Simulation-Based Iteration of Tree Transducers

Parosh Aziz Abdulla[1], Axel Legay[2*], Julien d'Orso[3], and Ahmed Rezine[1]

[1] Dept. of Information Technology
Uppsala University
P.O. Box 337
SE-751 05 Uppsala, Sweden
{parosh,rahmed}@it.uu.se
[2] Université de Liège
Institut Montefiore, B28
4000 Liège, Belgium
legay@montefiore.ulg.ac.be
[3] IRCCyN - UMR CNRS 6597
BP 92 101
44321 Nantes CEDEX 03, France
julien.dorso@irccyn.ec-nantes.fr

Abstract. *Regular model checking* is the name of a family of techniques for analyzing infinite-state systems in which states are represented by words, sets of states by finite automata, and transitions by finite-state transducers. The central problem is to compute the transitive closure of a transducer. A main obstacle is that the set of reachable states is in general not regular. Recently, regular model checking has been extended to systems with tree-like architectures. In this paper, we provide a procedure, based on a new implementable acceleration technique, for computing the transitive closure of a tree transducer. The procedure consists of incrementally adding new transitions while merging states which are related according to a pre-defined equivalence relation. The equivalence is induced by a *downward* and an *upward* simulation relation which can be efficiently computed. Our technique can also be used to compute the set of reachable states without computing the transitive closure. We have implemented and applied our technique to several protocols.

1 Introduction

Regular model checking is the name of a family of techniques for analyzing infinite-state systems in which states are represented by words, sets of states by finite automata, and transitions by finite automata operating on pairs of states, i.e. finite-state transducers. The central problem in regular model checking is to compute the transitive closure of a finite-state transducer. Such a representation allows to compute the set of reachable states of the system (which is useful to verify safety properties) and to detect loops between states (which is useful to verify

* This author is supported by a F.R.I.A grant.

liveness properties). However, computing the transitive closure is in general undecidable; consequently any method for solving the problem is necessarily incomplete. One of the goals of regular model checking is to provide semi-algorithms which terminate on many practical applications. Such semi-algorithms have already been successfully applied to parameterized systems with linear topologies, and to systems that operate on linear unbounded data structures such as queues, stacks, integers, reals, and hybrid systems [BJNT00,BLW03,BLW04,DLS01,BHV04].

This work aims at extending the paradigm of regular model checking to verify systems which operate on tree-like architectures. This includes several interesting protocols such as the percolate protocol ([KMM$^+$01]) or the Tree-arbiter protocol ([ABH$^+$97]).

To verify such systems, we use the extension of regular model checking called *tree regular model checking*, which was first introduced in [AJMd02], [KMM$^+$01], and [BT02]. In tree regular model checking, states of the systems are represented by trees, sets of states by tree automata, and transitions by tree automata operating on pairs of trees, i.e. tree transducers. As in the case of regular model checking, the central problem is to provide semi-algorithms for computing the transitive closure of a tree transducer. This problem was considered in [AJMd02,BT02]; however the proposed algorithms are most of the time inefficient or non-implementable.

In this work, we provide an efficient and implementable semi-algorithm for computing the transitive closure of a tree transducer. Starting from a tree transducer D, describing the set of transitions of the system, we derive a transducer, called the *history transducer* whose states are *columns* (words) of states of D. The history transducer characterizes the transitive closure of the rewriting relation corresponding to D. The set of states of the history transducer is infinite which makes it inappropriate for computational purposes. Therefore, we present a method for computing a finite-state transducer which is an abstraction of the history transducer. The abstract transducer is generated on-the-fly by a procedure which starts from the original transducer D, and then incrementally adds new transitions and merges equivalent states. To compute the abstract transducer, we define an equivalence relation on columns (states of the history transducer). We identify *good* equivalence relations, i.e., equivalence relations which can be used by our on-the-fly algorithm. An equivalence relation is considered to be *good* if it satisfies the following two conditions:

- *Soundness and completeness:* merging two equivalent columns must not add any traces which are not present in the history transducer. Consequently, the abstract transducer accepts the same language as the history transducer (and therefore characterizes exactly the transitive closure of D).
- *Computability of the equivalence relation:* This allows on-the-fly merging of equivalent states during the generation of the abstract transducer.

We present a methodology for deriving good equivalence relations. More precisely, an equivalence relation is induced by two simulation relations; namely a *downward* and an *upward* simulation relation, both of which are defined on tree automata. We provide sufficient conditions on the simulation relations which guarantee that the induced equivalence is good. Furthermore, we give examples of concrete simulations which satisfy the sufficient conditions. These simulations can be computed by efficient algorithms derived from those of Henzinger *et al.* ([HHK95]) for finite words.

We show that our technique can be directly adapted in order to compute the set of reachable states of a system without computing its entire transitive closure. When checking for safety properties, such an approach is often (but not always) more efficient.

We have implemented our algorithms in a tool which we have applied to a number of protocols including a Two-Way Token protocol, the Percolate Protocol ([KMM$^+$01]), a parametrized version of the Tree-arbiter protocol ([ABH$^+$97]), and a tree-parametrized version of a leader election protocol.

Related Work: There are several works on efficient computation of transitive closures for *word* transducers [BJNT00,DLS01,BLW03,BLW04,AJNd03,BHV04]. However, all current algorithms devoted to the computation of the transitive closure of a *tree* transducer are not efficient or not implementable. In [AJMd02], we presented a method for computing transitive closures of tree transducers. The method presented in [AJMd02] is very heavy and relies on several layers of expensive automata-theoretic constructions. The method of this paper is much more light-weight and efficient, and can therefore be applied to a larger class of protocols. The work in [BT02] also considers tree transducers, but it is based on *widening* rather than acceleration. The idea is to compute successive applications of the transducer relation, and detect *increments* in the produced transducers. Based on the detected increments, the method makes a guess of the transitive closure. One of the main advantages of this work is that it also allows non-structure preserving transformations (in contrast to the method provided in this paper). However, the widening procedure in [BT02] is not implemented. Furthermore, no efficient method is provided to detect the increments. This indicates that any potential implementation of the widening technique would be inefficient. In [DLS01], Dams, Lakhnech, and Steffen present an extension of the word case to trees. However, this is done for top-down tree automata which are not closed under determinization (and thus many other operations). In [DLS01], the authors consider several definitions of simulations and bisimulations between top-down tree automata without providing methods for computing them. Hence, it is not clear how to implement their algorithms.

2 Tree automata

In this section, we introduce some preliminaries on trees and tree automata (more details can be found in [CDG$^+$99]).

A *ranked alphabet* is a pair (Σ, ρ), where Σ is a finite set of symbols and ρ is a mapping from Σ to $\mathbb{N}$. For a symbol $f \in \Sigma$, we call $\rho(f)$ the *arity* of f. We let Σ_p denote the set of symbols in Σ with arity p. Intuitively, each node in a tree is labeled with a symbol in Σ with the same arity as the out-degree of the node.

Sometimes, we abuse notation and use Σ to denote the ranked alphabet (Σ, ρ).

Following [CDG$^+$99], the nodes in a tree are represented by words over $\mathbb{N}$. More precisely, the empty word ϵ represents the root of the tree, while a node $b_1 b_2 ... b_k$ is a child of the node $b_1 b_2 ... b_{k-1}$. Also, nodes are labeled by symbols from Σ.

Definition 1. [**Trees**]
A tree T over a ranked alphabet Σ is a pair (S, λ), where

- *S, called the tree structure, is a finite set of sequences over $\mathbb{N}$ (i.e, a finite subset of $\mathbb{N}^*$). Each sequence n in S is called a node of T. If S contains a node $n = b_1 b_2 ... b_k$, then S will also contain the node $n' = b_1 b_2 ... b_{k-1}$, and the nodes $n_r = b_1 b_2 ... b_{k-1} r$, for $r : 0 \leq r < b_k$. We say that n' is the parent of n, and that n is a child of n'. A leaf of T is a node n which does not have any child, i.e., there is no $b \in \mathbb{N}$ with $nb \in S$.*
- *λ is a mapping from S to Σ. The number of children of n is equal to $\rho(\lambda(n))$. Observe that if n is a leaf then $\lambda(n) \in \Sigma_0$.*

We use $T(\Sigma)$ to denote the set of all trees over Σ.

Sets of trees are recognized using tree automata. There exist various kinds of tree automata. In this paper, we use bottom-up tree automata since they are closed under all operations needed by the classical model checking procedure: intersection, union, minimization, determinization, inclusion test, complementation, etc. In the sequel, we will omit the term bottom-up.

Definition 2. [**Tree Automata and Languages**]
A tree language is a set of trees.
A tree automaton [CDG$^+$99,Tho90] over a ranked alphabet Σ is a tuple $A = (Q, F, \delta)$, where Q is a set of states, $F \subseteq Q$ is a set of final states, and δ is the transition relation, represented by a set of rules each of the form

$$(q_1, \ldots, q_p) \xrightarrow{f} q$$

where $f \in \Sigma_p$ and $q_1, \ldots, q_p, q \in Q$. Unless stated otherwise, we assume Q and δ to be finite.

We say that A is *deterministic* when δ does not contain two rules of the form $(q_1, \ldots, q_p) \xrightarrow{f} q$ and $(q_1, \ldots, q_p) \xrightarrow{f} q'$ with $q \neq q'$.

Intuitively, the automaton A takes a tree $T \in T(\Sigma)$ as input. It proceeds from the leaves to the root (that explains why it is called bottom-up), annotating states to the nodes of T. A transition rule of the form shown above tells us that if the children of a node n are already annotated from left to right with $q_1, \ldots, q_p$ respectively, and if $\lambda(n) = f$, then the node n can be annotated by q. As a special case, a transition rule of the form $\xrightarrow{f} q$ implies that a leaf labeled with $f \in \Sigma_0$ can be annotated by q.

Formally, a *run* r of A on a tree $T = (S, \lambda) \in T(\Sigma)$ is a mapping from S to Q such that for each node $n \in T$ with children $n_1, \ldots, n_k$ we have $\left((r(n_1), \ldots, r(n_k)) \xrightarrow{\lambda(n)} r(n) \right) \in \delta$.

For a state q, we let $T \overset{r}{\Longrightarrow}_A q$ denote that r is a run of A on T such that $r(\epsilon) = q$. We use $T \Longrightarrow_A q$ denote that $T \overset{r}{\Longrightarrow}_A q$ for some r. For a set $S \subseteq Q$ of states, we let $T \overset{r}{\Longrightarrow}_A S$ ($T \Longrightarrow_A S$) denote that $T \overset{r}{\Longrightarrow}_A q$ ($T \Longrightarrow_A q$) for some $q \in S$. We say that A *accepts* T if $T \Longrightarrow_A F$. We define $L(A) = \{T |\, T \text{ is accepted by } A\}$. A tree language K is said to be *regular* if there is a tree automaton A such that $K = L(A)$.

