nondeterministic. Consider the edge-labeled undirected graph $G = (V, \{E_u\}_{u \in \mathcal{C}})$ such that $V = \mathbb{Z}^+$ and

$$E_u = \{\{j, k\} : 0 < j < k \text{ and } u \text{ is a prefix of } \theta_{j,k}\}.$$

Notice that $\{E_u\}_{u \in \mathcal{C}}$ is a partition of $\{\{j, k\} : j \neq k, j, k > 0\}$, and so G is a complete graph. By (infinite) Ramsey theorem, G has a monochromatic complete infinite subgraph $H = (V', E_u)$ for some $u \in \mathcal{C}$. Set $r' := u$. Notice that if V' contains the elements $i_1 < i_2 < \dots$, then θ_{i_j, i_k} with $k > j \geq 1$ has u as a prefix. Therefore, we can discard all words s_i ($i > 0$) in σ such that $i \notin V'$ and by renaming indices call the resulting sequence $\sigma = \{s_i\}_{i \in \mathbb{N}}$. We also adjust the sequence $\{\eta_i\}_{i > 0}$ of accepting runs by omitting the appropriate runs and adjusting indices. We now set w'' to be the unique word v such that $s_1 = w'v$. It is easy to see that (1) and (2) are satisfied. Setting $r = \eta_1$, it is easy to check that $w'' \in \text{Rec}_2(\mathcal{N}^{q_1}, \mathcal{T}^{q_1})$ with witnessing sequence $\{t_i\}_{i > 0}$, where t_i is the unique word such that $s_i = w't_i$ for all $i > 0$. \square

Now it is not difficult to inductively construct the desired sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ by using lemma 5 at every induction step. The gist of the proof is that from the word $w'w''$ given by lemma 5 at induction step k , we will set $\beta_k = w'$, $\alpha_{k+1} = w''$, and extend the partial runs r and r' in lemma 4. Notice that we now have $w'' \in \text{Rec}_2(\mathcal{N}^{q_1}, \mathcal{T}^{q_1})$, which sets up the next induction step. See appendix for a detailed argument. This completes the proof of lemma 4.

Now we construct a Büchi automaton \mathcal{B} accepting ω -words of the form (*), where α_i 's and β_i 's are given by Lemma 4. We first give an informal description of how to implement \mathcal{B} . The automaton \mathcal{B} will attempt to guess the runs r and r' , while at the same time checking that the runs satisfy conditions 3–4 in Lemma 4. To achieve this, \mathcal{B} will run a copy of \mathcal{A} and \mathcal{M} , while simultaneously also running a few other copies of \mathcal{A} and \mathcal{M} to check that the runs r and r' guessed so far satisfy conditions 3) and 4) along the way. The automaton \mathcal{B}

consists of three components depicted as Boxes 1, 2, and 3 in Figure 1. The first box is used for reading the prefix of the input before the first occurrence of $\begin{bmatrix} \# \\ \# \end{bmatrix}$, while the other boxes are used for reading the remaining suffix. Boxes 2-3 are essentially identical, i.e., they have the same sets of states and essentially the same transition functions. When \mathcal{B} arrives in Box 2, it will read a single letter in Σ^2 and goes to Box 3 so as to make sure that $|\alpha_i| > 0$ for each $i > 0$. When \mathcal{B} is in Box 3, it will go to Box 2 upon reading the letter $\begin{bmatrix} \# \\ \# \end{bmatrix}$. We will set all states in Box 2 as the final states so as to make sure that infinitely many $\begin{bmatrix} \# \\ \# \end{bmatrix}$ is seen, i.e., the sequences $\{\alpha_i\}_i$ and $\{\beta_i\}_i$ are both infinite, and each words α_i and β_i are finite.




Fig. 1. A bird's eye view of the Büchi automaton \mathcal{B}

More formally, the automaton $\mathcal{B} = (\Sigma^2 \cup \{\begin{bmatrix} \# \\ \# \end{bmatrix}\}, Q, \delta, q_0, F)$ is defined as follows. We set $Q := (Q_1 \times Q_2 \times Q_2) \uplus (Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2 \times \{1, 2\})$, where $Q_1 \times Q_2 \times Q_2$ are the states in box #1, and $Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2 \times \{i\}$ are states in box $\#(i+1)$. The initial state is $q_0 := (q_0^1, q_0^2, q_0^2)$. The first and the last components in each state are meant for guessing the infinite runs r and r' . The second component of each state in box #1 is used for guessing a prefix of the accepting run of \mathcal{M} on $\alpha_0 \otimes \beta_0 \alpha_1$. The automaton \mathcal{B} will finish this guessing when it reaches box #3 upon the completion of parsing $\alpha_1 \otimes \beta_1$. When \mathcal{B} is presently in box #2 or #3 and reading $\alpha_i \otimes \beta_i$, where $i > 0$, the third and fourth components of the states are used for checking that $\beta_0 \dots \beta_{i-1} \alpha_i \in L(\mathcal{A})$ and $\beta_0 \dots \beta_{i-2} \alpha_{i-1} \otimes \beta_0 \dots \beta_{i-1} \alpha_i \in L(\mathcal{M})$, respectively. At the same time, the second component will be checking that $\beta_0 \dots \beta_{i-1} \alpha_i \otimes \beta_0 \dots \beta_i \alpha_{i+1} \in L(\mathcal{M})$, which will be completed in the next iteration. We now formally define the transition function. We set

$$\delta((q, q', q''), \begin{bmatrix} a \\ b \end{bmatrix}) := \begin{cases} \delta_1(q, b) \times \delta_2(q', \begin{bmatrix} a \\ b \end{bmatrix}) \times \delta_2(q'', \begin{bmatrix} b \\ b \end{bmatrix}) & , \text{ if } a, b \neq \# \\ (q, q'', q, q', q'', 1) & , \text{ if } a = b = \# \\ \emptyset & , \text{ otherwise.} \end{cases}$$

and, when \mathcal{B} is in a state in $Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2 \times \{i\}$, where $i = 1, 2$, we define

$$\begin{aligned}
 \delta((q_1, q_2, q'_1, q'_2, q''_2, i), \begin{bmatrix} a \\ b \end{bmatrix}) &:= \delta_1(q_1, b) \times \delta_2(q_2, \begin{bmatrix} a \\ b \end{bmatrix}) \times \delta_1(q'_1, a) \times \\
 &\delta_2(q'_2, \begin{bmatrix} \perp \\ a \end{bmatrix}) \times \delta_2(q''_2, \begin{bmatrix} b \\ b \end{bmatrix}) \times \{2\}
 \end{aligned}$$

if $a, b \neq \#$. If $q'_1 \in F_1$, and $q'_2 \in F_2$, then we set

$$\delta((q_1, q_2, q'_1, q'_2, q''_2, 2), \begin{bmatrix} \# \\ \# \end{bmatrix}) = (q_1, q''_2, q_1, q_2, q''_2, 1).$$

Finally, the set of final states are $F := Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2 \times \{1\}$. It is easy to see that \mathcal{B} , as claimed, recognizes exactly ω -words of the word of the form $(*)$ such that the sequences $\{\alpha_i\}_{i \in \mathbb{N}}$ and $\{\beta_i\}_{i \in \mathbb{N}}$ satisfy the conditions in Lemma 4.

