# Model-Checking Infinite Systems Generated by Ground Tree Rewriting 

Christof Löding<br>RWTH Aachen, Lehrstuhl Informatik VII, 52056 Aachen, Germany<br>loeding@informatik.rwth-aachen.de


#### Abstract

We consider infinite graphs that are generated by ground tree (or term) rewriting systems. The vertices of these graphs are trees. Thus, with a finite tree automaton one can represent a regular set of vertices. It is shown that for a regular set $T$ of vertices the set of vertices from where one can reach (respectively, infinitely often reach) the set $T$ is again regular. Furthermore it is shown that the problems, given a tree $t$ and a regular set $T$, whether all paths starting in $t$ eventually (respectively, infinitely often) reach $T$, are undecidable. We then define a logic which is in some sense a maximal fragment of temporal logic with a decidable model-checking problem for the class of ground tree rewriting graphs.


## 1 Introduction

Graphs play an important role in the behavioral description of programs or processes. The use of theoretically unbounded data structures (such as stacks, queues, etc.) in these programs requires the use of infinite (transition-)graphs. In order to allow algorithmic applications like verification of infinite-state systems, these systems have to be given effectively, i.e., by some finite object. In recent years many classes of finitely representable infinite graphs have been studied. A starting point for this research on infinite graphs was an analysis of the configuration graphs of pushdown automata by Muller and Schupp in [13]. The vertices of pushdown graphs are words (the control state followed by the stack content) and the edges are defined via the transitions which are prefix rewriting rules on these words. In 13 it is shown that pushdown graphs have a decidable monadic second-order (MSO) theory. Later, more efficient algorithms have been developed for solving reachability problems, model-checking temporal logics [8], and synthezising strategies for games 1014 on pushdown graphs.

Using rewriting as the basic process for generation of infinite graphs, a variety of possibilities arises to define classes of graphs different from pushdown graphs. Let us briefly review some of the approaches. The result on the decidability of the MSO theory generalizes from pushdown graphs to prefix recognizable graphs [3], which are also generated by prefix rewriting on words; but the rewriting rules refer to to regular languages instead of single words. Even more general classes of graphs can be obtained when using finite word transducers to define the edge relations of graphs. This leads to automatic graphs in case of synchronous
transducers (see e.g. [1) and to rational graphs [12] in case of asynchronous transducers. The formalisms to generate automatic and rational graphs are very strong and even simple reachability problems on these graphs are undecidable. In [11] model-checking problems for process rewriting graphs are studied (in process rewriting parallel composition of words is allowed in addition to the usual sequential composition).

In the present paper we use another generalization of word rewriting, namely tree rewriting, as already considered in [2]. In tree (or term) rewriting the basic objects in the rewriting systems are trees. For ground tree rewrite systems (GTRS) confluence [6], the first-order theory [7, and several reachability problems [5] have been shown to be decidable. As the example of the infinite grid (see Section 2 below) shows, the MSO theory of a GTRS graph can be undecidable. Our goal is to develop an expressive fragment of temporal logic with a decidable model-checking problem on GTRS graphs. Since reachability analysis of systems constitutes an important part of verification, we analyze different reachability problems. As a means of specification we will use finite tree automata to represent sets of vertices. Given such a regular set $T$ represented by a tree automaton we address the following problems for GTRS graphs.
(1) Compute the set of all vertices from where one can reach $T$.
(2) Compute the set of all vertices from where one can infinitely often visit $T$.
(3) Given a single tree $t$, do all paths starting in $t$ eventually (respectively, infinitely often) visit $T$ ?

We give algorithms to solve the Problems (1) and (2). The first problem was already solved in the context of term rewriting (see e.g. [4]), so the solution of the second problem is the main contribution of the present paper. The problems in (3) are shown to be undecidable and hence mark a boundary in the design of a fragment of temporal logic with decidable model-checking problem for GTRS graphs.

The paper is organized as follows. In Section 2 we introduce GTRS and tree automata. In Section 3 we give an algorithm to solve the reachability problem (1), in Section 4 we solve the recurrence problem (2), and in Section 5 we show the undecidability of the problems of universal reachability and universal recurrence from (3). Finally, in Section [6, we use the results of the former sections and present a logic with a decidable model-checking problem for GTRS graphs.

I would like to thank Wolfgang Thomas and the anonymous referees for many helpful comments.

## 2 Ground Tree Rewriting Systems and Tree Automata

A ranked alphabet $A$ is a family of sets $\left(A_{i}\right)_{i \in[k]}$, where $[k]=\{0, \ldots, k\}$. For simplicity we identify $A$ with the set $\bigcup_{i=0}^{k} A_{i}$. A ranked tree $t$ over $A$ is a mapping $t: D_{t} \rightarrow A$ with $D_{t} \subseteq[k-1]^{*}$ such that $D_{t} \neq \emptyset$ is prefix closed, for each $u i \in D_{t}$ we have $u j \in D_{t}$ for all $j \leq i$, and if $u 0, \ldots, u(i-1) \in D_{t}$, ui $\notin D_{t}$, then $t(u) \in A_{i} . D_{t}$ is called the domain of $t$. We will mainly be interested in finite
trees, except for Section 4 where we also use infinite trees. The set of all finite trees over $A$ is called $T_{A}$. By $\sqsubseteq$ we denote the prefix ordering on $\mathbb{N}^{*}$.

For $u \in \mathbb{N}^{*}$ we define $u D_{t}=\left\{u v \in \mathbb{N}^{*} \mid v \in D_{t}\right\}$ and $u^{-1} D_{t}=\left\{v \in \mathbb{N}^{*} \mid u v \in\right.$ $\left.D_{t}\right\}$. For $u \in D_{t}$ the subtree $t^{u}$ of $t$ at $u$ is the tree with domain $D_{t^{u}}=u^{-1} D_{t}$ and $t^{u}(v)=t(u v)$. For trees $t, t^{\prime} \in T_{A}$ and $u \in D_{t}$ we define $t\left(u \leftarrow t^{\prime}\right)$ to be the tree $s$ with domain $D_{s}=u D_{t^{\prime}} \cup\left(D_{t} \backslash u\left(u^{-1} D_{t}\right)\right)$ and $s(v)=t^{\prime}\left(u^{-1} v\right)$ for $v \in u D_{t^{\prime}}$ and $s(v)=t(v)$ for $v \in D_{t} \backslash u\left(u^{-1} D_{t}\right)$. This means we replace the subtree $t^{u}$ in $t$ by $t^{\prime}$.

To denote trees we will use the usual graphical notation (as e.g. in Figure 1) and the term notation, where $a\left(t_{0}, \ldots, t_{i-1}\right)$ denotes the tree with $a$ at the root and $t_{0}, \ldots, t_{i-1}$ as subtrees.