We now define the notion of context. Intuitively, a context is a tree with "holes" instead of leaves. Formally, we consider a special symbol $\square \notin \Sigma$ with arity 0. A *context* over Σ is a tree (S_C, λ_C) over $\Sigma \cup \{\square\}$ such that for all leaves $n_c \in S_C$, we have $\lambda_C(n_c) = \square$. For a context $C = (S_C, \lambda_C)$ with holes at leaves $n_1, \ldots, n_k \in S_C$, and trees $T_1 = (S_1, \lambda_1), \ldots, T_k = (S_k, \lambda_k)$, we define $C[T_1, \ldots, T_k]$ to be the tree (S, λ), where

- $S = S_C \cup \bigcup_{i \in \{1,\ldots,k\}} \{n_i \cdot n' |\, n' \in S_i\}$
- for each $n = n_i \cdot n'$ with $n' \in S_i$ for some $1 \leq i \leq k$, we have $\lambda(n) = \lambda_i(n')$.
- for each $n \in S_C - \{n_1, \ldots, n_k\}$, we have $\lambda(n) = \lambda_C(n)$.

Intuitively, $C[T_1, \ldots, T_k]$ is the result of appending the trees $T_1, \ldots, T_k$ to the holes of C. Consider a tree automaton $A = (Q, F, \delta)$ over a ranked alphabet Σ. We extend the notion of runs to contexts. Let $C = (S_C, \lambda_C)$ be a context with leaves $n_1, \ldots, n_k$. A *run* r of A on C from $(q_1, \ldots, q_k)$ is defined in a similar manner to a run except that for leaf n_i, we have $r(n_i) = q_i$. In other words, each leaf labeled with $\square$ is annotated by one q_i. We use $C[q_1, \ldots, q_k] \overset{r}{\Longrightarrow}_A q$ to denote that r is a run of A on C from $(q_1, \ldots, q_k)$ such that $r(\epsilon) = q$. The notation $C[q_1, \ldots, q_k] \Longrightarrow_A q$ and its extension to sets of states are explained in a similar manner to runs on trees.

Definition 3. *For an automaton $A = (Q, F, \delta)$, we define the suffix of a tuple of states $(q_1, \ldots, q_n)$ to be* $\mathrm{suff}(q_1, \ldots, q_n) = \{C : \textit{context} |\, C[q_1, \ldots, q_n] \Longrightarrow_A F\}$. *For a state $q \in Q$, its prefix is the set of trees* $\mathrm{pref}(q) = \{T : \textit{tree} |\, T \Longrightarrow_A q\}$.

Remark Our definition of a context coincides with the one of [BT03] where all leaves are holes. On the other hand, a context in [CDG+99] and [AJMd02] is a tree with a *single* hole.

3 Tree Relations and Transducers

In this section we introduce tree relations and transducers.

For a binary relation R, we use R^+ to denote the transitive closure of R.

For a ranked alphabet Σ and $m \geq 1$, we let $\Sigma^\bullet(m)$ be the ranked alphabet which contains all tuples $(f_1, \ldots, f_m)$ such that $f_1, \ldots, f_m \in \Sigma_p$ for some p. We define $\rho((f_1, \ldots, f_m)) = \rho(f_1)$. In other words, the set $\Sigma^\bullet(m)$ contains the m-tuples, where all the elements in the same tuple have equal arities. Furthermore, the arity of a tuple in $\Sigma^\bullet(m)$ is equal to the arity of any of its elements. For trees $T_1 = (S_1, \lambda_1)$ and $T_2 = (S_2, \lambda_2)$, we say that T_1 and T_2 are *structurally equivalent*, denoted $T_1 \cong T_2$, if $S_1 = S_2$.

Consider structurally equivalent trees $T_1, \ldots, T_m$ over an alphabet Σ, where $T_i = (S, \lambda_i)$ for $i : 1 \leq i \leq m$. We let $T_1 \times \cdots \times T_m$ be the tree $T = (S, \lambda)$ over $\Sigma^\bullet(m)$ such that $\lambda(n) = (\lambda_1(n), \ldots, \lambda_m(n))$ for each $n \in S$. An m-ary *relation* on the alphabet Σ is a set of tuples of the form $(T_1, \ldots, T_m)$, where $T_1, \ldots, T_m \in T(\Sigma)$ and $T_1 \cong \cdots \cong T_m$. A tree language K over $\Sigma^\bullet(m)$ characterizes an m-ary tree relation $[K]$ on $T(\Sigma)$ as follows: $(T_1, \ldots, T_m) \in [K]$ iff $T_1 \times \cdots \times T_m \in K$.

We use tree automata also to characterize tree relations: An automaton A over $\Sigma^\bullet(m)$ characterizes an m-ary relation on $T(\Sigma)$, namely the relation $[L(A)]$. A tree relation is said to be *regular* if it is equal to $[L(A)]$, for some tree automaton A. In such as case, we denote this relation by $R(A)$.

Definition 4. [**Tree Transducers**]
In the special case where D is a tree automaton over $\Sigma^\bullet(2)$, we call D a tree transducer over Σ.

Remark Our definition of tree transducers is a restricted version of the one considered in [BT02] in the sense that we only consider transducers that do not modify the structure of the tree. In [BT02], such transducers are called relabeling transducers.

4 Tree Regular Model Checking

We use the following framework known as *tree regular model checking* [AJMd02,BT02,KMM$^+$01]:

Definition 5. *A* program *is a triple $P = (\Sigma, \phi_I, D)$ where*

- *Σ is a ranked alphabet.*
- *ϕ_I is a set of initial configurations represented by a tree automaton over Σ.*
- *D is a transducer over Σ characterizing a transition relation $R(D)$.*

In a similar manner to the the case of words (see [BJNT00]), the problems we are going to consider are the following:

- *Computing the transitive closure of D:* The goal is to compute a new tree transducer D^+ representing the transitive closure of D, i.e., $R(D^+) = (R(D))^+$. Such a representation can be used for computing the reachability set of the program or for finding cycles between reachable program configurations.

- *Computing the reachable states:* The goal is to compute a tree automaton representing $R(D^+)(\phi_I)$. This set can be used for checking safety properties of the program.

We will first provide a technique for computing D^+. Afterwards, we will show the modifications needed for computing $R(D^+)(\phi_I)$ without computing D^+.

5 Computing the Transitive Closure

In this section we introduce the notion of *history transducer*. With a transducer D we associate a *history transducer* which corresponds to the transitive closure of D. Each state of H is a word of the form $q_1 \cdots q_k$ where $q_1, \ldots, q_k$ are states of D. For a word w, we let $w(i)$ denote the i-th symbol of w. Intuitively, for each $(T, T') \in D^+$, the history transducer H encodes the successive runs of D needed to derive T' from T. The term "history transducer" reflects the fact that the transducer encodes the histories of all such derivations.

Definition 6. *Consider a tree transducer $D = (Q, F, \delta)$ over a ranked alphabet Σ. The history (tree) transducer H for D is an (infinite) transducer (Q_H, F_H, δ_H), where $Q_H = Q^+$, $F_H = F^+$, and δ_H contains all rules of the form*

$$(w_1, \ldots, w_p) \xrightarrow{(f,f')} w$$

such that there is $k \geq 1$ where the following conditions are satisfied

- $|w_1| = \cdots = |w_p| = |w| = k$.
- *there are $f_1, f_2, \ldots, f_{k+1}$, with $f = f_1$, $f' = f_{k+1}$, and* $(w_1(i) \ldots, w_p(i)) \xrightarrow{(f_i, f_{i+1})} w(i)$ *belongs to δ, for each $i : 1 \leq i \leq k$.*

Observe that all the symbols $f_1, \ldots, f_{k+1}$ are of the same arity p. Also, notice that if $(T \times T') \stackrel{r}{\Longrightarrow}_H w$, then there is a $k \geq 1$ such that $|r(n)| = k$ for each $n \in (T \times T')$. In other words, any run of the history transducer assigns states (words) of the same length to the nodes. From the definition of H we derive the following lemma (proved in [AJMd02]) which states that H characterizes the transitive closure of the relation of D.

Lemma 1. *For a transducer D and its history transducer H, we have that $R(H) = R(D^+)$.*

The problem with H is that it has infinitely many states. Therefore, we define an *equivalence* $\simeq$ on the states of H, and construct a new transducer where equivalent states are merged. This new transducer will hopefully only have a finite number of states.

Given an equivalence relation $\simeq$, the symbolic transducer $D_\simeq$ obtained by merging states of H according to $\simeq$ is defined as $(Q/\simeq, F/\simeq, \delta_\simeq)$, where:

- $Q/\simeq$ is the set of equivalence classes of Q_H w.r.t. $\simeq$;

- $F/\simeq$ is the set of equivalence classes of F_H w.r.t. $\simeq$;
- $\delta_{\simeq}$ contains rules of the form $(x_1, \ldots, x_n) \xrightarrow{f}_{\simeq} x$ iff there are states $q_1 \in x_1, \ldots, q_n \in x_n, q \in x$ such that there is a rule $(q_1, \ldots, q_n) \xrightarrow{f} q$ of H.

Since H is infinite we cannot derive $D_{\simeq}$ by first computing H. Instead, we compute $D_{\simeq}$ on-the-fly collapsing states which are equivalent according to $\simeq$. In other words, we performing the following *procedure* (which need not terminate in general).

- The procedure computes successive reflexive powers of D: $D^{\leq 1}, D^{\leq 2}, D^{\leq 3}, \ldots$ (where $D^{\leq i} = \bigcup_{n=1}^{n=i} D^n$), and collapses states[4] according to $\simeq$. We thus obtain $D_{\simeq}^{\leq 1}, D_{\simeq}^{\leq 2}, \ldots$
- The procedure terminates when the relation R^+ is accepted by $D_{\simeq}^{\leq i}$. This can be tested by checking if the language $D_{\simeq}^{\leq i} \circ D$ is included in $D_{\simeq}^{\leq i}$.

6 Soundness, Completeness, and Computability

In this section, we describe how to derive equivalence relations on the states of the history transducer which can be used in the procedure given in Section 5. A *good* equivalence relation $\simeq$ satisfies the following two conditions:

- It is sound and complete, i.e., $R(D_{\simeq}) = R(H)$. This means that $D_{\simeq}$ characterizes the same relation as D^+.
- It is computable. This turns the procedure of Section 5 into an *implementable algorithm*, since it allows on-the-fly merging of equivalent states.

We provide a methodology for deriving good equivalence relations as follows: we define two simulation relations; namely a downward simulation relation $\preccurlyeq_{down}$ and an *upward simulation relation* $\preccurlyeq_{up}$, which together induce an equivalence relation $\simeq$. Then, we give sufficient conditions of the simulation relations which guarantee that the induced equivalence $\simeq$ is a good one.

6.1 Downward and Upward Simulation

We start by giving the definitions.

Definition 7 (Downward Simulation). *Let $A = (Q, F, \delta)$ be a tree automaton. A binary relation $\preccurlyeq_{down}$ is a downward simulation iff for any $n \geq 1$ and any symbol $f \in \Sigma_n$, for all states $q, q_1, \ldots, q_n, r$, the following holds:*
Whenever $q \preccurlyeq_{down} r$ and $(q_1, \ldots, q_n) \xrightarrow{f} q$, then there exist states $r_1, \ldots, r_n$ such that $q_1 \preccurlyeq_{down} r_1, \ldots, q_n \preccurlyeq_{down} r_n$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$.