Now, from \mathcal{B} we can easily compute the automaton $\mathcal{A}_2 = (Q', \Sigma, \delta', q'_0, F')$ that recognizes Rec_2 . Roughly speaking, the automaton \mathcal{A}_2 will accept the set of finite words α_0 such that there exist two sequences $\{\alpha_i\}_{i > 0}$ and $\{\beta_i\}_{i \geq 0}$ such that the ω -word $(*)$ is accepted by \mathcal{B} . Therefore, we will set the new set of states Q' to be $Q_1 \times Q_2 \times Q_2$, i.e., the first component of \mathcal{B} in Fig. 1. We apply projection operation on the transition function δ of \mathcal{B} to obtain δ' . More formally, if $a \in \Sigma$, we set

$$\delta'((q_1, q_2, q'_2), a) = \bigvee_{b \in \Sigma} \delta((q_1, q_2, q'_2), \begin{bmatrix} a \\ b \end{bmatrix}).$$

Finally, the new set F' of final states will those states in Q' from which \mathcal{B} can accept some ω -words of the form $\begin{bmatrix} \# \\ \# \end{bmatrix} w$ for some ω -word w . For this, we can apply the standard algorithm for testing nonemptiness for Büchi automata, which takes linear time. Theorem 1 is now immediate. \square

Application: Pushdown systems We shall use the definition of [10, 11, 22], which subsumes a more common definition of [15, 14, 27] based on configurations of pushdown automata and transitions between them. A *pushdown system* over the alphabet Σ is given by a finite set Δ of rules of the form $u \rightarrow v$ where $u, v \in \Sigma^*$. Let $\text{Dom}(\Delta)$ denote the set of words u for which there is a rule $u \rightarrow v$ in Δ . Then Δ generates a relation \rightarrow_Δ over Σ^* as follows: $(w, w') \in \mathcal{R}_\Delta$ iff there exist $x, u, v \in \Sigma^*$ such that $w = xu$, $w' = xv$, and $u \rightarrow v$ is a rule in Δ . We thus compute recurrent reachability over pushdown systems $\langle \Sigma^*, \rightarrow_\Delta, \lambda \rangle$.

The binary relation \rightarrow_Δ is regular, and can be given by a transducer \mathcal{R}_Δ whose size is linear in $\|\Delta\|$ (where $\|\Delta\|$ is the sum of the lengths of each word in Δ). Caucal [10] proved that, for each pushdown system Δ , the relation \rightarrow_Δ^* is a poly-time-computable rational transduction¹. Later in [11] he noted that the given transducer is also regular. For completeness, we sketch how his construction gives a regular transducer \mathcal{R}_Δ^* for \rightarrow_Δ^* in poly-time. Recall the following well-known proposition, which is proven using the standard “saturation” construction (e.g. see [8, 14, 10]).

Proposition 6. *Given a pushdown system Δ and a nondeterministic automaton \mathcal{A} , one can compute two automata \mathcal{A}_{pre^*} and \mathcal{A}_{post^*} for $pre^*(L(\mathcal{A}))$ and $post^*(L(\mathcal{A}))$ in poly-time.*

In fact, the algorithm given in [14] computes the automata in cubic time, and the sizes of \mathcal{A}_{pre^*} and \mathcal{A}_{post^*} are at most quadratic in $|\mathcal{A}|$. To construct \mathcal{R}_Δ^* using this proposition, we shall need the following easy lemma.

¹ Rational transducers are strictly more powerful than regular transducers.

Lemma 7 ([10]). *Given a pushdown system Δ and two words $u, v \in \Sigma^*$, then $u \xrightarrow{\Delta}^* v$ iff there exist words $x, y, z \in \Sigma^*$ and word $w \in \text{Dom}(\Delta) \cup \{\varepsilon\}$ such that $u = xy$, $v = xz$, $y \xrightarrow{\Delta}^* w$, and $w \xrightarrow{\Delta}^* z$.*

Now constructing \mathcal{R}_{Δ}^* is easy. First, we use Proposition 6 to compute the automata $\mathcal{A}_{pre^*}^w$ and $\mathcal{A}_{post^*}^w$ that recognize $pre^*(w)$ and $post^*(w)$ for every $w \in \text{Dom}(\Delta) \cup \{\varepsilon\}$. Then, on input $u \otimes v$, the transducer guesses a word $w \in \text{Dom}(\Delta) \cup \{\varepsilon\}$ and a position at which the prefix x in Lemma 7 ends, and then simultaneously runs the automata $\mathcal{A}_{pre^*}^w$ and $\mathcal{A}_{post^*}^w$ to verify that the top part y and the bottom part z of the remaining input word (preceding the \perp symbol) satisfy $y \in L(\mathcal{A}_{pre^*}^w)$ and $z \in L(\mathcal{A}_{post^*}^w)$. We thus obtain a transducer \mathcal{R}^* of size $O(\|\Delta\|^2)$. By taking a product, we compute a transducer \mathcal{R}^+ of size $O(\|\Delta\|^3)$ in poly-time. Therefore, Theorem 1 implies the following.

Theorem 8. *Both global and local model-checking for recurrent reachability over pushdown systems are solvable in poly-time.*

That is, for a pushdown system Δ and a nondeterministic automaton \mathcal{A} over an alphabet Σ , one can compute in polynomial time an NWA recognizing $\text{Rec}(L(\mathcal{A}), \xrightarrow{\Delta}^+)$.

4 Recurrent reachability: the tree case

Recall that in a tree-automatic transition system $\mathcal{S} = (S, \rightarrow, \lambda)$, the relation \rightarrow and the sets $\lambda(P_i)$'s are given as tree automata. As in the word case, such a transition system is said to be *transitive* if the relation \rightarrow^+ is regular. We now extend our results from Section 3 to transitive tree-automatic transition systems.

Theorem 9. *Given an NTA \mathcal{A} and a transitive tree-automatic transition system $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$, the set $\text{Rec}(L(\mathcal{A}))$ of states $T \in S$ such that $T \xrightarrow{\omega} L(\mathcal{A})$ is regular. Moreover, if the transducer \mathcal{R}^+ for \rightarrow^+ is computable in time $t(|\mathcal{R}|)$, then one can compute an NTA recognizing $\text{Rec}(L(\mathcal{A}))$ of size $O(|\mathcal{R}^+|^2 \times |\mathcal{A}|)$ in time $t(|\mathcal{R}|) + O(|\mathcal{R}^+|^6 \times |\mathcal{A}|^4)$.*

As in the word case, this implies two corollaries:

Corollary 10. *If \mathcal{S} is transitive and tree-automatic and \mathcal{R}^+ is poly-time computable, then both global and local model-checking for recurrent reachability are solvable in poly-time.*

Corollary 11. *If $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ is a transitive tree-automatic transition system such that the transducer \mathcal{R}^+ is computable, then for each formula φ of (EF+EGF)-logic, the set $\llbracket \varphi \rrbracket_{\mathcal{S}}$ is regular, and an NTA defining $\llbracket \varphi \rrbracket_{\mathcal{S}}$ can be effectively constructed.*

The proof follows the same basic steps as the proof of Theorem 1: we first show that it is sufficient to consider only infinite witnessing sequences that have a representation as an infinite tree over a finite labeling alphabet; we then construct a tree automaton (over infinite trees) with a Büchi acceptance condition that

recognizes such sequences; and from such an automaton we construct a NTA for $\text{Rec}(L(\mathcal{A}))$ by applying projection and checking nonemptiness for Büchi tree automata. As checking nonemptiness for Büchi tree automata is quadratic [26] instead of linear as in the word case, the degree of the polynomials in Theorem 9 doubles. While all steps are similar to those in the word case, there are many technical differences; in particular in the coding of an infinite sequence by a single infinite tree. See appendix for details of the proof.

Application: Ground tree rewrite systems Ground tree rewrite systems have been intensely studied in the rewriting, automata, and verification communities [13, 12, 18, 19]. We now show that a result by Löding [19] on poly-time model-checking for recurrent reachability and decidability of model checking (EF+EGF)-logic over such systems, which was proved with a specialized method for RGTRSs, can be obtained as a corollary of Theorem 9.

