Closure of Tree Automata Languages under Innermost Rewriting

Adrià Gascón{1}, Guillem Godoy{1} & Florent Jacquemard{2} [2]

{1} *Technical University of Catalonia, Jordi Girona 1, Barcelona, Spain.* [1]

{2} *INRIA Futurs & LSV, UMR CNRS/ENS Cachan, France.*

Abstract

Preservation of regularity by a term rewrite system (TRS) states that the set of reachable terms from a tree automata (TA) language (aka regular term set) is also a TA language. It is an important and useful property, and there have been many works on identifying classes of TRS ensuring it; unfortunately, regularity is not preserved for restricted classes of TRS like shallow TRS. Nevertheless, this property has not been studied for important strategies of rewriting like the innermost strategy – which corresponds to the *call by value* computation of programming languages.

We prove that the set of innermost-reachable terms from a TA language by a shallow TRS is not necessarily regular, but it can be recognized by a tree automaton with equality and disequality constraints between brothers. As a consequence we conclude decidability of regularity of the reachable set of terms from a TA language by innermost rewriting and shallow TRS. This result is in contrast with plain (not necessarily innermost) rewriting for which we prove undecidability. We also show that, like for plain rewriting, innermost rewriting with linear and right-shallow TRS preserves regularity.

Introduction

Finite representations of infinite sets of terms are useful in many areas of computer science. The choice of a formalism for this purpose depends on its expressiveness, but also on its computational properties. Finite-state Tree Automata (TA) [3] are a well studied formalism for representing term languages, due to their good computational and expressiveness properties. They are used in many fields of computer science, from a theoretical and a practical point of view. For instance, for the analysis of systems or programs, when configurations can be represented by trees (e.g. concurrent processes with parallel and sequential composition operators) TA provide a finite representation of possibly infinite sets of configurations.

Term rewriting is a general formalism for the symbolic evaluation of terms by replacement of some patterns by others, following oriented equations, or rewrite

[1] The first two authors were supported by Spanish Min. of Educ. and Science by the FORMALISM project (TIN2007-66523) and by the LOGICTOOLS-2 project (TIN2007-68093-C02-01)

[2] adriagascon@gmail.com, ggodoy@lsi.upc.edu, florent.jacquemard@lsv.ens-cachan.fr

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

rules, given in a finite set (a term rewrite system, or TRS). Plain rewriting is sometimes too general, and in many contexts rewriting is applied with specific strategies giving a finer representation of the system behaviour. This is the case of the innermost strategy, which corresponds to the *call by value* computation of programming languages, where arguments are fully evaluated before the application of the function.

In the above application to system verification, transitions in infinite state systems can usually be represented by rewrite rules. There have been many studies of the connections between tree automata and rewriting, and a central property in this domain is the *preservation of regularity*. It states that for any given regular language L (which means that L is accepted by a TA), the set of reachable terms from L by a TRS R, denoted $R^*(L)$ is also regular. Preservation of regularity has been widely studied. The first result of this kind was that preservation of regularity holds for every ground TRS, as shown in [16]. In [14] this property was established for linear (variables occur at most once in every left-hand and right-hand side of a rule) and right-flat (the right-hand sides of the rules have height 0 or 1) TRS. There have been several extensions of this result, e.g. [6,10,13,15,5], and [13] represents a breakthrough since the left-linearity condition (linearity of left-hand sides of rules of the TRS) was dropped. However, in all the above cases, the condition of right-linearity remains necessary and in fact, a rewrite rule like $g(x) \to f(x, x)$ does not preserve regularity. Moreover, only plain rewriting is considered in these works, except in [5] where the bottom-up strategy is considered; there have been (up to our knowledge) no studies of regularity preservation under the innermost strategy.

The aim of this work is to study the preservation of regularity for innermost rewriting, and to identify a class of TRS for which better results can be found under the innermost strategy than under plain rewriting. We consider the class of shallow (all variables occur at depth 0 or 1 in the terms of the rules) TRS. Although the shallow case seems restrictive, for plain rewriting, shallow TRS do not preserve regularity. Moreover, several interesting properties of TRS, like reachability, joinability, confluence [12] and termination [8], are undecidable for shallow TRS, while adding certain linearity restrictions allows the decidability of all these problems [13,15,9,8]. Hence, from a theoretical point of view, the shallow case draws a frontier for decidability when one considers classes of TRS defined by syntactic restrictions.

Our main result (Theorem 4.6, Section 4.2) is that, given a regular language L and a *shallow* TRS R, the set $R^*(L)$ of terms reachable from L using R with the innermost strategy is recognized by tree automata extended with equality and disequality constraints between brothers in their state transitions. This kind of automata, which we call BTTA, was introduced in [2] as an extension of TA, and it has also good closure and decidability properties, but with worst complexity than standard TA. This is in contrast with the situation with plain rewriting: $R^*(L)$ is in general neither a TA nor a BTTA language under the same hypotheses (Proposition 3.2, Section 3).

One of the classical techniques for proving results of preservation of regularity consists of adding transitions to the automaton recognizing the starting language L, in order to simulate rule applications of R and recognize also all the terms reachable from L. Apparently, this completion technique which worked well for standard TA (in all the regularity preservation results cited so far) does not work for general shallow TRS. Innermost rewriting cannot be simulated by TA transitions, despite it does operate almost in a bottom-up fashion for shallow TRS [5]. The reason follows from two other results of the paper:

- First, we show that innermost rewriting with flat TRS (TRS whose all left-hand-side and right-hand-sides of rules have depth at most one) does not preserve regularity (Proposition 4.2, Section 4). As a consequence, we need to consider BTTA instead of standard TA.
- Second, flat and linear TRS do neither preserve BTTA-recognizably (Proposition 4.3, Section 4.1). Consequently, TA completion cannot work in this case.

The main result is obtained in two steps. First, we reduce the problem of representing the reachable terms from a regular set to the reachable terms from a constant. Next, we give a direct construction of a BTTA recognizing the reachable terms from a constant. It is based on a representation of the set of reachable terms introduced in [7] using constrained terms. As an immediate consequence of the main result, we obtain from [1] that given a regular language L and a shallow TRS R, it is decidable whether $R^*(L)$ is regular for innermost rewriting. In contraposition, we prove undecidability of regularity of $R^*(L)$ for plain (not necessarily innermost) rewriting.

Another positive result (Theorem 5.3, Section 5.1) is that, like for plain rewriting, innermost rewriting with linear and right-shallow TRS preserves regular languages. This result has been independently obtained in [11] In our case it is proved with a non trivial adaptation of the tree automata completion technique of e.g. [14,10]. The cases of plain and innermost rewriting are different in essence to treat, and some subtle differences need to be introduced. We show in particular that even though TA completion permits to establish that right-linear and right-flat TRS (*i.e.* when left-hand sides of rules might be not linear) preserve regular languages under plain rewriting, we show that this property is no longer true for under innermost rewriting (Proposition 5.4, Section 5.2).

1 Preliminaries

We use standard notation from the term rewriting literature [4]. A *signature* Σ is a finite set of function symbols with arity. We write Σ_m for the subset of function symbols of Σ of arity m. Given an infinite set $\mathcal{V}$ of variables, the set of terms built over Σ and $\mathcal{V}$ is denoted $\mathcal{T}(\Sigma, \mathcal{V})$, and the subset of ground terms is denoted $\mathcal{T}(\Sigma)$. The set of variables occurring in a term $t \in \mathcal{T}(\Sigma, \mathcal{V})$ is denoted $vars(t)$. A *substitution* σ is a mapping from $\mathcal{V}$ to $\mathcal{T}(\Sigma, \mathcal{V})$. The application of a substitution σ to a term t is written $\sigma(t)$, and is the homomorphic extension of σ to $\mathcal{T}(\Sigma, \mathcal{V})$.

A term t is identified as usual to a function from its set of *positions* (strings of positive integers) $Pos(t)$ to symbols of $\mathcal{F}$ and $\mathcal{V}$. We note Λ the empty string (root position). The length of a position p is denoted $|p|$. The *height* of a term t, denoted $h(t)$, is the maximum of $\{|p| \mid p \in Pos(t)\}$. A subterm of t at position p is written $t|_p$, and the replacement in t of the subterm at position p by u denoted $t[u]_p$.

Rewriting and the innermost strategy.

A *term rewriting system* (TRS) over a signature Σ is a finite set of rewrite rules $\ell \to r$, where $\ell \in \mathcal{T}(\Sigma, \mathcal{V}) \setminus \mathcal{V}$ (it is called left-hand side of the rule) and $r \in \mathcal{T}(\Sigma, vars(\ell))$ (it is called right-hand side). A term $s \in \mathcal{T}(\Sigma, \mathcal{V})$ rewrites to t by a TRS R at a position p of s with a substitution σ, denoted $s \xrightarrow[R,p,\sigma]{} t$ (p and σ may be omitted in this notation) if there is a rewrite rule $\ell \to r \in R$ such that $s|_p = \sigma(\ell)$ and $t = s[\sigma(r)]_p$. In this case, s is said to be *reducible*. The set of irreducible terms, also called *R-normal-forms*, is denoted by NF_R. The transitive and reflexive closure of $\xrightarrow[R]{}$ is denoted $\xrightarrow[R]{*}$. Given $L \subseteq \mathcal{T}(\Sigma)$, we note $R^*(L) = \{t \mid \exists s \in L, s \xrightarrow[R]{*} t\}$. The above rewrite step is called *innermost* if all proper subterms of $s|_p$ are R-normal forms. In this case, we write $s \xrightarrow[R]{\imath} t$, and $\xrightarrow[R]{\curlywedge}$ for the the transitive and reflexive closure of this relation, and $R^{\curlywedge}(L)$ for $\{t \mid \exists s \in L, s \xrightarrow[R]{\curlywedge} t\}$. We shall also use the notations $\mathsf{NF}^*_R(s)$ and $\mathsf{NF}^{\curlywedge}_R(s)$ (with $s \in \mathcal{T}(\mathcal{F}, \mathcal{V})$) for resp. $R^*(\{s\}) \cap \mathsf{NF}_R$ and $R^{\curlywedge}(\{s\}) \cap \mathsf{NF}_R$.

A TRS is called *linear* (resp. *right-linear*, *left-linear*) if every variable occurs at most once in each term (resp. right-hand side, left-hand side) of the rules. It is called *shallow* (resp. *right-shallow*, *left-shallow*) if variables occur at depth 0 or 1 in the terms (resp. in the right-hand sides, in the left-hand sides) of the rules and *flat* (resp. *right-flat*, *left-flat*) if the terms (resp. the right-hand sides, the left-hand sides) in the rules have height at most 1. A rule $\ell \to r$ is called *collapsing* if r is a variable.

