

From Outermost to Context-Sensitive Rewriting

Jörg Endrullis and Dimitri Hendriks

Vrije Universiteit Amsterdam
{joerg,diem}@few.vu.nl

Abstract. We define a transformation from term rewriting systems (TRSs) to context-sensitive TRSs in such a way that termination of the target system implies outermost termination of the original system. For the class of left-linear TRSs the transformation is complete. Thereby state-of-the-art termination methods and automated termination provers for context-sensitive rewriting become available for proving termination of outermost rewriting. The translation has been implemented in Jambox, making it the most successful tool in the category of outermost rewriting of the last edition of the annual termination competition.

1 Introduction

Termination is a key aspect of program correctness, and therefore a widely studied subject in term rewriting and program verification. While termination is undecidable in general, various automated techniques have been developed for proving termination. One of the most powerful techniques is the method of dependency pairs [2]. Recently [1], this method has been generalized to context-sensitive TRSs, thereby significantly extending the class of context-sensitive TRSs for which termination can be shown automatically. Context-sensitive rewriting [6] is a restriction on term rewriting where rewriting in some fixed arguments of function symbols is disallowed. It offers a flexible paradigm to analyze properties of rewrite strategies, in particular of (lazy) evaluation strategies employed by functional programming languages.

In this paper context-sensitive rewriting is the target formalism for a transformational approach to the problem of outermost termination, that is, termination with respect to outermost rewriting. Outermost rewriting is a rewriting strategy where a redex may be contracted as long as it is not a proper subterm of another redex occurrence. The main reason for studying outermost termination is its practical relevancy: lazy functional programming languages like Miranda, Haskell or Clean, are based on outermost rewriting as an evaluation strategy, and in implementations of rewrite logic such as Maude and CafeOBJ, outermost rewriting can be specified. Consider the TRS R_0 consisting of the following rules:

$$a \rightarrow f(a) \qquad f(f(x)) \rightarrow b \qquad (R_0)$$

Clearly, this system is not terminating as witnessed by the infinite rewrite sequence $a \rightarrow f(a) \rightarrow f(f(a)) \rightarrow f(f(f(a))) \rightarrow \dots$, but it is outermost terminating. Indeed, the third step in the infinite sequence is not an outermost step, since

the contraction takes place inside another redex. The only (maximal) outermost rewrite sequence the term a admits is $a \rightarrow f(a) \rightarrow f(f(a)) \rightarrow b$.

Our contribution is a transformation of arbitrary TRSs into context-sensitive TRSs (μ TRSs) in such a way that rewriting in the μ TRS corresponds to outermost rewriting in the original TRS. As a result advanced termination techniques for μ TRSs become applicable for proving outermost termination, and automated termination provers for μ TRSs can directly (without modification, only preprocessing) be used for proving outermost termination. Our transformation is complete for the class of quasi left-linear TRSs (a generalized form of left-linear TRSs, see [8]), that is, termination of the resulting μ TRS is equivalent to outermost termination of the original system.

The transformation is comprised of a variant of semantic labeling [12]. In semantic labeling the function symbols in a term are labeled by the interpretation of their arguments (or a label depending on these values) according to some given semantics. We employ semantic labeling in order to mark symbols at redex positions, and we obtain a μ TRS by defining a replacement map that disallows rewriting inside arguments of marked symbols.

We illustrate our use of semantic labeling by the TRS R_0 from the first page. We choose the algebra $\mathcal{A}_0 = \langle \{0, 1\}, [\cdot] \rangle$ where the interpretation indicates the presence of the symbol f , thus $[a] = [b] = 0$, and $[f](x) = 1$ for $x \in \{0, 1\}$. Then we write f^* if the value of its argument is 1, and just f if the value is 0. The symbol a is always marked, and b never is. If f is marked it corresponds to a redex position with respect to the rule $f(f(x)) \rightarrow b$. For example the term $f(f(f(a)))$ is labeled as $f^*(f^*(f(a^*)))$. We obtain a μ TRS by forbidding rewriting inside the argument of the symbol f^* ; since a^* is a constant, there is nothing to be forbidden. Then for correctly labeled terms, rewriting inside redex positions is disallowed, corresponding to the outermost rewriting strategy. In order to rewrite labeled terms we have to label the rules of the TRS. By labeling R_0 with the algebra \mathcal{A}_0 we obtain the μ TRS \bar{R}_0 :

$$a^* \rightarrow f(a^*) \qquad f^*(f(x)) \rightarrow b \qquad f^*(f^*(x)) \rightarrow b \qquad (\bar{R}_0)$$

which has two instances of the second rule, one for each possible value of x .

Now, despite the fact that the original TRS is outermost terminating, the transformed μ TRS \bar{R}_0 admits an infinite rewrite sequence:

$$a^* \rightarrow f(a^*) \rightarrow f(f(a^*)) \rightarrow f(f(f(a^*))) \rightarrow \dots \qquad (1)$$

The reason is that the term $f(f(a^*))$ is not correctly labeled, as the root symbol f should have been marked. In [12] this problem is avoided by allowing labeling only with ‘models’. Roughly, an algebra is a model if left- and right-hand sides of all rewrite rules have equal interpretations. However, this requirement is too strict for the purpose of marking redexes, because contraction of a redex at a position p may create a redex above p in the term tree, as witnessed by (1). In fact, for R_0 there exists *no* model which is able to distinguish between redex and non-redex positions. The rewrite step $f(a) \rightarrow f(f(a))$ creates a redex at the top. The term $f(a)$ is not a redex, and therefore its root symbol f should not

be marked. On the other hand $f(f(a))$ is a redex and so the outermost f has to be marked. The change of the labeling of a context (here $f(\square)$) implies that the interpretation of its arguments a and $f(a)$ cannot be the same. Therefore we cannot require the rule $a \rightarrow f(a)$ to preserve the interpretation.

For this reason, we generalize the concept of model, and relax the requirement $[\ell] = [r]$ to $\exists n. [C[\ell]] = [C[r]]$ for all contexts C of depth n . Thus rules are allowed to change the interpretation as long as the effect is limited to contexts of a bounded depth. Algebras satisfying this weaker requirement, which we call ‘ C -models’, are strong enough to recognize redex positions. In particular, the algebra \mathcal{A}_0 given above is a C -model for the system R_0 . As demonstrated by the rewrite sequence (1) in the μ TRS \bar{R}_0 , for C -models it is no longer sufficient to simply label the rules: an application of the rule $a^* \rightarrow f(a^*)$ in the term $f(a^*)$ creates the invalid labeled term $f(f(a^*))$. In order to preserve correct labeling, we sometimes need to extend the rewrite rules by putting contexts around their left- and right-hand sides. Using the C -model \mathcal{A}_0 , our algorithm transforms R_0 into the μ TRS $\Delta_{\mathcal{A}_0}^\pi(R_0)$, which truthfully simulates outermost rewriting of R_0 :

$$\begin{array}{ll} f(a^*) \rightarrow f^*(f(a^*)) & \text{top}(f^*(f(x))) \rightarrow \text{top}(b) \\ \text{top}(a^*) \rightarrow \text{top}(f(a^*)) & \text{top}(f^*(f^*(x))) \rightarrow \text{top}(b) \end{array} \quad (\Delta_{\mathcal{A}_0}^\pi(R_0))$$

We explain this magical transformation. The rule $f(a^*) \rightarrow f^*(f(a^*))$ is obtained from prepending the context $f(\square)$ to $a \rightarrow f(a)$. This enables correct updating of the labeling of the context during rewriting. Because we still have to allow rewrite steps $a \rightarrow f(a)$ of the original TRS at the top of a term, we extend the signature with a unary function symbol top which represents the top of a term. Thus when prepending contexts we include $\text{top}(\square)$, giving rise to the rule $\text{top}(a^*) \rightarrow \text{top}(f(a^*))$. The necessity of the symbol top becomes especially apparent when considering the rule $f(f(x)) \rightarrow b$. Here prepending the context $f(\square)$ is not even an option since $f(f(f(x))) \rightarrow f(b)$ is not an outermost rewrite step; this rule can only be applied at the top of a term. Hence we get the two rules displayed on the right, one for each possible interpretation of the variable x .