A *ground tree rewrite system* (GTRS) over Σ -labeled \mathcal{Y} -trees is a finite set Δ of transformation rules of the form $t \rightarrow t'$ where $t, t' \in \text{TREE}_{\mathcal{Y}}(\Sigma)$. If we permit rules of the form $L \rightarrow L'$, where L and L' are tree languages given by some NTAs, then we call Δ a *regular ground tree rewrite system* (RGTRS). Obviously, RGTRSs generalize GTRSs. We define $\|\Delta\|$ as the sum of the sizes of automata in Δ . The RGTRS Δ also generates a binary relation \rightarrow_{Δ} over $\text{TREE}_{\mathcal{Y}}(\Sigma)$: For a tree T and a node u in it, let T_u be the subtree of T rooted at u . Given two trees T and T' , we let $T \rightarrow_{\Delta} T'$ iff there exists a node u in T and a rule $L \rightarrow L'$ in Δ such that $T_u \in L$ and $T' = T[t'/u]$ for some $t' \in L'$, where $T[t'/u]$ is the tree obtained from T by replacing the node u by the tree t' .

Given Δ , it is easy to compute a tree transducer \mathcal{R}_{Δ} for \rightarrow_{Δ} in time $O(\|\Delta\|)$; it guesses a node u in the input tree $T \otimes T'$ and a rule in Δ to apply at u in T to obtain T' . The following has been proven in [13] and [12, chapter 3].

Proposition 12. *Given a RGTRS Δ , the transitive closure relation \rightarrow_{Δ}^+ is regular, and a transducer defining it can be computed in time polynomial in $|\mathcal{R}_{\Delta}|$.*

In fact, the proof for the above proposition constructs “ground tree transducers”, which are a subclass of the notion of transducers we are considering in this paper (e.g. see [12, chapter 3]).

Combining this proposition with corollaries 10 and 11, we obtain:

Corollary 13. *(Löding [19]) Both global and local model checking for recurrent reachability over RGTRSs are solvable in poly-time. Model checking (EF+EGF)-logic over RGTRSs with regular atomic predicates is decidable.*

Application: PA-processes PA [4, 22] is a well-known process algebra allowing sequential and parallel compositions, but no communication. It generalizes basic parallel processes (BPP), and context-free processes (BPA), but is incomparable to pushdown processes and Petri nets (e.g. see [22]). PA has found applications in the interprocedural dataflow analysis of parallel programs [16].

We review the basic definitions, following the presentation of [20]: we initially distinguish terms that are equivalent up to simplification laws. The definition of PA usually includes transition labels, which we omit to simplify our presentation (however, the results easily hold when we incorporate transition labels). Fix a finite set $Var = \{X, Y, Z, \dots\}$ of process variables. *Process terms* over Var , denoted by \mathcal{F}_{Var} , are generated by the grammar:

$$t, t' := 0 \mid X, \quad X \in Var \mid t.t' \mid t \parallel t'$$

where 0 denotes a “nil” process, and $t.t'$ and $t \parallel t'$ are sequential and parallel compositions, resp. Process terms can be viewed as Σ -labeled binary trees, where $\Sigma = Var \cup \{0, \parallel, \cdot\}$. In particular, inner nodes are always labeled by ‘.’ or ‘ \parallel ’, while leaves are labeled by elements in $Var \cup \{0\}$. A PA *declaration* over \mathcal{F}_{Var} is a finite set Δ of rewrite rules of the form $X \rightarrow t$, where $X \in Var$, and $t \in \mathcal{F}_{Var}$. We set $\text{Dom}(\Delta) = \{X : (X \rightarrow t) \in \Delta, \text{ for some } t \in \mathcal{F}_{Var}\}$, and $Var_\emptyset = Var - \text{Dom}(\Delta)$. The set Δ generates a transition relation \rightarrow_Δ on process terms defined by:

$\frac{t_1 \rightarrow t'_1}{t_1 \parallel t_2 \rightarrow t'_1 \parallel t_2}$	$\frac{t_1 \rightarrow t'_1}{t_1.t_2 \rightarrow t'_1.t_2}$	$\frac{}{X \rightarrow t} \quad (X \rightarrow t) \in \Delta$
$\frac{t_2 \rightarrow t'_2}{t_1 \parallel t_2 \rightarrow t_1 \parallel t'_2}$	$\frac{t_2 \rightarrow t'_2}{t_1.t_2 \rightarrow t_1.t'_2}$	$t_1 \in \text{IsNil}$

Here IsNil is the set of “terminated” process terms, i.e., those in which all variables are in Var_\emptyset . It is easy to see that there is a regular transducer \mathcal{R}_Δ over process terms for \rightarrow_Δ , whose size is linear in the size $\|\Delta\|$ of Δ . It is defined in the same way as for GTRSs, except that when it guesses a leaf node at which a rule is applied, it must further ensure that v has no ‘.’-labeled ancestor u such that v is a descendant $u1$ and that T_{u0} is not a terminated process term.

Theorem 14 ([20, 21, 16]). *Given a PA declaration Δ and a NTA \mathcal{A} describing a set of process terms over Var , the sets $\text{pre}^*(L(\mathcal{A}))$ and $\text{post}^*(L(\mathcal{A}))$ are regular, for which NTAs can be computed in time $O(\|\Delta\| \times |\mathcal{A}|)$, and one can construct a regular transducer \mathcal{R}^+ for \rightarrow^+ in poly-time².*

We consider only languages and atomic propositions that are interpreted as regular subsets of \mathcal{F}_{Var} . This poses no problem as \mathcal{F}_{Var} is easily seen a regular subset of $\text{TREE}_2(\Sigma)$ and no tree $t \in \mathcal{F}_{Var}$ is related by \rightarrow to a tree $t' \in \text{TREE}_2(\Sigma) - \mathcal{F}_{Var}$. From Theorem 14 and Corollaries 10 and 11, we obtain:

Theorem 15. *Both global and local model checking for recurrent reachability over PA are solvable in poly-time. Model checking (EF+EGF)-logic over PA is decidable.*

In the study of PA processes, it is common to use a structural equivalence on process terms. We now extend our results to PA modulo structural equivalence.

² Lugiez and Schnoebelen first proved this in [20] for a more general notion of transducers, but later in [21] realized that regular transducers suffice.

Let \equiv be the smallest equivalence relation on \mathcal{F}_{Var} that satisfies the following:

$$\begin{array}{cccc} t.0 \equiv t & 0.t \equiv t & t\|0 \equiv t & t\|t' \equiv t'\|t \\ (t\|t')\|t'' \equiv t\|(t'\|t'') & & (t.t').t'' \equiv t.(t'.t'') & \end{array}$$

We let $[t]_{\equiv}$ stand for the equivalence class of t and $[L]_{\equiv}$ for $\bigcup_{t \in L} [t]_{\equiv}$. We write L/\equiv for $\{[t]_{\equiv} \mid t \in L\}$. It was shown in [20] that, for each $t \in \mathcal{F}_{Var}$, $[t]_{\equiv}$ is a regular tree language, although the set $[L]_{\equiv}$ need not be regular even for regular L . Given a PA declaration Δ , the equivalence \equiv generates a transition relation $[t]_{\equiv} \Rightarrow [u]_{\equiv}$ over \mathcal{F}_{Var}/\equiv which holds iff there exist $t' \in [t]_{\equiv}$ and $u' \in [u]_{\equiv}$ such that $t' \rightarrow u'$. We need the following result:

Lemma 16 ([20]). *The relation \equiv is bisimulation: for all $t, t', u \in \mathcal{F}_{Var}$, if $t \equiv t'$ and $t \rightarrow u$, then there exists $u' \in \mathcal{F}_{Var}$ such that $t' \rightarrow u'$ and $u \equiv u'$.*

Now it is not hard to show that, for every NTA \mathcal{A} , the set $Rec(L(\mathcal{A}))$ is closed under \equiv , if $L(\mathcal{A})$ is closed under \equiv . This also implies that $Rec(L(\mathcal{A})) = [Rec(L(\mathcal{A}))]_{\equiv} = \{t : t \in Rec(L(\mathcal{A}))/\equiv, \Rightarrow^+\}$. In the following, we consider only languages that are closed under \equiv .