2 Tree automata with constraints between brothers

A *tree automaton* (TA) A on a signature Σ is a tuple $(Q, Q^{\mathsf{f}}, \Delta)$ where Q is a finite set of nullary state symbols, disjoint from Σ, $Q^{\mathsf{f}} \subseteq Q$ is the subset of final states and Δ is a set of ground rewrite rules of the form: $f(q_1, \ldots, q_m) \to q$, or $q_1 \to q$ (ε-transition) where $f \in \Sigma_m$, and $q_1, \ldots, q_m, q \in Q$ (q is called the *target* state of the rule).

A *Bogaert-Tison tree automaton* (BTTA, or tree automaton with constraints between brothers) is defined like a TA except that its states are unary and its transitions are constrained rewrite rules of the form $f\big(q_1(x_1), \ldots, q_m(x_m)\big) \to q\big(f(x_1, \ldots, x_m)\big)\, [\![c]\!]$, or ε-transitions $q_1(x_1) \to q(x_1)$, where $x_1, \ldots, x_m$ are distinct variables and the constraint c is a Boolean combination of equalities $x_i = x_j$. Equivalently, the constraint c can be defined as a partition P of $\{1, \ldots, m\}$ with the same meaning as a conjunction of equalities $x_i = x_j$ for the indexes i, j such that $i \equiv_P j$, and disequalities $x_i \neq x_j$ for the indexes i, j such that $i \not\equiv_P j$. Following the notations of [2,3], the above transitions are written $f(q_1, \ldots, q_m) \xrightarrow{c} q$ (or $f(q_1, \ldots, q_m) \to q$ when c is *true*) and $q_1 \to q$, and every equality $x_i = x_j$ (resp. disequality $x_i \neq x_j$) in the constraint c is written $i = j$ (resp. $i \neq j$). Note that every TA is the special case of BTTA whose constraints are all equal to *true*.

The *language* $L(A, q)$ of a BTTA A in state q is the set of ground terms *accepted* in state q by A, i.e. the terms t such that $t \xrightarrow[\Delta]{*} q(t)$. The language $L(A)$ of A is $\bigcup_{q \in Q^{\mathsf{f}}} L(A, q)$ and a set of ground terms is called *regular* (resp. *BT-regular*) if it is the language of a TA (resp. BTTA).

A BTTA A is called *deterministic* (resp. *complete*) if for every term $t \in \mathcal{T}(\Sigma)$, there is at most (resp. at least) one state q such that $t \in L(A, q)$. If A is deterministic and complete, this unique state is denoted $A(t)$. A BTTA A is *normalized* if it does not contain ϵ-transitions, constraints in transitions are defined using partitions, and for every function symbol f with arity m, states $q_1, \ldots, q_m$ and partition P of $\{1, \ldots, m\}$, A contains exactly one rule of the form $f(q_1, \ldots, q_m) \xrightarrow{P} q$. A normalized BTTA A is deterministic and complete, and any BTTA can be transformed into a normalized one recognizing the same language. If A is normalized, we write $A(t, P)$, for a flat term $t \in \mathcal{T}(\Sigma \cup Q)$, to denote the unique state q such that $t \xrightarrow{P} q$ is a transition of A. BTTA are useful for representing the set of normal forms of certain classes of TRS, like flat TRS, see *e.g.* [3].

Lemma 2.1 *[3] Let R be a flat TRS over Σ. There exists a normalized BTTA $B = (Q_B, Q^{\mathsf{f}}_B, \Delta_B)$ on Σ which recognizes the set of ground R-normal forms. Moreover $|Q_B \setminus Q^{\mathsf{f}}_B| = 1$.*

3 Closure under plain rewriting with shallow TRS

Right-(shallow and linear) TRS preserve regularity [13]. It is well known that right-linearity cannot be omitted, as the following example shows.

Example 3.1 Let us consider the TRS $R := \{g(x) \to f(x, x)\}$ and the regular language $L = \{g^n(a) \mid n \geq 0\} = \{a, g(a), g(g(a)), \ldots\}$. The set $R^*(L)$ is not regular because its intersection with the regular set $\mathcal{T}(\{f, a\})$ is the non-regular set *Bin* of complete binary trees whose internal nodes are labeled by f and whose leaves are labeled by a, and the class of regular tree languages is closed under intersection. $\diamond$

We show below that considering BTTA does not help in this case.

Proposition 3.2 *In general, $R^*(L)$ is not BT-regular when L is a regular tree language and R a flat TRS.*

Proof. Let us consider R, L and *Bin* as in Example 3.1. The set $R^*(L)$ is not BT-regular. Indeed, its intersection with the regular (hence BT-regular) set $L_2 := \{f(s, t) \mid s \in \mathcal{T}(\{g, a\}), t \in \mathcal{T}(\{f, a\})\}$ is the subset L' of terms $f(s, t) \in L_2$ with $t \in Bin$ and $h(s) = h(t)$. This latter set is not BT-regular, as shown below. It follows that $R^*(L)$ is not BT-regular because the class of BT-regular tree languages is closed under intersection [2].

Let us now show that L' is not BT-regular. Assume that it is recognized by a BTTA $A = (Q, Q^{\mathsf{f}}, \Delta)$ on Σ with n states, and for all $i \geq 1$ let $f(s_i, t_i)$ be the term of L' with $h(s) = h(t) = i$. For each i, there exists a reduction sequence $f(s_i, t_i) \xrightarrow[\Delta]{*} q(f(s_i, t_i))$ with $q \in Q^{\mathsf{f}}$, and we consider the last rule ρ_i of Δ applied in this reduction sequence. There exist two distinct indexes $i_1, i_2 \geq 1$ such that $\rho_{i_1} = \rho_{i_2}$. Let $f(q_1(x_1), q_2(x_2)) \xrightarrow{c} q(f(x_1, x_2))$ be this unique rule of Δ. Note that the constraint c does not contain the equality $x_1 = x_2$, actually c may be $x_1 \neq x_2$ or *true*. In both cases, it follows that $f(s_{i_1}, t_{i_2}) \xrightarrow[\Delta]{*} f(q_1(s_{i_1}), q_2(t_{i_2})) \xrightarrow[\rho_{i_1}]{} q(f(s_{i_1}, t_{i_2}))$. This is contradiction with the fact that $f(s_{i_1}, t_{i_2}) \notin L'$ because $h(s_{i_1}) \neq h(t_{i_2})$. $\square$

4 Closure under innermost rewriting with shallow TRS

The essential problem in Example 3.1 and Proposition 3.2 relies on the fact that after an application of the rule $g(x) \to f(x, x)$ on a term $g(t)$, producing $f(t, t)$, the following application of rewrite rules can change the two occurrences of t in different ways, producing terms $f(t_1, t_2)$ with $t_1 \neq t_2$. The equality constraints of BTTA have not the expressive power to capture the relation relating t_1 and t_2 (i.e. that both are reachable from a common term). The situation is getting better when the innermost strategy is applied.

Example 4.1 When we apply the rule $g(x) \to f(x, x)$ (Example 3.1) with the innermost strategy, the subterm where it is applied must be $g(t)$ for a R-normal form t. Hence, in the term $f(t, t)$ obtained, t cannot be modified by rewriting. Hence $R^{\curlywedge}(L) = \{g^n(t) \mid t \in Bin\}$ is BT-regular. ◇

Note however that $R^{\curlywedge}(L)$ is not regular in the above example.

Proposition 4.2 *In general, $R^{\curlywedge}(L)$ is not regular when L is a regular tree language L and R a flat TRS.*

4.1 Closure of BTTA languages with flat TRS

Linear and flat TRS preserve regularity [14]. This result cannot be extended to BT-regularity, neither for plain nor innermost rewriting.

Proposition 4.3 *In general, $R^*(L)$ and $R^{\curlywedge}(L)$ are not BT-regular when L is BT-regular and R is a flat and linear TRS.*

Proof. The tree language $L = \{h(f^n(0), f^n(0)) \mid n \geq 0\}$ is recognized by the following BTTA, with one equality constraint tested at the root position: $\big(\{q, q^{\mathsf{f}}\}, \{q^{\mathsf{f}}\}, \{0 \to q, f(q) \to q, h(q, q) \xrightarrow{1=2} q^{\mathsf{f}}\}\big)$. Note that L is not regular. Let us consider the flat and linear TRS $R = \{f(x) \to g(x)\}$ and the regular tree language $L' = \{h(f^n(0), g^m(0)) \mid n \geq 0\}$. The closure $R^*(L) \cap L' = \{h(f^n(0), g^n(0))\}$ is not BT-recognizable, hence $R^*(L)$ is neither BT-recognizable. This is also true if we consider innermost rewriting. □

4.2 Closure of TA languages with shallow TRS

The classical approach for proving preservation of regularity [10,13,15] consists of completing a TA recognizing the original language L with new rules inferred using R. This method cannot be generalized to BT-regular languages, according to Proposition 4.3. Therefore, we follow a different approach. We first prove that given L regular and R flat, we can generate a new TRS R_c over an extended signature including a new constant c such that $R_c^{\curlywedge}(\{c\})$ coincides with $R^{\curlywedge}(L)$ on the given signature. This simple and enabling result permits to represent the set of terms reachable from a regular term set as the set of terms reachable from a constant. Later, we show how to compute a BTTA recognizing the reachable terms from a constant. To this end we make use of some results in [7] on innermost rewriting with shallow TRS.

Simplifying assumptions on the signature and the TRS.

From now on in this section, we assume that all terms are built from a given fixed signature Σ which contains several constant symbols and only one non-constant symbol f of arity m. We assume moreover that the TRS R is flat. Such assumptions, already used *e.g.* in [7], can be made without loss of generality for the problem considered here.

Reduction to reachable terms from constants.

Our intention is to reduce the effort of characterizing the reachable terms from a regular language L to just characterizing the reachable terms from a single constant. This is possible using the common idea of adding the inverse rules of an automaton A recognizing L to the rewrite system. The generation of the terms of L starting from the final states of A before any rewrite step application is ensured by the innermost strategy.

Lemma 4.4 *For every flat TRS R and regular language L, over a signature Σ, there exists an extension $\Sigma' \supset \Sigma$, a constant $c \in \Sigma' \setminus \Sigma$ and a flat TRS R_c over Σ' such that $R_c^{\curlywedge}(\{c\}) \cap \mathcal{T}(\Sigma) = R^{\curlywedge}(L)$.*

Weak normal forms and constrained terms.

From [7] we have the following definitions and results. A term t is a *weak normal form* if it is either a constant or a term of the form $t = f(t_1, \ldots, t_m)$ such that every t_i is either a constant or a normal form.