Semantic labeling has the nice property that it does not complicate termination proofs. Although semantic labeling increases the search space, termination proofs for the unlabeled system carry over to the labeled one. That is, whenever R' is a labeling of a TRS R and $\mathcal{A} = \langle A, [\cdot], \succ, \sqsubseteq \rangle$ is a monotone Σ -algebra [4] which proves termination of R , then extending $[\cdot]$ to the labeled signature Σ' by $[f^\lambda] = [f]$ for every $f \in \Sigma$ and label λ , yields a monotone Σ' -algebra which proves termination of R' . Consequently our transformation, based on a variant of semantic labeling, also does not complicate termination proofs. On the contrary, the labeled systems often allow for simpler proofs arising from the enriched signature which provides more freedom for the choice of interpretations, see [12].

Semantic labeling possibly creates extra copies of rules, and our extension might even create more copies: one for each context that has to be prepended. Despite of this fact, the implementation of our transformation in the termination prover Jambox performs efficiently on the set of examples from the Termination Problem Database (TPDB [10]).

Jambox was best in proving termination in the category of outermost rewriting of the termination competition of 2008 [10], see Figure 1. With an average time of 4.1 seconds per termination proof, Jambox was also faster than the other participants, providing empirical evidence for the efficiency of our transformation. Not listed in Figure 1 is TTT2, which did not prove outermost termination, but performed best in disproving outermost termination.

	score	average time
Jambox	72 (93.5%)	4.1s
TrafO	46 (59.7%)	8.1s
AProVE	27 (35.0%)	10.8s

Fig. 1. Results of proving outermost termination in the competition of 2008.

The secret behind the efficiency of Jambox is threefold: First, we construct and minimize the algebras employed for marking redex positions, see Sections 4 and 5. Secondly, we try two labeling strategies: minimal and maximal. Minimal labeling is very efficient and contributes to 75% of the success of Jambox. In order to have a complete transformation we also employ maximal labeling. Both labelings are described in Section 6. Thirdly, we combine labeling and context extension into what we call ‘dynamic labeling’, where contexts are prepended depending on the interpretation of the variables. This is formalized in Section 3. All these optimizations minimize the number of rules and their size in the transformed μ TRS, which is important to keep a manageable search space.

Related work. Cariboo [5] deals with outermost termination using a stand-alone approach based on induction. The very idea for a transformational approach to outermost termination comes from [8]. There the signature is enriched with unary symbols `top`, `up`, and `down` and the TRS is extended with ‘anti-matching’ rules such that `down(t)` is a redex if and only if t is not a redex with respect to the original TRS. The idea is that the symbol `down` is moved down in the term tree as long as no redex is encountered. Once a redex is encountered, a rewrite step is performed, and the symbol `down` is replaced by `up`, which then moves upwards again to the top of the term, marked by `top`. This transformation is implemented in TrafO. Based on a similarly elegant idea, Thiemann [11] defines a complete transformation from outermost to innermost rewriting, which is implemented in AProVE. For traversal to the redex positions, rules of the form $\text{down}(\text{isRedex}(f(\dots))) \rightarrow f(\dots, \text{down}(\text{isRedex}(\dots)), \dots)$ are used. In order to simulate outermost rewriting and to prevent from moving inside redexes, rules $\text{isRedex}(\ell) \rightarrow \text{up}(r)$ are added for every rule $\ell \rightarrow r$ of the original TRS. Then, by the innermost rewriting strategy, the latter rules have priority over the traversal rules, whenever an original redex is encountered. The simplicity of both approaches is attractive, but the yo-yoing effect in the resulting TRSs makes that the original outermost rewrite steps are ‘hidden’ among a vast amount of auxiliary steps. This increases derivational complexity, and makes it hard for automated termination provers to find proofs for the transformed systems.

2 Preliminaries

For a general introduction to term rewriting and to context-sensitive rewriting, the reader is referred to [9] and [6], respectively. Here we repeat some of the main definitions, for the sake of completeness, and to fix notations.

A *signature* Σ is a non-empty set of symbols each having a fixed *arity*, given by a mapping $\sharp : \Sigma \rightarrow \mathbb{N}$. Given Σ and a set \mathcal{X} of variables, the *set* $Ter(\Sigma, \mathcal{X})$ of *terms over* Σ is the smallest set satisfying: $\mathcal{X} \subseteq Ter(\Sigma, \mathcal{X})$, and $f(t_1, \dots, t_n) \in Ter(\Sigma, \mathcal{X})$ if $f \in \Sigma$ of arity n and $t_i \in Ter(\Sigma, \mathcal{X})$ for all $1 \leq i \leq n$. We use x, y, z, \dots to range over variables, and write $Var(t)$ for the set of variables occurring in a term t . Usually we leave \mathcal{X} implicit and write $Ter(\Sigma)$ for the set of terms over Σ and a fixed, countably infinite set of variables \mathcal{X} . The set of positions $Pos(t) \subseteq \mathbb{N}^*$ of a term $t \in Ter(\Sigma)$ is defined as follows: $Pos(x) = \{\epsilon\}$ for variables $x \in \mathcal{X}$ and $Pos(f(t_1, \dots, t_n)) = \{\epsilon\} \cup \{ip \mid 1 \leq i \leq \sharp f, p \in Pos(t_i)\}$. We write $t(p)$ to denote the root symbol of $t|_p$, the subterm of t rooted at p , and we write $root(t)$ for the root symbol of t , that is $root(t) = t(\epsilon)$.

A *substitution* σ is a map $\sigma : \mathcal{X} \rightarrow Ter(\Sigma, \mathcal{X})$ from variables to terms. For terms $t \in Ter(\Sigma, \mathcal{X})$ and substitutions σ , $t\sigma$ is inductively defined by $x\sigma = \sigma(x)$ if $x \in \mathcal{X}$, and $f(t_1, \dots, t_n)\sigma = f(t_1\sigma, \dots, t_n\sigma)$ otherwise. Let \square be a fresh symbol, i.e. $\square \notin \Sigma \cup \mathcal{X}$. A *context* C is a term from $Ter(\Sigma, \mathcal{X} \cup \{\square\})$ containing precisely one occurrence of \square . By $C[s]$ we denote the term $C\sigma$ where $\sigma(\square) = s$ and $\sigma(x) = x$ for all $x \in \mathcal{X}$. The *depth* of a context C is defined as the length $|p|$ of the position p at which \square resides, that is, the position p such that $C(p) = \square$.