Theorem 17. *Given an NTA \mathcal{A} such that $L(\mathcal{A})$ is closed under \equiv and a process term $t \in \mathcal{F}_{Var}$, it is possible to decide whether $[t]_{\equiv} \Rightarrow^{\omega} L(\mathcal{A})/\equiv$ in PTIME.*

Since $Rec(L(\mathcal{A})) = [Rec(L(\mathcal{A}))]_{\equiv}$, we need only compute an NTA for $Rec(L(\mathcal{A}))$ and test whether $t \in Rec(L(\mathcal{A}))$. These can be done in PTIME by theorem 15.

We now move to model checking (EF+EGF)-logic over PA modulo \equiv . Suppose $\mathcal{S} = \langle S, \rightarrow, \lambda \rangle$ is a transition system generated by some PA-declaration and that each $\lambda(P)$ is closed under \equiv . In fact, the standard atomic propositions for PA-processes include sets of process terms of the form $[t]_{\equiv}$ and *action-based predicates*, i.e., sets of all terms t in which some transitions in Δ can be applied (and these are obviously closed under \equiv and of size $O(\|\Delta\|)$). Now Lemma 16 implies that $\llbracket \varphi \rrbracket_{\mathcal{S}}$ is closed under \equiv for (EF+EGF)-formulae φ , and we obtain:

Theorem 18. *The problem of model checking for (EF+EGF)-logic over PA modulo \equiv is decidable whenever all atomic propositions are closed under \equiv .*

5 Future work

We mention some possible future work. We would like to further study algorithmic improvements of our general technique, e.g., in its current form it gives a polynomial of degree higher than the specialized technique of [19] for RGTRSs. We would also like to investigate stronger but nonrestrictive conditions that ensure decidability of stronger logics (e.g. CTL*) within our framework; it is easy to show that our current condition is insufficient. Finally, we would like to study when our technique could generate elementary complexity algorithms for (EF+EGF)-logic, or just EF-logic alone. This problem is still open even for PA-processes and GTRSs [19, 23].

Acknowledgements We thank Richard Mayr and anonymous referees for their helpful comments. The authors were supported by EPSRC grant E005039, the second author also by EC grant MEXC-CT-2005-024502.

References

1. P. A. Abdulla, and B. Jonsson. Undecidable verification problems for programs with unreliable channels. In *Inf. Comput.*, 130(1):71–90, 1996.
2. P. A. Abdulla, B. Jonsson, P. Mahata, and J. d’Orso. Regular tree model checking. In *CAV ’02*, pages 555–568.
3. P. A. Abdulla, B. Jonsson, M. Nilsson, and M. Saksena. A survey of regular model checking. In *CONCUR’04*, pages 35–48.
4. J. Baeten and W. Weijland. *Process Algebra*. CUP, 1990.
5. M. Benedikt, L. Libkin, and F. Neven. Logical definability and query languages over ranked and unranked trees. *ACM Trans. Comput. Logic*, 8(2):11, 2007.
6. A. Blumensath and E. Grädel. Automatic structures. In *LICS ’00*, pages 51–60.
7. B. Boigelot, A. Legay, and P. Wolper. Iterating transducers in the large. In *CAV’03*, pages 223–235.
8. A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR ’97*, pages 135–150.
9. A. Bouajjani, A. Legay, and P. Wolper. Handling liveness properties in (ω -)regular model checking. In *INFINITY*, pages 101–115, 2004.
10. D. Caucal. On the regular structure of prefix rewriting. In *CAAP’90*, pages 61–86.
11. D. Caucal. On the regular structure of prefix rewriting. *Theor. Comput. Sci.*, 106(1):61–86, 1992.
12. H. Comon et al. *Tree Automata: Techniques and Applications*. Available on: <http://www.grappa.univ-lille3.fr/tata>, 2007.
13. M. Dauchet and S. Tison. The theory of ground rewrite systems is decidable. In *LICS’90*, pages 242–248.
14. J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient algorithms for model checking pushdown systems. In *CAV ’00*, pages 232–247.
15. J. Esparza, A. Kucera, and S. Schwoon. Model checking LTL with regular valuations for pushdown systems. *Inf. Comput.*, 186(2):355–376, 2003.
16. J. Esparza and A. Podelski. Efficient algorithms for pre* and post* on interprocedural parallel flow graphs. In *POPL ’00*, pages 1–11. USA, 2000. ACM.
17. B. Jonsson and M. Nilsson. Transitive closures of regular relations for verifying infinite-state systems. In *TACAS’00*, pages 220–234.
18. C. Löding. Model-checking infinite systems generated by ground tree rewriting. In *FoSSaCS’02*, pages 280–294.
19. C. Löding. Reachability problems on regular ground tree rewriting graphs. *Theory Comput. Syst.*, 39(2):347–383, 2006.
20. D. Lugiez and P. Schnoebelen. The regular viewpoint on PA-processes. *Theor. Comput. Sci.*, 274(1-2):89–115, 2002.
21. D. Lugiez and P. Schnoebelen. Decidable first-order transition logics for PA-processes. *Inf. Comput.*, 203(1):75–113, 2005.
22. R. Mayr. Process rewrite systems. *Inf. Comput.*, 156(1-2):264–286, 2000.
23. R. Mayr. Decidability of model checking with the temporal logic EF. *Theor. Comp. Sci.*, 256(1-2):31–62, 2001.
24. W. Thomas. Constructing infinite graphs with a decidable MSO-theory. In *MFCS*, pages 113–124, 2003.
25. M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Proc. Banff Higher-Order Workshop*, pages 238–266.
26. M. Y. Vardi, P. Wolper. Automata-theoretic techniques for modal logics of programs. *JCSS* 32(2): 183–221 (1986).
27. I. Walukiewicz. Model checking CTL properties of pushdown systems. In *FSTTCS’00*, pages 127–138.

APPENDIX

A Lemma 5 implies necessity in Lemma 4

We now complete the proof of necessity in lemma 4 by inductively constructing the desired sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ by using lemma 5 at every induction step. In the following, a sequence $\{\eta_i\}_{i \geq 0}$ of run segments of \mathcal{A} is said to be *good* if $\eta_0(0) = q_0^1$ and $\mathbf{last}(\eta_i) = \mathbf{first}(\eta_{i+1})$ for $i \in \mathbb{N}$. In other words, the sequence of run segments is good if it starts with the initial state of \mathcal{A} and two consecutive run segments can be concatenated. The same notion can similarly be defined for sequences of run segments of \mathcal{M} .