A *constraint* C is a partial function $C : \mathcal{V} \to \mathcal{P}(\Sigma_0)$ ($\mathcal{P}(\Sigma_0)$ denotes the powerset of Σ_0 minus the empty set) i.e. an assignment from variables to non-empty sets of constants. We say that a substitution σ is a *solution* of a constraint C (with respect to a TRS R) if for all x in $dom(C)$, $\sigma(x) \in \mathsf{NF}_R^{\curlywedge}(C(x)) \setminus \Sigma_0$. A *constrained term* is a pair denoted $t|C$, where t is a flat term and C is a constraint, with $dom(C) = vars(t)$. A term $\sigma(t)$ is called an *instance* of $t|C$ if σ is a solution of C. Note that every instance of a constrained term is a weak normal form.

In [7] it is shown how to compute for every flat TRS R and for every constant $c \in \Sigma_0$ two sets of constrained terms r_c and $\overline{r}_c$ satisfying the following properties.

- (a) for every $t|C \in r_c$, there exists at least one solution of C, and all instances of $t|C$ are innermost-reachable from c.
- (b) for every $t|C \in \overline{r}_c$, all non-constant normal form instances of $t|C$ are innermost-reachable from c.
- (c) for every weak normal form s innermost-reachable from c, there exists some constrained term $t|C \in r_c$ such that s is an instance of $t|C$.
- (d) for every non-constant normal form s innermost-reachable from c, there exists some constrained term $f(t_1, \ldots, t_m)|C \in \overline{r}_c$ such that s is an instance of $f(t_1, \ldots, t_m)|C$.

Recognizing terms reachable from constants.

We assume some sets r_c and $\overline{r_c}$ as above and we construct a normalized BTTA A_R which recognizes terms reachable from constants in Σ_0 using R. For this purpose, we shall use the BTTA B of Lemma 2.1 recognizing the ground normal forms of R. Let $Q_0 = \{q \in Q_B \mid \exists d \in \Sigma_0, B(d) = q\}$, and $Q_1 = Q^{\mathsf{f}}_B \setminus Q_0$. Without loss of generality we assume that only constants lead to states in Q_0. Thus, the states of Q_1 characterize the set of non-constant R-normal-form. The states of A_R are pairs $\langle S, q\rangle$, where $S \subseteq \Sigma_0$ and $q \in Q_B$. The intuitive idea is that a term t will lead to $\langle S, q\rangle$ with A_R if it leads to q with B and S is the set of all constants that R-reach t. To this end, the set of transition rules contains:

- $b \to \big\langle \{d \mid b|\emptyset \in r_d\}, B(b)\big\rangle$, for every constant b.
- $f\big(\langle S_1, q_1\rangle, \dots, \langle S_m, q_m\rangle\big) \xrightarrow{P} \big\langle S, B\big(f(q_1, \dots, q_m), P\big)\big\rangle$, for every $S_1, \dots, S_m \subseteq \Sigma_0$, $q_1, \dots, q_m \in Q_B$, and partition P of $\{1, \dots, m\}$, and where S is the set of constants $c \in \Sigma_0$ such that there exists $f(\alpha_1, \dots, \alpha_m)|C \in (r_c \cup \overline{r_c})$ with:
  - *i.* $\forall 1 \le i \le m$, if $\alpha_i \in \Sigma_0$ then $\alpha_i \in S_i$ and if $\alpha_i \in \mathcal{V}$ then $C(\alpha_i) \subseteq S_i$ and $q_i \in Q_1$,
  - *ii.* $\forall 1 \le i < j \le m$, if $\alpha_i = \alpha_j \in \mathcal{V}$ then $i \equiv_P j$ and if $f(\alpha_1, \dots, \alpha_m)|C \in \overline{r_c} \setminus r_c$ then $B\big(f(q_1 \dots q_m), P\big) \in Q_1$ and every $\alpha_i \in \Sigma_0$ $(1 \le i \le m)$ is in NF_R.

By construction, the automaton A_R is normalized. The following lemma states the correctness of its construction.

Lemma 4.5 *For all $t \in \mathcal{T}(\Sigma)$ $A_R(t) = \big\langle \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t\}, B(t)\big\rangle$.*

Proof. Let $A_R(t) = \langle S, q\rangle$. It is straightforward by construction that $B(t) = q$. We prove $S = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t\}$ by induction on the size of t.
We first consider the case where t is a constant. By definition of A_R, $S = \{d \mid (t|\emptyset) \in r_d\}$. It suffices to see that the conditions $d \xrightarrow[R]{\curlywedge} t$ and $(t|\emptyset) \in r_d$ are equivalent when t is a constant. This is a consequence of conditions (a) and (d) above.

Now, assume that t is not a constant. Then, t is of the form $f(t_1, \dots, t_m)$. Let $\langle S_1, q_1\rangle = A_R(t_1), \dots, \langle S_m, q_m\rangle = A_R(t_m)$. By induction hypothesis, $B(t_i) = q_i$ and $S_i = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t_i\}$, for every $i \in \{1, \dots, m\}$. Let $f(\langle S_1, q_1\rangle, \dots, \langle S_m, q_m\rangle) \xrightarrow{P} \langle S, q\rangle$ be the rule fired in the last applied transition of $A_R(t)$. Then, it holds that $i \equiv_P j$ iff $t_i = t_j$. We prove the two inclusions of $S = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t\}$ separately.

Direction $\subseteq$. Let $c \in S$. By construction of A_R there exists a constrained term $f(\alpha_1, \dots, \alpha_m)|C \in (r_c \cup \overline{r_c})$ for which the above conditions i and ii hold.

We obtain an instance of $f(\alpha_1, \dots, \alpha_m)|C$ by defining a substitution σ for the α_i's that are variables in $\mathit{vars}(f(\alpha_1, \dots, \alpha_m))$ with $\sigma(\alpha_i) := t_i$. The substitution σ is well defined: if for different i and j we have $\alpha_i = \alpha_j \in \mathcal{V}$, by condition (ii), it implies that $i \equiv_P j$ and then $t_i = t_j$. It holds that every of such $\sigma(\alpha_i)$ is a non-constant normal form, because $\alpha_i \in \mathcal{V}$ implies $q_i = B(t_i) \in Q_1$ due to condition (i). Moreover, $\sigma(\alpha_i)$ is reachable from $C(\alpha_i)$ because $\alpha_i \in \mathcal{V}$ implies $C(\alpha_i) \subseteq S_i$ and $S_i = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t_i\}$. Altogether, it follows that $\sigma(f(\alpha_1, \dots, \alpha_m))$ is an instance of $f(\alpha_1, \dots, \alpha_m)|C \in (r_c \cup \overline{r_c})$.
We know that either $f(\alpha_1, \dots, \alpha_m)|C \in r_c$ or $f(\alpha_1, \dots, \alpha_m)|C \in \overline{r_c} \setminus r_c$. On the one hand, if $f(\alpha_1, \dots, \alpha_m)|C \in r_c$ then, by condition a, $c \xrightarrow[R]{\curlywedge} \sigma(f(\alpha_1, \dots, \alpha_m))$. On the other hand, if $f(\alpha_1, \dots, \alpha_m)|C \in \overline{r_c} \setminus r_c$ then our conditions imply that

$B(f(q_1,\ldots,q_m),P) \in Q_1$, and hence $q \in Q_1$ and t is a non-constant normal form. Moreover, the α_i's that are constants are also normal forms. For every one of these constants α_i we know that $\alpha_i \in S_i$, and hence we also have $\alpha_i \xrightarrow[R]{\curlywedge} t_i$. But since this α_i is a normal form it follows that $\alpha_i = t_i$. This implies that $\sigma(f(\alpha_1,\ldots,\alpha_m)) = t$, and hence, that t is a non-constant normal form that is an instance of $f(\alpha_1,\ldots,\alpha_m)|C \in \overline{r_c}$, and by condition b, $c \xrightarrow[R]{\curlywedge} \sigma(f(\alpha_1,\ldots,\alpha_m))$.

Once we know that $c \xrightarrow[R]{\curlywedge} \sigma(f(\alpha_1,\ldots,\alpha_m))$, it suffices to show that $\sigma(f(\alpha_1,\ldots,\alpha_m)) \xrightarrow[R]{\curlywedge} t$ in order to conclude. By the definition of σ, terms $\sigma(f(\alpha_1,\ldots,\alpha_m))$ and t can only differ in the positions i such that α_i is a constant. But in such cases we know that $\alpha_i \in S_i$, and using $S_i = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t_i\}$ we obtain $\alpha_i \xrightarrow[R]{\curlywedge} t_i$. Hence, $\sigma(f(\alpha_1,\ldots,\alpha_m)) \xrightarrow[R]{\curlywedge} t$ follows.

Direction $\supseteq$. Let c be such that there exists a rewrite sequence $c \xrightarrow[R]{\curlywedge} t$. Since t is not a constant, the previous derivation can be written by making explicit the last rewrite step at position Λ as ($>\Lambda$ represents any position other than Λ):

$$c \xrightarrow[R]{\curlywedge}\xrightarrow[R,\Lambda]{} f(s_1,\ldots,s_m) \xrightarrow[R,>\Lambda]{\curlywedge} t = f(t_1,\ldots,t_m)$$

Hence, there exist (sub-)derivations $s_i \xrightarrow[R]{\curlywedge} t_i$. The term $s = f(s_1,\ldots,s_m)$ is a weak normal form, and hence, by condition c, there exists a constrained term $u|C \in r_c$ such that s is an instance of $u|C$. At this point, either there exists such a u of the form $f(\alpha_1,\ldots,\alpha_m)$, or every u satisfying this condition is a variable. In the second case, s is necessarily a normal form, and hence, by condition d, there exists a constrained term $f(\alpha_1,\ldots,\alpha_m)|C$ in $\overline{r}_c$ such that s is an instance of $f(\alpha_1,\ldots,\alpha_m)|C$. For proving that $c \in S$, it suffices to show that the conditions i and ii hold.

If a certain α_i is a constant, then it coincides with s_i, which R-reaches t_i. Since $S_i = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t_i\}$, it necessarily contains α_i.

If a certain α_i is a variable, then s_i coincides with t_i and is a non-constant normal form reachable from $C(\alpha_i)$. Hence, $q_i = B(t_i)$ is in Q_1, and again since $S_i = \{d \in \Sigma_0 \mid d \xrightarrow[R]{\curlywedge} t_i\}$, it necessarily includes $C(\alpha_i)$.

If $\alpha_i = \alpha_j \in \mathcal{V}$ then $s_i = s_j$ and since both are normal forms we also have $t_i = t_j$, from which $i \equiv_P j$ follows.

In the case where $f(\alpha_1,\ldots,\alpha_m)|C$ belongs to $\overline{r}_c \setminus r_c$, $f(s_1,\ldots,s_m)$ is a non-constant normal form. Therefore, $q = B(f(q_1,\ldots,q_m),P) \in Q_1$ and all the constants α_i are also normal forms. $\square$

Given a flat TRS R and a regular L, the BTTA A_{R_c} (for R_c associated to R as in Lemma 4.4), restricted to the signature Σ of R, recognizes $R^{\curlywedge}(L)$ by marking as accepting states the pairs $\langle S, q\rangle$ such that $c \in S$, according to Lemmas 4.4 and 4.5. This permits to conclude the proof of Theorem 4.6.