A ground tree rewriting system (GTRS) is a tuple $S=\left(A, \Sigma, P, t_{I}\right)$, where $A=\left(A_{i}\right)_{i \in[k]}$ is a ranked alphabet, $\Sigma$ is an alphabet, $P$ is a finite set of rules $t_{1} \rightarrow^{\sigma} t_{2}$ with $t_{1}, t_{2} \in T_{A}, \sigma \in \Sigma$, and $t_{I} \in T_{A}$ is the initial tree. For two trees $t, t^{\prime} \in T_{A}$ we write $t \rightarrow{ }_{S}^{\sigma} t^{\prime}$ iff there exists a rule $t_{1} \rightarrow^{\sigma} t_{2}$ in $P$ and $u \in D_{t}$ such that $t^{u}=t_{1}$ and $t^{\prime}=t\left(u \leftarrow t_{2}\right)$. We write $t \rightarrow_{S} t^{\prime}$ iff there is a $\sigma \in \Sigma$ with $t \rightarrow{ }_{S}^{\sigma} t^{\prime}$. By $\rightarrow_{S}^{+}$we denote the transitive closure of $\rightarrow_{S}$ and by $\rightarrow_{S}^{*}$ we denote the transitive and reflexive closure of $\rightarrow_{S}$.

The tree language generated by $S$ is $T(S)=\left\{t \in T_{A} \mid t_{I} \rightarrow_{S}^{*} t\right\}$. The edge labeled graph $G_{S}=\left(V_{S}, E_{S}, \Sigma\right)$ generated by $S$ is defined by $V_{S}=T(S)$, and $\left(t, \sigma, t^{\prime}\right) \in E_{S}$ iff $t \rightarrow{ }_{S}^{\sigma} t^{\prime}$.

Example 1. The GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ given by $\Sigma=\{0,1\}$, $A=\left(A_{i}\right)_{i \in[2]}$ with $A_{0}=\{a, b\}, A_{1}=\{c\}$, and $A_{2}=\{d\}, P=\left\{b \rightarrow^{0} c(b), a \rightarrow^{1} c(a)\right\}$, and $t_{I}=d(a, b)$ generates the infinite grid as shown in Figure 1 .

As usual a path $\pi$ through a graph $G=(V, E)$ is a (possibly infinite) sequence of vertices such that two successive vertices on this path form an edge. If we consider a path $\pi$ through a graph $G_{S}$ for a GTRS $S$, then the edges are generated by rewriting rules. We will sometimes refer to these rewritings as the


Fig. 1. The infinite grid generated by a GTRS.
rewritings on $\pi$. Given a path $\pi$ and $i \in \mathbb{N}$ we denote the $i$ th element on $\pi$ by $\pi(i)$ and the suffix of $\pi$ starting at $\pi(i)$ by $\pi^{i}$.

If $T \subseteq T_{A}$, then $t \rightarrow{ }_{S}^{\omega} T$ means that there is a path through $G_{S}$ starting in $t$ that infinitely often visits the set $T$ and if $\pi$ is a path through $G_{S}$, then $\pi: t \rightarrow{ }_{S}^{\omega} T$ means that $\pi$ is a path starting in $t$ that infinitely often visits $T$. Furthermore we use similar notations as, e.g., $t \rightarrow{ }_{S}^{*} T_{1} \rightarrow_{S}^{*} T_{2}$ if there is a path that starts in $t$ and first visits $T_{1}$ and then $T_{2}$.

For a set $T \subseteq T_{A}$ of trees let $\operatorname{sub}(T)=\left\{t^{\prime} \in T_{A} \mid \exists t \in T\right.$ and $u \in$ $D_{t}$ with $\left.t^{u}=t^{\prime}\right\}$ be the set of all subtrees of trees in $T$. For later use we define the sets $P_{L}=\left\{t \in T_{A} \mid \exists t^{\prime} \in T_{A}\right.$ with $\left.t \rightarrow t^{\prime} \in P\right\}, P_{R}=\left\{t \in T_{A} \mid \exists t^{\prime} \in\right.$ $T_{A}$ with $\left.t \rightarrow t^{\prime} \in P\right\}$, and $\operatorname{sub}(P)=\operatorname{sub}\left(P_{L}\right) \cup \operatorname{sub}\left(P_{R}\right)$.

A nondeterministic tree automaton (NTA) is a tuple $\mathcal{A}=(Q, A, \Delta, F)$, where $Q$ is a finite set of states, $A=\left(A_{i}\right)_{i \in[k]}$ is a ranked alphabet, $F \subseteq Q$ is a set of final states and $\Delta \subseteq\left(\bigcup_{i=0}^{k} Q^{i} \times A_{i}\right) \times Q$ is the transition relation. Given a tree $t \in T_{A}$, a tree $\rho \in T_{Q}$ is a run of $\mathcal{A}$ on $t$ iff $D_{\rho}=D_{t}$, and for each $u \in D_{t}$ with $t(u) \in A_{i}$ we have $(\rho(u 0), \ldots, \rho(u(i-1)), t(u), \rho(u)) \in \Delta$. We extend $\Delta$ to trees in the natural way: $(t, q) \in \Delta$ iff $\rho(\varepsilon)=q$ for a run $\rho$ of $\mathcal{A}$ on $t$. A tree $t$ is accepted by $\mathcal{A}$ iff $(t, q) \in \Delta$ for some $q \in F$ and the language $T(\mathcal{A})$ accepted by $\mathcal{A}$ is $T(\mathcal{A})=\left\{t \in T_{A} \mid(t, q) \in \Delta(t)\right.$ for some $\left.q \in F\right\}$. A language of trees is called regular iff it can be recognized by an NTA.

For $t \in T_{A}$, a run $\rho$ of an NTA on $t$, and a vertex $u \in D_{t}$ it is clear that the subtree $\rho^{u}$ of $\rho$ is a run of the NTA on $t^{u}$.

The notion of a run can also be applied to infinite trees. This will be used in Section 4

In the next section we also consider a more general type of automata, namely nondeterministic tree automata with $\varepsilon$-transitions ( $\varepsilon$-NTA). An $\varepsilon$-NTA is a tuple $\mathcal{A}=(Q, A, \Delta, F)$, where $Q, A$, and $F$ are as in NTA and $\Delta \subseteq\left(\left(\bigcup_{i=0}^{k} Q^{i} \times A_{i}\right) \cup\right.$ $Q) \times Q$ is the transition relation. The transitions of $\Delta$ from the set $Q \times Q$ are called $\varepsilon$-transitions because in a run these $\varepsilon$-transitions do not depend on the input tree. For an $\varepsilon$-NTA $\mathcal{A}=(Q, A, \Delta, F)$ one can easily define an equivalent $\varepsilon$-free NTA $\mathcal{A}^{-}=\left(Q, A, \Delta^{-}, F\right)$ by $\left(q_{1}, \ldots, q_{i}, a, q\right) \in \Delta^{-}$if there are $p_{0}, \ldots, p_{j} \in Q$ with $p_{j}=q,\left(q_{1}, \ldots, q_{i}, a, p_{0}\right) \in \Delta$, and $\left(p_{l}, p_{l+1}\right) \in \Delta$ for all $l \in\{0, \ldots, j-1\}$. Then we can define $(t, q) \in \Delta$ iff $(t, q) \in \Delta^{-}$.

## 3 Reachability

In this section we give an algorithm to solve the following reachability problem:
(Reachability): Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and a regular set of trees $T \subseteq T_{A}$, compute the set of all trees in $T_{A}$ from which one can reach a tree of $T$ in the graph $G_{S}$ associated to $S$.