A *term rewriting system (TRS)* over Σ is a set R of pairs $\langle \ell, r \rangle \in Ter(\Sigma, \mathcal{X})^2$, called *rewrite rules* and written as $\ell \rightarrow r$, for which the *left-hand side* ℓ is not a variable ($\ell \notin \mathcal{X}$) and all variables in the *right-hand side* r occur in ℓ : $Var(r) \subseteq Var(\ell)$. For a TRS R we define \rightarrow_R , the *rewrite relation* induced by R as follows. For terms $s, t \in Ter(\Sigma, \mathcal{X})$ we write $s \rightarrow_R t$, or just $s \rightarrow t$ if R is clear from the context, if there exists a rule $\ell \rightarrow r \in R$, a substitution σ and a context $C \in Ter(\Sigma, \mathcal{X} \cup \{\square\})$ such that $s = C[\ell\sigma]$ and $t = C[r\sigma]$; we sometimes write $s \rightarrow_{R,p} t$ to explicitly indicate the rewrite position p , i.e. when $C(p) = \square$. Then s *outermost rewrites to* t at a position $p \in Pos(s)$, denoted by $s \xrightarrow{out}_{R,p} t$, if $s \rightarrow_{R,p} t$ and for all positions p' that are a proper prefix of p : $s|_{p'}$ is not a redex.

A mapping $\mu : \Sigma \rightarrow \mathbf{2}^{\mathbb{N}}$ is called a *replacement map (for Σ)* if for all $f \in \Sigma$ we have $\mu(f) \subseteq \{1, \dots, \sharp f\}$. A *context-sensitive term rewriting system (μ TRS)* is a pair $\langle R, \mu \rangle$ consisting of a TRS R and a replacement map μ . The set of μ -*replacing positions* $Pos^\mu(t)$ of a term $t \in Ter(\Sigma, \mathcal{X})$ is defined by $Pos^\mu(x) = \{\epsilon\}$ for $x \in \mathcal{X}$ and $Pos^\mu(f(t_1, \dots, t_n)) = \{\epsilon\} \cup \{ip \mid i \in \mu(f), p \in Pos^\mu(t_i)\}$. In context-sensitive term rewriting only redexes at μ -replacing positions are contracted: s μ -*rewrites to* t , denoted $s \rightarrow_{R,\mu} t$, whenever $s \rightarrow_{R,p} t$ with $p \in Pos^\mu(s)$.

A Σ -*algebra* $\langle A, [\cdot] \rangle$ consists of a non-empty set A and for each n -ary $f \in \Sigma$ a function $[f] : A^n \rightarrow A$, called the *interpretation of f* . Given an *assignment* $\alpha : \mathcal{X} \rightarrow A$, the interpretation of terms $t \in Ter(\Sigma)$ is defined by: $[x, \alpha] = \alpha(x)$ and $[f(t_1, \dots, t_n), \alpha] = [f]([t_1, \alpha], \dots, [t_n, \alpha])$. For substitutions $\sigma : \mathcal{X} \rightarrow Ter(\Sigma, \mathcal{X})$, we write $[\sigma, \alpha]$ for the function $\lambda x. [\sigma(x), \alpha]$. For ground terms $t \in Ter(\Sigma, \emptyset)$ and substitutions $\sigma : \mathcal{X} \rightarrow Ter(\Sigma, \emptyset)$ we write $[t]$ and $[\sigma]$ for short, respectively.

3 Transformation by Dynamic Labeling

In outermost rewriting the only redexes which are allowed to be rewritten are those which are not nested within any other redex occurrence. We represent this strategy by context-sensitive rewriting by using semantic labeling: we mark the symbols which are the root of a redex in order to disallow rewriting within that redex. We first recall the definition of semantic labeling and models from [12], and then generalize these to fit our purpose.

Definition 3.1 ([12]). Let Σ be a signature. A *semantic labeling* $\langle A, [\cdot], \pi \rangle$ consists of a Σ -algebra $\langle A, [\cdot] \rangle$ and a family $\pi = \{\pi_f\}_{f \in \Sigma}$ of functions $\pi_f : A^{\sharp f} \rightarrow A_f$ where, for each $f \in \Sigma$, A_f is a finite and non-empty set of labels. For a term $t \in \text{Ter}(\Sigma)$ and $\alpha : \text{Var}(t) \rightarrow A$, an interpretation of its variables, we define the *labeling* $\text{lab}(t, \alpha)$ of t with respect to α inductively as follows:

$$\begin{aligned} \text{lab}(x, \alpha) &= x, \\ \text{lab}(f(t_1, \dots, t_n), \alpha) &= f^{\pi_f([t_1, \alpha], \dots, [t_n, \alpha])}(\text{lab}(t_1, \alpha), \dots, \text{lab}(t_n, \alpha)). \end{aligned}$$

For ground terms $t \in \text{Ter}(\Sigma, \emptyset)$ we just write $\text{lab}(t)$. Let R be a TRS over Σ . The *semantic labeling of R* is the TRS $\text{lab}(R)$ over the labeled signature $\text{lab}(\Sigma) = \{f^\lambda \mid f \in \Sigma, \lambda \in A_f\}$, defined by:

$$\text{lab}(R) = \{\text{lab}(\ell, \alpha) \rightarrow \text{lab}(r, \alpha) \mid \ell \rightarrow r \in R, \alpha : \text{Var}(\ell) \rightarrow A\}.$$

Term labeling satisfies the following useful property; see [12] for a proof.

Lemma 3.2 ([12]). *Let $\langle A, [\cdot] \rangle$ be a Σ -algebra, $\alpha : \mathcal{X} \rightarrow A$, $\sigma : \mathcal{X} \rightarrow \text{Ter}(\Sigma)$, and $\bar{\sigma}(x) = \text{lab}(\sigma(x), \alpha)$. Then $\text{lab}(t\sigma, \alpha) = \text{lab}(t, [\sigma, \alpha])\bar{\sigma}$, for all $t \in \text{Ter}(\Sigma)$.*

The Σ -algebra of a semantic labeling has to satisfy certain constraints in order to obtain that a TRS is terminating if and only if its labeled version is. In [12] the algebra has to be a ‘model’: A Σ -algebra $\langle A, [\cdot] \rangle$ is called a *model* of a TRS R if for all rules $\ell \rightarrow r \in R$ and assignments of variables $\alpha : \text{Var}(\ell) \rightarrow A$ we have that $[\ell, \alpha] = [r, \alpha]$. In the introduction we explained why this notion of model is too restrictive for our purpose. In order to be able to distinguish between redex and non-redex positions we introduce C -models, a generalization of models.