Theorem 4.6 *$R^{\curlywedge}(L)$ is BT-regular when L is regular and R is shallow.*

A term t is *reachable* from another term s if there exists a rewrite sequence that transforms s into t. Two terms s and t are *joinable* if there exists a term u reachable from s and t. Ground reachability and joinability (the restriction of these problems to ground terms) are undecidable for flat TRS [12]. Theorem 4.6 show that they become decidable when the innermost strategy is applied.

Corollary 4.7 *Ground reachability and joinability are decidable for ground terms for innermost rewriting with shallow TRS.*

Proof. When restricting to innermost rewriting, t is reachable from s iff $t \in R^{\curlywedge}(\{s\})$. Since $\{s\}$ is a regular language when s is ground, $R^{\curlywedge}(\{s\})$ is BT-regular by Theorem 4.6. Therefore ground reachability reduces to the membership problem for BTTA, which is decidable.

Similarly, s and t are joinable iff $R^{\curlywedge}(\{s\}) \cap R^{\curlywedge}(\{t\}) \neq \emptyset$. By Theorem 4.6 and closure of BT-languages under Boolean operations [3], we obtain a reduction of ground joinability to the emptiness problem for BTTA, which is also decidable. $\Box$

In [1] the decidability of the regularity of a BTTA was shown. Combining this result with Theorem 4.6 we obtain the following corollary.

Corollary 4.8 *Given a regular language L and a shallow TRS R, it is decidable whether $R^{\curlywedge}(L)$ is regular.*

This result does not hold when we deal with plain rewriting. In [12] it has been proved that reachability of flat TRS is undecidable by reducing the Post correspondence problem into $0 \xrightarrow[R]{*} 1$. We show below how to extend R into R_0 such that ${R_0}^*(0)$ is regular iff $0 \xrightarrow[R]{*} 1$.

Theorem 4.9 *Given a regular language L and a flat TRS R, it is undecidable whether $R^*(L)$ is regular.*

Proof. In [12] it is proved that reachability of flat TRS is undecidable by reducing a PCP instance P into a TRS R over a signature including $\{0, 1\}$ such that P has a solution iff $0 \xrightarrow[R]{*} 1$. The reduction in [12] also satisfies that if P has no solution, the 0 does not reach any term containing 1 nor any term containing 0 properly.

This reduction can be modified by adding new symbols $\{f, h, g, a, b, c\}$ to the current signature Σ, and adding two new sets of rules to R: $R_1 = \{0 \to f(a, b), a \to g(a), b \to g(b), a \to c, b \to c, f(x, x) \to h(x, x)\}$ and R_2 containing all the necessary rules for making $R_2^*(1)$ to be $\mathcal{T}(\Sigma \cup \{f, h, a, b, c\})$. The rules of R ensure $(R \cup R_1 \cup R_2)^*(0)$ to be a non-regular language, unless $0 \xrightarrow[R]{*} 1$. Note that if P has solution, then $0 \xrightarrow[R]{*} 1$, and hence $(R \cup R_1 \cup R_2)^*(0)$ is $\mathcal{T}(\Sigma \cup \{f, h, a, b, c\})$, which is regular. Otherwise, if P has no solution, then 0 does not reach any term containing 1, nor containing 0 properly, and hence $(R \cup R_1 \cup R_2)^*(0) \cap \mathcal{T}(\{h, c\})$ is the set $\{h(g^n(c), g^n(c)) \mid n > 0\}$, which is not regular. $\Box$

5 Innermost rewriting and right-shallow TRS

In this section, we study the closure of regular languages under innermost rewriting with TRS whose right-hand sides of rules are shallow. We show that regularity is preserved by innermost rewriting with linear right-shallow TRS (Subsection 5.1), but not by innermost rewriting with right-(linear and flat) (non left-linear) TRS (Subsection 5.2). The first result was proved independently in [11]. ** both with techniques of TA completion, with some differences however:

REV.2 Transforming right-shallow TRS to flat TRS (in the first and second paragraph of Sec 5.1). This makes the construction and the proof little bit simpler, which costs almost equivalently to the argument of the transformation.

REV.2 Arguing the number of states in the constructed automata (in 3 lines between Lem 5.2 and Th 5.3) **

5.1 TA languages and linear and right-shallow TRS

First, we observe that every right-shallow TRS R can be transformed into a right-flat TRS R' (on an extended signature) such that for all $s, t \in \mathcal{T}(\Sigma)$, $s \xrightarrow[R]{\curlywedge} t$ iff $s \xrightarrow[R']{\curlywedge} t$. The idea is to add a new constant c_r and a rule $c_r \to r$ for every ground proper subterm r of a right-hand side of a rule of R, and to replace r by c_r in all the right-hand sides of R.

Let $A = (Q, Q^{\mathsf{f}}, \Delta)$ be a deterministic and complete TA on Σ recognizing a tree language L, and let R be a linear and right-flat TRS. For all $c \in \Sigma_0$ we denote as q_c the unique state of Q such that $c \to q_c \in \Delta$. We assume moreover *wlog* that $L(A, q_c) = \{c\}$.

We construct a finite sequence of TA $A_0, A_1, \ldots$ whose last element recognizes $R^{\curlywedge}(L)$. The construction of the sequence is incremental. Every A_{k+1} is obtained from A_k by the addition of some new transitions, such that if some term s is recognized by A_k and s rewrites (in one step of innermost rewriting) to t, then t is recognized by A_{k+1}.

In order to restrict to innermost rewriting, we shall use a complete and deterministic TA $B = (Q_B, Q_B^{\mathsf{f}}, \Delta_B)$ (without ε-transitions) recognizing the ground R-normal forms (see *e.g.* [3] for its construction). As in Lemma 2.1, we can assume that B has only one non-accepting state q_{reject}. Let A_0 be a TA recognizing $L(A)$: $A_0 := \big(Q \times Q_B, Q^{\mathsf{f}} \times Q_B, \Delta_0\big)$ where Δ_0 is the set of transitions $f\big(\langle q_1, q'_1\rangle, \ldots, \langle q_m, q'_m\rangle\big) \to \langle q, q'\rangle$ such that $f(q_1, \ldots, q_m) \to q \in \Delta$ and $f(q'_1, \ldots, q'_m) \to q' \in \Delta_B$.

The addition of transition rules to A_k, giving A_{k+1}, is defined by the superposition of rules of R into a sequence of transitions of Δ_k. More precisely, $A_{k+1} \setminus A_k$ contains all the transitions which can be constructed from a rewrite rule $\ell \to r$ of R (we let $\ell = f(\ell_1, \ldots, \ell_m)$) and a substitution θ of the variables of ℓ into states of $Q \times Q_B^{\mathsf{f}}$ whose accepted language wrt A_k is not empty, such that: $\theta(\ell) \xrightarrow[\Delta_k]{*} \langle q_0, q_{\mathsf{reject}}\rangle$, and the last step of the above reduction is $f\big(\langle q_1, q'_1\rangle, \ldots, \langle q_m, q'_m\rangle\big) \xrightarrow[\Delta_k]{} \langle q_0, q_{\mathsf{reject}}\rangle$ and for all $i \leq m$, $q'_i \neq q_{\mathsf{reject}}$. There are two cases for the transitions of $A_{k+1} \setminus A_k$:

- case 1: r is a variable. In this case, $r \in \mathit{vars}(\ell)$. Let $\langle \overline{q}, \overline{q}'\rangle = \theta(r)$, we add the ε-transition $\langle \overline{q}, \overline{q}'\rangle \to \langle q_0, \overline{q}'\rangle$.
- case 2: $r = g(r_1, \ldots, r_m)$. We add all the transitions $g\big(\langle \overline{q}_1, \overline{q}'_1\rangle, \ldots, \langle \overline{q}_m, \overline{q}'_m\rangle\big) \to \langle q_0, \overline{q}'\rangle$ such that $g(\overline{q}'_1, \ldots, \overline{q}'_m) \to \overline{q}' \in \Delta_B$ and for each $i \leq m$, if r_i is a variable then $\langle \overline{q}_i, \overline{q}'_i\rangle := \theta(r_i)$, otherwise, if r_i is a constant then $\overline{q}_i$ is q_{r_i} and there is no restriction for $\overline{q}'_i$.

All the TAs have the same state set, hence the construction terminates with a fixpoint denoted A^*. We can now show that $L(A^*) = R^{\curlywedge}(L(A))$, more precisely, that for all $t \in \mathcal{T}(\Sigma)$, $t \in L\big(A^*, \langle q, q'\rangle\big)$ iff $t \in L(B, q')$ and there exists $s \in L(A, q)$

such that $s \xrightarrow[R]{\curlywedge} t$. To this end we follow the principle of the proofs given *e.g.* in [14,10,13], but some technical difficulties appear when we try to replace a subterm by another subterm while preserving an execution with Δ^*. They are solved thanks to the following technical Lemma 5.2.

The if direction is proved by induction on the number of rewrite steps in $s \xrightarrow[R]{\curlywedge} t$, using Lemma 5.2. The other direction is proved by an induction on the multiset associated to the derivation $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$ by mapping each transition rule ρ used to the least index i of the A_i to which ρ belongs.

Lemma 5.2 *For all $t \in \mathcal{T}(\Sigma \cup Q_{A^*})$, if $t\big[\langle q_0, q_{\mathsf{reject}} \rangle\big]_p \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ then, for all $q' \in Q_B$, there exists $q'' \in Q_B$ such that $t[\langle q_0, q' \rangle]_p \xrightarrow[\Delta^*]{*} \langle q, q'' \rangle$.*

Proof. First of all, we note that no ε-transition is applied in $t \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ at any position $p' \leq p$ *wrt* the prefix ordering. This is due to the fact that the ε-transitions are always of the form $\langle \_, q_{nr} \rangle \to \langle \_, q_{nr} \in Q_B \rangle$ for some q_{nr} different from q_{reject}, since this kind of transitions are derived in case 1 using rules where the right-hand side is a variable, and hence, it is necessarily instantiated by a normal form.