Claim. There exist two sequences $\{\alpha_i\}_{i \geq 0}$ and $\{\beta_i\}_{i \geq 0}$ of words, a good sequence $\{\eta_i\}_{i \geq 0}$ of run segments of \mathcal{A} , and a good sequence $\{\theta_i\}_{i \geq 0}$ of run segments of \mathcal{M} such that $\alpha_0 = w$, $\eta_0 = q_0^1$, $\theta_0 = q_0^2$, and for all $k \in \mathbb{N}$:

1. for all $0 < i \leq k$, $|\alpha_i| > 0$,
2. for all $0 \leq i < k$, $|\beta_i| = |\alpha_i|$,
3. for all $0 \leq i \leq k$, $r_i := \eta_0 \odot \dots \odot \eta_i$ is a run of \mathcal{A} on $\beta_0 \dots \beta_{i-1}$,
4. for all $0 \leq i \leq k$, $r'_i := \theta_0 \odot \dots \odot \theta_i$ is a run of \mathcal{M} on $(\beta_0 \otimes \beta_0) \dots (\beta_{i-1} \otimes \beta_{i-1})$,
5. for all $0 < i \leq k$, \mathcal{A}^{q_i} accepts α_i , where $q_i = \mathbf{last}(r_i)$,
6. for all $0 \leq i < k$, $\mathcal{M}^{q'_i}$ accepts $\alpha_i \otimes \beta_i \alpha_{i+1}$, where $q'_i = \mathbf{last}(r'_i)$,
7. for all $0 \leq i \leq k$, $\alpha_i \in \mathit{Rec}_2(\mathcal{A}^{q_i}, \mathcal{M}^{q'_i})$, where $q_i = \mathbf{last}(r_i)$ and $q'_i = \mathbf{last}(r'_i)$.

Observe that this claim immediately implies Lemma 4 as we may simply define $r = \eta_0 \odot \eta_1 \odot \dots$ and $r' = \theta_0 \odot \theta_1 \odot \dots$. To prove this claim, we shall define these four sequences inductively. For each $k \in \mathbb{N}$, we shall define four partial sequences $\{\alpha_i\}_{0 \leq i \leq k}$, $\{\beta_i\}_{0 \leq i < k}$, $\{\eta_i\}_{0 \leq i \leq k}$, and $\{\theta_i\}_{0 \leq i \leq k}$ satisfying the conditions in the claim. We shall first deal with the base case $k = 0$. We define $\alpha_0 = w$, $\eta_0 = q_0^1$, and $\theta_0 = q_0^2$. It is easy to see that statements (1),(2),(5), and (6) are vacuous. Statements (3)–(4) are also true because q_0^1 (resp. q_0^2) is a run of \mathcal{A} (resp. \mathcal{M}) on ε . Statement (7) is true by assumption that $w \in \mathit{Rec}_2(\mathcal{A}, \mathcal{M})$. Assume now that $k > 0$ and the four partial sequences have been defined satisfying the seven conditions in the claim for all natural numbers up to k . We shall now extend these partial sequences by defining α_{k+1} , β_k , η_{k+1} , and θ_k . By induction, we have $\alpha_k \in \mathit{Rec}_2(\mathcal{A}^{q_k}, \mathcal{M}^{q'_k})$ and so Lemma 5 gives us a word $w'w''$. We may set $\beta_k = w'$ and $\alpha_{k+1} = w''$. It is immediate that condition (1) and (2) are satisfied. We define η_{k+1} to be the prefix of length $|w'|$ of the run r of \mathcal{A}^{q_k} on $w'w''$ given by Lemma 5. We define θ_k to be the run r' of $\mathcal{M}^{q'_k}$ of length $|w'|$ given by Lemma 5. It is easy to see now that condition (3)–(5) hold whenever $i = k + 1$ and condition (6) hold whenever $i = k$. Lemma 5 also implies that $\alpha_{k+1} \in \mathit{Rec}_2(\mathcal{A}^{q_{k+1}}, \mathcal{M}^{q'_{k+1}})$, where $q_{k+1} := \mathbf{last}(r_{k+1})$ and $q'_{k+1} := \mathbf{last}(r'_{k+1})$.

Finally, conditions (3)–(7) hold for other smaller values of i by induction. This completes our proof for the claim and therefore the proof of Lemma 4.

B Proofs from Section 4

We will first fix some definitions and notations that we will use in the proof of theorem 9. For a tree $T = (D, \tau)$ and a nonempty $D' \subseteq D$, we define $T|_{D'}$ to be the tree T whose domain is restricted to D' . We write $T' \preceq T$ iff T' is the “prefix” of T , i.e. the domain of T' is contained in the domain of T , and the labeling of T and T' coincide on the domain of T' .

Next we need some definitions about contexts. A *context tree* with variables x_1, \dots, x_n is a $(\Sigma \cup \{x_1, \dots, x_n\})$ -labeled \mathcal{Y} -tree $T' = (D', \tau')$ such that for each $i = 1, \dots, n$, there is exactly one node $u_i \in D'$ with $\tau'(u_i) = x_i$; furthermore, u_i is a leaf node in T' . The leaf nodes u_1, \dots, u_n are also called *context leaves*. To emphasize which variables are in T' , we write $T'[x_1, \dots, x_n]$ for T' . Whenever $n = 0$, T' is just a normal tree (a.k.a. *ground tree*). Given ground trees t_1, \dots, t_n , the tree $T'[t_1/x_1, \dots, t_n/x_n]$ is the ground tree obtained by replacing all context leaves u_1, \dots, u_n by the ground trees t_1, \dots, t_n , respectively. We also define $T' \otimes T''$ just as we defined the operator \otimes for ground trees, but we replace the label $\begin{bmatrix} x_i \\ x_i \end{bmatrix}$ by x_i . Given an NTA \mathcal{A}' over Σ -labeled \mathcal{Y} -trees, we want to extend the notion of runs of $\mathcal{A}' = (\Sigma', Q', \delta', q'_0, F')$ to context trees. We first define $\mathbf{virt}(T')$ to be the $(\Sigma_{\perp} \cup \{x_1, \dots, x_n\})$ -labeled tree $T'' = (D'', \tau'')$ such that $D'' := D' \cup \{vj : v \in D' - \{u_1, \dots, u_n\}, j \in \mathcal{Y}\}$ and, if $u \in D'$, we set $\tau''(u) := \tau'(u)$; if $u \notin D'$, then $\tau''(u) := \perp$. A *run* of \mathcal{A}' on T' is a mapping $\rho : D'' \rightarrow Q'$ that is defined in the same way as for ground trees. The run r is said to be *potentially accepting* if $\rho(u) \in F$ for each leaf $u \in D'' - \{u_1, \dots, u_n\}$. In other words, potentially accepting runs *might* become accepting after we replace the context leaves with some ground trees. For a context tree $T' = (D', \tau')$ and a tree $T = (D, \tau)$, we write $T' \preceq T$ if $D' \subseteq D$ and $\tau'(u) = \tau(u)$ whenever $u \in D'$ and u is not a context leaf. A tree $T = (D, \tau)$ is said to be *unranked* if its direction alphabet is \mathbb{Z}^+ , and for each $u \in D$ there is $n \in \mathbb{N}$ such that $uj \in D$ for all $1 \leq j \leq n$ and $u(j+1) \notin D$. Notice that unranked trees are finitely branching.