This problem was already solved in [5] and an algorithm similar to the one presented here can be found in [4] for more general rewrite systems. Thus, we
only provide the algorithm for the sake of completeness and without correctness proof.

The main idea of the algorithm is to simulate the rewriting rules within the tree automaton $\mathcal{A}=(Q, A, \Delta, F)$ that accepts the regular set $T=T(\mathcal{A})$. If the automaton contains for each rewriting rule $t_{1} \rightarrow t_{2}$ a unique state $q_{t_{1}}$ that identifies $t_{1}$, i.e., $\left(t, q_{t_{1}}\right) \in \Delta$ iff $t=t_{1}$, then such the rule $t_{1} \rightarrow t_{2}$ can be simulated by adding an $\varepsilon$-transition $\left(q_{t_{1}}, q\right)$ to $\Delta$, where $q$ is a state with $\left(t_{2}, q\right) \in \Delta$. So the algorithm adds a new part to $\mathcal{A}$ such that subtrees from $P_{L}$ can be uniquely identified and then starts adding $\varepsilon$-transitions as described above.

First we define the new part that is added to $\mathcal{A}$. Given a set $P$ of rewriting rules over a ranked alphabet $A=\left(A_{i}\right)_{i \in[k]}$, let

- $Q^{P}=\left\{q_{t} \mid t \in \operatorname{sub}\left(P_{L}\right)\right\}$, and
$-\Delta^{P}=\left\{\left(q_{t_{1}}, \ldots, q_{t_{i}}, a, q_{a\left(t_{1}, \ldots, t_{i}\right)} \mid a\left(t_{1}, \ldots, t_{i}\right) \in \operatorname{sub}\left(P_{L}\right)\right\}\right.$.
If these new states and transitions are added to $\mathcal{A}$, then the accepted language does not change and for each $t \in P_{L}$ the state $q_{t}$ can only be reached via $t$.

Figure 2 shows an algorithm solving the reachability problem. Note that the algorithm always terminates since only finitely many transitions can be added to the automaton $\mathcal{A}_{0}$.

```
INPUT: GTRS \(S=\left(A, \Sigma, P, t_{I}\right), \varepsilon\)-NTA \(\mathcal{A}=(Q, A, \Delta, F)\)
\(Q_{0}=Q \dot{\cup} Q^{P}, \Delta_{0}=\Delta \dot{\cup} \Delta^{P}, F_{0}=F\)
\(\mathcal{A}_{0}:=\left(Q_{0}, A, \Delta_{0}, F_{0}\right)\)
\(i:=0\)
while \(\exists t_{1} \rightarrow t_{2} \in P, q \in Q_{0}\) with \(\left(t_{2}, q\right) \in \Delta_{i}\) and \(\left(t_{1}, q\right) \notin \Delta_{i}\) do
    \(i:=i+1\)
    \(\Delta_{i}:=\Delta_{i-1} \cup\left\{\left(q_{t_{1}}, q\right)\right\}\)
    \(\mathcal{A}_{i}:=\left(Q_{0}, A, \Delta_{i}, F_{0}\right)\)
end
\(m:=i\)
OUTPUT: \(\varepsilon\)-NTA \(\mathcal{A}_{m}\)
```

Fig. 2. Algorithm to solve the reachability problem.

For a proof of the following theorem see e.g. [4] (or [8] for the special case of pushdown systems).
Theorem 1. Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and an $\varepsilon-N T A \mathcal{A}$, the algorithm from Figure 圆 computes an $\varepsilon-N T A$ with $|Q|+\left|\operatorname{sub}\left(P_{L}\right)\right|$ states accepting the set $\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{*} T\right\}$.

## 4 Recurrence

After having solved the reachability problem from the former section the next step is to deal with repeated reachability or recurrence.
(Recurrence): Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and a regular set of trees $T \subseteq T_{A}$, compute the set $\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\omega} T\right\}$.

For this section we fix a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and an NTA $\mathcal{A}=(Q, A, \Delta, F)$. By $\mathcal{A}_{q}$ we denote the automaton $\mathcal{A}_{q}=(Q, A, \Delta,\{q\})$.

There are two main steps in the construction of an NTA for the set $\left\{t \in T_{A} \mid\right.$ $\left.t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})\right\}$.

Step 1. We reduce the recurrence problem for $T(\mathcal{A})$ to the reachability problem for a set $R(\mathcal{A})$ (defined below) by showing $\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})\right\}=\{t \in$ $\left.T_{A} \mid t \rightarrow{ }_{S}^{*} R(\mathcal{A})\right\}$. The set $R(\mathcal{A})$ is regular and can be constructed if we can decide for all trees $t \in P_{L}$ and all states $q \in Q$ whether $t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$.
Step 2. We give a procedure for deciding for $t \in P_{L}$ and $q \in Q$ if $t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$.
Step 1. Informally speaking, the set $R(\mathcal{A})$ (for a fixed GTRS $S$ ) contains all the trees $t$ such that there is a state $q \in Q$ and vertex $u \in D_{t}$ with $t^{u} \in P_{L}$, $t^{u} \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$, and if we assume that there is a run of $\mathcal{A}$ on $t$ that labels $u$ with $q$, then there is an accepting run of $\mathcal{A}$ on $t$. The last condition ensures that $t\left(u \leftarrow t^{\prime}\right) \in T(\mathcal{A})$ for each $t^{\prime} \in T\left(\mathcal{A}_{q}\right)$. Since the second requirement says that from $t^{u}$ we can infinitely often reach a tree from $T\left(\mathcal{A}_{q}\right)$, it is clear that from $t$ we can infinitely often reach $T(\mathcal{A})$. On the other hand, if we are given a tree $t$ with $t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$, then we can also reach $R(\mathcal{A})$ from $t$. Before we prove this in Lemma we give a formal definition of $R(\mathcal{A})$.

Given $t \in T_{A}, u \in D_{t}$, and $q \in Q$, we say that $t(u \leftarrow q)$ is in $T(\mathcal{A})$ if there is a tree $t_{1} \in T_{A}$ such that there is an accepting run $\rho$ of $\mathcal{A}$ on $t\left(u \leftarrow t_{1}\right)$ with $\rho(u)=q$. For this definition we assume that all states of $\mathcal{A}$ are reachable, i.e., for each $q \in Q$ there is a $t \in T_{A}$ with $(t, q) \in \Delta$. We define the set

$$
R(\mathcal{A})=\left\{t \in T_{A} \mid \exists u \in D_{t}, q \in Q: \begin{array}{c}
t^{u} \in P_{L}, t^{u} \rightarrow_{S}^{\omega} T\left(\mathcal{A}_{q}\right), \\
\text { and } t(u \leftarrow q) \in T(\mathcal{A})
\end{array}\right\}
$$