[4] The states of $D^{\leq i}$ are by construction states of the history transducer.

Definition 8 (Upward Simulation). *Let $A = (Q, F, \delta)$ be a tree automaton. Given a downward simulation $\preccurlyeq_{down}$, a binary relation $\preccurlyeq_{up}$ is an upward simulation w.r.t. $\preccurlyeq_{down}$ iff for any $n \geq 1$ and any symbol $f \in \Sigma_n$, for all states $q, q_1, \ldots, q_i, \ldots, q_n, r_i \in Q$, the following holds:*
Whenever $q_i \preccurlyeq_{up} r_i$ and $(q_1, \ldots, q_n) \xrightarrow{f} q$, then there exist states $r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n, r \in Q$ such that $q \preccurlyeq_{up} r$ and $\forall j \neq i : q_j \preccurlyeq_{down} r_j$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$.

While the notion of a downward simulation is a straightforward extension of the word case, the notion of an upward simulation is not as obvious. This comes from the asymmetric nature of trees. If we follow the execution of a tree automaton downwards, it is easy to see that all respective children of two nodes related by simulation should continue to be related pairwise. If we now consider how a tree automaton executes when going upwards, we are confronted to the problem that the parent of the current node may have several children. The question is then how to characterize the behavior of such children. The answer lies in constraining their prefixes, i.e. using a downward simulation.

We state some elementary properties of the simulation relations

Lemma 2. *A downward simulation $\preccurlyeq_{down}$ can be closed under reflexivity and transitivity. Furthermore, there is a unique maximal downward simulation.*

Lemma 3. *Let $\preccurlyeq_{down}$ be a reflexive (transitive) downward simulation. An upward simulation $\preccurlyeq_{up}$ w.r.t. $\preccurlyeq_{down}$ can be closed under reflexivity (transitivity). Let $\preccurlyeq_{down}$ be a downward simulation. There exists a unique maximal upward simulation w.r.t. $\preccurlyeq_{down}$.*

Observe that both for downward simulations, and upward simulations, maximality implies transitivity and reflexivity.

We now define an equivalence relation derived from two simulation relations.

Definition 9. *Two binary relations $\preceq_1$ and $\preceq_2$ are said to be* independent *iff whenever $q \preceq_1 r$ and $q \preceq_2 r'$, there exists s such that $r \preceq_2 s$ and $r' \preceq_1 s$.*

Definition 10. *The relation $\simeq$* induced *by two binary relations $\preceq_1$ and $\preceq_2$ is defined as:*

$$\preceq_1 \circ \preceq_2^{-1} \cap \preceq_2 \circ \preceq_1^{-1}$$

The following Lemma gives sufficient conditions for two relations to induce an equivalence relation.

Lemma 4. *Let $\preceq_1$ and $\preceq_2$ be two binary relations. If $\preceq_1$ and $\preceq_2$ are reflexive, transitive, and independent, then their induced relation $\simeq$ is an* equivalence relation.

6.2 Sufficient Conditions for Soundness and Completeness

We give sufficient conditions for the two simulation relations to induce a sound and complete equivalence relation on states of a tree automaton.

We assume a tree automaton $A = (Q, F, \delta)$. We now define a relation $\simeq$ induced by the two relations $\preceq$ and $\preccurlyeq_{down}$ satisfying the conditions:

1. $\preccurlyeq_{down}$ is a downward simulation;
2. $\preceq$ is a reflexive and transitive relation included in $\preccurlyeq_{up}$ which is an upward simulation w.r.t. $\preccurlyeq_{down}$;
3. $\preccurlyeq_{down}$ and $\preceq$ are independent;
4. whenever $x \in F$ and $x \preccurlyeq_{up} y$, then $y \in F$;
5. F is a union of equivalence classes w.r.t. $\simeq$;
6. whenever $\xrightarrow{f} x$ and $x \preccurlyeq_{down} y$, then $\xrightarrow{f} y$;

$\square$

We first obtain the following Lemma which shows that if the simulations satisfy the sufficient conditions, then the induced relation is indeed an equivalence.

Lemma 5. *Let $A = (Q, F, \delta)$ be a tree automaton. Consider two binary relations $\preccurlyeq_{down}$ and $\preceq$ which satisfies the above sufficient conditions, as well as their induced relation $\simeq$. We have that $\simeq$ is an equivalence relation on states of A.*

The above Lemma holds since Conditions 1 through 3 imply directly that $\preccurlyeq_{down}$ and $\preceq$ satisfy the premises needed by Lemma 4.

Next, we state that such an equivalence relation is sound and precise.

Theorem 1. *Let $A = (Q, F, \delta)$ be a tree automaton. Consider two binary relations $\preccurlyeq_{down}$ and $\preceq$ satisfying the above sufficient conditions, and let $\simeq$ be their induced relation. Let $A_{\simeq} = (Q/\simeq, F/\simeq, \delta_{\simeq})$ be the automaton obtained by merging the states of A according to $\simeq$. Then, $L(A_{\simeq}) = L(A)$.*

Theorem 1 can be used to relate the languages of H and $D_{\simeq}$.

We are now ready to prove the soundness and the completeness of our on-the-fly algorithm (assuming a computable equivalence relation $\simeq$).

Theorem 2. *Consider two binary relations on the states of H $\preccurlyeq_{down}$ and $\preceq$, satisfying the hypothesis of Theorem 1. Let $\simeq$ be their induced equivalence relation. If the algorithm terminates at step i, then the transducer $D_{\simeq}^{\leq i}$ accepts the same relation as $D_{\simeq}$.*

6.3 Sufficient Condition for Computability

The next step is to give conditions on the simulations which ensure that the induced equivalence relation is computable.

Definition 11. [**Effective relation:**] *A relation $\preceq$ is said to be effective if the image of a regular set w.r.t. $\preceq$ and w.r.t. $\preceq^{-1}$ is regular and computable*

Effective relations induce an equivalence relation which is also computable.

Theorem 3. *Let $\preceq_1$ and $\preceq_2$ be both reflexive, transitive, effective and independent. Let $\simeq$ be their induced equivalence. Then for any state x of H, we can compute its equivalence class $[x]$ w.r.t. $\simeq$.*

The Theorem follows by definition of $\simeq$, and effectiveness [5] of $\preceq_1$ and $\preceq_2$. $\square$

An equivalence relation that satisfies hypothesis of Theorem 1 and Theorem 3 can be used in the on-the-fly algorithm of Section 5 to compute the transitive closure of a tree transducer. The next step is to provide a concrete example of such an equivalence. Because we are not able to compute the infinite representation of H, the equivalence will be directly computed from the powers of D provided by the on-the-fly algorithm.

7 Good Equivalence Relation

In this section, we provide concrete relations satisfying Theorem 1 and Theorem 3. We first introduce prefix- and suffix-copying states.

Definition 12. Prefix-copying state: *Given a transducer D, and a state q, we say that q is a prefix-copying state if for any tree $T = (S, \lambda) \in \mathrm{pref}(q)$, then for any node $n \in S$, $\lambda(n) = (f, f)$ for some symbol $f \in \Sigma$.*

Definition 13. Suffix-copying state: *Given a transducer D, and a state q, we say that q is a suffix-copying state if for any context $C = (S_C, \lambda_C) \in \mathrm{suff}(q)$, then for any node $n \in S_C$ with $\lambda_C(n) \neq \square$, we have $\lambda_C(n) = (f, f)$ for some symbol $f \in \Sigma$.*

We let Q_{pref} (resp. Q_{suff}) denote the set of prefix-copying states (resp. the set of suffix-copying states) of D and we assume that $Q_{pref} \cap Q_{suff} = \emptyset$. We let $Q_N = Q - Q_{pref} \cup Q_{suff}$.

We now define relations by the means of rewriting relation on the states of the history transducer.

Definition 14. Generated relation: *Given a set S of pairs of states of H, we define the relation $\mapsto$* generated *by S to be the smallest reflexive and transitive relation such that $\mapsto$ contains S, and $\mapsto$ is a congruence w.r.t. concatenation (i.e. if $x \mapsto y$, then for any w_1, w_2, we have $w_1 \cdot x \cdot w_2 \mapsto w_1 \cdot y \cdot w_2$).*

Next, we find relations $\preceq$ and $\preccurlyeq_{down}$ that satisfy the sufficient conditions for computability (Theorem 3) and conditions for exactness of abstraction (Lemma 6.2).

Definition 15. Simulation relations

- *We define $\preccurlyeq_{down}$ to be the downward simulation generated by all pairs of the form $(q_{pref} \cdot q_{pref}, q_{pref})$ and $(q_{pref}, q_{pref} \cdot q_{pref})$, where $q_{pref} \in Q_{pref}$.*

[5] A state x of the history transducer is a word. The set $\{x\}$ is regular.

– *Let* $\preccurlyeq^1_{up}$ *be the maximal upward simulation computed on* $D \cup D^2$*. Then, we define* $\preceq$ *to be the relation generated by the maximal set* $S \subseteq \preccurlyeq^1_{up}$ *such that*
  • $(q_{suff} \cdot q_{suff}, q_{suff}) \in S$ *iff* $(q_{suff}, q_{suff} \cdot q_{suff}) \in S$
  • $(q \cdot q_{suff}, q) \in S$ *iff* $(q, q \cdot q_{suff}) \in S$
  • $(q_{suff} \cdot q, q) \in S$ *iff* $(q, q_{suff} \cdot q) \in S$

  where $q_{suff} \in Q_{suff}$*, and* $q \in Q_N$*.*

Let us state that the simulations of Definition 15 satisfy the hypothesis needed by Theorems 1 and 3.

Lemma 6. *The following properties of* $\preccurlyeq_{down}$ *hold:*

1. $\preccurlyeq_{down}$ *is a downward simulation;*
2. $\preccurlyeq_{down}$ *is effective.*

Lemma 7. *The following properties of* $\preceq$ *hold:*

1. $\preceq$ *is included in an upward simulation;*
2. $\preceq$ *is effective.*

We now state that $\preceq$ and $\preccurlyeq_{down}$ are independent.

Lemma 8. $\preceq$ *and* $\preccurlyeq_{down}$ *are independent.*

Lemma 9. *The following holds:*

– *whenever* $x \in F_H$ *and* $x \preccurlyeq_{up} y$*, then* $y \in F_H$*;*
– F_H *is a union of equivalence classes w.r.t.* $\simeq$*;*
– *whenever* $\xrightarrow{f} x$ *and* $x \preccurlyeq_{down} y$*, then* $\xrightarrow{f} y$*;*

We conclude that $\preceq$ and $\preccurlyeq_{down}$ satisfy the hypothesis of Theorem 1 and Theorem 3 and can thus be used by the on-the-fly procedure presented in Section 5.

8 Computing Reachable Configurations

We now sketch the modifications needed to compute $R(D^+)(\phi_I)$ without computing D^+. When checking for safety properties, such a computation is known to be sufficient (see [VW86]). Computing $R(D^+)(\phi_I)$ rather than D^+, can simply be done by lightly modifying the definition of the history transducer. Assume that we have constructed a tree automaton A_{ϕ_I} for ϕ_I, we replace the transducer run in the first "row" of the history transducer by a transducer that only accept trees from A_{ϕ_I} in input. Such a transducer can easily by constructed. Let D be the transducer representing the transition of the system, the restricted transducer is obtained by taking the intersection between D and $A_{\phi_I} \times T(\Sigma)$ where Σ is the ranked alphabet of the system. Computing $R(D^+)(\phi_I)$ is often less expensive than computing D^+ because it only considers reachable sets of states (see Section 9 for a time comparison). We have an example for which our technique can compute $R(D^+)(\phi_I)$ but cannot compute D^+.