We now prove theorem 9. To simplify notations, we shall make an assumption that $\mathcal{Y} = \{1, 2\}$. This is not necessary for the proof and it is easy to adapt the proof to the case when \mathcal{Y} is any finite nonempty direction alphabet. For the rest of this section, we write $\mathcal{A} = (Q_1, \delta_1, q_0^1, q_F^1)$ and $\mathcal{M} = \mathcal{R}^+ = (Q_2, \delta_2, q_0^2, q_F^2)$. As our automata are nondeterministic, we may assume without loss of generality that there are no transitions from q_F^1 and q_F^2 in the automata \mathcal{A} and \mathcal{M} , respectively. By definition, for every tree T , we have $T \in \mathit{Rec}(\mathcal{A})$ iff there exists an infinite sequence $\{T_i\}_{i \in \mathbb{N}}$ of such that $T_0 = T$, $T_{i-1} \rightarrow^+ T_i$ and $T_i \in L(\mathcal{A})$ for all $i > 0$. As in the case of words, we shall prove that it is sufficient to consider only infinite sequences of trees of a special form that can be recognized by a

Büchi tree automaton \mathcal{B} , after which constructing the desired automaton \mathcal{A}' for $\text{Rec}(\mathcal{A})$ is easy. Unlike in the word case, we shall not treat separately trees with looping and non-looping witnessing sequences. We shall now prove an analogue of lemma 4. For the rest of the section, we shall need the following definition. For any context tree $T'[x_1, \dots, x_n] = (D', \tau')$ and a tree $T = (D, \tau)$, we write $T'[x_1, \dots, x_n] \sqsubseteq T$ (or just $T' \sqsubseteq T$) if, whenever u_1, \dots, u_n are leaf nodes in T' labeled by x_1, \dots, x_n , respectively, it is the case that

- $u_i = v_i r_i$ for some $r_i \in \mathcal{T}$ and $v_i \in D \cap D'$
- for $1 \leq i \leq n$, we have $u_i \notin D$ and $D' - \{u_1, \dots, u_n\} \subseteq D$.

In other words, $T' \sqsubseteq T$ iff all the nodes in T' are in T except for the context leaves. In the word case, an analogous notion that we used in the previous section is the relation between words w' and w such that $|w'| = |w|$. The use of equality as opposed to the inequality ' \leq ' is owing to the looping and the non-looping case that we considered separately.

Lemma 19. *For every tree T , it is the case that $T \in \text{Rec}(\mathcal{A})$ iff there exists a Γ -labeled unranked tree $\mathfrak{T} = (D_{\mathfrak{T}}, \tau_{\mathfrak{T}})$, where $\Gamma := \{(t \otimes t'[x_1, \dots, x_r], q, q') : q \in Q_1, q' \in Q_2, r \in \mathbb{N}, t \text{ is a tree, } t' \text{ is a context tree, and } t' \sqsubseteq t\}$, and $\tau_{\mathfrak{T}}(u) = (\alpha_u \otimes \beta_u[x_1, \dots, x_{r_u}], q_u, q'_u)$ for all $u \in D_{\mathfrak{T}}$, such that the following conditions hold:*

1. $\tau_{\mathfrak{T}}(\varepsilon) = (T \otimes \beta_{\varepsilon}[x_1, \dots, x_{r_{\varepsilon}}], q_0^1, q_0^2)$ for some context tree $\beta_{\varepsilon}[x_1, \dots, x_{r_{\varepsilon}}]$ and some $r_{\varepsilon} \in \mathbb{N}$ such that $\beta_{\varepsilon} \sqsubseteq \alpha_{\varepsilon}$,
2. for all $u \in D_{\mathfrak{T}}$ we have
 - (a) the number of children of u is the same as r_u ,
 - (b) $\alpha_u \otimes \beta_u[\alpha_{u_1}, \dots, \alpha_{u_{r_u}}] \in L(\mathcal{M}^{q_u})$,
 - (c) if v_1, \dots, v_{r_u} are the nodes of β_u labeled by x_1, \dots, x_{r_u} respectively, then there exist an accepting run ρ_u of \mathcal{A}^{q_u} on $\beta_u[\alpha_{u_1}, \dots, \alpha_{u_{r_u}}]$ and a potentially accepting run ρ'_u of $\mathcal{M}^{q'_u}$ on $\beta_u \otimes \beta_u$ such that, for each $i = 1, \dots, r_u$, it is the case that $q_{ui} = \rho_u(v_i)$ and $q'_{ui} = \rho'_u(v_i)$.

Notice that, if we restrict the value of r in the definition of Γ in lemma 19 to be at most 1, we obtain a lemma for the word case, which is similar to lemma 4. Observe that the tree \mathfrak{T} might be finite in which case each leaf node is labeled by some $(t \otimes t'[x_1, \dots, x_r], q, q') \in \Gamma$ with $r = 0$. It is also possible that some (or all) branches are infinite. We shall first show that lemma 19 is sufficient. For this, we shall construct a witnessing sequence $\{T_i\}_{i \geq 0}$ out of the tree \mathfrak{T} . We shall inductively define $\{T_i\}_{i \geq 0}$ together with a sequence $\{C_i\}_{i \geq 0}$ of context trees as follows. We set $T_0 := \alpha_{\varepsilon} = T$, $C_0 := x$, $T_1 := \beta_{\varepsilon}[\alpha_1, \dots, \alpha_{r_{\varepsilon}}]$, and $C_1 := \beta_{\varepsilon}[x_1^{\varepsilon}, \dots, x_{r_{\varepsilon}}^{\varepsilon}]$. Suppose that C_i with $i \geq 1$ has been defined to be the context tree $T'[x_1^{u_1}, \dots, x_{r_{u_1}}^{u_1}, \dots, x_1^{u_n}, \dots, x_{r_{u_n}}^{u_n}]$ for some nodes u_1, \dots, u_n in \mathfrak{T} of the same height, where $n \in \mathbb{N}$ and $r_{u_1}, \dots, r_{u_n} \in \mathbb{N}$. We define C_{i+1}

to be $T'[\sigma]$, where σ replaces $x_k^{u_j}$ by $\beta_{u_j k}[x_1^{u_j k}, \dots, x_{r_{u_j k}}^{u_j k}]$. Similarly, we define T_{i+1} to be $T'[\sigma]$, where σ replaces $x_k^{u_j}$ by $\alpha_{u_j k}$. Notice that if C_i is ground, then $T_{i+1} = T_i$ and $C_{i+1} = C_i$. By induction, the sequence $\{T_i\}_{i \geq 0}$ together with $\{C_i\}_{i \geq 0}$ have been defined. It is not difficult to prove by induction that $T_i \in L(\mathcal{A})$ and $T_{i-1} \otimes T_i \in L(\mathcal{M})$ for all $i \in \mathbb{Z}^+$. Therefore, we conclude that $T \in \text{Rec}(\mathcal{A})$.

We shall now prove the converse. As we did for the word case, we first extend the definition of $\text{Rec}(\mathcal{N}, \mathcal{T})$ in the same way to allow not necessarily transitive transducers \mathcal{T} . We need the following intermediate lemma.

Lemma 20. *Suppose \mathcal{N} and \mathcal{T} are, respectively, an automaton and a transducer over Σ . For every tree $T = (D, \tau)$, if $T \in \text{Rec}(\mathcal{N}, \mathcal{T})$, then one of the following is true:*

1. *there exists a tree $T' = (D', \tau')$ such that $D' \subseteq D$, $T' \in L(\mathcal{N})$, $(T, T') \in L(\mathcal{T})$, and $(T', T') \in L(\mathcal{T})$.*
2. *There exist a context tree $T'[x_1, \dots, x_n] = (D', \tau')$ and trees t_1, \dots, t_n such that*
 - (a) $T' \sqsubseteq T$,
 - (b) $(T, T'[t_1, \dots, t_n]) \in L(\mathcal{T})$,
 - (c) *there exist an accepting run $r = (D_r, \tau_r)$ of \mathcal{N} on $T'[t_1, \dots, t_n]$ and a potentially accepting run $r' = (D_{r'}, \tau_{r'})$ of \mathcal{T} on $T' \otimes T'$ such that, whenever $1 \leq i \leq n$, it is the case that $t_i \in \text{Rec}(\mathcal{N}^{q_i}, \mathcal{T}^{q'_i})$ where $q_i = \tau_r(u_i)$ and $q'_i = \tau_{r'}(u_i)$.*