Definition 3.3. A C -model for a TRS R is a Σ -algebra $\langle A, [\cdot] \rangle$ such that for every rule $\ell \rightarrow r \in R$ there exists $n \in \mathbb{N}$ such that for every context C of depth n and assignment $\alpha : \mathcal{X} \rightarrow A$ we have $[C[\ell], \alpha] = [C[r], \alpha]$. When $n \in \mathbb{N}$ is minimal for a rule $\ell \rightarrow r$ with respect to this property, we call n the C -depth for $\ell \rightarrow r$.

As this definition suggests, it is possible to use a C -model \mathcal{A} to transform a TRS R by prepending contexts to its rules in such a way that \mathcal{A} becomes a model for the transformed system \tilde{R} , and then perform semantic labeling. But then, from the labeled version $\text{lab}(\tilde{R})$, every rule that contains a marked (redex) symbol within the context that was prepended to the rule in the construction of \tilde{R} has to be removed again, as it would enable a rewrite step which is not outermost.

We choose for a different approach which we call ‘dynamic labeling’. We step-wise extend rules by contexts, only when needed and dependent on the variable

interpretation used for the semantic labeling. For different interpretations of the variables usually different context depths are necessary for achieving equal interpretations of left- and right-hand side. In each extension step we check whether a candidate symbol is a redex symbol, and, if it is, this symbol is excluded from prepending. Here, by a redex symbol we mean a labeled symbol which indicates the presence of a redex in the original system (at the same position). Dynamic labeling is more efficient in that both the number and the size of the rules of the resulting μ TRS are smaller than in the ‘static’ version. We explain dynamic labeling by means of the TRS R_1 consisting of the rules:

$$f(g(x)) \rightarrow f(f(g(x))) \qquad f(f(f(x))) \rightarrow x \qquad (R_1)$$

and the algebra $\mathcal{A}_1 = \langle A_1, [\cdot] \rangle$ where $A_1 = \{g, f_0, f_1, f_2\}$ and where the interpretation of the symbols is defined by $[g](x) = g$ for all $x \in A_1$, $[f](g) = f_1$ and $[f](f_i) = f_{\min(i+1,2)}$ for $i = 0, 1, 2$. Let further $\langle \mathcal{A}_1, \pi \rangle$ be the semantic labeling where π labels the symbols with the interpretations of their arguments. Then the symbols f^g and f^{f_2} are redex symbols, corresponding to redex positions with respect to the first and the second rule of R_1 , respectively. The algebra \mathcal{A}_1 is a C -model where for the first rule the C -depth is 1, and for the second rule it is 2.

We iteratively construct sets P_0, P_1, \dots , until $P_{i+1} = P_i$ for some i . The initial set P_0 consists of pairs $\langle \ell \rightarrow r, \alpha \rangle$ for each rule $\ell \rightarrow r$, and each interpretation $\alpha : \text{Var}(\ell) \rightarrow A_1$ of the variables. Then, in each step, P_{i+1} is obtained from P_i by replacing every pair $\langle \ell \rightarrow r, \alpha \rangle$ of P_i for which the interpretation of the left-hand side differs from the right-hand side ($[\ell, \alpha] \neq [r, \alpha]$), by the pairs $\langle C[\ell] \rightarrow C[r], \alpha' \rangle$ for every flat context C (see (2) on the next page) and every extension $\alpha' : \text{Var}(C[\ell]) \rightarrow A_1$ of α , such that the root of the labeled, extended left-hand side $\text{lab}(C[\ell], \alpha')$ is not a redex symbol. Among the flat contexts to prepend we include $\text{top}(\square)$ to cater for the case that the rule is applied at the top of the term. For R_1 the initial set P_0 is:

$$P_0 = \{ \langle f(g(x)) \rightarrow f(f(g(x))), \lambda x.a \rangle, \langle f(f(f(x))) \rightarrow x, \lambda x.a \rangle \mid a \in A_1 \}.$$

The only element $\langle \ell \rightarrow r, \alpha \rangle$ of P_0 such that $[\ell, \alpha] = [r, \alpha]$ is $\langle f(f(f(x))) \rightarrow x, \lambda x.f_2 \rangle$. For this pair no context needs to be prepended. The other pairs have to be replaced by their context extensions, and thus P_1 consists of:

$$\begin{aligned} &\langle C[f(g(x))] \rightarrow C[f(f(g(x))), \lambda x.a \rangle, \quad \text{for all } a \in A_1, C \in \{\text{top}(\square), f(\square), g(\square)\}, \\ &\quad \langle f(f(f(x))) \rightarrow x, \lambda x.f_2 \rangle, \\ &\langle C[f(f(f(x)))] \rightarrow C[x], \lambda x.a \rangle, \quad \text{for all } a \in A_1 \setminus \{f_2\}, C \in \{\text{top}(\square), g(\square)\}. \end{aligned}$$

In the last line the context $f(\square)$ is excluded, because the labeled left-hand side of the rule would contain the redex symbol f^{f_2} within the prepended context, and thus the step would not be outermost. Due to the outermost strategy, the original rule is only applicable under a context $C[g(\square)]$ (where C does not contain any redexes) or at the top of a term. Now for all $(4 \cdot 3 + 1 + 3 \cdot 2 = 19)$ pairs of P_1 left- and right-hand side have equal interpretations, and the iterative construction is ended. We define $\Delta_{\mathcal{A}_1}(R_1) = P_1$, and call this set the ‘dynamic context extension’ of R_1 with respect to \mathcal{A}_1 .

Secondly, the dynamic extension $\Delta_{\mathcal{A}_1}(R_1)$ is labeled using the family π which returns for each symbol f a label consisting of the interpretation of its arguments, i.e., $\pi_f(a_1, \dots, a_{\#f}) = \langle a_1, \dots, a_{\#f} \rangle$. Then the desired μ TRS $\Delta_{\mathcal{A}_1}^\pi(R_1)$, which we call the *dynamic labeling of R_1* , consists of the rules $lab(\ell, \alpha) \rightarrow lab(r, \alpha)$ for every $\langle \ell \rightarrow r, \alpha \rangle \in \Delta_{\mathcal{A}_1}(R_1)$, with the replacement map μ defined by $\mu(f) = \emptyset$ if $f \in \{f^g, f^{f^2}\}$, and $\mu(f) = \{1, \dots, \#f\}$ otherwise, for all $f \in lab(\Sigma)$.

We now formalize dynamic labeling. For the remainder of this section we fix an arbitrary TRS R over Σ , and let $\mathcal{A} = \langle A, [\cdot] \rangle$ be a C -model for R . We assume \mathbf{top} to be a fresh symbol ($\mathbf{top} \notin \Sigma$) representing the top of a term, and abbreviate $\Sigma_{\mathbf{top}} = \Sigma \cup \{\mathbf{top}\}$. We extend \mathcal{A} to a $\Sigma_{\mathbf{top}}$ -algebra by choosing an arbitrary but fixed element $a \in A$ and defining the interpretation of \mathbf{top} as the constant function $[\mathbf{top}] = \lambda x.a$. Furthermore, we let $\langle \mathcal{A}, \pi \rangle$ be a semantic labeling, and $\Sigma^{red} \subseteq lab(\Sigma)$ a subset of the labeled signature, called the set of *reder symbols*. The definition makes use of *flat contexts fresh for $t \in Ter(\Sigma_{\mathbf{top}})$* :

$$\mathcal{C}_t^b = \{ f(x_1, \dots, x_{j-1}, \square, x_{j+1}, \dots, x_{\#f}) \mid f \in \Sigma_{\mathbf{top}}, j \in \{1, \dots, \#f\} \} \quad (2)$$

where $x_1, x_2, \dots \in \mathcal{X}$ such that $x_i \notin Var(t)$. Furthermore, for partial functions $f, g : S \rightarrow T$ with disjoint domains, we write $f+g$ for the *union of f and g* , defined by $(f+g)(x) = f(x)$ if $x \in dom(f)$, and $(f+g)(x) = g(x)$ if $x \in dom(g)$.