9 Experimental Results

The techniques presented in this paper have been applied on several case studies using a prototype implementation that relies in part on the regular model checking tool (see www.regularmodelchecking.com).

In Table 1 we report the result of running our implementation on a number of parametrized protocols for which we have computed the set of reachable states as well as the transitive closure of their transition relation. A full description of the protocols is given in the Appendix.

Relation	$\|R\|$	$\|R^+\|$	max size	time $\|R^+\|$	$\|R^+(\phi_I)\|$	max size	time $\|R^+(\phi_I)\|$
Simple Token Protocol	3	4	15	1,6s	3	17	0.75s
Two-Way Token Protocol	4	6	28	6.2s	3	26	1.01s
Percolate Protocol	4	6	40	16.76s	3	21	1.87s
Tree-arbiter Protocol	8	-	-	-	10	246	4mn18s
Leader Election Protocol	6	9	105	3mn29s	10	150	6mn2s

Table 1. Results

In our previous work [AJMd02], we were able to handle the first three protocols of the table. Moreover, for those protocols, we were only able to compute the transitive closure for individual *actions* representing one class of statements in the protocol, sometimes with manual intervention. Here we compute automatically the transitive closure of the tree transducer representing the *entire* transition relation of the protocol which is necessary to check for repeated reachability.

In order to compute the set of reachable states, we have used the technique presented in Section 8. In [AJMd02,BT02], the reachability computation was done by first computing the transitive closure for each individual action, and then applying a classical forward reachability algorithm using these results. However, such an approach requires manual intervention: to make the reachability analysis terminate, it is often necessary to combine actions in a certain order, or even to accelerate combinations of individual actions.

Observe that we are not able to compute the transitive closure of the transition relation of the tree-arbiter protocol. However, we are already able to compute transitive closure of individual actions for this protocol as well as the reachable set of states with the technique of Section 8.

10 Conclusions and Future Work

In this paper, we have presented a technique for computing the transitive closure of a tree transducer.

This technique has been implemented and successfully tested on a number of protocols, several of which are beyond the capabilities of existing tree regular model checking techniques.

We believe that substantial efficiency improvement can be achieved by considering more general equivalence relations than the one defined in Section 7, and by refining our algorithms for computing simulation relations.

The restriction to structure-preserving tree transducers might be seen as a weakness of our approach. However, structure-preserving tree transducers can model the relation of many interesting parametrized network protocols. In the future, we plan to investigate the case of non structure-preserving tree transducers. One possible solution would be to use *padding* to simulate a structure-preserving behaviour. This would allow us to extend our method to work on such systems as Process Rewrite Systems (PRS). PRS are useful when modeling systems with a dynamic behavior [BT03,KRS04].

It would also be interesting to see if one can extend our simulations, as well as the algorithms for computing them, in order to efficiently implement the technique presented in [BT02] (the detection of an increment can be done by isolating part of the automaton with the help of (bi)simulations).

Finally, we intend to extend our framework to check for liveness properties on tree-like architecture systems (as done for words in [AJN+04]).

References

[ABH+97] R. Alur, R.K. Brayton, T.A. Henzinger, S. Qadeer, and S.K. Rajamani. Partial-order reduction in symbolic state space exploration. In O. Grumberg, editor, *Proc. 9^{th} Int. Conf. on Computer Aided Verification*, volume 1254, pages 340–351, Haifa, Israel, 1997. Springer Verlag.

[AJMd02] Parosh Aziz Abdulla, Bengt Jonsson, Pritha Mahata, and Julien d'Orso. Regular tree model checking. In *Proc. 14^{th} Int. Conf. on Computer Aided Verification*, volume 2404 of *Lecture Notes in Computer Science*, 2002.

[AJN+04] P. A. Abdulla, B. Jonsson, M. Nilsson, J. d'Orso, and M. Saksena. Regular model checking for s1s + ltl. In *CAV04*, Lecture Notes in Computer Science, Boston, July 2004. Springer-Verlag.

[AJNd03] Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, and Julien d'Orso. Algorithmic improvements in regular model checking. In *Proc. 15^{th} Int. Conf. on Computer Aided Verification*, volume 2725 of *Lecture Notes in Computer Science*, pages 236–248, 2003.

[BHV04] A. Bouajjani, P. Habermehl, and T. Vojnar. Abstract regular model checking. In *CAV04*, Lecture Notes in Computer Science, Boston, July 2004. Springer-Verlag.

[BJNT00] A. Bouajjani, B. Jonsson, M. Nilsson, and T. Touili. Regular model checking. In Emerson and Sistla, editors, *Proc. 12^{th} Int. Conf. on Computer Aided Verification*, volume 1855 of *Lecture Notes in Computer Science*, pages 403–418. Springer Verlag, 2000.

[BLW03] Bernard Boigelot, Axel Legay, and Pierre Wolper. Iterating transducers in the large. In *Proc. 15^{th} Int. Conf. on Computer Aided Verification*, volume 2725 of *Lecture Notes in Computer Science*, pages 223–235, 2003.

[BLW04] Bernard Boigelot, Axel Legay, and Pierre Wolper. Omega regular model checking. In *Proc. TACAS '04, 10^{th} Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science, 2004.

[BT02] Ahmed Bouajjani and Tayssir Touili. Extrapolating Tree Transformations. In *Proc. 14^{th} Int. Conf. on Computer Aided Verification*, volume 2404 of *Lecture Notes in Computer Science*, 2002.

[BT03] A. Bouajjani and T. Touili. Reachability analysis of process rewrite systems. In *Proc. Int. Conf. on Foundations of Software Technology and Theoritical Computer Science (FSTTCS'03)*, Lecture Notes in Computer Science, 2003.

[CDG+99] H. Common, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. *Tree Automata Techniques and Applications.* not yet published, October 1999.

[DLS01] D. Dams, Y. Lakhnech, and M. Steffen. Iterating transducers. In G. Berry, H. Comon, and A. Finkel, editors, *Computer Aided Verification*, volume 2102 of *Lecture Notes in Computer Science*, 2001.

[HHK95] M. Henzinger, T. Henzinger, and P. Kopke. Computing simulations on finite and infinite graphs. In *Proc. 36^{th} Annual Symp. Foundations of Computer Science*, pages 453–463, 1995.

[KMM+01] Y. Kesten, O. Maler, M. Marcus, A. Pnueli, and E. Shahar. Symbolic model checking with rich assertional languages. *Theoretical Computer Science*, 256:93–112, 2001.

[KRS04] P. Kretinsky, V. Rehak, and J. Strejcek. Extended process rewrite systems: Expressiveness and reachability. In *CONCUR04*, volume 3170 of *Lecture Notes in Computer Science*, pages 355–370, London, august 2004. Springer-Verlag.

[Tho90] W. Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science, Volume B: Formal Methods and Semantics*, pages 133–192, 1990.

[VW86] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. LICS '86, 1^{st} IEEE Int. Symp. on Logic in Computer Science*, pages 332–344, June 1986.

A An example for Section 2

In this section, we illustrate the tree regular model checking framework introduced in Section 4.

We consider the simple token tree protocol. Roughly speaking, the protocol consists of processes that are connected in a binary tree-like fashion. Each process stores a single bit which reflects whether the process has a token or not. The token tree passes a token from a leaf to the root. The system is represented by a tree transducer over an alphabet consisting of $t, n \in \Sigma_0$ representing processes at the leaves, and $N, T \in \Sigma_2$ representing processes at the inner nodes of the tree. Processes labeled by $\{n, N\}$ are those which do not have a token, while those labeled by $\{t, T\}$ are those which do have a token.

A configuration of the system for seven processes can be given by the following tree (denoted by α) represented by the pair (S, λ) and defined over the alphabet $\Sigma \cup \Sigma_2$ where:

- $\Sigma \cup \Sigma_2 = \{n, t, N, T\}$,
- $S = \{\epsilon, 0, 1, 00, 01, 10, 11\}$,
- λ is defined by :
  - $\lambda(00) = \lambda(01) = \lambda(10) = n \in \Sigma_0$,
  - $\lambda(11) = t \in \Sigma_0$,
  - $\lambda(\epsilon) = \lambda(1) = \lambda(2) = N \in \Sigma_2$.

Graphically, we have

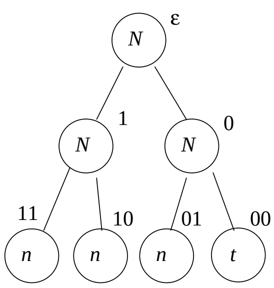

The set of initial configurations can be encoded by a tree automaton which recognizes all possible configurations in which one leaf has the token, i.e. $A = (Q, F, \delta)$, where:

– $\Sigma = \{n, t, N, T\}$,
– $Q = \{q_0, q_1\}$,
– $F = \{q_1\}$,

and δ is defined by the following set of rules:

$$\xrightarrow{n} q_0 \qquad \xrightarrow{t} q_1 \qquad (q_0, q_0) \xrightarrow{T} q_1$$
$$(q_0, q_0) \xrightarrow{N} q_0 \qquad (q_0, q_1) \xrightarrow{N} q_1 \qquad (q_1, q_0) \xrightarrow{N} q_1$$

The simple token passing protocol consists in passing, the token from the leaf having the token to the root. The transducer D that models the behavior of the protocol is given by :

$$\begin{array}{llllll} & & \xrightarrow{(n,n)} q_0 & & \xrightarrow{(t,n)} q_1 & \\ (q_0,q_0) \xrightarrow{(T,N)} q_1 & (q_0,q_1) \xrightarrow{(N,T)} q_2 & (q_1,q_0) \xrightarrow{(N,T)} q_2 \\ (q_0,q_0) \xrightarrow{(N,N)} q_0 & (q_0,q_2) \xrightarrow{(N,N)} q_2 & (q_2,q_0) \xrightarrow{(N,N)} q_2 \end{array}$$

Where the states correspond to the following :
q_0 : the node is idle, i.e., the token is not in the node, nor in the subtree below the node;
q_1 : the node is releasing the token to the node above it in the tree;
q_2 : the token is either in the node or in a subtree below the node (accepting state).

Graphically, one and two applications of D on α give:

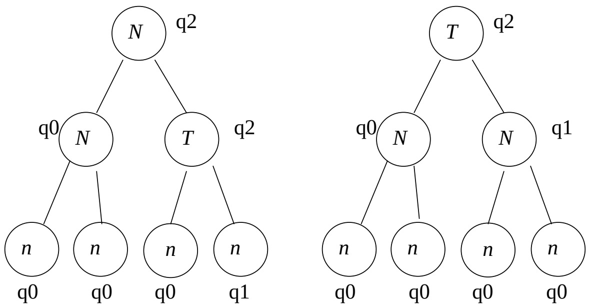

B Proofs of Some Lemmas

Lemma 2

Proof. We consider the three points.