Proof. Suppose that $T = (D, \tau) \in \text{Rec}(\mathcal{N}, \mathcal{T})$, but statement (1) on RHS is false. Then, there exists an infinite sequence $\sigma = \{T_i\}_{i \in \mathbb{N}}$ of trees such that $T_0 = T$, $T_j \neq T_k$ for all distinct indices j, k , and it is the case that, for all $i > 0$, $T_i \in L(\mathcal{N})$ with accepting run $\eta_i = (D_{\eta_i}, \tau_{\eta_i})$, and for all distinct pair of indices $0 \leq i < i'$, $T_i \otimes T_{i'} \in L(\mathcal{T})$. Now, for every tree T_i , where $i > 0$, there exists a unique context tree $C_i[x_1, \dots, x_{n_i}] = (D_i, \tau_i)$, for some $n_i \in \mathbb{N}$, such that $C_i \sqsubseteq T$, and $T_i = C_i[t_1^i, \dots, t_{n_i}^i]$ for some (ground) trees $t_1^i, \dots, t_{n_i}^i$. Let $H = \{C_i[x_1, \dots, x_{n_i}] : i > 0\}$. For infinitely many $i > 0$, it is the case that $n_i > 0$, i.e., there exists a node in T_i that is not in T ; for, otherwise, there are infinitely many indices i such that $D_i \subseteq D$ where $T_i = (D_i, \tau_i)$ and, since there are only finitely many different such trees, pigeonhole principle tells us that one of these trees must repeat in σ , which contradicts our assumption that statement (1) is false. On the other hand, it is easy to see that the number of nodes in any context tree C_i in H is bounded by $|\mathcal{Y}| \times |D|$. Therefore, the set H is finite and so is the number of different potentially accepting runs of \mathcal{N} on context trees in H . So, if we define $\eta'_i := (\eta_i)|_D$, i.e., the part of the run tree η_i restricted to the domain D of T , then by pigeonhole principle there exists $k > 0$ such that $C_k[x_1, \dots, x_{n_k}] = C_j[x_1, \dots, x_{n_j}]$ and $\eta_k = \eta_j$ for infinitely many indices j s. Let $n := n_k$, $T'[x_1, \dots, x_n] := C_k[x_1, \dots, x_n]$, and $\eta' = \eta'_k$. We remove all elements

T_i ($i > 0$) from σ such that $C_i \neq T'$ or $\eta'_i \neq \eta'$ and, by renaming indices, call the resulting sequence $\sigma = \{T_i\}_{i \in \mathbb{N}}$ where $T_0 = T$. The same is done for the sequence $\{\eta_i\}$ of runs so that η_i is an accepting run \mathcal{N} on T_i ($i \geq 1$) such that $\eta' \preceq \eta_i$. Notice that σ is still a witness for $T \in \text{Rec}(\mathcal{N}, \mathcal{T})$. Now let $\theta(j, k)$, where $0 \leq j < k$, be an accepting run of \mathcal{T} on $T_j \otimes T_k$. Let \mathcal{C} be the *finite* set of all potentially accepting runs of \mathcal{T} on $T' \otimes T'$. The set \mathcal{C} is nonempty as $T' \otimes T' \preceq T_j \otimes T_k$ and $T_j \otimes T_k \in L(\mathcal{T})$. Consider the edge-labeled undirected graph $G = (V, \{E_\rho\}_{\rho \in \mathcal{C}})$ such that $V = \mathbb{Z}^+$ and

$$E_\rho := \{\{j, k\} : 0 < j < k \text{ and } \rho \preceq \theta(j, k)\}.$$

Notice that $\{E_u\}_{u \in \mathcal{C}}$ is a partition of $\{\{j, k\} : j \neq k \in \mathbb{Z}^+\}$, and so G is a complete graph. By (infinite) Ramsey theorem, G has a monochromatic complete infinite subgraph $H = (V', E_\rho)$ for some $\rho \in \mathcal{C}$. Set $r' := \rho$. Notice that if V' contains the elements $j_1 < j_2 < \dots$ then $r' \preceq \theta(j_k, j_{k'})$ for all $k' > k \geq 1$. We now remove all T_i ($i \geq 1$) in σ with $i \notin V'$ and, again, rename indices. Notice that σ is still a witness for $T \in \text{Rec}(\mathcal{N}, \mathcal{T})$. Recall that for each $i \geq 1$, we have $T_i = T'[t_i^1, \dots, t_i^n]$ for some ground trees t_i^1, \dots, t_i^n . Set $r := \eta_1$ and $t_k := t_1^k$ for each $k = 1, \dots, n$. Letting $\sigma_k = \{t_i^k\}_{i \geq 1}$ for each $k = 1, \dots, n$, it is easy now to check that $t_k \in \text{Rec}(\mathcal{N}^{q_k}, \mathcal{T}^{q'_k})$ with witnessing sequence σ_k , where $q_k = \tau_r(u_k)$ and $q'_k = \tau_{r'}(u_k)$ if u_k is the leaf node of T' labeled by x_k . So, condition (2c) holds. That (2b) holds is also immediate. As we already saw that $T' \sqsubseteq T$, the lemma is proven. \square

In the same way we used lemma 5 to complete the proof of necessity in lemma 4, we can now finish off the proof of necessity in lemma 19 by constructing the tree \mathfrak{T} inductively and adding nodes of height n at step $n \in \mathbb{N}$ by using lemma 20. Therefore, the proof of lemma 19 is complete.

We now construct the Büchi automaton \mathcal{B} . For any tree $T = (D, \tau)$, we write \widehat{T} for the tree T with root marked by the new symbol $\#$, i.e., $\widehat{T} := (1D, \widehat{\tau})$ with $\widehat{\tau}(\varepsilon) := \#$ and, whenever $u \in D$, $\widehat{\tau}(1u) := \tau(u)$. Given an unranked tree $\mathfrak{T} = (D_{\mathfrak{T}}, \tau_{\mathfrak{T}})$ satisfying conditions in lemma 19 on the right, we can inductively define a Ω -labeled Υ -tree H_v for every $v \in D_{\mathfrak{T}}$, where $\Omega := \Sigma_{\perp}^2 \cup \{\#\}$. We set $H_v := (\alpha_v \otimes \beta_v)[\widehat{H}_{v_1}, \dots, \widehat{H}_{v_{r_v}}]$. Note that H_v might be infinite for some $v \in D_{\mathfrak{T}}$. If $H_v = (D, \tau)$, we also denote by $\mathbf{full}(H_v)$ the full infinite tree (Υ^*, τ') such that if $u \in D$, then $\tau'(u) := \tau(u)$; if $u \notin D$, then $\tau'(u) := \perp$ where $\perp := \begin{bmatrix} \perp \\ \perp \end{bmatrix}$. In other words, the tree $\mathbf{full}(T)$ is the tree T made full by padding finite branches by \perp .