Lemma 1. Let $t \in T_{A}$. Then $t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ iff $t \rightarrow{ }_{S}^{*} R(\mathcal{A})$.
Proof. " $\Rightarrow$ ": Assume that $t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ and let $\pi=t_{0} \rightarrow_{S} t_{1} \rightarrow_{S} t_{2} \rightarrow_{S} t_{3} \rightarrow_{S} \cdots$ be a path starting at $t=t_{0}$ that infinitely often visits $T(\mathcal{A})$. For $i \in \mathbb{N}$ let $u_{i} \in D_{t_{i}}$ and $t_{i}^{\prime} \in T_{A}$ such that $t_{i+1}=t_{i}\left(u_{i} \leftarrow t_{i}^{\prime}\right)$ (i.e. the rewriting rules are applied to the $t_{i}^{u_{i}}$ ). There is a $j \in \mathbb{N}$ such that no $u_{i}$ is a proper prefix of $u_{j}$ for all $i>j$, and there is an infinite number of $i>j$ such that $u_{j}$ is a prefix of $u_{i}$. In particular, this means $u_{j} \in D_{t_{i}}$ for all $i>j$. Let $t_{L}=t_{j}^{u_{j}}$. We know that $t_{L} \in P_{L}$. Since $\pi$ infinitely often visits $T(\mathcal{A})$, there must be $q \in Q$ that is infinitely often at $u_{j}$ in accepting runs on trees from $T(\mathcal{A})$ on $\pi$. We get $t_{L} \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ since $u_{i}$ is not a proper prefix of $u_{j}$ for $i>j$ and because there are infinitely many substitutions "below" $u_{j}\left(u_{j}\right.$ is a prefix of $u_{i}$ for infinitely many $\left.i>j\right)$.

Let $k>j$ be such that there is an accepting run of $\mathcal{A}$ on $t_{k}$ that labels $u_{j}$ with $q$. Define $t^{\prime}=t_{k}\left(u_{j} \leftarrow t_{L}\right)$ and $u=u_{j}$. Then $\left(t^{\prime}\right)^{u}=t_{L}$. From the accepting run on $t_{k}$ that labels $u$ with $q$ we get $t^{\prime}(u \leftarrow q) \in T(\mathcal{A})$. Thus, $t^{\prime} \in R(\mathcal{A})$.

For all $i>j$ we have that $u_{i}$ is not a proper prefix of $u_{j}$. Therefore we get $t_{j} \rightarrow{ }_{S}^{*} t^{\prime}$ and as a consequence $t \rightarrow{ }_{S}^{*} t^{\prime}$.
" $\Leftarrow$ ": Let $t^{\prime} \in R(\mathcal{A})$ such that $t \rightarrow{ }_{S}^{*} t^{\prime}$. It is sufficient to show that $t^{\prime} \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$. Let $q \in Q$ and $u \in D_{t^{\prime}}$ be such that $\left(t^{\prime}\right)^{u} \in P_{L}$ and $\left(t^{\prime}\right)^{u} \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$. Let $t_{0}=\left(t^{\prime}\right)^{u}$ and let $t_{0} \rightarrow_{S} t_{1} \rightarrow_{S} t_{2} \rightarrow_{S} t_{3} \rightarrow_{S} \cdots$ be a path starting at $t_{0}$ that infinitely often visits $T\left(\mathcal{A}_{q}\right)$. Then $t \rightarrow_{S} t\left(u \leftarrow t_{1}\right) \rightarrow_{S} t\left(u \leftarrow t_{2}\right) \rightarrow_{S} t\left(u \leftarrow t_{3}\right) \rightarrow_{S} \cdots$ is a path starting in $t$ that infinitely often visits $T(\mathcal{A})$.

Since the set $P_{L}$ is finite, it is not difficult to verify that the defining properties of trees in $R(\mathcal{A})$ can be checked by a finite tree automaton. But to construct such an automaton we have to decide for each $t \in P_{L}$ and $q \in Q$ whether $t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$. This is the second step mentioned above and will be done in the remainder of this section.

Step 2. There are two possibilities on how to visit the set $T(\mathcal{A})$ infinitely often along a path $\pi$.
(a) There is a single tree $t \in T(\mathcal{A})$ that is visited infinitely often along $\pi$.
(b) There are infinitely many different trees from $T(\mathcal{A})$ on $\pi$.

To distinguish the two cases we analyze how the trees evolve along a path $\pi$ by defining the limit of $\pi$. This is the tree consisting of all the vertices of trees on $\pi$ that are eventually not involved in the rewriting steps any more (i.e. the vertices that are fixed from a certain point onwards). In case (a) this limit will be a finite tree. If the limit of $\pi$ is infinite, then we have to deal with case (b). In both cases we can find a normal form for paths infinitely often visiting $T(\mathcal{A})$. That is, if there is a path $\pi^{\prime}: t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ for $t \in P_{L}$, then there is also a path $\pi$ in normal form with $\pi: t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$. Then we describe a procedure to decide for $t \in P_{L}$ and $q \in Q$ if there exists a path $\pi$ in normal form with $\pi: t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$.

The normal form for case (a) is based on the same ideas as in Step 1 from above and is stated without proof in Lemma 2. The normal form for case (b) (Lemma 3) is more complicated and will be described in more detail after the formal definition of the limit of a path.

A branch $\beta$ of a tree $t$ is a maximal prefix closed subset of $D_{t}$ such that for each $u \in \beta$ there is at most one $i \in \mathbb{N}$ with $u i \in \beta$ (in other words: a maximal subset of $D_{t}$ that is linearly ordered by $\sqsubseteq$ ). Given a branch $\beta$ of a tree $t$ and a vertex $v \in D_{t}$ we define $\beta(v)$ to be the maximal prefix $u$ of $v$ with $u \in \beta$.

Let $\pi$ be an infinite path through $G_{S}$. We say that $u \in \mathbb{N}^{*}$ is stable on $\pi$ if $u \in D_{\pi(i)}$ for each $i \in \mathbb{N}$ and none of the vertices used in the rewritings on $\pi$ is a prefix of $u$. This simply means that $u$ is never involved in any of the substitutions on the path $\pi$ (note that $u$ is a prefix of itself). The $\operatorname{limit} \lim (\pi)$ of $\pi$ is a (possibly infinite) tree with domain

$$
D_{\lim (\pi)}=\left\{u \in \mathbb{N}^{*} \mid \exists j \in \mathbb{N}: u \text { stable on } \pi^{j}\right\} .
$$

If a vertex $u$ is eventually stable on $\pi$ then it eventually has a fixed label $a$. We take this label to be the label of $u$ in $\lim (\pi)$, i.e., $\lim (\pi)(u)=a$.

The path $\pi$ is called stable iff for all $u \notin D_{\lim (\pi)}$ there is a $j \in \mathbb{N}$ such that $u \notin D_{\pi(i)}$ for all $i \geq j$.

If $\lim (\pi)$ is finite for a path $\pi: t \rightarrow T(\mathcal{A})$, then it is not difficult to see that it is possible to visit $T(\mathcal{A})$ infinitely often in the following way. From $t$ one can reach a tree $t^{\prime \prime}$ which has a subtree $t^{\prime}=\left(t^{\prime \prime}\right)^{u} \in P_{L}$. From $t^{\prime}$ we can reach a tree $r$ and from $r$ one can reach $t^{\prime}$ again such that $t^{\prime \prime}(u \leftarrow r) \in T(\mathcal{A})$. This is formalized in the next lemma. The proof is similar to the construction in the proof of Lemma
Lemma 2. If there exists a path $\pi: t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ with $\lim (\pi)$ finite, then there are $t^{\prime} \in P_{L}, t^{\prime \prime} \in T_{A}, u \in D_{t^{\prime \prime}}$, and $q \in Q$ such that $t \rightarrow_{S}^{*} t^{\prime \prime},\left(t^{\prime \prime}\right)^{u}=t^{\prime}$, $t^{\prime} \rightarrow_{S}^{*} T\left(\mathcal{A}_{q}\right) \rightarrow_{S}^{+} t^{\prime}$, and $t^{\prime \prime}(u \leftarrow q) \in T(\mathcal{A})$.