Reflexivity
Let $\preccurlyeq^1_{down} = \preccurlyeq_{down} \cup Id$. We show that $\preccurlyeq^1_{down}$ is also a downward simulation. Assume $(q_1, \ldots, q_n) \xrightarrow{f} q$ and $q \preccurlyeq^1_{down} r$. We find states $r_1, \ldots, r_n$ such that $q_1 \preccurlyeq^1_{down} r_1, \ldots, q_n \preccurlyeq^1_{down} r_n$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q = r$. Then we choose $r_1 = q_1, \ldots, r_n = q_n$. Observe that since $\preccurlyeq^1_{down} \supseteq Id$, we have $q_1 \preccurlyeq^1_{down}, r_1, \ldots, q_n \preccurlyeq^1_{down} r_n$. Thus, the claim holds.
- **Case 2:** $q \preccurlyeq_{down} r$. Then we apply the hypothesis that $\preccurlyeq_{down}$ is a downward simulation, and conclude that there exist $r_1, \ldots, r_n$ such that $(r_1, \ldots, r_n) \xrightarrow{f}$

r and $q_1 \preccurlyeq_{down} r_1, \ldots, q_n \preccurlyeq_{down} r_n$. From $\preccurlyeq^1_{down} \supseteq \preccurlyeq_{down}$, we conclude that the claim holds.

Transitivity

Let $\preccurlyeq^1_{down}$ be the transitive closure of $\preccurlyeq_{down}$. We show that $\preccurlyeq^1_{down}$ is also a downward simulation. Assume $(q_1, \ldots, q_n) \xrightarrow{f} q$ and $q \preccurlyeq^1_{down} r$. We find states $r_1, \ldots, r_n$ such that $q_1 \preccurlyeq^1_{down} r_1, \ldots, q_n \preccurlyeq^1_{down} r_n$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q \preccurlyeq_{down} r$. Then the claim trivially holds.
- **Case 2:** there is s such that $q \preccurlyeq_{down} s \preccurlyeq_{down} r$. We apply the hypothesis that $\preccurlyeq_{down}$ is a downward simulation with $q \preccurlyeq_{down} s$, and find states $s_1, \ldots, s_n$ with $(s_1, \ldots, s_n) \xrightarrow{f} s$ and $q_1 \preccurlyeq_{down} s_1, \ldots, q_n \preccurlyeq_{down} s_n$. Now, we apply this a second step using $s \preccurlyeq_{down} r$, and find states $r_1, \ldots, r_n$ such that $(r_1, \ldots, r_n) \xrightarrow{f} r$ and $s_1 \preccurlyeq_{down} r_1, \ldots, s_n \preccurlyeq_{down} r_n$. By transitivity, we get $q_1 \preccurlyeq^1_{down} r_1, \ldots, q_n \preccurlyeq^1_{down} r_n$. Thus, the claim holds.

Observe that in case 2 above, we only treat the case of one step transitivity, but by induction on the number of steps, we can get arbitrary transitivity.

Uniqueness

Assume two maximal downward simulations $\preccurlyeq^1_{down}$ and $\preccurlyeq^2_{down}$.
Let $\preccurlyeq_{down} = \preccurlyeq^1_{down} \cup \preccurlyeq^2_{down}$. We show that $\preccurlyeq_{down}$ is also a simulation.

Assume $(q_1, \ldots, q_n) \xrightarrow{f} q$ and $q \preccurlyeq_{down} r$. We find states $r_1, \ldots, r_n$ such that $q_1 \preccurlyeq_{down} r_1, \ldots, q_n \preccurlyeq_{down} r_n$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q \preccurlyeq^1_{down} r$. Since $\preccurlyeq^1_{down}$ is a simulation, and $\preccurlyeq_{down} \supseteq \preccurlyeq^1_{down}$, the claim holds.
- **Case 2:** $q \preccurlyeq^2_{down} r$. Since $\preccurlyeq^2_{down}$ is a simulation, and $\preccurlyeq_{down} \supseteq \preccurlyeq^2_{down}$, the claim holds.

We have $\preccurlyeq_{down} \supseteq \preccurlyeq^1_{down}$ and $\preccurlyeq_{down} \supseteq \preccurlyeq^2_{down}$. Now, if we assume $\preccurlyeq^1_{down} \neq \preccurlyeq^2_{down}$, we get either $\preccurlyeq_{down} \supset \preccurlyeq^1_{down}$ or $\preccurlyeq_{down} \supset \preccurlyeq^2_{down}$. This violates the maximality of either $\preccurlyeq^1_{down}$ or $\preccurlyeq^2_{down}$. □

Lemma 3

Proof. We consider the three points.

Reflexivity

Let $\preccurlyeq^1_{up} = \preccurlyeq_{up} \cup Id$. We show that $\preccurlyeq^1_{up}$ is also an upward simulation.

Assume $(q_1, \ldots, q_i, \ldots, q_n) \xrightarrow{f} q$ and $q_i \preccurlyeq^1_{up} r_i$. We find states $r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n, r$ such that $q \preccurlyeq^1_{up} r$ and $\forall j \neq i : q_j \preccurlyeq_{down} r_j$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q_i \preccurlyeq_{up} r_i$. Since $\preccurlyeq_{up}$ is an upward simulation, and $\preccurlyeq_{up}^1 \supseteq \preccurlyeq_{up}$, the claim trivially holds.
- **Case 2:** $q_i = r_i$. Then we choose $r = q$ and $r_j = q_j$ for each $j \neq i$. By reflexivity of $\preccurlyeq_{down}$, we have $q_j \preccurlyeq_{down} r_j$ for each $j \neq i$. Since $\preccurlyeq_{up}^1 \supset Id$, we also have $q \preccurlyeq_{up}^1 r$. Hence, the claim holds.

□

Transitivity

Let $\preccurlyeq_{up}^1$ be the transitive closure of $\preccurlyeq_{up}$. We show that $\preccurlyeq_{up}^1$ is also an upward simulation.

Assume $(q_1, \ldots, q_i, \ldots, q_n) \xrightarrow{f} q$ and $q_i \preccurlyeq_{up}^1 r_i$.

We find states $r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n, r$ such that $q \preccurlyeq_{up}^1 r$ and $\forall j \neq i : q_j \preccurlyeq_{down} r_j$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q_i \preccurlyeq_{up} r_i$. Since $\preccurlyeq_{up}$ is an upward simulation, and $\preccurlyeq_{up}^1 \supseteq \preccurlyeq_{up}$, the claim trivially holds.
- **Case 2:** there is s_i with $q_i \preccurlyeq_{up} s_i \preccurlyeq_{up} r_i$. Since $q_i \preccurlyeq_{up} s_i$ we apply the hypothesis that $\preccurlyeq_{up}$ is an upward simulation. We get states $s, s_1, \ldots, s_n$ with $(s_1, \ldots, s_i, \ldots, s_n) \xrightarrow{f} s$ and $q \preccurlyeq_{up} s$ and for each $j \neq i$, $q_j \preccurlyeq_{down} s_j$. With $s_i \preccurlyeq_{up} r_i$, we use simulation a second time, and get states $r, r_1, \ldots, r_n$ with $(r_1, \ldots, r_i, \ldots, r_n) \xrightarrow{f} r$ and $s \preccurlyeq_{up} r$ and for each $j \neq i$, $s_j \preccurlyeq_{down} r_j$. By transitivity of $\preccurlyeq_{down}$, we get for each $j \neq i$, $q_j \preccurlyeq_{down} r_j$. By transitivity of $\preccurlyeq_{up}^1$, we also get $q \preccurlyeq_{up}^1 r$. Hence, the claim holds.

Observe that in case 2 above, we only treat the case of one step transitivity, but by induction on the number of steps, we can get arbitrary transitivity.

Uniqueness

Assume two maximal upward simulations $\preccurlyeq_{up}^1$ and $\preccurlyeq_{up}^2$.

Let $\preccurlyeq_{up} = \preccurlyeq_{up}^1 \cup \preccurlyeq_{up}^2$. We show that $\preccurlyeq_{up}$ is also a simulation.

Assume $(q_1, \ldots, q_i, \ldots, q_n) \xrightarrow{f} q$ and $q_i \preccurlyeq_{up} r_i$.

We find states $r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n, r$ such that $q \preccurlyeq_{up} r$ and $\forall j \neq i : q_j \preccurlyeq_{down} r_j$ and $(r_1, \ldots, r_n) \xrightarrow{f} r$ as follows:

- **Case 1:** $q_i \preccurlyeq_{up}^1 r_i$. Since $\preccurlyeq_{up}^1$ is a simulation, and $\preccurlyeq_{up} \supseteq \preccurlyeq_{up}^1$, the claim holds.
- **Case 2:** $q_i \preccurlyeq_{up}^2 r_i$. Since $\preccurlyeq_{up}^2$ is a simulation, and $\preccurlyeq_{up} \supseteq \preccurlyeq_{up}^2$, the claim holds.

We have $\preccurlyeq_{up} \supseteq \preccurlyeq_{up}^1$ and $\preccurlyeq_{up} \supseteq \preccurlyeq_{up}^2$. Now, if we assume $\preccurlyeq_{up}^1 \neq \preccurlyeq_{up}^2$, we get either $\preccurlyeq_{up} \supset \preccurlyeq_{up}^1$ or $\preccurlyeq_{up} \supset \preccurlyeq_{up}^2$. This violates the maximality of either $\preccurlyeq_{up}^1$ or $\preccurlyeq_{up}^2$. □

Lemma 4

Proof. 1. If $\preceq_1$ and $\preceq_2$ are reflexive, then for any q we have $q \preceq_1 q \preceq_2^{-1} q$ and $q \preceq_2 q \preceq_1^{-1} q$. Thus, we have $q \simeq q$.
2. If $q \simeq r$, then $q \preceq_1 \circ \preceq_2^{-1} r$ and $q \preceq_2 \circ \preceq_1^{-1} r$. We can rewrite this $r \preceq_1 \circ \preceq_2^{-1} q$ and $r \preceq_2 \circ \preceq_1^{-1} q$. Hence, $r \simeq q$.

3. Assume $q \simeq r \simeq s$. Then by definition of $\simeq$, we can find t', t'' such that $q \preceq_1 t' \preceq_2^{-1} r$ and $r \preceq_1 t'' \preceq_2^{-1} s$. Since $\preceq_1$ and $\preceq_2$ are independent, there is t such that $t' \preceq_1 t$ and $t'' \preceq_2 t$. By transitivity, we get $q \preceq_1 t \preceq_2^{-1} s$. Hence, $q \preceq_1 \circ \preceq_2^{-1} s$. Similarly, we can get $q \preceq_2 \circ \preceq_1^{-1} s$, and finally conclude $q \simeq s$ □

Theorem 1

We need a property of $\simeq$ which follows from independence. This is induced by two lemmas.