Hence, it suffices to prove the statement of the lemma for the case where t is of the form $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_m, q'_m \rangle)$, p is a certain position i, and there is just one rewrite step derivation with a rule of the form $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_i, q'_i \rangle, \ldots, \langle q_m, q'_m \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$. The global statement is then obtained inductively. In this case, q_i is q_0, and q'_i is q_{reject}. If the rule $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_i, q'_i \rangle, \ldots, \langle q_m, q'_m \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$ is in Δ_0, by the definition of Δ_0 we have that for any $q' \in Q_B$ there exists also an alternative rule of the form $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_i, q' \rangle, \ldots, \langle q_m, q'_m \rangle) \to \langle q, q'' \rangle$, for some q'', and we are done. Otherwise, assume that this rule $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_i, q'_i \rangle, \ldots, \langle q_m, q'_m \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$ is not in Δ_0. Then it has been derived by case 2 using a rule $\ell \to f(r_1, \ldots, r_m) \in R$. Moreover, by the conditions of case 2 and the fact that q'_i is q_{reject} we know that r_i is not a variable, and hence, it is a constant. Again by the conditions of case 2, for any $q' \in Q_B$ there exists also an alternative rule of the form $f(\langle q_1, q'_1 \rangle, \ldots, \langle q_i, q' \rangle, \ldots, \langle q_m, q'_m \rangle) \to \langle q, q'' \rangle$, for some q'', and we are done. $\square$

The number $|A^*|$ of states of A^* is at most $|A| \times |B|$, and the number of rules of A^* is polynomial in the same measure [3], if we assume as usual that the maximum arity of a function symbol is fixed for the problem.

Theorem 5.3 *$R^{\curlywedge}(L)$ is regular when L is regular and R is linear and right-shallow.*

Proof.

- Direction $\Leftarrow$ We prove it by induction on the number of rewrite steps of $s \xrightarrow[R]{\curlywedge} t$. For 0 rewrite steps we have $s = t$, and hence, $t \xrightarrow[\Delta]{*} q$. From the construction of Δ_0, $t \xrightarrow[\Delta_0]{*} \langle q, q' \rangle$ follows, and since $\Delta_0 \subseteq \Delta^*$ we also have $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$, and we are done. Hence, assume that there is at least one rewrite step in $s \xrightarrow[R]{\curlywedge} t$. We can write this derivation by making explicit the last rewrite step as $s \xrightarrow[R]{\curlywedge} t[\sigma(l)]_p \xrightarrow[\ell \to r, p, \sigma]{\imath} t[\sigma(r)]_p = t$. By induction hypothesis, $t[\sigma(\ell)]_p \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ (the q_{reject} is due to the fact that $t[\sigma(\ell)]_p$ is not a normal form). By re-ordering the rule applications of Δ^*, we can assume that this derivation is of the form

[3] Note however that $|B|$ can be exponential in the size of R in worst case.

$$t[\sigma(\ell)]_p \xrightarrow[\Delta^*]{*} t[\theta(\ell)]_p \xrightarrow[\Delta^*]{*} t[\langle q_0, q_{\mathsf{reject}}\rangle]_p \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}}\rangle$$

where θ is a mapping from variables to states of Δ^*. By the conditions for adding new rules into the Δ_j's, we also have $\theta(r) \xrightarrow[\Delta^*]{*} \langle q_0, q'_0\rangle$ for some $q'_0 \in Q_B$, and since the variables of r occur also in ℓ, there exists a derivation $\sigma(r) \xrightarrow[\Delta^*]{*} \theta(r) \xrightarrow[\Delta^*]{*} \langle q_0, q'_0\rangle$, and hence, by Lemma 5.2,

$$t[\sigma(r)]_p \xrightarrow[\Delta^*]{*} t[\theta(r)]_p \xrightarrow[\Delta^*]{*} t[\langle q_0, q'_0\rangle]_p \xrightarrow[\Delta^*]{*} \langle q, q''\rangle$$

for some q''. But this q'' is necessarily q' since Δ^* simulates B on the second component of the pairs, and we are done.

- Direction $\Rightarrow$ The fact that $t \xrightarrow[B]{*} q'$ follows directly from the construction of B: note that all the Δ_i's always simulate the execution of B for the second component of the pair of states. For proving that $\exists s : (s \xrightarrow[\Delta]{*} q \wedge s \xrightarrow[R]{\curlywedge} t)$, we do an induction on the measure defined as follows. For a concrete derivation $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$, we call $\mathsf{Rules}(t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle)$ the multiset of all rules used in it, and define $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle)$ as the multiset obtained by replacing every occurrence of a rule $\ell \to r$ in $\mathsf{Rules}(t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle)$ by the least index i such that $\ell \to r \in \Delta_i$. We compare indexes of derivations by the multiset extension of the usual ordering on natural numbers, and prove the statement of the lemma by induction on this ordering.

  Now, we distinguish cases depending on whether the last rewrite step of $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$ is an ε-transition or not.

  Assume first that it is an ε-transition. Then, this derivation can be written of the form $t \xrightarrow[\Delta^*]{*} \langle q_0, q'\rangle \xrightarrow[\Delta^*]{} \langle q, q'\rangle$ for a certain state $\langle q_0, q'\rangle$ of Q_{Δ^*}. The rule $\langle q_0, q'\rangle \to \langle q, q'\rangle$ is not in Δ_0, and has been necessarily added into some Δ_i with $i > 0$ by case 1, using a collapsing rule $\ell \to x \in R$ and a substitution θ such that $\theta(x) = \langle q_0, q'\rangle$ and $\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}}\rangle$. Moreover, for the rest of variables y of ℓ different from x, there exist terms $t_y \in \mathcal{T}(\Sigma)$ such that $t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y)$.

  We define a substitution σ such that $\sigma(y) = t_y$ for every of such y, and $\sigma(x) = t$. The substitution σ is well defined because R is right-linear.

  Hence, we have a one rewrite step derivation $\sigma(\ell) \xrightarrow[\ell \to x, \sigma]{\imath} t$, which is innermost due to the conditions for the addition of the rule to Δ_i (condition $q'_i \neq q_{\mathsf{reject}}$ there). Now, note that the multiset $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q_0, q'\rangle)$ has just one less i than $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle)$, and that the multiset $\mathsf{Ind}(\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}}\rangle)$ and the respective multisets $\mathsf{Ind}(t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y))$ contain numbers smaller than i. Therefore, composing the previous derivations we can construct a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}}\rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$. Hence, by induction hypothesis, there exists a term s such that $s \xrightarrow[R]{\curlywedge} \sigma(\ell)$ and $s \xrightarrow[\Delta]{*} q$, and, since t is R-reachable from s, we are done.

  Now, assume that the last rewrite step of $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$ is not an ε-transition. Hence, this derivation can be written by making explicit the last step as $t \xrightarrow[\Delta^*]{*} f(\langle q_1, q'_1\rangle, \ldots, \langle q_m, q'_m\rangle) \xrightarrow[\Delta^*]{} \langle q, q'\rangle$. Thus, if we write t of the form $f(t_1, \ldots, t_m)$, we can take (sub-)derivations $t_1 \xrightarrow[\Delta^*]{*} \langle q_1, q'_1\rangle$, $\ldots$,$t_m \xrightarrow[\Delta^*]{*} \langle q_m, q'_m\rangle$. Every $\mathsf{Ind}(t_i \xrightarrow[\Delta^*]{*} \langle q_i, q'_i\rangle)$ is smaller than $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle)$, and thus, by induction hypothesis there exist terms $s_1, \ldots s_m$ and derivations $s_i \xrightarrow[\Delta]{*} q_i$ and $s_i \xrightarrow[R]{\curlywedge} t_i$.

  At this point we distinguish cases depending on whether the rule

$f(\langle q_1, q_1' \rangle, \ldots, \langle q_m, q_m' \rangle) \rightarrow \langle q, q' \rangle$ belongs to Δ_0 or not. Assume first that it belongs to Δ_0. Then, by the definition of Δ_0, there exists also a rule $f(q_1, \ldots, q_m) \rightarrow q$ in Δ. Therefore, by composing this and the previous derivations, we have $f(s_1, \ldots, s_m) \xrightarrow[\Delta]{*} q$, but also $f(s_1, \ldots, s_m) \xrightarrow[R]{\curlywedge} f(t_1, \ldots, t_m) = t$, and we are done.

Now, assume that $f(\langle q_1, q_1' \rangle, \ldots, \langle q_m, q_m' \rangle) \rightarrow \langle q, q' \rangle$ has been added by case 2 into some Δ_i with $i > 0$ using a rule of the form $\ell \rightarrow f(r_1, \ldots, r_m)$ and a substitution θ satisfying the following properties. On the one side, $\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}} \rangle$. On the other side, for every r_i, if r_i is a variable then $\langle q_i, q_i' \rangle$ is $\theta(r_i)$ and $q_i' \neq q_{\mathsf{reject}}$, and otherwise, if r_i is a constant then q_i is q_{r_i} and there is no restriction on q_i'. Moreover, for the variables y occurring in ℓ and not in $f(r_1, \ldots, r_m)$ there exist terms $t_y \in \mathcal{T}(\Sigma)$ such that $t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y)$. We define a substitution σ such that $\sigma(y) = t_y$ for every of such y, and $\sigma(r_i) = t_i$ for the r_i that are variables. As above, σ is well defined because R is right-linear. Hence, we have a one rewrite step derivation $\sigma(\ell) \xrightarrow[\ell \rightarrow f(r_1, \ldots, r_m), \sigma]{} \sigma(f(r_1, \ldots, r_m))$, which is innermost due to the conditions for the addition of the rule to Δ_i, again. The terms $\sigma(f(r_1, \ldots, r_m))$ and $t = f(t_1, \ldots, t_m)$ can differ only in the positions i such that r_i is a constant. For an i of this kind, q_i is q_{r_i}, and the only term of $\mathcal{T}(\Sigma)$ that can be derived into q_{r_i} with Δ is r_i. Therefore, s_i is r_i, and hence, there exists a derivation $r_i \xrightarrow[R]{\curlywedge} t_i$.

This implies that $\sigma(f(r_1, \ldots, r_m)) \xrightarrow[R]{\curlywedge} t$. To conclude, it suffices to see that there exists a term s such that $s \xrightarrow[R]{\curlywedge} \sigma(\ell)$ and $s \xrightarrow[\Delta]{*} q$. To this end, we will prove that there exists a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$, and the statement will follow by induction hypothesis. But this is identical to what we have done in a previous case. Note that the multiset $\mathsf{Ind}(\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}} \rangle)$ and the respective multisets $\mathsf{Ind}(t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y))$ contain numbers smaller than i. Moreover, for the variables x that occur in $f(r_1, \ldots, r_m)$, left-linearity ensures that the indexes are preserved in these cases. Therefore, composing the previous derivations we can construct a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$, and we are done.

□

5.2 *Closure of TA languages with right-(linear and flat) TRS*

When we drop the restriction that R is left-linear in Theorem 5.3, we lose regularity preservation with innermost rewriting. This is in contrast with plain rewriting, where regularity is preserved for right-linear and right-shallow TRS [13].