Definition 3.4. The *dynamic context extension of R* , denoted by $\Delta_{\mathcal{A}}(R)$, is defined as the fixed point of the following construction of sets P_0, P_1, \dots , that is, $\Delta_{\mathcal{A}}(R) = P_i$ as soon as $P_{i+1} = P_i$ for some i . The initial set is defined by:

$$P_0 = \{ \langle \ell \rightarrow r, \alpha \rangle \mid \ell \rightarrow r \in R, \alpha : Var(\ell) \rightarrow A \},$$

and for $i = 0, 1, \dots$ the set P_{i+1} is obtained from P_i by replacing every pair $\langle \ell \rightarrow r, \alpha \rangle$ such that $[\ell, \alpha] \neq [r, \alpha]$, or $r \in \mathcal{X}$,¹ by all pairs in $\Delta(\ell \rightarrow r, \alpha)$ where:

$$\Delta(\ell \rightarrow r, \alpha) = \{ \langle C[\ell] \rightarrow C[r], \alpha + \beta \rangle \mid C \in \mathcal{C}_\ell^b, \beta : Var(C) \rightarrow A, \\ root(lab(C[\ell], \alpha + \beta)) \notin \Sigma^{red} \}.$$

Then, the *dynamic labeling of R* is the μ TRS $\langle \Delta_{\mathcal{A}}^\pi(R), \mu \rangle$ consisting of:

$$\Delta_{\mathcal{A}}^\pi(R) = \{ lab(\ell, \alpha) \rightarrow lab(r, \alpha) \mid \langle \ell \rightarrow r, \alpha \rangle \in \Delta_{\mathcal{A}}(R) \},$$

and the replacement map μ , defined by $\mu(f) = \emptyset$ if $f \in \Sigma^{red}$, and $\mu(f) = \{1, \dots, \#f\}$ otherwise, for all $f \in lab(\Sigma_{\mathbf{top}})$. Whenever the set Σ^{red} , which determines the replacement map, is clear from the context, we write $\Delta_{\mathcal{A}}^\pi(R)$ as a shorthand for $\langle \Delta_{\mathcal{A}}^\pi(R), \mu \rangle$.

Notice that the construction of $\Delta_{\mathcal{A}}(R)$ is guaranteed to terminate because of the assumption that \mathcal{A} is a C -model.

We come to the first main theorem, stating that outermost ground termination of R is implied by termination of the transformed system $\Delta_{\mathcal{A}}^\pi(R)$.

¹ The condition $r \notin \mathcal{X}$ eliminates collapsing rules. This is used in the proof of Thm. 6.5, which states completeness. Without this condition, the transformation is still sound (Thm. 3.7). Nota bene: in the TRS R_1 worked out before, $r \notin \mathcal{X}$ is not used.

For the remainder of this section, we assume that the set $\Sigma^{red} \subseteq \text{lab}(\Sigma)$ contains redex symbols only, that is, for all ground terms t and $p \in \text{Pos}(t)$:

$$\text{lab}(t)(p) \in \Sigma^{red} \text{ implies that } t|_p \text{ is a redex with respect to } R \quad (\ddagger)$$

Lemma 3.5. *Let $s, t \in \text{Ter}(\Sigma, \emptyset)$ be ground terms and $p \in \text{Pos}(s)$ such that $s \xrightarrow{\text{out}}_{R,p} t$. Then for all proper prefixes q of $1p$ we have $\text{lab}(\text{top}(s))(q) \notin \Sigma^{red}$.*

Proof. If $q = \epsilon$, this follows from $\text{top}^\lambda \notin \Sigma^{red}$ for any label λ . If $q \neq \epsilon$, then $\text{lab}(\text{top}(s))(q) = \text{lab}(s)(q')$ with q' a proper prefix of p , and if $\text{lab}(s)(q') \in \Sigma^{red}$, then by assumption (\ddagger) the term s contains a redex at position q' , quod non. \square

The following lemma states that any outermost ground rewrite step in R can be transformed into a rewrite step in $\Delta_{\mathcal{A}}^\pi(R)$. For ground substitutions $\sigma : \mathcal{X} \rightarrow \text{Ter}(\Sigma, \emptyset)$ define $\bar{\sigma}(x) = \text{lab}(x\sigma)$.

Lemma 3.6. *Let $s, t \in \text{Ter}(\Sigma, \emptyset)$ be ground terms such that $s \xrightarrow{\text{out}}_R t$. Then:*

$$\text{lab}(\text{top}(s)) \rightarrow_{\Delta_{\mathcal{A}}^\pi(R)} \text{lab}(\text{top}(t)) .$$

Proof. Assume $s \xrightarrow{\text{out}}_{R,p} t$ for some position $p \in \text{Pos}(s)$. Then there exists a rule $\ell \rightarrow r \in R$, a context C with $C(p) = \square$ and a ground substitution σ such that $s = C[\ell\sigma]$ and $t = C[r\sigma]$. We consider the construction of the dynamic context extension from Definition 3.4, and prove by induction that for all $i = 0, 1, \dots$ there exists a context C_i which is a prefix of $\text{top}(C)$, a ground substitution σ_i , and terms ℓ_i, r_i such that $\text{top}(s) = C_i[\ell_i\sigma_i]$, $\text{top}(t) = C_i[r_i\sigma_i]$ and $\langle \ell_i \rightarrow r_i, [\sigma_i] \rangle \in P_i$. For the base case we have $\langle \ell_0 \rightarrow r_0, [\sigma_0] \rangle \in P_0$ with $\ell_0 = \ell$, $r_0 = r$, $\sigma_0 = \sigma$, and $C_0 = \text{top}(C)$. For the induction step we assume the existence of C_i , σ_i , and $\langle \ell_i \rightarrow r_i, [\sigma_i] \rangle \in P_i$ with the above properties. If $[\ell_i, [\sigma_i]] = [r_i, [\sigma_i]]$ and $r_i \notin \mathcal{X}$ then by definition $\langle \ell_i \rightarrow r_i, [\sigma_i] \rangle \in P_{i+1}$, and so we are done. For the remaining cases $[\ell_i, [\sigma_i]] \neq [r_i, [\sigma_i]]$ and $r_i \in \mathcal{X}$, we first show that $C_i \neq \square$. If $[\ell_i, [\sigma_i]] \neq [r_i, [\sigma_i]]$ and $C_i = \square$, then $\ell_i\sigma_i = \text{top}(s)$ and $r_i\sigma_i = \text{top}(t)$, and hence $\text{root}(\ell_i) = \text{root}(r_i) = \text{top}$, contradicting $[\ell_i, [\sigma_i]] \neq [r_i, [\sigma_i]]$ (recall that the interpretation of top is constant). Furthermore, we have $r_i \in \mathcal{X}$ only if $i = 0$, and then $C_i = \text{top}(C) \neq \square$. Thus we have $C_i = D[D'\sigma']$ for some context D , flat context $D' \in \mathcal{C}_{\ell_i}^b$ and substitution σ' . We choose $C_{i+1} = D$, $\ell_{i+1} = D'[\ell_i]$, $r_{i+1} = D'[r_i]$, and $\sigma_{i+1} = \sigma_i + \sigma'$. It remains to be shown that $\langle \ell_{i+1} \rightarrow r_{i+1}, [\sigma_{i+1}] \rangle \in P_{i+1}$. For this it suffices to prove that $\text{root}(\text{lab}(\ell_{i+1}, [\sigma_{i+1}])) \notin \Sigma^{red}$. We have $C_{i+1}[\ell_{i+1}\sigma_{i+1}] = \text{top}(s)$. Let q be the position such that $C_{i+1}(q) = \square$. Then, by Lemma 3.2 we obtain $\text{root}(\text{lab}(\ell_{i+1}, [\sigma_{i+1}])) = \text{root}(\text{lab}(\ell_{i+1}\sigma_{i+1})) = \text{top}(\text{lab}(s))(q)$. Note that q is a proper prefix of $1p$, and so $\text{lab}(s)(q) \notin \Sigma^{red}$ by Lemma 3.5.