Lemma 10. *Let $\preceq_1$ and $\preceq_2$ be both reflexive and transitive. Furthermore, $\preceq_1$ and $\preceq_2$ are independent. Whenever $x \simeq y$ and $x \preceq_1 z$, then there exists t such that $y \preceq_1 t$ and $z \preceq_2 t$.*

Proof. Assume $x \simeq y$ and $x \preceq_1 z$. By definition of $\simeq$, we know that there is u with $x \preceq_2 u$ and $y \preceq_1 u$. We apply the definition of independence to x, u, z, and conclude that there is a state t such that $z \preceq_2 t$ and $u \preceq_1 t$. By transitivity of $\preceq_1$, we have $y \preceq_1 t$. □

We first show that the tree automaton $A_\simeq$ has the same traces as A.

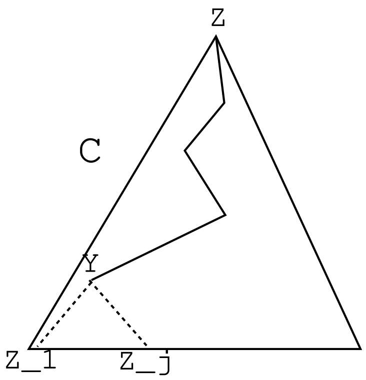

Fig. 1. Lemma 11: a path within context C.

Lemma 11. *For any states $Z_1, \ldots, Z_k, Z$ of $A_\simeq$ and context C, if $C[Z_1, \ldots, Z_k] \Longrightarrow_\simeq Z$, then there exist states $z_1, \ldots, z_k, z$ and states $t_1, \ldots, t_k, t$ of A such that $C[t_1, \ldots, t_k] \Longrightarrow t$ and $z_1 \in Z_1, \ldots, z_k \in Z_k, z \in Z$ and $z_1 \preccurlyeq_{down} t_1, \ldots, z_k \preccurlyeq_{down} t_k, z \preccurlyeq_{up} t$.*

Proof. We show the claim by induction on the structure of C.

Base case: C contains only a hole. We choose a $z \in Z$. By reflexivity of $\preccurlyeq_{down}$ and $\preccurlyeq_{up}$, the claim obviously holds.

Induction case: Consider a run r of $A_{\simeq}$ on $C = (S_C, \lambda_C)$ satisfying $C[Z_1, \ldots, Z_k] \Longrightarrow_{\simeq} Z$. Let $n_1, \ldots, n_j$ be the left-most leaves of C with a common parent. Let n be the parent of $n_1, \ldots, n_j$. Let $Z_1 = r(n_1), \ldots, Z_j = r(n_j)$, and let $Y = r(n)$. Refer to Figure 1.

We let C' be the context C, with the leaves $n_1, \ldots, n_j$ deleted. In other words $C' = (S'_C, \lambda'_C)$ where $S'_C = S_C - \{n_1, \ldots, n_j\}$, $\lambda'_C(n') = \lambda_C(n')$ if $n' \in S_C - \{n, n_1, \ldots, n_j\}$, and $\lambda'_C(n) = \square$. Since C' is smaller than C, we can apply the induction hypothesis. Let $u, z_{j+1}, \ldots, z_k, y$ and $v, t'_{j+1}, \ldots, t'_k, t'$ be states of A such that $C'\left[v, t'_{j+1}, \ldots, t'_k\right] \Longrightarrow t'$ and $u \in Y, z_{j+1} \in Z_{j+1}, \ldots, z_k \in Z_k, y \in Z$ and $u \preccurlyeq_{down} v, z_{j+1} \preccurlyeq_{down} t'_{j+1}, \ldots, z_k \preccurlyeq_{down} t'_k, y \preccurlyeq_{up} t'$.

By definition of $A_{\simeq}$, there are states $z \in Y, z_1 \in Z_1, \ldots, z_j \in Z_j$ such that $(z_1, \ldots, z_j) \xrightarrow{f} z$ for some f.

We now use Lemma 10 with premise $u \simeq z$ and $u \preccurlyeq_{down} v$. We thus find state w such that $z \preccurlyeq_{down} w$ and $v \preceq w$. Note that this implies $v \preccurlyeq_{up} w$.

By definition of a downward simulation, and premises $z \preccurlyeq_{down} w$ and $(z_1, \ldots, z_j) \xrightarrow{f} z$, we find states $t_1, \ldots, t_j$ with $z_1 \preccurlyeq_{down} t_1, \ldots, z_j \preccurlyeq_{down} t_j$ and $(t_1, \ldots, t_j) \xrightarrow{f} w$.

By definition of an upward simulation and premises $v \preccurlyeq_{up} w$ and $C'\left[v, t'_{j+1}, \ldots, t'_k\right] \Longrightarrow t'$, we find states $t, t_{j+1}, \ldots, t_k$ with $t' \preccurlyeq_{up} t$, $t'_{j+1} \preccurlyeq_{down} t_{j+1}, \ldots, t'_k \preccurlyeq_{down} t_k$ and $C'[w, t_{j+1}, \ldots, t_k] \Longrightarrow t$.

The claim thus holds. $\square$

Now, we are ready to prove Theorem 1 itself.

Proof. Let T be a tree accepted by $A_{\simeq}$. We construct a context C by replacing all leaves in T by holes. We follow the construction of Lemma 11 for the context C. We now have a run of A on C. The Sufficient Conditions 4 and 5 of Section 6 ensure that this run is accepting. Sufficient Condition 6 of Section 6 ensures that we can complete the run on C to a run on T. $\square$

Theorem 2

Proof. We can easily see that by construction, $D_{\cong}^{\leq i}$ is a sub-automaton of $D_{\simeq}$.

Conversely, let (T_1, T_2) be a pair accepted by $D_{\simeq}$. We use Theorem 1, and let r be the corresponding run in H. Let $w_0, w_1, \ldots, w_n$ be the states in r. Let k be the length $k = |w_0| = |w_1| = \ldots = |w_n|$. Note that (T_1, T_2) is accepted by $D^{\leq k}$.

If $k \leq i$, then by construction, states $[w_0], [w_1], \ldots, [w_n]$ are in $D_{\cong}^{\leq i}$, and there is an accepting run in $D_{\cong}^{\leq i}$ for the pair (T_1, T_2).

If $k > i$, then we let T be such that (T_1, T) is recognized by $D^{\leq i}$ and (T, T_2) is recognized by $D^{\leq k-i}$. By the reasoning above, we know that (T_1, T) is recognized by $D_{\cong}^{\leq i}$.

Hence, we can write $(T_1, T_2) \in D_{\simeq}^{\leq i} \circ D^{\leq k-i}$. Using the termination condition $D_{\simeq}^{\leq i} \circ D \subseteq D_{\simeq}^{\leq i}$, we get that (T_1, T_2) is recognized by $D_{\simeq}^{\leq i}$.

□

Lemma 6

Proof. 1. Let $x \cdot q_{pref} \cdot y$ be a state of H. Any transition rule leading to that state will be of the form:

$$\left(x^1 \cdot q^1_{pref} \cdot y^1, \ldots, x^n \cdot q^n_{pref} \cdot y^n\right) \xrightarrow{f}_H x \cdot q_{pref} \cdot y$$

Suppose $x \cdot q_{pref} \cdot y \preccurlyeq_{down} z$. Then by definition, we know that z is of the form $x \cdot q_{pref} \cdot q_{pref} \cdot y$. Observe that for each $i : 1 \leq i \leq n$, we have $x^i \cdot q^i_{pref} \cdot y^i \preccurlyeq_{down} x^i \cdot q^i_{pref} \cdot q^i_{pref} \cdot y^i$. We also have a rule:

$$\left(x^1 \cdot q^1_{pref} \cdot q^1_{pref} \cdot y^1, \ldots, x^n \cdot q^n_{pref} \cdot q^n_{pref} \cdot y^n\right) \xrightarrow{f}_H x \cdot q_{pref} \cdot q_{pref} \cdot y$$

Conversely, we consider the state $x \cdot q_{pref} \cdot q_{pref} \cdot y$ of H. We notice that since D is deterministic, it follows that a state of form $q^1_{pref} \cdot q^2_{pref}$ is not reachable in H unless $q^1_{pref} = q^2_{pref}$. We can thus ignore states in which $q^1_{pref} \neq q^2_{pref}$. Then, any transition rule leading to state $x \cdot q_{pref} \cdot q_{pref} \cdot y$ will be of the form:

$$\left(x^1 \cdot q^1_{pref} \cdot q^1_{pref} \cdot y^1, \ldots, x^n \cdot q^n_{pref} \cdot q^n_{pref} \cdot y^n\right) \xrightarrow{f}_H x \cdot q_{pref} \cdot q_{pref} \cdot y$$

Suppose $x \cdot q_{pref} \cdot q_{pref} \cdot y \preccurlyeq_{down} z$. Then we have z of the form $x \cdot q_{pref} \cdot y$. Observe that we also have for each i: $x^i \cdot q^i_{pref} \cdot q^i_{pref} \cdot y' \preccurlyeq_{down} x^i \cdot q^i_{pref} \cdot y^i$. We also have a rule:

$$\left(x^1 \cdot q^1_{pref} \cdot y^1, \ldots, x^n \cdot q^n_{pref} \cdot y^n\right) \xrightarrow{f}_H x \cdot q_{pref} \cdot y$$

2. If we consider a regular set of states of H given by a word automaton, then its image w.r.t. $\preccurlyeq_{down}$ or $\preccurlyeq_{down}^{-1}$ can be expressed by adding edges to this automaton: for each transition $x \xrightarrow{q_{pref}} y$, add an edge $y \xrightarrow{q_{pref}} y$; similarly, for each consecutive edges $x \xrightarrow{q_{pref}} y$ and $y \xrightarrow{q_{pref}} z$, add an edge $x \xrightarrow{q_{pref}} z$.

□

Lemma 7

Proof. 1. We know that $\preccurlyeq^1_{up}$ is an upward simulation. If we let S be the relation generated by $\preccurlyeq^1_{up}$, then S is also an upward simulation. Furthermore, we have $\preceq \subseteq S$.
2. Similar to lemma 6 above.

□

Lemma 8

Proof. Assume $x \preceq y$ and $x \preccurlyeq_{down} z$. Then $x = x_1 \cdot q_{pref} \cdot x_2$ and $z = x_1 \cdot z' \cdot x_2$, with $z' \in \{\epsilon, q_{pref} \cdot q_{pref}\}$. Since the left-hand side of each pair generating $\preceq$ does not contain any prefix-copying state, we conclude that $y = y_1 \cdot q_{pref} \cdot y_2$, where either $x_1 \preceq y_1$ and $x_2 = y_2$, or $x_1 = y_1$ and $x_2 \preceq y_2$. In either case, we have $z = x_1 \cdot z' \cdot x_2 \preceq y_1 \cdot z' \cdot y_2$. Furthermore, we also have $y = y_1 \cdot q_{pref} \cdot y_2 \preccurlyeq_{down} y_1 \cdot z' \cdot y_2$.

We have show independence of *single steps* of $\preceq$ and $\preccurlyeq_{down}$. This is sufficient for proving that independence also holds for the transitive closure w.r.t. concatenation. Hence the claim holds. □

Lemma 9

Proof. We observe that all states in F are either in Q_{suff} or in Q_N. Therefore, the first and second claim hold.