The normal form for paths $\pi^{\prime}$ with $\lim \left(\pi^{\prime}\right)$ infinite is obtained by removing unnecessary rewritings. By unnecessary we mean that these rewritings are not essential for visiting $T(\mathcal{A})$ infinitely often. This results in a new path $\pi$ that is stable such that $\lim (\pi)$ only has one infinite branch. Furthermore we can ensure that there are infinitely many trees on $\pi$ that are accepted by $\mathcal{A}$ such that the accepting runs agree on growing initial segments of the infinite branch of $\lim (\pi)$.

Lemma 3. If there exists a path $\pi^{\prime}: t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ with $\lim \left(\pi^{\prime}\right)$ infinite, then there is a path $\pi: t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})$ with
(i) $\pi$ is stable,
(ii) $\lim (\pi)$ has exactly one infinite branch $\beta$, and
(iii) there is a run $\rho_{\lim }$ of $\mathcal{A}$ on $\lim (\pi)$ with the following property. For each $u \in \beta$ there are infinitely many $i \in \mathbb{N}$ such that there is an accepting run $\rho_{i}$ of $\mathcal{A}$ on $\pi(i)$ that agrees with $\rho_{\lim }$ on $\{v \mid \beta(v) \sqsubseteq u\}$.

Proof. For simplicity we will only speak of runs instead of runs of $\mathcal{A}$.
Let $\beta$ be an infinite branch of $\lim \left(\pi^{\prime}\right)$. We first define a mapping $\rho_{\beta}$ on $\beta$ that will be extended to the run $\rho_{\mathrm{lim}}$. In the definition of $\rho_{\beta}$ we make use of König's Lemma to be able to satisfy property (iii).

For $u \in \beta$ let $\beta_{u}=\{v \in \beta \mid v \sqsubseteq u\}$ be the initial segment of $\beta$ up to $u$ and $V_{u}=\left\{\rho: \beta_{u} \rightarrow Q \mid \exists^{\omega} i: \pi^{\prime}(i)\right.$ is accepted by a run that agrees with $\rho$ on $\left.\beta_{u}\right\}$ ( $\exists^{\omega}$ means "there are infinitely many"). Define a graph $G=(V, E)$ with $V=$ $\left(\bigcup_{u \in \beta} V_{u}\right) \cup\{\emptyset\}$. The edge relation $E$ is defined as follows.
$-\left(\rho, \rho^{\prime}\right) \in E$ if there are $u, u^{\prime} \in \beta$ with $u^{\prime}=u i$ for some $i \in \mathbb{N}$ such that $\rho: \beta_{u} \rightarrow Q, \rho^{\prime}: \beta_{u^{\prime}} \rightarrow Q$, and $\rho, \rho^{\prime}$ agree on $\beta_{u}$.
$-(\emptyset, \rho) \in E$ if $\rho: \beta_{\varepsilon} \rightarrow Q$.
Then $G$ is an infinite tree (in the graph theoretic sense) of finite degree. Therefore, by König's Lemma, there is an infinite path in $G$. On this path growing initial segments of $\beta$ are labelled consistently. This gives us our mapping $\rho_{\beta}$.

Now we will modify the path $\pi^{\prime}$ to obtain the path $\pi$. For each $u \in \beta$ let $i_{u} \in \mathbb{N}$ be minimal such that the successor $s(u)$ of $u$ on $\beta$ is stable on $\left(\pi^{\prime}\right)^{i_{u}}$ and $\pi^{\prime}\left(i_{u}\right)$ is accepted by a run $\rho_{u}$ that agrees with $\rho_{\beta}$ on $\beta_{s(u)}$. To obtain $\pi$ remove
for each $u \in \beta$ all substitutions on $\left(\pi^{\prime}\right)^{i_{u}}$ at a vertex $v$ with $\beta(v)=u$. Then $\pi$ is still an infinite path and $\pi$ is stable. Furthermore $\lim (\pi)$ only contains the infinite branch $\beta$. It remains to define $\rho_{\lim }$ and verify (iii). For each $u \in D_{\lim (\pi)}$ let $\rho_{\lim }(u)=\rho_{\beta(u)}(u)$. Then one can verify that $\rho_{\text {lim }}$ is indeed a run of $\mathcal{A}$ on $\lim (\pi)$. Since $\pi$ is obtained from $\pi^{\prime}$ by removing rewritings we get a natural mapping $\varphi: \mathbb{N} \rightarrow \mathbb{N}$ associating $\pi^{\prime}(i)$ with $\pi(\varphi(i))$. This mapping is inductively defined as $\varphi(0)=0$ and

$$
\varphi(i+1)= \begin{cases}\varphi(i) & \text { if the rewriting generating the edge from } \pi^{\prime}(i) \text { to } \\ \varphi(i)+1 & \pi^{\prime}(i+1) \text { otherwise } .\end{cases}
$$

By the definition of $\pi$ the tree $\pi\left(\varphi\left(i_{u}\right)\right)$ is accepted for each $u \in \beta$ by a run that agrees with $\rho_{\mathrm{lim}}$ on $\{v \mid \beta(v) \sqsubseteq u\}$.

Lemmas 2 and 3 provide us with normal forms for paths infinitely often visiting $T(\mathcal{A})$ (or any other regular set of trees). To decide for $t \in P_{L}$ and $q \in Q$ whether $t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ we construct a finite graph $G$ with vertex set $\operatorname{sub}(P) \times Q \times Q$ and edges labelled with 0 or 1 such that there is a path $\pi$ in normal form with $\pi: t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ iff there is a path through $G$ with infinitely many edges labelled 1. We will motivate the construction of $G$ by the normal form obtained in Lemma 3 but it is not difficult to see that we can also use $G$ to find paths with finite limit.

In the first component of the vertices of $G$ the rewriting steps at vertices on the infinite branch $\beta$ (according to Lemma 3) will be simulated (usually more than one at a time). The second component will keep track of the state labelling $\rho_{\mathrm{lim}}$ of $\beta$. As we have seen in Lemma 3, there are accepting runs for trees on the infinite path $\pi$ that agree with $\rho_{\lim }$ on growing initial segments of $\lim (\pi)$. But it is not guaranteed that these runs completely agree with $\rho_{\mathrm{lim}}$. So, in the third component of the vertices we will simulate the actual accepting runs on trees that are on $\pi$. The property that we will find runs that agree with $\rho_{\mathrm{lim}}$ on growing initial segments allow us to reset the third component to the second one whenever we find an accepting run.