Let $i \in \mathbb{N}$ be such that $P_{i+1} = P_i$. By the result above we have $\langle \ell_i \rightarrow r_i, [\sigma_i] \rangle \in \Delta_{\mathcal{A}}^\pi(R)$, with $[\ell_i\sigma_i] = [r_i\sigma_i]$, and then $\text{lab}(\ell_i, [\sigma_i]) \rightarrow \text{lab}(r_i, [\sigma_i]) \in \Delta_{\mathcal{A}}^\pi(R)$ by definition. Let τ and v be defined by $\tau(\square) = \ell_i\sigma_i$, $v(\square) = r_i\sigma_i$, and $\tau(x) = v(x) = x$ for $x \in \mathcal{X}$. Then we have that $\text{lab}(C_i, [\tau]) = \text{lab}(C_i, [v])$ since $[\tau] = [v]$. Let $E = \text{lab}(C_i, [\tau])$. We get $\text{lab}(\text{top}(s)) = \text{lab}(C_i[\ell_i\sigma_i]) = \text{lab}(C_i\tau) = E\bar{\tau} = E[\text{lab}(\ell_i\sigma_i)] = E[\text{lab}(\ell_i, [\sigma_i])\bar{\sigma}_i]$ and $\text{lab}(\text{top}(t)) = \dots = E[\text{lab}(r_i, [\sigma_i])\bar{\sigma}_i]$, by Lemma 3.2. Then, by Lemma 3.5 we have $\text{lab}(\text{top}(s)) \rightarrow_{\Delta_{\mathcal{A}}^\pi(R)} \text{lab}(\text{top}(t))$. \square

Theorem 3.7. *R is outermost ground terminating if $\Delta_{\mathcal{A}}^{\pi}(R)$ is terminating.*

Proof. Assume R admits an infinite outermost rewrite sequence $t_1 \xrightarrow{\text{out}}_R t_2 \xrightarrow{\text{out}}_R t_3 \xrightarrow{\text{out}}_R \dots$. Then from Lemma 3.6 it follows that $\Delta_{\mathcal{A}}^{\pi}(R)$ allows an infinite rewrite sequence: $\text{lab}(\text{top}(t_1)) \rightarrow_{\Delta_{\mathcal{A}}^{\pi}(R)} \text{lab}(\text{top}(t_2)) \rightarrow_{\Delta_{\mathcal{A}}^{\pi}(R)} \text{lab}(\text{top}(t_3)) \rightarrow_{\Delta_{\mathcal{A}}^{\pi}(R)} \dots$ \square

Theorem 3.7 is about outermost ground termination. This is not a restriction because by adding a fresh constant 0 and a fresh unary symbol s , outermost ground termination coincides with outermost termination:

Lemma 3.8. *The TRS R over the signature Σ is outermost terminating if and only if R over the signature $\Sigma \cup \{s, 0\}$ is outermost ground terminating.* \square

4 Constructing Suitable Algebras

In this section we construct C -models that are able to recognize redex positions with respect to left-linear rules. The construction of C -models is similar to the construction of a deterministic tree automaton (DTA) for recognizing left-linear redexes. A DTA is a Σ -algebra $\langle A, [\cdot] \rangle$ with a distinguished set $A_F \subseteq A$ of final states. A term t is accepted by the automaton whenever $[t] \in A_F$. The difference with the construction of a DTA is that for the construction of a C -model we do not distinguish final and non-final states, but instead have a family of functions $\text{isRedex}_f : A^{\sharp f} \rightarrow \text{Bool}$ for indicating the presence of a redex.

Definition 4.1. A *redex-algebra* $\mathcal{A} = \langle A, [\cdot], \text{isRedex} \rangle$ is a Σ -algebra $\langle A, [\cdot] \rangle$ with a family $\{\text{isRedex}_f\}_{f \in \Sigma}$ of functions $\text{isRedex}_f : A^{\sharp f} \rightarrow \text{Bool}$. The *language* of \mathcal{A} is the set $\mathcal{L}(\mathcal{A}) = \{f(t_1, \dots, t_n) \in \text{Ter}(\Sigma, \emptyset) \mid \text{isRedex}_f([t_1], \dots, [t_n]) = \text{true}\}$.

By this separation of tasks our approach allows for smaller algebras, because, intuitively, the algebra needs to ‘remember’ only the subterms t_1, \dots, t_n and not $f(t_1, \dots, t_n)$ itself. To see this, consider the single-rule system:

$$f(g(x)) \rightarrow a.$$

A tree automaton recognizing redex positions for this TRS needs at least three states: one for indicating a redex $f(g(\dots))$, one for $g(\dots)$, and one garbage state. For redex-algebras two states suffice: one state for $g(\dots)$ and one for garbage. Then $\text{isRedex}_f(g(\dots)) = \text{true}$ and false , otherwise.

The ‘core’ of a redex-algebra consists of all interpretations of ground terms:

Definition 4.2. Let $\mathcal{A} = \langle A, [\cdot], \text{isRedex} \rangle$ be a redex-algebra over Σ . The *core* of \mathcal{A} is the redex-algebra $\mathcal{A}_c = \langle A_c, [\cdot]_c, \text{isRedex}_c \rangle$ where A_c is the smallest set such that $[f](a_1, \dots, a_n) \in A_c$ whenever $f \in \Sigma$ and $a_1, \dots, a_n \in A_c$, and where $[\cdot]_c$ and isRedex_c are the restrictions of $[\cdot]$ and isRedex to A_c , respectively. Furthermore we say that \mathcal{A} is *core* whenever $\mathcal{A} = \mathcal{A}_c$.