Proposition 5.4 *In general, $R^{\curlywedge}(L)$ is not BT-regular when L is regular and R is right-linear and right-flat.*

Proof. Let $L = \{f(f(a, a), c)\}$, and $R = \{f(x, c) \rightarrow x, f(g(x), x) \rightarrow h(x), h(x) \rightarrow h(x), a \rightarrow g(a), a \rightarrow b\}$. The intersection of $R^{\curlywedge}(L)$ with the language of all terms containing only the symbols f, g, b is the set $\{f(g^n(b), g^m(b)) \mid n \neq m + 1\}$, which is not BT-regular. □

6 Conclusion and further work

We have covered much of the cases of closure of TA and BTTA languages by innermost rewriting, providing results for each case. The positive results are that the set of terms innermost-reachable from a regular language with a shallow TRS is BT-regular, and it is regular when the TRS is linear and right-shallow. Moreover, given a shallow TRS, regularity of the reachable terms from a regular language is decidable. Other consequences are the decidability of the problems of *ground reachability*, *ground joinability* and *regular tree model checking* (given two regular languages L_{init} and L_{bad} and the TRS R, do we have $R^{\curlywedge}(L_{\mathsf{init}}) \cap L_{\mathsf{bad}} = \emptyset$?) for innermost rewriting with TRS in the above classes.

As future work, it could be interesting to consider other variants of tree automata with more general or different constraints, and to consider other strategies of rewriting different from innermost.

References

[1] B. Bogaert, F. Seynhaeve, and S. Tison. The recognizability problem for tree automata with comparisons between brothers. In Wolfgang Thomas, editor, *Foundations of Software Science and Computation Structure, Second International Conference, FoSSaCS'99*, volume 1578 of *Lecture Notes in Computer Science*, pages 150–164. Springer, March 1999.

[2] B. Bogaert and S. Tison. Equality and disequality constraints on direct subterms in tree automata. In *International Symposium on Theoretical Aspects of Computeter Science*, pages 161–171, 1992.

[3] H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. Available on: `http://www.grappa.univ-lille3.fr/tata`, 2007.

[4] N. Dershowitz and J. P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science (Vol. B: Formal Models and Semantics)*, pages 243–320, Amsterdam, 1990. North-Holland.

[5] I. Durand and G. Sénizergues. Bottom-up rewriting is inverse recognizability preserving. In *18th Int. Conf. Term Rewriting and Applications (RTA'07)*, volume 4533 of *LNCS*, pages 107–121. Springer, 2007.

[6] R. Gilleron and S. Tison. Regular tree languages and rewrite systems. *Fundam. Inform.*, 24(1/2):157–174, 1995.

[7] G. Godoy and E. Huntingford. Innermost-reachability and innermost-joinability are decidable for shallow term rewrite systems. In F. Baader, editor, *Term Rewriting and Applications, 18th International Conference, RTA*, volume 4533 of *LNCS*, pages 184–199. Springer, June 2007.

[8] G. Godoy, E. Huntingford, and A.Tiwari. Termination of rewriting with right-flat rules. In F. Baader, editor, *Term Rewriting and Applications, 18th International Conference, RTA*, volume 4533 of *LNCS*, pages 200–213. Springer, June 2007.

[9] G. Godoy and A. Tiwari. Confluence of shallow right-linear rewrite systems. In C.-H. Luke Ong, editor, *19th International Workshop of Computer Science Logic, CSL*, volume 3634 of *LNCS*, pages 541–556. Springer, August 2005.

[10] F. Jacquemard. Decidable approximations of term rewriting systems. In *Rewriting Techniques and Applications, 7th International Conference*, pages 362–376, 1996.

[11] Y. Kojima and M. Sakai. Innermost reachability and context sensitive reachability properties are decidable for linear right-shallow term rewriting systems. In *Rewriting Techniques and Applications, 19th International Conference, 2008*, to appear.

[12] I. Mitsuhashi, M. Oyamaguchi, and F. Jacquemard. The confluence problem for flat TRSs. In *Proc. 8th Intl. Conf. on Artificial Intelligence and Symbolic Computation (AISC'06)*, volume 4120 of *LNAI*, pages 68–81. Springer, 2006.

[13] T. Nagaya and Y. Toyama. Decidability for left-linear growing term rewriting systems. In *Proc. 10th Intl. Conf. RTA*, volume 1631 of *LNCS*, pages 256–270. Springer, 1999.

[14] K. Salomaa. Deterministic tree pushdown automata and monadic tree rewriting systems. *J. Comput. Syst. Sci.*, 37:367–394, 1988.

[15] T. Takai, Y. Kaji, and H. Seki. Right-linear finite path overlapping term rewriting systems effectively preserve recognizability. In *Rewriting Techniques and Applications, RTA*, volume 1833 of *LNCS*, pages 246–260, 2000.

[16] W.S.Brainerd. Tree generating regular systems. *Information and Control*, 4:217–231, 1969.

Appendix

A Proof of Lemma 2.1

We present here the construction of a BTTA B recognizing R-normal forms for a flat TRS R.

Lemma 2.1 *[3] Let R be a flat TRS over Σ. There exists a normalized BTTA $B = (Q_B, Q^{\mathsf{f}}_B, \Delta_B)$ on Σ which recognizes the set of ground R-normal forms. Moreover $|Q_B \setminus Q^{\mathsf{f}}_B| = 1$.*

Proof. The construction of $B = (Q_B, Q_B \setminus \{q_{\mathsf{reject}}\}, \Delta_B)$ on Σ is as follows. Its set of states Q_B is $\{q_c \mid c \in \Sigma_0\} \cup \{q, q_{\mathsf{reject}}\}$ where all of them except q_{reject} are accepting states. Its set of rules Δ_B contains:

- the rules $c \to q_c$ for every constant c that is a R-normal form.
- the rules $c \to q_{\mathsf{reject}}$ for every constant c that is not a R-normal form.
- the rules $f(q_1, \dots q_m) \xrightarrow{c} q_{\mathsf{reject}}$ such that either some q_i is q_{reject} and c is *true*, or every q_i is different from q_{reject} and c is

$$\bigvee_{\substack{f(\ell_1 \dots \ell_m) \to r \in R \\ \forall 1 \le i \le m\ \ell_i \in \Sigma_0 \Rightarrow q_i = q_{\ell_i}}} \bigwedge_{1 \le i < j \le m\ \ell_i = \ell_j \in \mathcal{V}} i = j$$

- the rules $f(q_1 \dots q_m) \xrightarrow{c} q$ such that every q_i is different from q_{reject} and c is

$$\bigwedge_{\substack{f(\ell_1, \dots, \ell_m) \to r \in R \\ \forall 1 \le i \le m\ \ell_i \in \Sigma_0 \Rightarrow q_i = q_{\ell_i}}} \bigvee_{1 \le i < j \le m\ \ell_i = \ell_j \in \mathcal{V}} i \neq j$$

$\square$

B Simplifying assumptions

We show that the assumptions made in Section 4.2 are without loss of generality. Such assumptions have appeared in some previous works, specifically in [7].

First assumption: signature.

We assume that all terms are constructed over a given fixed signature Σ that contains several constants and only one non-constant function symbol f. If this was not the case, we can define a transformation T from terms over Σ into terms over a new signature Σ' as follows. Let m be the maximum arity of a symbol in Σ plus 1. We chose a new function symbol f with arity m and define the new signature $\Sigma' = \Sigma'_0 \cup \Sigma'_m$ as $\Sigma'_0 = \Sigma$ and $\Sigma'_m = \{f\}$. Note that all symbols of Σ appear also in Σ' but with arity 0. Now, we recursively define $T : \mathcal{T}(\Sigma, \mathcal{V}) \to \mathcal{T}(\Sigma', \mathcal{V})$ as $T(c) = c$ and $T(x) = x$ for constants $c \in \Sigma_0$ and variables $x \in \mathcal{V}$, and $T(gt_1 \dots t_k) = f(T(t_1), \dots, T(t_k), g, \dots, g)$ for terms headed with $g \in \Sigma \setminus \Sigma_0$. We denote $\{T(l) \to T(r) \mid l \to r \in R\}$ as $T(R)$ for a given TRS R. Note that the size of $T(R)$ is at most m times the size of R, and hence, this transformation can be easily obtained in polynomial time.

Lemma B.1 *Let s, t be terms in $\mathcal{T}(\Sigma, \mathcal{V})$ and R a TRS with all its terms in $\mathcal{T}(\Sigma, \mathcal{V})$. Then, $s \xrightarrow[R]{\imath} t$ iff $T(s) \xrightarrow[T(R)]{\imath} T(t)$, and $s \xrightarrow[R]{\curlywedge} t$ iff $T(s) \xrightarrow[T(R)]{\curlywedge} T(t)$. Moreover, for all $u \in \mathcal{T}(\Sigma, \mathcal{V})$ and all $v' \in \mathcal{T}(\Sigma', \mathcal{V})$, if $T(u) \xrightarrow[T(R)]{\curlywedge} v'$, then there exists $v \in \mathcal{T}(\Sigma, \mathcal{V})$ such that $T(v) = v'$.*

Proof. By the recursive definition of T it is easy to see that $T(u[v]_p) = T(u)[T(v)]_p$ for all terms $u, v \in \mathcal{T}(\Sigma, \mathcal{V})$. Moreover, for any substitution $\sigma : \mathcal{V} \to \mathcal{T}(\Sigma, \mathcal{V})$, $T(u\sigma) = T(u)T(\sigma)$, where we understand $T(\sigma)$ as the substitution $\{x \mapsto T(x\sigma) \mid x \in dom(\sigma)\}$, i.e. the substitution holding $xT(\sigma) = T(x\sigma)$. Hence, for every rule $l \to r \in R$, it holds that $u \xrightarrow[l \to r, p, \sigma]{\imath} v$ iff $T(u) \xrightarrow[T(l) \to T(r), p, T(\sigma)]{\imath} T(v)$. Moreover, all rule applications with $T(R)$ on terms of the form $T(u)$ are always of the form $T(u) \xrightarrow[T(l) \to T(r), p]{\imath} T(v)$. The lemma follows from these considerations.

□

Recall that our goal is to compute a BTTA for $R^{\curlywedge}(L)$ for a given regular language L and TRS R. After the previous change the TRS R' works on a different signature. Thus we need to adapt the TA A for L. This can be done easily with the transformation $T(g(q_1, \ldots, q_n) \to q) \mapsto (f(q_1, \ldots, q_n, q_g, \ldots, q_g) \to q)$ for every rule $g(q_1, \ldots, q_n) \to q$, and adding a rule $g \to q_g$ for every new constant g. If L' is the language recognized by A', after computing a BTTA A'_2 for $R'^{\curlywedge}(L')$ we can retrieve an automaton for $R^{\curlywedge}(L)$ as follows. Without loss of generality we can assume that for every new constant g, there is a state in A'_2, which we call q'_g, such that g is the only term that reaches q'_g with R'. Note also that, by previous lemma, the reachable terms from L' are images of terms of the original signature by T. Hence, we can just invert the transformation for the rules of the form $f(q_1, \ldots, q_n, q'_g, \ldots, q'_g) \xrightarrow{P \cup \{\{n+1, \ldots, m\}\}} q$ by converting these rules into $g(q_1, \ldots, q_n) \xrightarrow{P} q$.