The third claim holds since $x \preccurlyeq_{down} y$ only involves prefix-copying states. For a prefix-copying state q_{pref}, an arity 0 rule will be of the form $\xrightarrow{(f,f)} q_{pref}$, which means that the claim holds. □

C Simulation Algorithms

In the present section, we give algorithms for compute downwrad and upward simulations on states of a finite tree automaton. The algorithms are extensions of the one presented in [HHK95] for word automata.

Definition 16. *Let $A = (Q, F, \delta)$. For $p \in Q$, we let $pre_f(p)$ be the set $\left\{(r_1, \ldots, r_n) \mid (r_1, \ldots, r_n) \xrightarrow{f} p\right\}$. We let $pre_{f,i}(p)$ be the set $\left\{r \mid \exists r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n.\ (r_1, \ldots, r_{i-1}, r, r_{i+1}, \ldots, r_n) \xrightarrow{f} p\right\}$. For $P \subseteq Q$, we let $pre_f(P)$ to be the set $\bigcup_{p \in P} pre_f(p)$. We extend $pre_{f,i}$ to sets of states in a similar manner. Also, we define $post_f$ and $post_{f,i}$ in a similar manner.*

C.1 Computing Downward Simulations

We present an algorithm for computing the maximal downward simulation on states of a tree automaton. The result will be stored in the variable *sim* after termination.

Input: a tree automaton $A = (Q, F, \delta)$
Variables: sim, oldsim $\subseteq Q \times Q$

– **INITIALIZATION**
for each $q \in Q$ **do**
  $oldsim(q) := Q$;
  $sim(q) := \{r \in Q \mid \forall f \in \Sigma.\ pre_f(q) \neq \emptyset \Rightarrow pre_f(r) \neq \emptyset\}$;
od;

– **COMPUTATION**

While $\exists q \in Q$ such that $sim(q) \neq oldsim(q)$ **do**

{**Invariant 1:** $\forall q \in Q$, we have $sim(q) \subseteq oldsim(q)$}

{**Invariant 2:** $\forall q, q_1, \ldots, q_n, r$, if $(q_1, \ldots, q_n) \xrightarrow{f} q$ and $r \in sim(q)$, then $\exists r_1 \in oldsim(q_1), \ldots, \exists r_n \in oldsim(q_n)$ such that $(r_1, \ldots, r_n) \xrightarrow{f} r$}

For each symbol f **do**

For each i:$1 \leq i \leq rank(f)$ **do**

$remove := post_{f,i}(oldsim(q)) - post_{f,i}(sim(q))$;

For each $r \in post_{f,i}(q)$ **do**

$sim(r) := sim(r) - remove$;

od

od

od

$oldsim(q) := sim(q)$;

od

Observe that when the algorithm terminates, we get $sim = oldsim$, and hence invariant 2 implies that sim is a downward simulation.

The simulation returned by the algorithm is maximal: whenever a pair is removed in our algorithm, then this pair doesn't belong to any downward simulation. A state r is removed from $sim(q)$ when it is present in $remove$. This may have occurred because information was propagated, but ultimately, this is done because some pair was in $oldsim$ but not in sim after the initialization procedure. This means that if we follow the path through which pairs were put in the $remove$ sets, we can construct an execution of length $k+1$, showing that q is simulated by r up to depth k but not depth $k+1$. We thus conclude that there is no simulation relating q and r.

C.2 Computing Upward Simulations

We present an algorithm for computing the maximal upward simulation between states of a tree automaton.

Input: a tree automaton $A = (Q, F, \delta)$ and a downward simulation $\preccurlyeq_{down}$

Variables: sim, oldsim $\subseteq Q \times Q$

– **INITIALIZATION**

for each $q \in Q$ **do**

$oldsim(q) := Q$;

$sim(q) := \{r \in Q | \forall f, \forall i : 1 \leq i \leq rank(f)$, if $\exists q_1, \ldots, q_{i-1}, q_{i+1}, \ldots, q_n$ s.t $post_f\,(q_1, \ldots, q_{i-1}, q, q_{i+1}, \ldots, q_n) \neq \emptyset$, then $\exists r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n$ s.t $post_f\,(r_1, \ldots, r_{i-1}, r, r_{i+1}, \ldots, r_n) \neq \emptyset$ and $\forall j \neq i . q_j \preccurlyeq_{down} r_j\}$;

od;

– **COMPUTATION**

While $\exists q \in Q$ such that $sim(q) \neq oldsim(q)$ **do**

{**Invariant 1:** $\forall q \in Q$, we have $sim(q) \subseteq oldsim(q)$}

{**Invariant 2:** $\forall q, q_1, \ldots, q_n$ and r_i, if $(q_1, \ldots, q_n) \xrightarrow{f} q$ and $r_i \in sim(q_i)$ then $\exists r, r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n$ s.t. $q_1 \preccurlyeq_{down} r_1, \ldots, q_{i-1} \preccurlyeq_{down} r_{i-1}$, $q_{i+1} \preccurlyeq_{down} r_{i+1}, \ldots, q_n \preccurlyeq_{down} r_n$ and $(r_1, \ldots, r_{i-1}, r, r_{i+1}, \ldots, r_n) \xrightarrow{f} r$ and $r \in oldsim(q)$}

```
For each symbol f do
    For each i : 1 ≤ i ≤ rank(f) do
        remove:=pre_{f,i}(oldsim(q)) − pre_{f,i}(sim(q));
        For each r ∈ pre_{f,i}(q) do
            sim(r):=sim(r) − remove;
        od
    od
od
oldsim(q):=sim(q);
od;
```

Correctness of this algorithm is similar to the case of the downward simulation above. The only notable difference is in the use of $\preccurlyeq_{down}$ to further constrain the pairs we consider as potential candidates in the initialization step. This is due to the definition of an upward simulation, which depends on the given downward simulation. When we compare two states q_i and r_i, with respect to a transition rule $(q_1, \ldots, q_n) \xrightarrow{f} q$, we only consider other rules $(r_1, \ldots, r_n) \xrightarrow{f} r$ for which it holds that $\forall j \neq i. q_j \preccurlyeq_{down} r_j$. Hence, we don't need to redo that check later in the algorithm.

Remark In [HHK95], less straightforward but more efficient variants of the above algorithms are presented for the word case. The algorithms presented here can be modified in a similar manner.

D Description of the Protocols

We want to analyze parametrized systems whose processes are arranged in a tree-like architecture (examples: many network protocols). To do this, we will encode each configuration of the system by a tree. Each node of the tree will represent a process of the system.

D.1 Simple Token Protocol

This protocol has been described as an example in Appendix A. We have been able to compute the set of reachable states, and the transitive closure of the entire relation of this protocol.

D.2 Two Way Token Protocol

This protocol is a generalization of the simple token protocol in the sense that the token can also move downwards. We thus use the same alphabet as the one for the simple token protocol.

The transition relation is given by:

$$\begin{array}{lll} \xrightarrow{(n,n)} q_0 & \xrightarrow{(t,n)} q_1 & \xrightarrow{(n,t)} q_3 \\ (q_0,q_0) \xrightarrow{(T,N)} q_1 & (q_0,q_1) \xrightarrow{(N,T)} q_2 & (q_1,q_0) \xrightarrow{(N,T)} q_2 \\ (q_0,q_0) \xrightarrow{(N,N)} q_0 & (q_0,q_2) \xrightarrow{(N,N)} q_2 & (q_2,q_0) \xrightarrow{(N,N)} q_2 \\ (q_3,q_0) \xrightarrow{(T,N)} q_2 & (q_0,q_3) \xrightarrow{(T,N)} q_2 & (q_0,q_0) \xrightarrow{(N,T)} q_3 \end{array}$$

where,

q_2 is the accepting state;

q_0 the node is idle, i.e., the token is not in the node, nor in the subtree below the node;

q_1 the node is releasing the token to the node above it in the tree;

q_2 the token is either in the node or in a subtree below the node;

q_3 the node is receiving the token from the node above it in the tree;

We consider as initial configurations all trees with just one token. These configurations are the ones accepted by the following transducer, where q_5 is the only accepting state:

$$\begin{array}{lll} \xrightarrow{(n)} q_4 & \xrightarrow{(t)} q_5 & (q_4,q_4) \xrightarrow{(T)} q_5 \\ (q_4,q_4) \xrightarrow{(N)} q_4 & (q_4,q_5) \xrightarrow{(N)} q_5 & (q_5,q_4) \xrightarrow{(N)} q_5 \end{array}$$

We have been able to compute the set of reachable states, and the transitive closure of this protocol.

D.3 The Percolate Protocol

The protocol PERCOLATE, described in [KMM+01], operates on a tree of processes.

Each process has a local variable with values {0,1} for the leaf nodes and {U,0,1} for the internal nodes[6], (U is interpreted as "undefined yet").

The system percolates the disjunction of values in the leaves up to the root.

The states we use are $Q = \{q_0, q_1, q_u, q_d\}$. Intuitively, these states correspond to the following

- q_0: all nodes below (and including) the current node are labeled with 0; do not make any change to them;
- q_1: the current node and at least one node below is labeled with 1. No node below is labeled with U. Do not change the nodes below;

[6] To simplify the notation, we do not distinguish between the nullary and binary versions of the symbols 0 and 1.

q_u: all nodes above (and including) the current node have not yet been changed (they are still undefined);
q_d: a single change has occurred in the current node or below (accepting state);

The transition relation δ is given below (we use q_m to denote any member of $\{q_u, q_1, q_0\}$)

$$
\begin{array}{ccc}
\xrightarrow{(0,0)} q_0 & \xrightarrow{(1,1)} q_1 & (q_0, q_0) \xrightarrow{(0,0)} q_0 \\
(q_0, q_1) \xrightarrow{(1,1)} q_1 & (q_1, q_0) \xrightarrow{(1,1)} q_1 & (q_1, q_1) \xrightarrow{(1,1)} q_1 \\
(q_m, q_m) \xrightarrow{(U,U)} q_u & (q_m, q_d) \xrightarrow{(U,U)} q_d & (q_d, q_m) \xrightarrow{(U,U)} q_d \\
(q_0, q_0) \xrightarrow{(U,0)} q_d & (q_0, q_1) \xrightarrow{(U,1)} q_d & (q_1, q_0) \xrightarrow{(U,1)} q_d \\
(q_1, q_1) \xrightarrow{(U,1)} q_d & &
\end{array}
$$

We consider as initial configurations the ones accepted by the following tree automaton (i.e. a binary tree with leaves labeled with 0 and 1, the rest with u). q_{iu} is the only accepting state:

$$
\begin{array}{ccc}
\xmapsto{(0)} q_{i0} & \xmapsto{(1)} q_{i1} & (q_{i0}, q_{i0}) \xmapsto{(u)} q_{iu} \\
(q_{i1}, q_{i0}) \xmapsto{(N)} q_{iu} & (q_{i0}, q_{i1}) \xmapsto{(N)} q_{iu} & (q_{i1}, q_{i1}) \xmapsto{(N)} q_{iu}
\end{array}
$$

We have been able to compute the set of reachable states, and the transitive closure of this protocol.

D.4 The Tree-Arbiter Protocol

The protocol operates on a set of processes arranged in a tree-like architecture, and aims at preserving a mutual exclusion property between them.