Lemma 4.3. *Let $\mathcal{A} = \langle A, [\cdot], \text{isRedex} \rangle$ be a redex-algebra over Σ . Then for every $a \in A_c$ there exists a ground term $t \in \text{Ter}(\Sigma, \emptyset)$ with $[t] = a$.* \square

We now describe a syntactical construction of redex-algebras. The idea is to build Σ -algebras that ‘remember’ proper subterms of left-hand sides. Given this interpretation, the *isRedex* functions decide whether a redex is present. We first define some auxiliary functions.

Let \perp be a fresh symbol, $\perp \notin \Sigma$, and define $\mathcal{T} = \text{Ter}(\Sigma \cup \{\perp\}, \emptyset)$. The function $\text{cut} : \text{Ter}(\Sigma, \mathcal{X}) \rightarrow \mathcal{T}$ is defined such that $\text{cut}(t)$ is the result of replacing all variables in t by \perp . We define $\text{match} : \mathcal{T} \times \mathcal{T} \rightarrow \text{Bool}$ such that $\text{match}(s, t) = \text{true}$ if s can be obtained from t by replacing subterms of t by \perp , and $\text{match}(s, t) = \text{false}$, otherwise. Let further $\text{merge}(s, t)$ be the ‘most general common instance’ of s and t , that is, $\text{merge} : \mathcal{T} \times \mathcal{T} \rightarrow \mathcal{T}$ is defined by: $\text{merge}(\perp, t) = t$, $\text{merge}(t, \perp) = t$, and $\text{merge}(f(s_1, \dots, s_n), f(t_1, \dots, t_n)) = f(\text{merge}(s_1, t_1), \dots, \text{merge}(s_n, t_n))$, and undefined whenever there exists a position $p \in \text{Pos}(s)$ such that $s(p) \in \Sigma$, $t(p) \in \Sigma$, and $s(p) \neq t(p)$. Finally we define $\text{shrink} : \mathcal{T} \times \mathbf{2}^T \rightarrow \mathcal{T}$ such that $\text{shrink}(s, T)$ is the largest $t \in T$ (with respect to the number of symbols) such that $\text{match}(t, s)$ is *true*. Note that $\text{shrink}(s, T)$ is well-defined whenever T is closed under merge and $\perp \in T$: whenever two terms $t_1 \neq t_2$ of equal size match s then $\text{merge}(t_1, t_2)$ is larger and matches s .

Definition 4.4. Let R be a TRS. The *redex-algebra* for R is the core of the redex-algebra $\langle A, [\cdot], \text{isRedex} \rangle$, where A is the smallest set such that $\perp \in A$ and

- $t \in A$ for every proper subterm t of $\text{cut}(\ell)$ with ℓ a linear left-hand side of R ,
- $\text{merge}(s, t) \in A$ whenever $s, t \in A$ and $\text{merge}(s, t)$ is defined.

Then $[\cdot]$ is defined by $[f](t_1, \dots, t_n) = \text{shrink}(f(t_1, \dots, t_n), A)$. And, for every $f \in \Sigma$ we define $\text{isRedex}_f(t_1, \dots, t_n) = \text{true}$ if $f(t_1, \dots, t_n)$ is an instance of a linear left-hand side of R , and $\text{isRedex}_f(t_1, \dots, t_n) = \text{false}$, otherwise.

Example 4.5. Consider the term rewriting system R consisting of the rules:

$$c(c(c(x))) \rightarrow a, \quad c(c(a)) \rightarrow c(c(c(c(a)))) .$$

By Definition 4.4 we obtain $A = \{c(c(\perp)), c(\perp), \perp, c(a), a\}$. Here $[a] = a$, $[c](a) = c(a)$, $[c](c(a)) = c(c(\perp))$ and $[c](c(c(\perp))) = c(c(\perp))$. Thus the elements \perp and $c(\perp)$ are not part of the core and hence not of the redex-algebra for the TRS. Here we have $\text{isRedex}_c(c(a)) = \text{isRedex}_c(c(c(\perp))) = \text{true}$, and *false* otherwise.

Example 4.6. We compute the domain of the redex-algebra for the TRS:

$$\begin{aligned} f(x, y) &\rightarrow a(f(c(x), y)), & a(f(c(c(x)), y)) &\rightarrow e, \\ f(x, y) &\rightarrow b(f(x, c(y))), & b(f(x, c(c(y)))) &\rightarrow e. \end{aligned}$$

The subterms of $\text{cut}(\ell)$ of linear left-hand sides ℓ are: $S = \{\perp, f(c(c(\perp)), \perp), f(\perp, c(c(\perp))), c(c(\perp)), c(\perp)\}$. The closure of S under merge yields the domain: $A = S \cup \{f(c(c(\perp)), c(c(\perp)))\}$.

Lemma 4.7. *Let R be a TRS over Σ and \mathcal{A} the redex-algebra for R . Then for all ground terms $t \in \text{Ter}(\Sigma, \emptyset)$ we have $t \in \mathcal{L}(\mathcal{A})$ if and only if t is a redex with respect to a left-linear rule in R .*

The proof proceeds by induction over the term structure. The ‘only if’-part is crucial for soundness of our transformation, whereas the ‘if’-part is needed for completeness for left-linear TRSs.

5 Minimizing Algebras

In this section we are concerned with the minimization of redex-algebras. The algorithm is similar to the minimization of deterministic tree automata, see [3]. For the set of 291 TRSs of the outermost termination competition of 2008 [10], the redex-algebras constructed according to Definition 4.4 have an average size of 4.6 elements. After an application of the minimization algorithm described here, the average size falls to 3.4, a reduction of 27%. This reduction has a polynomial influence on the number of rules of the transformed system.

Definition 5.1. Core redex-algebras $\mathcal{A}_1, \mathcal{A}_2$ are *equivalent* if $\mathcal{L}(\mathcal{A}_1) = \mathcal{L}(\mathcal{A}_2)$.

For a given core redex-algebra we now construct a minimal equivalent algebra. The difference to the minimization of tree automata from [3] lies in the initial equivalence E_0 . For tree automata E_0 consists of two partitions, the final and the non-final states. In our setting two states are initially equivalent if they cannot be distinguished using the *isRedex* functions. In general this can yield any number of partitions between 1 and $|A|$.

Definition 5.2. Let $\mathcal{A} = \langle A, [\cdot], isRedex \rangle$ be a core redex-algebra over Σ . We define equivalence relations E_i for $i \in \mathbb{N}$ on the elements of A . Initially two elements $a, b \in A$ are equivalent, $a E_0 b$, if $isRedex_f(\mathbf{x}, a, \mathbf{y}) = isRedex_f(\mathbf{x}, b, \mathbf{y})$ for all symbols $f \in \Sigma$, $j \in \{1, \dots, \#f\}$, $\mathbf{x} \in A^{j-1}$, and $\mathbf{y} \in A^{\#f-j}$. Then for $i = 0, 1, \dots$ and $a, b \in A$ we define $a E_{i+1} b$ to be the conjunction of $a E_i b$ and $[f](\mathbf{x}, a, \mathbf{y}) E_i [f](\mathbf{x}, b, \mathbf{y})$ for all $f \in \Sigma$, $j \in \{1, \dots, \#f\}$, $\mathbf{x} \in A^{j-1}$, and $\mathbf{y} \in A^{\#f-j}$. We stop when $E_{i+1} = E_i$ for some $i \in \mathbb{N}$. Then we define $E = E_i$.