Traversing one edge in $G$ corresponds to going one vertex along the branch $\beta$. If we are at vertex $u$ on $\beta$ and $u i$ is the successor of $u$ on $\beta$, then going from $(t, q, p)$ to $\left(t^{\prime}, q^{\prime}, p^{\prime}\right)$ in $G$ means that $t$ is the tree at $u$ before the first substitution takes place at $u$ and $t^{\prime}$ is the tree at $u i$ after the last substitution at $u$. Thus, one edge of $G$ corresponds to the sequence of substitutions starting at the first substitution at a certain vertex on $\beta$ and ending after the last substitution at this vertex. The edge of $G$ is labelled with 1 if during this sequence of substitutions a tree from $T\left(\mathcal{A}_{q}\right)$ occurs. This is checked using the states in the second and third component of the vertices of $G$.

Formally we define $G=(V, E)$, where $V=\operatorname{sub}(P) \times Q \times Q$ and $E$ contains the following edges.
(1) $(t, q, p) \xrightarrow{0}\left(t^{\prime}, q^{\prime}, p^{\prime}\right) \in E$ if there exist $t_{1}, \ldots, t_{l} \in T_{A}, q_{1}, \ldots, q_{l}, p_{1}, \ldots, p_{l} \in$ $Q, i \in\{1, \ldots, l\}$ such that
$-a\left(t_{1}, \ldots, t_{l}\right) \in \operatorname{sub}(P)$ and $t \rightarrow{ }_{S}^{*} a\left(t_{1}, \ldots, t_{l}\right)$,
$-t^{\prime}=t_{i}, q^{\prime}=q_{i}, p^{\prime}=p_{i}$,
$-\left(q_{1}, \ldots, q_{l}, a, q\right),\left(p_{1}, \ldots, p_{l}, a, p\right) \in \Delta$, and
$-t_{j} \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{p_{j}}\right) \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{q_{j}}\right)$ for all $j \in\{1, \ldots, l\}$ with $j \neq i$.
(2) $(t, q, p) \xrightarrow{1}\left(t^{\prime}, q^{\prime}, p^{\prime}\right) \in E$ if there exist $t_{1}, \ldots, t_{l} \in T_{A}, q_{1}, \ldots, q_{l}, p_{1}, \ldots, p_{l} \in$ $Q, i \in\{1, \ldots, l\}$ such that
$-a\left(t_{1}, \ldots, t_{l}\right) \in \operatorname{sub}(P)$ and $t \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{p}\right) \rightarrow_{S}^{*} a\left(t_{1}, \ldots, t_{l}\right)$,
$-t^{\prime}=t_{i}, q^{\prime}=q_{i}, p^{\prime}=p_{i}$,
$-\left(q_{1}, \ldots, q_{l}, a, q\right),\left(p_{1}, \ldots, p_{l}, a, q\right) \in \Delta$, and
$-t_{j} \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{p_{j}}\right) \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{q_{j}}\right)$ for all $j \in\{1, \ldots, l\}$ with $j \neq i$.
The only differences between (1) and (2) are that in (2) we require $a\left(t_{1}, \ldots, t_{l}\right)$ to be reachable from $t$ via a tree from $T\left(\mathcal{A}_{p}\right)$ and that we reset the simulation of a run in the third component.

Lemma 4. Let $t \in P_{L}$ and $q \in Q$.
(i) There is a path $\pi$ in $G_{S}$ with $\pi: t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ and $\lim (\pi)$ infinite iff there is an infinite path in $G$ with infinitely many edges labelled 1 starting from $(t, q, q)$.
(ii) There is a path $\pi$ in $G_{S}$ with $\pi: t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ and $\lim (\pi)$ finite iff there is a path in $G$ from $(t, q, q)$ to $\left(t^{\prime}, q^{\prime}, q^{\prime}\right)$ such that $t^{\prime} \rightarrow_{S}^{*} T\left(\mathcal{A}_{q^{\prime}}\right) \rightarrow_{S}^{+} t^{\prime}$.

Proof. The proof for (ii) is rather simple with the use of Lemma 2 The idea for the proof of (i) is as follows.
" $\Leftarrow$ ": By induction on $n$ one can show that if there is a path from $(t, q, p)$ to ( $t^{\prime}, q^{\prime}, p^{\prime}$ ) in $G$ starting with $n 0$-edges and ending with one 1-edge, then there is a tree $t^{\prime \prime}$ and $u \in D_{t^{\prime \prime}}$ such that $t \rightarrow{ }_{S}^{*} T\left(\mathcal{A}_{p}\right) \rightarrow_{S}^{*} t^{\prime \prime},\left(t^{\prime \prime}\right)^{u}=t^{\prime}$, and $t^{\prime \prime}\left(u \leftarrow q^{\prime}\right) \in T\left(\mathcal{A}_{q}\right)$. So, from an infinite path in $G$ starting in $(t, q, q)$ and having infinitely many 1-edges one can construct a path in $G_{S}$ that infinitely often visits $T\left(\mathcal{A}_{q}\right)$.
$" \Rightarrow$ ": For this direction the normal form of Lemma 3 is used. The first component in the vertices of $G$ is used to simulate the rewritings on the branch $\beta$ at the point where the vertices on $\beta$ become stable. The second component guesses the labelling of $\beta$ by $\rho_{\text {lim }}$ and the third component is used to simulate the labelling of $\beta$ by accepting runs on trees that are on $\pi$. Property (iii) from Lemma 3 ensures that there are accepting runs that agree with $\rho_{\text {lim }}$ on growing initial segments of $\lim (\pi)$.

The graph $G$ can be constructed effectively since the only nontrivial conditions in the construction of the edges are instances of the reachability problem. Furthermore condition (i) of the previous lemma is decidable by standard algorithms on finite graphs. Condition (ii) of the previous lemma is decidable because the condition $t^{\prime} \rightarrow_{S}^{*} T\left(\mathcal{A}_{q^{\prime}}\right) \rightarrow_{S}^{+} t^{\prime}$ can also be formulated as an instance of the reachability problem. Therefore we get the following lemma.

Lemma 5. For $t \in P_{L}$ and $q \in Q$ it is decidable whether $t \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$.

Summarizing the results from this section we get the following theorem.
Theorem 2. Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and an $N T A \mathcal{A}=(Q, A, \Delta, F)$, one can construct an $\varepsilon$-NTA with $\mathcal{O}\left(|Q|+\left|\operatorname{sub}\left(P_{L}\right)\right|\right)$ states accepting the set $\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})\right\}$.

Proof. An automaton for $R(\mathcal{A})$ on an input $t \in T_{A}$ has to guess a subtree $t^{u} \in P_{L}$ and $q \in Q$ with $t^{u} \rightarrow{ }_{S}^{\omega} T\left(\mathcal{A}_{q}\right)$ and then verify that $t(u \leftarrow q)$ is in $T(\mathcal{A})$. This can be done by an automaton with $\mathcal{O}\left(|Q|+\left|\operatorname{sub}\left(P_{L}\right)\right|\right)$ states (and by Lemma 5 we can also construct this automaton). Then we apply Theorem 1 and obtain an automaton for $\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\omega} T(\mathcal{A})\right\}$ with $\mathcal{O}\left(|Q|+\left|\operatorname{sub}\left(P_{L}\right)\right|\right)$ states.