Second assumption: flat TRS.

Finally we do a last transformation that allows us to assume that all rules in R are flat. If this was not the case, we can modify R as follows. First remove all rules $l \to r$ such that l has a proper subterm that is not a normal form: note that in such a case this rule is useless for innermost rewriting. Second, we proceed by applying several times the following transformation step a), until the obtained TRS is left-flat.

step a) If there is a non-constant term t that is a proper subterm of a left-hand side of a rule[4] in R, then create a new constant c, replace all occurrences of t in the left-hand sides of the rules of R by c, and add the rule $t \to c$ to R.

In the following two lemmas and proofs we assume that R' is obtained from R by the previous transformation with t and c, Σ is the original signature, and Σ' is

[4] Note that the subterm t is a normal form.

$\Sigma \cup \{c\}$.

Lemma B.2 *Let s_1, s_2 be terms in $\mathcal{T}(\Sigma, \mathcal{V})$. Then, $s_1 \xrightarrow[R]{\curlywedge} s_2$ iff $s_1 \xrightarrow[R']{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$.*

Proof.

- Direction $\Rightarrow$

  We first show that if $s_1 \xrightarrow[R]{\imath} s_2$, then $\mathsf{NF}_{\{t\to c\}}(s_1) \xrightarrow[R]{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$. This is because if $l \to r \in R$ is the used rule in $s_1 \xrightarrow[R]{\imath} s_2$, then, using the corresponding $l[c]\ldots[c] \to r$ obtained by the transformation from R into R' we have $\mathsf{NF}_{\{t\to c\}}(s_1) \xrightarrow[l[c]\ldots[c]\to r]{\imath} \xrightarrow[t\to c]{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$.

  Now, from any derivation $s_1 \xrightarrow[R]{\curlywedge} s_2$ can be obtained a derivation $s_1 \xrightarrow[t\to c]{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_1) \xrightarrow[R']{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$ by applying several times the previous observation.

- Direction $\Leftarrow$

  On the other side, consider any rewrite step $u \xrightarrow[R']{\imath} v$. It holds that $\mathsf{NF}_{\{c\to t\}}(u) \xrightarrow[R]{0,1} \mathsf{NF}_{\{c\to t\}}(v)$. Therefore, since s_1 and s_2 are terms without the constant c, $s_1 \xrightarrow[R']{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$ implies $s_1 = \mathsf{NF}_{\{c\to t\}}(s_1) \xrightarrow[R]{\curlywedge} \mathsf{NF}_{\{c\to t\}}(\mathsf{NF}_{\{t\to c\}}(s_2)) = s_2$.

$\square$

The new R' does not necessarily preserve the language of reachable terms from a given regular set. Since our global goal is to justify that this reachable set is BT-regular, we will need to retrieve it by, essentially, inverting the transformation. This can be done using Lemma **??** below.

At this point we assume that L is a regular language and A' is a BTTA recognizing $R'^{\curlywedge}(L)$. Without loss of generality, we assume that A' has two states q_c and q_t such that c and t are the only terms that run to q_c and q_t, respectively, by A'. Moreover, for the case of c we assume the existence of a rule $c \to q_c$. We construct a BTTA A for $R^{\curlywedge}(L)$ by removing $c \to q_c$ from A' and adding $q_t \to q_c$.

Lemma B.3 *$L(A) = R^{\curlywedge}(L)$.*

Proof.

$\supseteq$) Let $s_2 \in R^{\curlywedge}(L)$. Then, there exists a term $s_1 \in L$ such that $s_1 \xrightarrow[R]{\curlywedge} s_2$. By Lemma **??** $s_1 \xrightarrow[R']{\curlywedge} \mathsf{NF}_{\{t\to c\}}(s_2)$. Therefore, A' recognizes $s_2' = \mathsf{NF}_{\{t\to c\}}(s_2)$, and by the construction of A, $s_2 \in L(A)$.

$\subseteq$) Let $s_2 \in L(A)$. Then, there exists a derivation from s_2 into an accepting state of A using the rules of A. Let s_2'' be obtained from s_2 by replacing the subterms t that are rewritten to q_c in the previous derivation by c. Then, $s_2'' \in L(A')$, and hence, there exists $s_1 \in L$ such that $s_1 \xrightarrow[R']{\curlywedge} s_2''$. Let s_2' be obtained from s_2'' (or from s_2, it does not matter) by replacing all the occurrences of t by c. Since the rule $t \to c$ is in R', we also have $s_1 \xrightarrow[R']{\curlywedge} s_2'$, and since $s_2' = \mathsf{NF}_{t\to c}(s_2)$, by Lemma **??**, $s_1 \xrightarrow[R]{\curlywedge} s_2$ and hence $s_2 \in R^{\curlywedge}(L)$.

$\square$

Finally, we apply several times the following transformation step until the TRS is right-flat.

step b) If there is a non-constant term t that is a proper subterm of a right-hand side of a rule in R, then create a new constant c, replace all occurrences of t in the right-hand sides of the rules of R by c, and add the rule $c \to t$ to R.

Lemma B.4 *Let s_1, s_2 be terms in $\mathcal{T}(\Sigma, \mathcal{V})$ and let R' be obtained from R by the previous transformation step b). Then, $s_1 \xrightarrow[R]{\curlywedge} s_2$ iff $s_1 \xrightarrow[R']{\curlywedge} s_2$.*

Since in step b) the reachable terms over the original signature are preserved, a simple intersection allows to retrieve the desired BTTA.

Every step (a or b) decreases the total sum of the number of positions at depth more than one in all the left-hand and right-hand sides of R. Moreover, also at every step, the total size of the TRS increases by the size occupied by a constant. Hence, this process terminates in polynomial time.

C Proof of Lemma 4.4

Lemma 4.4 *For every flat TRS R and regular language L, over a signature Σ, there exists an extension $\Sigma' \supset \Sigma$, a constant $c \in \Sigma' \setminus \Sigma$ and a flat TRS R_c over Σ' such that $R_c^{\curlywedge}(\{c\}) \cap \mathcal{T}(\Sigma) = R^{\curlywedge}(L)$.*

Proof. Let $A = (Q, Q^{\mathsf{f}}, \Delta)$ be a tree automaton on Σ recognizing L. Without loss of generality we assume that every state $q \in Q$ is the target of a rule of Δ. Let Σ' be $\Sigma \cup Q \cup \{c\}$, where c is a new constant not in $\Sigma \cup Q$ and let R_c be $R \cup \Delta^{-1} \cup \{c \to q \mid q \in Q^{\mathsf{f}}\}$. We prove that $R_c^{\curlywedge}(\{c\}) \cap \mathcal{T}(\Sigma) = R^{\curlywedge}(L)$.

Direction $\supseteq$.

Let $t \in R^{\curlywedge}(L)$ and let $s \in L$ such that $s \xrightarrow[R]{\curlywedge} t$. We have $s \xrightarrow[\Delta]{\curlywedge} q \in Q^{\mathsf{f}}$. Therefore $c \xrightarrow[c \to q]{\imath} q \xrightarrow[\Delta^{-1}]{\curlywedge} s \xrightarrow[R]{\curlywedge} t$ and hence $c \xrightarrow[R_c]{\curlywedge} t$.

Direction $\subseteq$.

Let $t \in R_c^{\curlywedge}(\{c\}) \cap \mathcal{T}(\Sigma)$. We consider a concrete derivation $c \xrightarrow[R_c]{\curlywedge} t$. Since $c \notin \Sigma$, there is at least one rewrite step in this derivation, which is necessarily of the form $c \xrightarrow[R_c]{\imath} q$ for $q \in Q^{\mathsf{f}}$. The subderivation $q \xrightarrow[R_c]{\curlywedge} t$ can contain alternate rewrite steps using rules of Δ^{-1} or of R. We want to show that these rewrite steps can be commuted such that a $(\Delta^{-1} \cup R)$-innermost derivation of the form $q \xrightarrow[\Delta^{-1}]{*} s \xrightarrow[R]{*} t$ is possible. To this end it suffices to see that any $(\Delta^{-1} \cup R)$-innermost subderivation $u \xrightarrow[R,p_1]{} v \xrightarrow[\Delta^{-1},p_2]{} w$ can be commuted to a $(\Delta^{-1} \cup R)$-innermost subderivation $u \xrightarrow[\Delta^{-1},p_2]{} v' \xrightarrow[R,p_1]{} w$, which will be straightforward if we prove that the positions p_1 and p_2 are disjoint. Note that the term $u|_{p_1}$ does not contain any symbol $q \in Q$: otherwise the rewrite step $u \xrightarrow[R,p_1]{} v$ would not be $(\Delta^{-1} \cup R)$-innermost since there exists a rule of the form $q \to r$ in Δ^{-1} due to our assumptions. Hence, $v|_{p_1}$ does not contain any symbol $q \in Q$ neither. Therefore, the rewrite step $v \xrightarrow[\Delta^{-1},p_2]{} w$ is produced at a position p_2 disjoint with p_1, and we are done. Hence, there exists a $(\Delta^{-1} \cup R)$-innermost derivation $q \xrightarrow[\Delta^{-1}]{*} s \xrightarrow[R]{*} t$. In order to prove that $t \in R^{\curlywedge}(L)$ it suffices to see that $s \in L$. Since $s \xrightarrow[\Delta]{\curlywedge} q \in Q^{\mathsf{f}}$, it is enough to prove that $s \in \mathcal{T}(\Sigma)$, i.e. that s does not contain any symbol $q \in Q$. Suppose that $s|_p$ is a certain $q \in Q$.

Similarly to before, no rule in R can be applied to a position $p' < p$, since otherwise this rewrite step would not be innermost. Hence, repeated applications of rules of R under innermost rewriting do not remove q. Therefore q is a symbol occurring in t, a contradiction.

$\square$

D Proof of Lemma 5.2

Proof. First of all, we note that no ε-transition is applied in $t \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ at any position $p' \leq p$ *wrt* the prefix ordering. This is due to the fact that the ε-transitions are always of the form $\langle \_, q_{nr} \rangle \to \langle \_, q_{nr} \in Q_B \rangle$ for some q_{nr} different from q_{reject}, since this kind of transitions are derived in case 1 using rules where the right-hand side is a variable, and hence, it is necessarily instantiated by a normal form.