In fact, processes encoded as leaf nodes try to access a shared resource, while the interior nodes are used to manage this resource. The access to the resource is represented by a token, i.e. the process which has the token is supposed to own the resource. The protocol must ensure that at maximum one process will have the shared resource.

In our model of the protocol, any process can be labeled as follows:
idle: the process does not do anything;
requesting: the process wants to access the shared resource;
token: the process has been granted the shared resource.
Furthermore, an interior process can be labeled as follows:
idle: the process together with all the process below are idle;
below: the token is somewhere in one subtree below this node (but not in the node itself).

The alphabet we use is $\{i, r, t, b\}$ for respectively *idle*, *requesting*, *token*, and *below*.

When a leaf is in state *requesting* the request is propagated upwards until it encounters a node which is aware of the presence of the token (i.e. a node in state token or below). If a node has the token it can always pass it upwards, or pass it downwards to a child which is requesting. Each time the token moves a step, the propagation moves a step or there is no move; the request and the token can't propagate at the same time.

Let us now describe the transition relation. The states are $\{q_i, q_r, q_t, q_{req}, q_{rel}, q_{granted}, q_m, q_{rt}\}$, where:
q_i : every node up to the current one is idle;
q_r : every node up to the current node are either idle or requesting, with at least one requesting. there was no move of the propagation below;
q_t : The token is either in this node or below; token hasn't moved (this is an accepting state);
q_{req} : The current node is requesting the token for itself or on behalf of a child : The request is being propagated;
q_{rel} : The token is moving upwards from the current node;
$q_{granted}$: The token is moving downwards to the current node;
q_m : The token is in this node or below; the token has moved (this is an accepting state);
q_{rt} : The token is either in this node or below it, i.e, nothing happens above the current node (this is an accepting state).
Using these states, the transition relation is given by:

$$
\begin{array}{ccc}
\xmapsto{(i,i)} q_i & \xmapsto{(i,r)} q_{req} & \xmapsto{(r,r)} q_r \\
\xmapsto{(r,t)} q_{grant} & \xmapsto{(t,t)} q_t & \xmapsto{(t,i)} q_{rel} \\
(q_i, q_i) \xmapsto{(i,i)} q_i & (q_r, q_i) \xmapsto{(i,i)} q_i & (q_i, q_r) \xmapsto{(i,i)} q_i \\
(q_r, q_r) \xmapsto{(i,i)} q_r & (q_{req}, q_i) \xmapsto{(i,i)} q_{req} & (q_i, q_{req}) \xmapsto{(i,i)} q_{req} \\
(q_{req}, q_r) \xmapsto{(i,i)} q_{req} & (q_r, q_{req}) \xmapsto{(i,i)} q_{req} & (q_r, q_r) \xmapsto{(i,r)} q_{req} \\
(q_r, q_i) \xmapsto{(i,r)} q_{req} & (q_i, q_r) \xmapsto{(i,r)} q_{req} & (q_r, q_r) \xmapsto{(r,r)} q_r \\
(q_r, q_i) \xmapsto{(r,r)} q_r & (q_i, q_r) \xmapsto{(r,r)} q_r & (q_{req}, q_r) \xmapsto{(r,r)} q_{req} \\
(q_r, q_{req}) \xmapsto{(r,r)} q_{req} & (q_r, q_i) \xmapsto{(r,t)} q_{grant} & (q_i, q_r) \xmapsto{(r,t)} q_{grant} \\
(q_r, q_r) \xmapsto{(r,t)} q_{grant} & (q_i, q_i) \xmapsto{(t,t)} q_t & (q_i, q_r) \xmapsto{(t,t)} q_t \\
(q_r, q_i) \xmapsto{(t,t)} q_t & (q_r, q_r) \xmapsto{(t,t)} q_t & (q_i, q_{req}) \xmapsto{(t,t)} q_{rt} \\
(q_{req}, q_i) \xmapsto{(t,t)} q_{rt} & (q_r, q_{req}) \xmapsto{(t,t)} q_{rt} & (q_{req}, q_r) \xmapsto{(t,t)} q_{rt} \\
(q_i, q_{grant}) \xmapsto{(t,b)} q_m & (q_{grant}, q_i) \xmapsto{(t,b)} q_m & (q_{grant}, q_r) \xmapsto{(t,b)} q_m \\
(q_r, q_{grant}) \xmapsto{(t,b)} q_m & (q_t, q_i) \xmapsto{(b,b)} q_t & (q_t, q_r) \xmapsto{(b,b)} q_t \\
(q_r, q_t) \xmapsto{(b,b)} q_t & (q_i, q_t) \xmapsto{(b,b)} q_t & (q_m, q_r) \xmapsto{(b,b)} q_m \\
(q_m, q_i) \xmapsto{(b,b)} q_m & (q_r, q_m) \xmapsto{(b,b)} q_m & (q_i, q_m) \xmapsto{(b,b)} q_m \\
(q_t, q_{req}) \xmapsto{(b,b)} q_{rt} & (q_{req}, q_t) \xmapsto{(b,b)} q_{rt} & (q_r, q_{rt}) \xmapsto{(b,b)} q_{rt} \\
(q_{rt}, q_r) \xmapsto{(b,b)} q_{rt} & (q_i, q_{rt}) \xmapsto{(b,b)} q_{rt} & (q_{rt}, q_i) \xmapsto{(b,b)} q_{rt} \\
(q_i, q_{rel}) \xmapsto{(b,t)} q_m & (q_{rel}, q_i) \xmapsto{(b,t)} q_m & (q_r, q_{rel}) \xmapsto{(b,t)} q_m \\
(q_{rel}, q_r) \xmapsto{(b,t)} q_m & (q_i, q_i) \xmapsto{(t,i)} q_{rel} & (q_i, q_r) \xmapsto{(t,i)} q_{rel} \\
(q_r, q_i) \xmapsto{(t,i)} q_{rel} & (q_r, q_r) \xmapsto{(t,i)} q_{rel} &
\end{array}
$$

Unfortunately, the reflexive closure used too much memory. We nevertheless succeded in computing the set of recheable configurations from the ones accepted by the below automaton.

$$
\begin{array}{rl}
 & \xmapsto{(i)} qi_i \\
(qi_i, qi_i) & \xmapsto{(i)} qi_i \\
(qi_i, qi_i) & \xmapsto{(t)} qi_t
\end{array}
$$

This initial automaton accepts binary trees where all the nodes are idle except the root holding the token. qi_t is the accepting state.

D.5 The Leader Election Protocol

The protocol operates on a binary tree of processes to elect a leader among the leaf processes that are possible candidates. A leaf process can be labeled as candidate or not candidate. An internal node can be labeled as candidate if at least one of its children is candidate; not candidate if none of his children is candidate; undefined if not defined yet; just elected if it just changed from candidate to elected; or elected.
There are six states:
q_c : There is at least a candidate in the tree below (this is an accepting state);
q_n : no candidates below;
q_{el} : the candidate to be elected is below (this is an accepting state);
q_u: undefined yet;
q_{jel} : just elected;
q_{ch} : something changed below.

$$
\begin{array}{rlrlrl}
 & \xmapsto{c,c} q_{(c)} & & \xmapsto{c,el} q_{(jel)} & & \xmapsto{n,n} q_{(n)} \\
 & \xmapsto{el,el} q_{(el)} & (q_{(c)}, q_{(c)}) & \xmapsto{c,c} q_{(c)} & (q_{(c)}, q_{(n)}) & \xmapsto{c,c} q_{(c)} \\
(q_{(n)}, q_{(c)}) & \xmapsto{c,c} q_{(c)} & (q_{(c)}, q_{(n)}) & \xmapsto{c,el} q_{(jel)} & (q_{(c)}, q_{(c)}) & \xmapsto{c,el} q_{(jel)} \\
(q_{(n)}, q_{(c)}) & \xmapsto{c,el} q_{(jel)} & (q_{(n)}, q_{(n)}) & \xmapsto{n,n} q_{(n)} & (q_{(u)}, q_{(u)}) & \xmapsto{u,u} q_{(u)} \\
(q_{(n)}, q_{(u)}) & \xmapsto{u,u} q_{(u)} & (q_{(u)}, q_{(n)}) & \xmapsto{u,u} q_{(u)} & (q_{(u)}, q_{(c)}) & \xmapsto{u,u} q_{(u)} \\
(q_{(c)}, q_{(u)}) & \xmapsto{u,u} q_{(u)} & (q_{(c)}, q_{(c)}) & \xmapsto{u,u} q_{(u)} & (q_{(n)}, q_{(n)}) & \xmapsto{u,u} q_{(u)} \\
(q_{(n)}, q_{(c)}) & \xmapsto{u,u} q_{(u)} & (q_{(c)}, q_{(n)}) & \xmapsto{u,u} q_{(u)} & (q_{(ch)}, q_{(u)}) & \xmapsto{u,u} q_{(ch)} \\
(q_{(ch)}, q_{(n)}) & \xmapsto{u,u} q_{(ch)} & (q_{(ch)}, q_{(c)}) & \xmapsto{u,u} q_{(ch)} & (q_{(u)}, q_{(ch)}) & \xmapsto{u,u} q_{(ch)} \\
(q_{(n)}, q_{(ch)}) & \xmapsto{u,u} q_{(ch)} & (q_{(c)}, q_{(ch)}) & \xmapsto{u,u} q_{(ch)} & (q_{(c)}, q_{(c)}) & \xmapsto{u,c} q_{(ch)} \\
(q_{(n)}, q_{(c)}) & \xmapsto{u,c} q_{(ch)} & (q_{(c)}, q_{(n)}) & \xmapsto{u,c} q_{(ch)} & (q_{(n)}, q_{(n)}) & \xmapsto{u,n} q_{(ch)} \\
(q_{(n)}, q_{(el)}) & \xmapsto{el,el} q_{(el)} & (q_{(c)}, q_{(el)}) & \xmapsto{el,el} q_{(el)} & (q_{(el)}, q_{(n)}) & \xmapsto{el,el} q_{(el)} \\
(q_{(el)}, q_{(c)}) & \xmapsto{el,el} q_{(el)} & (q_{(n)}, q_{(jel)}) & \xmapsto{el,el} q_{(el)} & (q_{(c)}, q_{(jel)}) & \xmapsto{el,el} q_{(el)} \\
(q_{(jel)}, q_{(n)}) & \xmapsto{el,el} q_{(el)} & (q_{(jel)}, q_{(c)}) & \xmapsto{el,el} q_{(el)} & &
\end{array}
$$

The set of initial configurations is given by the following tree automaton. This automaton recognizes all binary trees whith leaves labeled with n and c, the rest with u. q_0 is the only accepting state:

$$\xrightarrow{(c)} q_0 \qquad \xrightarrow{(n)} q_1$$
$$(q_0, q_0) \xrightarrow{(u)} q_0 \ (q_0, q_1) \xrightarrow{(u)} q_0 \ (q_1, q_0) \xrightarrow{(u)} q_0$$
$$(q_1, q_1) \xrightarrow{(u)} q_0$$

We have been able to compute the set of reachable states, and the transitive closure of this protocol.

$(T^*(I) \cap T2) == \emptyset$, where $T2$ models the configurations where more than one leaf is elected.