## 5 Universal Reachability and Universal Recurrence

In the previous sections we asked for the existence of a path with a certain property. In this section we address the dual problems.
(Universal Reachability): Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and a regular set of trees $T \subseteq T_{A}$, does every maximal path in $G_{S}$ starting in $t_{I}$ visit $T$ ?
(Universal Recurrence): Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and a regular set of trees $T \subseteq T_{A}$, does every infinite path in $G_{S}$ starting in $t_{I}$ visit $T$ infinitely often?

We will show the undecidability of these two problems using the same ideas as in [9] where the undecidability of a similar property for Basic Parallel Processes is shown. The general idea is to use reductions from undecidable problems for Turing machines (TM) by defining for a given TM $M$ a GTRS $S_{M}$ that can simulate computations of $M$. A TM configuration $a_{1} \cdots a_{k} q b_{l} \cdots b_{1}$, where the word $a_{1} \cdots a_{k} b_{l} \cdots b_{1}$ is on the tape and the TM is in state $q$ with the reading head on $b_{l}$, will be represented by a tree with two branches with $X, a_{1}, \ldots, a_{k}$ on the left branch and $X, b_{1}, \ldots, b_{l}, q$ on the right branch. The $X$ symbols are used to model the infinite blank parts of the tape to the left and to the right. For example the first tree in Figure 3 represents the configuration qa. For a configuration $\kappa$ of $M$ we will denote the corresponding tree as $t(\kappa)$.

With this coding of configurations it is not possible to exactly simulate the transitions of $M$ by $S_{M}$. The GTRS $S_{M}$ can simulate the correct behavior of $M$ as well as incorrect behavior (i.e., in $S_{M}$ one can reach $t\left(\kappa^{\prime}\right)$ from $t(\kappa)$ although $M$ cannot reach $\kappa^{\prime}$ when started in $\kappa$ ). But $S_{M}$ will be constructed in such a way that for every path $\pi$ through $G_{S_{M}}$ that leads from $t(\kappa)$ to $t\left(\kappa^{\prime}\right)$ there is a tree from a regular set $T_{\text {err }}$ on $\pi$ iff $M$ cannot reach $\kappa^{\prime}$ when started $\kappa$. This property is obtained by adding auxiliary symbols to the tree alphabet and simulating a transition of $M$ in more than one rewriting step with intermediate trees that do not code configurations.

Let $M=\left(Q, B, \Gamma, q_{0}, q_{s}, \delta\right)$ be a deterministic TM with state set $Q$, input alphabet $B$, tape alphabet $\Gamma$, starting state $q_{0}$, stop state $q_{s}$, and transition function $\delta: Q \times \Gamma \rightarrow Q \times \Gamma \times\{L, R\}$. Furthermore let $Q \cap \Gamma=\emptyset$, and $\sqcup$ be the blank symbol. If $\kappa$ and $\kappa^{\prime}$ are configurations of $M$, then we write as usual $\kappa \vdash_{M} \kappa^{\prime}$
if $\kappa^{\prime}$ is the successor configuration of $\kappa$, and $\kappa \vdash_{M}^{*} \kappa^{\prime}$ if $\kappa^{\prime}$ is a configuration that is reachable from $\kappa$. The initial configuration of $M$ on the empty tape is denoted by $\kappa_{\varepsilon}$.

The GTRS $S_{M}=\left(A, \Sigma, P, t_{I}\right)$ is defined as follows. The tree alphabet is $A=A_{0} \cup A_{1} \cup A_{2}$ with $A_{2}=\{\bullet\}, A_{1}=Q \cup \Gamma \cup\{X\}$, and $A_{0}=A_{1} \cup$ $\left\{a d d_{1}, a d d_{2}, a d d_{3}\right.$, rem $_{1}$, rem $\left._{2}, e r r\right\}$. We do not need the transition labels $\Sigma$, thus we assume that $\Sigma$ contains one symbol and omit this symbol in the rest of the proof.

The initial tree is $t_{I}=t\left(\kappa_{\varepsilon}\right)=\begin{aligned} \bullet & \begin{array}{l}X \\ q_{0} \\ q_{0}\end{array} .\end{aligned}$
The set $P$ contains the following rewriting rules.



4. $a d d_{1} \rightarrow a d d_{2}, a d d_{2} \rightarrow a d d_{3}$, rem $_{1} \rightarrow$ rem $_{2}$.



Figure 3shows the correct simulation of the transition $\delta(q, a)=(p, b, L)$ (the TM is in the configuration $q a$ ). The symbols $a$ and $q$ are replaced by $b$ and $p$. Since the TM moves the head to the left, it has to guess which symbol is on the left hand side of the head. Here it is the blank symbol. The symbol $a d d_{1}$ indicates that the blank symbol should be added to the right branch of the tree. Now the left branch has to confirm with rem $_{1}$. Then the two branches alternatingly increase their $a d d$ and rem symbols until they reach $a d d_{3}$ and rem 2 . Now the blank symbol can be removed from the left branch and then it is added to the right branch.


Fig. 3. Example for the simulation of a TM-transition by a GTRS.

The set $T_{\text {err }}$ contains all trees that result from a violation of the protocol and all the trees that encode a stop configuration of the TM. From the description below it is easy to see that $T_{\text {err }}$ can be defined by an NTA. A tree $t$ is in $T_{\text {err }}$ iff

1. $t$ contains the err symbol,
2. $t$ contains more than one $a d d$ or more than one rem symbol,
3. $t$ contains a remi symbol and no $a d d_{i}$ or $a d d_{i+1}$ symbol,
4. $t$ contains an $a d d_{2}$ symbol and no rem symbol, or
5. $t$ contains subtrees of the form $\underset{\text { rem }_{i}}{\stackrel{a}{\mid}}$ and $\underset{a d d_{j}}{b}$, with $a \neq b$.

By a tedious distinction of cases one can check the following lemma.
Lemma 6. (i) For each configuration $\kappa$ of $M: t\left(\kappa_{\varepsilon}\right) \rightarrow_{S}^{*} t(\kappa)$.
(ii) For all configurations $\kappa, \kappa^{\prime}$ of $M$ there is a path from $t(\kappa)$ to $t\left(\kappa^{\prime}\right)$ in $G_{S_{M}}$ not visiting $T_{\text {err }}$ iff $\kappa \vdash_{M}^{*} \kappa^{\prime}$.

Theorem 3. The problem (Universal Reachability) is undecidable.
Proof. Let $T_{\text {stop }}$ be the set of all trees encoding a stop or deadlocking configuration. Then, by Lemma (ii), it is clear that a TM $M$ does not stop on the empty tape iff there is a path through $G_{S_{M}}$ starting in $t_{I}$ that never visits $T_{\text {err }} \cup T_{\text {stop }}$.

Theorem 4. The problem (Universal Recurrence) is undecidable.
Proof. We use a reduction from a problem similar to the halting problem which can easily be shown to be undecidable:

Given a Turing machine $M$, does there exist a configuration $\kappa$ of $M$ such that $M$ does not stop when started in $\kappa$ ?