Hence, it suffices to prove the statement of the lemma for the case where t is of the form $f(\langle q_1, q_1' \rangle, \ldots, \langle q_m, q_m' \rangle)$, p is a certain position i, and there is just one rewrite step derivation with a rule of the form $f(\langle q_1, q_1' \rangle, \ldots, \langle q_i, q_i' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$. The global statement is then obtained inductively. In this case, q_i is q_0, and q_i' is q_{reject}. If the rule $f(\langle q_1, q_1' \rangle, \ldots, \langle q_i, q_i' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$ is in Δ_0, by the definition of Δ_0 we have that for any $q' \in Q_B$ there exists also an alternative rule of the form $f(\langle q_1, q_1' \rangle, \ldots, \langle q_i, q' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q'' \rangle$, for some q'', and we are done. Otherwise, assume that this rule $f(\langle q_1, q_1' \rangle, \ldots, \langle q_i, q_i' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q_{\mathsf{reject}} \rangle$ is not in Δ_0. Then it has been derived by case 2 using a rule $\ell \to f(r_1, \ldots, r_m) \in R$. Moreover, by the conditions of case 2 and the fact that q_i' is q_{reject} we know that r_i is not a variable, and hence, it is a constant. Again by the conditions of case 2, for any $q' \in Q_B$ there exists also an alternative rule of the form $f(\langle q_1, q_1' \rangle, \ldots, \langle q_i, q' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q'' \rangle$, for some q'', and we are done. $\square$

E Proof of Theorem 5.3

The proof of the correctness of the construction of A^* (page 11) relies on Lemma 5.1.

Lemma 5.1 *For all $t \in \mathcal{T}(\Sigma)$, $t \in L\big(A^*, \langle q, q' \rangle\big)$ iff $t \xrightarrow[B]{*} q'$ and there exists $s \in L(A, q)$ such that $s \xrightarrow[R]{\lambda} t$.*

Proof. *Direction* $\Leftarrow$. We prove it by induction on the number of rewrite steps of $s \xrightarrow[R]{\lambda} t$. For 0 rewrite steps we have $s = t$, and hence, $t \xrightarrow[\Delta]{*} q$. From the construction of Δ_0, $t \xrightarrow[\Delta_0]{*} \langle q, q' \rangle$ follows, and since $\Delta_0 \subseteq \Delta^*$ we also have $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$, and we are done. Hence, assume that there is at least one rewrite step in $s \xrightarrow[R]{\lambda} t$. We can write this derivation by making explicit the last rewrite step as $s \xrightarrow[R]{\lambda} t[\sigma(l)]_p \xrightarrow[\ell \to r, p, \sigma]{\imath} t[\sigma(r)]_p = t$. By induction hypothesis, $t[\sigma(\ell)]_p \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ (the q_{reject} is due to the fact that $t[\sigma(\ell)]_p$ is not a normal form). By re-ordering the rule applications of Δ^*, we can assume that this derivation is of the form

$$t[\sigma(\ell)]_p \xrightarrow[\Delta^*]{*} t[\theta(\ell)]_p \xrightarrow[\Delta^*]{*} t[\langle q_0, q_{\mathsf{reject}} \rangle]_p \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$$

where θ is a mapping from variables to states of Δ^*. By the conditions for adding new rules into the Δ_j's, we also have $\theta(r) \xrightarrow[\Delta^*]{*} \langle q_0, q_0' \rangle$ for some $q_0' \in Q_B$, and since the variables of r occur also in ℓ, there exists a derivation $\sigma(r) \xrightarrow[\Delta^*]{*} \theta(r) \xrightarrow[\Delta^*]{*} \langle q_0, q_0' \rangle$, and hence, by Lemma 5.2,

$$t[\sigma(r)]_p \xrightarrow[\Delta^*]{*} t[\theta(r)]_p \xrightarrow[\Delta^*]{*} t[\langle q_0, q_0' \rangle]_p \xrightarrow[\Delta^*]{*} \langle q, q'' \rangle$$

for some q''. But this q'' is necessarily q' since Δ^* simulates B on the second component of the pairs, and we are done.

Direction $\Rightarrow$. The fact that $t \xrightarrow[B]{*} q'$ follows directly from the construction of B: note that all the Δ_i's always simulate the execution of B for the second component of the pair of states. For proving that $\exists s : (s \xrightarrow[\Delta]{*} q \wedge s \xrightarrow[R]{\curlywedge} t)$, we do an induction on the measure defined as follows. For a concrete derivation $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$, we call $\mathsf{Rules}(t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle)$ the multiset of all rules used in it, and define $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle)$ as the multiset obtained by replacing every occurrence of a rule $\ell \to r$ in $\mathsf{Rules}(t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle)$ by the least index i such that $\ell \to r \in \Delta_i$. We compare indexes of derivations by the multiset extension of the usual ordering on natural numbers, and prove the statement of the lemma by induction on this ordering.

Now, we distinguish cases depending on whether the last rewrite step of $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$ is an ε-transition or not.

Assume first that it is an ε-transition. Then, this derivation can be written of the form $t \xrightarrow[\Delta^*]{*} \langle q_0, q' \rangle \xrightarrow[\Delta^*]{} \langle q, q' \rangle$ for a certain state $\langle q_0, q' \rangle$ of Q_{Δ^*}. The rule $\langle q_0, q' \rangle \to \langle q, q' \rangle$ is not in Δ_0, and has been necessarily added into some Δ_i with $i > 0$ by case 1, using a collapsing rule $\ell \to x \in R$ and a substitution θ such that $\theta(x) = \langle q_0, q' \rangle$ and $\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}} \rangle$. Moreover, for the rest of variables y of ℓ different from x, there exist terms $t_y \in \mathcal{T}(\Sigma)$ such that $t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y)$.

We define a substitution σ such that $\sigma(y) = t_y$ for every of such y, and $\sigma(x) = t$. The substitution σ is well defined because R is right-linear.

Hence, we have a one rewrite step derivation $\sigma(\ell) \xrightarrow[\ell \to x, \sigma]{\imath} t$, which is innermost due to the conditions for the addition of the rule to Δ_i (condition $q_i' \neq q_{\mathsf{reject}}$ there). Now, note that the multiset $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q_0, q' \rangle)$ has just one less i than $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle)$, and that the multiset $\mathsf{Ind}(\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}} \rangle)$ and the respective multisets $\mathsf{Ind}(t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y))$ contain numbers smaller than i. Therefore, composing the previous derivations we can construct a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}} \rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$. Hence, by induction hypothesis, there exists a term s such that $s \xrightarrow[R]{\curlywedge} \sigma(\ell)$ and $s \xrightarrow[\Delta]{*} q$, and, since t is R-reachable from s, we are done.

Now, assume that the last rewrite step of $t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle$ is not an ε-transition. Hence, this derivation can be written by making explicit the last step as $t \xrightarrow[\Delta^*]{*} f(\langle q_1, q_1' \rangle, \ldots, \langle q_m, q_m' \rangle) \xrightarrow[\Delta^*]{} \langle q, q' \rangle$. Thus, if we write t of the form $f(t_1, \ldots, t_m)$, we can take (sub-)derivations $t_1 \xrightarrow[\Delta^*]{*} \langle q_1, q_1' \rangle$, $\ldots$, $t_m \xrightarrow[\Delta^*]{*} \langle q_m, q_m' \rangle$. Every $\mathsf{Ind}(t_i \xrightarrow[\Delta^*]{*} \langle q_i, q_i' \rangle)$ is smaller than $\mathsf{Ind}(t \xrightarrow[\Delta^*]{*} \langle q, q' \rangle)$, and thus, by induction hypothesis there exist terms $s_1, \ldots s_m$ and derivations $s_i \xrightarrow[\Delta]{*} q_i$ and $s_i \xrightarrow[R]{\curlywedge} t_i$.

At this point we distinguish cases depending on whether the rule $f(\langle q_1, q_1' \rangle, \ldots, \langle q_m, q_m' \rangle) \to \langle q, q' \rangle$ belongs to Δ_0 or not. Assume first that it belongs to Δ_0. Then, by the definition of Δ_0, there exists also a rule $f(q_1, \ldots, q_m) \to q$ in Δ. Therefore, by composing this and the previous derivations, we have $f(s_1, \ldots, s_m) \xrightarrow[\Delta]{*} q$, but also $f(s_1, \ldots, s_m) \xrightarrow[R]{\curlywedge} f(t_1, \ldots, t_m) = t$, and we are done.

Now, assume that $f(\langle q_1, q'_1\rangle, \ldots, \langle q_m, q'_m\rangle) \to \langle q, q'\rangle$ has been added by case 2 into some Δ_i with $i > 0$ using a rule of the form $\ell \to f(r_1, \ldots, r_m)$ and a substitution θ satisfying the following properties. On the one side, $\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}}\rangle$. On the other side, for every r_i, if r_i is a variable then $\langle q_i, q'_i\rangle$ is $\theta(r_i)$ and $q'_i \neq q_{\mathsf{reject}}$, and otherwise, if r_i is a constant then q_i is q_{r_i} and there is no restriction on q'_i. Moreover, for the variables y occurring in ℓ and not in $f(r_1, \ldots, r_m)$ there exist terms $t_y \in \mathcal{T}(\Sigma)$ such that $t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y)$. We define a substitution σ such that $\sigma(y) = t_y$ for every of such y, and $\sigma(r_i) = t_i$ for the r_i that are variables. As above, σ is well defined because R is right-linear. Hence, we have a one rewrite step derivation $\sigma(\ell) \xrightarrow[\ell \to f(r_1,\ldots,r_m),\sigma]{} \sigma(f(r_1, \ldots, r_m))$, which is innermost due to the conditions for the addition of the rule to Δ_i, again. The terms $\sigma(f(r_1, \ldots, r_m))$ and $t = f(t_1, \ldots, t_m)$ can differ only in the positions i such that r_i is a constant. For an i of this kind, q_i is q_{r_i}, and the only term of $\mathcal{T}(\Sigma)$ that can be derived into q_{r_i} with Δ is r_i. Therefore, s_i is r_i, and hence, there exists a derivation $r_i \xrightarrow[R]{\curlywedge} t_i$.

This implies that $\sigma(f(r_1, \ldots, r_m)) \xrightarrow[R]{\curlywedge} t$. To conclude, it suffices to see that there exists a term s such that $s \xrightarrow[R]{\curlywedge} \sigma(\ell)$ and $s \xrightarrow[\Delta]{*} q$. To this end, we will prove that there exists a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}}\rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$, and the statement will follow by induction hypothesis. But this is identical to what we have done in a previous case. Note that the multiset $\mathsf{Ind}(\theta(\ell) \xrightarrow[\Delta_{i-1}]{*} \langle q, q_{\mathsf{reject}}\rangle)$ and the respective multisets $\mathsf{Ind}(t_y \xrightarrow[\Delta_{i-1}]{*} \theta(y))$ contain numbers smaller than i. Moreover, for the variables x that occur in $f(r_1, \ldots, r_m)$, left-linearity ensures that the indexes are preserved in these cases. Therefore, composing the previous derivations we can construct a derivation $\sigma(\ell) \xrightarrow[\Delta^*]{*} \langle q, q_{\mathsf{reject}}\rangle$ smaller than $t \xrightarrow[\Delta^*]{*} \langle q, q'\rangle$, and we are done. $\square$