Our Büchi tree automaton \mathcal{B} accepts precisely all Ω -labeled full infinite binary tree $\mathbf{full}(H_\varepsilon)$, where H_ε is generated by some unranked tree \mathfrak{T} satisfying lemma 19. Constructing a Büchi tree automaton $\mathcal{B} = (\Omega, Q, \delta, q_0, F)$ is now easy. The construction is very similar to the word case. We define

$$Q := (Q_1 \times Q_2 \times Q_2) \uplus (Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2),$$

where \uplus means disjoint union. The start state is $q_0 := (q_0^1, q_0^2, q_0^2)$. The states in $(Q_1 \times Q_2 \times Q_2)$ are meant to handle the cases when no $\#$ has thus far been

seen by \mathcal{B} . On the other hand, when \mathcal{B} is in $Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2$, at least one $\#$ has been seen. We now formally define the transition function δ . We first define how \mathcal{B} behaves when it is in $Q_1 \times Q_2 \times Q_2$. Suppose that $q_1 \in Q_1$, and $q_2, q'_2 \in Q_2$. For all $a, b \in \Sigma$, we set

$$\delta((q_1, q_2, q'_2), \begin{bmatrix} a \\ b \end{bmatrix}) := \left\{ \begin{array}{l} ((q_{L1}, q_{L2}, q'_{L2}), (q_{R1}, q_{R2}, q'_{R2})) \mid \\ (q_{L1}, q_{R1}) \in \delta_1(q_1, b), \\ (q_{L2}, q_{R2}) \in \delta_2(q_2, \begin{bmatrix} a \\ b \end{bmatrix}), \\ (q'_{L2}, q'_{R2}) \in \delta_2(q'_2, \begin{bmatrix} b \\ b \end{bmatrix}) \end{array} \right\}.$$

For all $a \in \Sigma$, we set

$$\delta((q_F^1, q_2, q_F^2), \begin{bmatrix} a \\ \perp \end{bmatrix}) := \{(q_F^1, q_{L2}, q_F^2), (q_F^1, q_{R2}, q_F^2) : (q_{L2}, q_{R2}) \in \delta_2(q_2, \begin{bmatrix} a \\ \perp \end{bmatrix})\}$$

and we set

$$\delta((q_F^1, q_F^2, q_F^2), \perp) := \{(q_F^1, q_F^2, q_F^2), (q_F^1, q_F^2, q_F^2)\}.$$

Recall that there are no transitions from q_F^1 (resp. q_F^2) in the automaton \mathcal{A} (resp. \mathcal{M}) and so, for all $a, b \in \Sigma$, we have $\delta((q_1, q_2, q'_2), \begin{bmatrix} a \\ b \end{bmatrix}) = \emptyset$ if at least one of the following holds: $q_1 = q_F^1$, $q_2 = q_F^2$, or $q'_2 = q_F^2$. Likewise, we have $\delta((q_F^1, q_F^2, q_F^2), \begin{bmatrix} a \\ \perp \end{bmatrix}) = \emptyset$ unless $a = \perp$. Finally, if $q_1 \neq q_F^1$ and $q_2, q'_2 \neq q_F^2$, we then set

$$\delta((q_1, q_2, q'_2), \#) := ((q_1, q'_2, q_1, q_2, q'_2), (q_F^1, q_F^2, q_F^2)).$$

Notice that the state sent to the right child is (q_F^1, q_F^2, q_F^2) as the right child of every $\#$ -labeled node in $\mathbf{full}(H_\varepsilon)$ is \perp -labeled.

We now proceed with our definition of δ when \mathcal{B} is in $Q_1 \times Q_2 \times Q_1 \times Q_2 \times Q_2$. Suppose that $q_1, q'_1 \in Q_1$ and $q_2, q'_2, q''_2 \in Q_2$. For all $a, b \in \Sigma$, we define

$$\delta((q_1, q_2, q'_1, q'_2, q''_2), \begin{bmatrix} a \\ b \end{bmatrix})$$

as

$$\left\{ \begin{array}{l} ((q_{L1}, q_{L2}, q'_{L1}, q'_{L2}, q''_{L2}), (q_{R1}, q_{R2}, q'_{R1}, q'_{R2}, q''_{R2})) \mid \\ (q_{L1}, q_{R1}) \in \delta_1(q_1, b), \\ (q_{L2}, q_{R2}) \in \delta_2(q_2, \begin{bmatrix} a \\ b \end{bmatrix}), \\ (q'_{L1}, q'_{R1}) \in \delta_1(q'_1, a), \\ (q'_{L2}, q'_{R2}) \in \delta_2(q'_2, \begin{bmatrix} \perp \\ a \end{bmatrix}), \\ (q''_{L2}, q''_{R2}) \in \delta_2(q''_2, \begin{bmatrix} b \\ b \end{bmatrix}) \end{array} \right\}$$

We also define

$$\delta((q_F^1, q_2, q'_1, q'_2, q_F^2), \begin{bmatrix} a \\ \perp \end{bmatrix})$$

as

$$\left\{ \left((q_F^1, q_{L2}, q'_{L1}, q'_{L2}, q_F^2), (q_F^1, q_{R2}, q'_{R1}, q'_{R2}, q_F^2) \right) \left| \begin{array}{l} (q_{L2}, q_{R2}) \in \delta_2(q_2, \begin{bmatrix} a \\ \perp \end{bmatrix}), \\ (q'_{L1}, q'_{R1}) \in \delta_1(q'_1, a), \\ (q'_{L2}, q'_{R2}) \in \delta_2(q'_2, \begin{bmatrix} \perp \\ a \end{bmatrix}) \end{array} \right. \right\}$$

We set $\delta((q_F^1, q_F^2, q_F^1, q_F^2, q_F^2), \perp) := (q_F^1, q_F^2, q_F^1, q_F^2, q_F^2)$. Finally, if $q_1 \neq q_F^1$ and $q_2, q_2'' \neq q_F^2$, then $\delta((q_1, q_2, q_F^1, q_F^2, q_2''), \#) := (q_1, q_2'', q_1, q_2, q_2'')$. In this case, notice that $\delta((q_1, q_2, q_F^1, q_F^2, q_2''), l) = \emptyset$ unless $l = \#$.

We now set

$$F := \{(q_F^1, q_F^2, q_F^2), (q_F^1, q_F^2, q_F^1, q_F^2, q_F^2)\} \cup \{(q_1, q_2, q_F^1, q_F^2, q_2'') : q_1 \neq q_F^1 \text{ and } q_2, q_2'' \neq q_F^2\}$$

It is easy to see that \mathcal{B} recognizes precisely all trees $\mathbf{full}(H_\varepsilon)$, where H_ε is generated by some unranked tree \mathfrak{T} satisfying lemma 19. We now show how to construct the automaton $\mathcal{A}' = (\Sigma, Q', \delta', q'_0, F')$ that recognizes $\mathit{Rec}(\mathcal{A})$ out of \mathcal{B} . The intuitive idea is similar to the word case: given a tree T , the automaton \mathcal{A}' guesses a tree $\mathbf{full}(H_\varepsilon)$, where H_ε is generated by an unranked tree $\mathfrak{T} = (D_{\mathfrak{T}}, \tau_{\mathfrak{T}})$ satisfying lemma 19 such that $\tau_{\mathfrak{T}}(\varepsilon) = (T \otimes \beta_\varepsilon[x_1, \dots, x_{r_\varepsilon}], q_0^1, q_0^2)$ for some context tree $\beta_\varepsilon[x_1, \dots, x_{r_\varepsilon}]$. More formally, we set $Q' := (Q_1 \times Q_2 \times Q_2)$ and $q'_0 = (q_0^1, q_0^2, q_0^2)$. The transition function is defined as follows:

$$\delta'((q_1, q_2, q'_2), a) = \bigvee_{b \in \Sigma_\perp} \delta((q_1, q_2, q'_2), \begin{bmatrix} a \\ b \end{bmatrix}).$$

Finally, we set

$$F' := \{(q_F^1, q_F^2, q_F^2)\} \cup \{(q_1, q_2, q'_2) \in Q_1 \times Q_2 \times Q_2 \mid \mathcal{B}^{(q_1, q_2, q'_2)} \text{ accepts some } \Omega\text{-labeled binary tree of the form } \widehat{T}\}$$

Observe that F' can be computed by using the algorithm for checking emptiness for Büchi tree automata, which runs in quadratic time $O(|\mathcal{B}|^2) = O(|\mathcal{A}|^4 \times |\mathcal{M}|^6)$ (e.g. see [26]). Theorem 9 is now immediate.