Suppose that such a configuration of $M$ exists. then $t\left(\kappa_{\varepsilon}\right) \rightarrow_{S}^{*} t(\kappa)$ by Lemma 6 (i), and by Lemma 6 (ii) there is an infinite path starting in $t(\kappa)$ not visiting $T_{\text {err }}$.

If there is an infinite path $\pi$ that visits $T_{\text {err }}$ only finitely often, then one can pick a tree of the form $t(\kappa)$ on $\pi$ such that there is no tree from $T_{\text {err }}$ on the suffix of $\pi$ starting in $t(\kappa)$ (it is not difficult to see that such a tree exists). Then the suffix of $\pi$ starting in $t(\kappa)$ corresponds to an infinite computation of $M$ starting in $\kappa$, again by Lemma 6 (ii).

## 6 A Logic for Model-Checking over GTRS Graphs

For a fixed ranked alphabet $A$ and an alphabet $\Sigma$ formulas of our logic are defined by the grammar (in CTL-like syntax)

$$
\phi:=\top|\perp| T|\neg \phi| \phi \vee \phi|\langle\sigma\rangle \phi| E F \phi \mid E G F \phi
$$

where $\sigma \in \Sigma$ and $T \subseteq T_{A}$ regular. Given a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ define the semantics $\|\phi\|_{S}$ (or simply $\|\phi\|$ ) of a formula $\phi$ as follows.

$$
\begin{aligned}
& -\|\top\|=T_{A},\|\perp\|=\emptyset,\|T\|=T,\|\neg \phi\|=T_{A} \backslash\|\phi\|,\left\|\phi_{1} \vee \phi_{2}\right\|=\left\|\phi_{1}\right\| \cup\left\|\phi_{2}\right\|, \\
& -\|\langle\sigma\rangle \phi\|=\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\sigma}\|\phi\|\right\}, \\
& -\|E F \phi\|=\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{*}\|\phi\|\right\}, \text { and } \\
& -\|E G F \phi\|=\left\{t \in T_{A} \mid t \rightarrow{ }_{S}^{\omega}\|\phi\|\right\} .
\end{aligned}
$$

Then $S$ is a model of $\phi$, denoted by $S \models \phi$, iff $t_{I} \in\|\phi\|$. With the results from the previous sections we get the following theorem.

Theorem 5. For a GTRS $S=\left(A, \Sigma, P, t_{I}\right)$ and a formula $\phi$ one can decide whether $S \models \phi$.

Proof. We can construct automata for $\|\top\|,\|\perp\|,\|T\|,\|\neg \phi\|,\left\|\phi_{1} \vee \phi\right\|$ with standard tree automaton constructions. Given an automaton for $\|\phi\|$, the automata for $\|E F \phi\|$ and $\|E G F \phi\|$ can be constructed according to Theorems 1 and 2 , An automaton for $\|\langle\sigma\rangle \phi\|$ has to guess a subtree and a rewriting rule such that rewriting this subtree with this rule yields a tree from $\|\phi\|$. We do not give the details of this straightforward construction here. Thus, to decide if $S \models \phi$ we construct the automaton for $\|\phi\|$ and then check if $t_{I} \in\|\phi\|$.

The fragment of temporal logic presented here is maximal in the sense that including the operators $E G$ or $E F G$ leads to an undecidable model-checking problem as shown by the results from the previous section. Furthermore, untiloperators (as known from temporal logic) also lead to undecidability, because in the GTRS constructed in the proof of Theorem 3 the fact that the TM stops on the empty tape can easily expressed by the until formula $E\left(\neg T_{\text {err }}\right) U T_{\text {stop }}$ because the existence of a path starting in the initial tree and remaining in the complement of $T_{\text {err }}$ until it eventually reaches $T_{\text {stop }}$ is equivalent to the existence of a halting computation of the TM on the empty tape.

## Conclusion

In the present paper we have analyzed various decision problems for infinite graphs generated by ground tree rewriting. The problems of reachability and recurrence as stated in Sections 3 and 4 were solved whereas the problems of universal reachability and universal recurrence from Section 5 were shown to be undecidable. These considerations lead to a rather expressive and in some sense maximal fragment of temporal logic with a decidable model-checking problem for GTRS graphs.

A drawback of the decision algorithm is that the size of the constructed automaton is non-elementary in the number of nested negations in the formula since complementing NTA gives an exponential blow-up. Note that we cannot simply push the negations inwards to the atomic formulas because this would lead to universal path quantifiers ( $A$ instead of $E$ ) which are difficult to handle with nondeterministic tree automata. The use of alternating tree automata may help to obtain a better complexity.

## References

1. Achim Blumensath and Erich Grädel. Automatic structures. In Proceedings of LICS '00, pages 51-62. IEEE Computer Society Press, 2000.
2. Walter S. Brainerd. Tree generating regular systems. Information and Control, 14:217-231, 1969.
3. Didier Caucal. On infinite transition graphs having a decidable monadic theory. In Proceedings ICALP '96, volume 1099 of LNCS. Springer-Verlag, 1996.
4. J.L. Coquidé, M. Dauchet, R. Gilleron, and S. Vágvölgyi. Bottom-up tree pushdown automata: Classification and connection with rewrite systems. Theoretical Computer Science, 127(1):69-98, 1994.
5. J.L. Coquidé and R. Gilleron. Proofs and reachability problem for ground rewrite systems. In Aspects and Prospects of Theoretical Computer Science, volume 464 of $L N C S$, pages 120-129. Springer, 1990.
6. Max Dauchet, Thierry Heuillard, Pierre Lescanne, and Sophie Tison. Decidability of the confluence of finite ground term rewrite systems and of other related term rewrite systems. Information and Computation, 88(2):187-201, October 1990.
7. Max Dauchet and Sophie Tison. The theory of ground rewrite systems is decidable. In Proceedings LICS '90, pages 242-248. IEEE Computer Society Press, 1990.
8. Javier Esparza, David Hansel, Peter Rossmanith, and Stefan Schwoon. Efficient algorithms for model checking pushdown systems. In Proceedings of $C A V$ 2000, volume 1855 of $L N C S$, pages 232-247. Springer-Verlag, 2000.
9. Javier Esparza and Astrid Kiehn. On the model checking problem for branching time logics and Basic Parallel Processes. In Proceedings of CAV '95, volume 939 of $L N C S$, pages 353-366, 1995.
10. Orna Kupferman and Moshe Y. Vardi. An automata-theoretic approach to reasoning about infinite-state systems. In E. A. Emerson and A. P. Sistla, editors, Proceedings of CAV 2000, volume 1855 of LNCS. Springer-Verlag, 2000.
11. Richard Mayr. Process rewrite systems. Information and Computation, 156(1-2):264-286, 2000.
12. Christophe Morvan. On rational graphs. In Proceedings of FoSSaCS '99, volume 1784 of $L N C S$, pages 252-266. Springer, 1999.
13. David E. Muller and Paul E. Schupp. The theory of ends, pushdown automata, and second-order logic. Theoretical Computer Science, 37:51-75, 1985.
14. Igor Walukiewicz. Pushdown processes: Games and model checking. In Rajeev Alur and Thomas A. Henzinger, editors, Proceedings of CAV '96, volume 1102 of $L N C S$, pages 62-74. Springer-Verlag, July-August 1996.