For $a \in A$ we use $[a]$ to denote the equivalence class of a with respect to E . The *minimized redex-algebra* is defined by $\mathcal{A}^{min} = \langle E, [\cdot]^E, isRedex^E \rangle$ where for every $f \in \Sigma$ we define $[f]^E : E^{\#f} \rightarrow E$ by $[f]^E([a_1], \dots, [a_n]) = [f(a_1, \dots, a_n)]$, and $isRedex_f^E : E^{\#f} \rightarrow Bool$ by $isRedex_f^E([a_1], \dots, [a_n]) = isRedex_f(a_1, \dots, a_n)$.

Lemma 5.3. Let \mathcal{A} be a core redex-algebra, then \mathcal{A} is equivalent to \mathcal{A}^{min} . \square

Example 5.4. We consider the TRS R consisting of the following three rules:

$$f(i(a)) \rightarrow a, \quad f(j(a)) \rightarrow a, \quad f(a) \rightarrow a.$$

The redex-algebra for R is $A = \{a, i(a), j(a), \perp\}$ with the interpretation $[a] = a$, $[i](a) = i(a)$, $[j](a) = j(a)$, and the interpretation is \perp in all non-listed cases; $isRedex_f(x) = true$ for all $x \neq \perp$, and *false*, otherwise.

The minimization algorithm starts with $E_0 = \{\{a, i(a), j(a)\}, \{\perp\}\}$ as initial equivalence, since \perp can be distinguished from the other elements due to $[f](\perp) = false$. The first iteration of the algorithm yields $E_1 = \{\{a\}, \{i(a), j(a)\}, \{\perp\}\}$ as

$[i](a) = i(a)$ whereas $[i](i(a)) = [i](j(a)) = \perp$. The elements $i(a)$ and $j(a)$ are indistinguishable, and so in the second iteration we obtain $E_2 = E_1$. Thus the elements $i(a)$ and $j(a)$ are identified and we obtain an algebra that has one element less than the algebra we started with.

6 Two Versions of Dynamic Labeling

In the previous sections we have constructed and minimized redex-algebras for recognizing redex positions. For completing the transformation we still need to define how the symbols are labeled. In this section we introduce two labelings that arise naturally: minimal and maximal labeling. In minimal labeling symbols are marked with a \star if they correspond to redex positions and stay unlabeled otherwise. This labeling creates a small signature and thereby results in a small number of rules of the transformed system.

In the sequel we fix R to be a TRS over Σ and $\langle \mathcal{A}, isRedex \rangle$ with $\mathcal{A} = \langle A, [\cdot] \rangle$ the minimized redex-algebra for R .

Definition 6.1. The *minimal labeling* for R is the semantic labeling $\langle \mathcal{A}, \pi \rangle$ defined for every $f \in \Sigma_{\text{top}}$ by $\pi_f(a_1, \dots, a_{\#f}) = \star$ if $isRedex_f(a_1, \dots, a_{\#f}) = true$, and $\pi_f(a_1, \dots, a_{\#f}) = \epsilon$, otherwise; the redex symbols are $\Sigma^{red} = \{f^\star \mid f \in \Sigma\}$.

Theorem 6.2. Let $\langle \mathcal{A}, \pi \rangle$ and Σ^{red} be the minimal labeling for R . Then R is outermost ground terminating if $\Delta_R^\pi(\mathcal{A})$ is terminating.

Proof. An application of Theorem 3.7 together with Lemmas 4.7 and 5.3. \square

Minimal labeling is sound and efficient, but it is not complete:

Example 6.3. The following term rewriting system is outermost terminating:

$$inf(x) \rightarrow cons(x, inf(s(x))) \qquad cons(s(x), y) \rightarrow nil \qquad (R_2)$$

The minimized redex-algebra for R_2 is $\mathcal{A}_2 = \langle A_2, [\cdot] \rangle$ with $A_2 = \{s, \perp\}$, $[s](x) = s$, $[inf](x) = [cons](x, y) = \perp$, and $[nil] = \perp$. With minimal labeling we have $\pi_{cons}(s, x) = \star$ and $\pi_{inf}(x) = \star$ for all $x \in A_2$, and unmarked (ϵ) otherwise. The dynamic labeling $\Delta_{\mathcal{A}_2}^\pi(R_2)$ of R_2 w.r.t. $\langle \mathcal{A}_2, \pi \rangle$ consists of the following rules, the first two of which arise from the *inf*-rule, with $\alpha(x) = \perp$, and $\alpha(x) = s$ resp.:

$$\begin{aligned} inf^\star(x) &\rightarrow cons(x, inf^\star(s(x))), \\ inf^\star(x) &\rightarrow cons^\star(x, inf^\star(s(x))), \\ cons^\star(s(x), y) &\rightarrow nil. \end{aligned}$$

with $\mu(inf^\star) = \mu(cons^\star) = \emptyset$. But now $\Delta_{\mathcal{A}_2}^\pi(R_2)$ admits an infinite derivation:

$$inf^\star(x) \rightarrow cons(x, inf^\star(s(x))) \rightarrow cons(x, cons(s(x), inf^\star(s(s(x)))))) \rightarrow \dots$$

The third term is labeled incorrectly, as the inner *cons* should be marked. The reason is that in the second step, instead of the first *inf*[∗]-rule, the second should have been applied; however, the left-hand side *inf*[∗](*x*) contains too little information to ‘decide’ what the labeling of the right-hand side should be.

This motivates the use of maximal labeling for which correct labeling is preserved under rewriting. Symbols are labeled with the interpretation of their arguments:

Definition 6.4. The *maximal labeling* for R is the semantic labeling $\langle \mathcal{A}, \pi \rangle$ defined for every $f \in \Sigma_{\text{top}}$ by: $A_f = A^{\#f}$, $\pi_f(a_1, \dots, a_{\#f}) = \langle a_1, \dots, a_{\#f} \rangle$ together with the redex symbols $\Sigma^{\text{red}} = \{f^{(a_1, \dots, a_{\#f})} \mid \text{isRedex}_f(a_1, \dots, a_{\#f}) = \text{true}\}$.

Maximal labeling is sound for all TRSs, and complete for quasi-left linear TRSs. A TRS R is called *quasi left-linear* if every non-linear left-hand side of a rule in R is an instance of a linear left-hand side from R .

Theorem 6.5. Let $\langle \mathcal{A}, \pi \rangle$ with Σ^{red} be the maximal labeling for R . Then R is outermost ground terminating if $\Delta_R^\pi(\mathcal{A})$ is terminating. Moreover, if R is quasi left-linear, the reverse direction holds as well.

Proof. The first claim is a consequence of Theorem 3.7 and Lemmas 4.7 and 5.3. For the second claim, let R be quasi left-linear. Assume that R is outermost terminating but $\Delta_{\mathcal{A}}^\pi(R)$ is not terminating. We introduce types for $\Delta_{\mathcal{A}}^\pi(R)$ over the sorts $A \cup \{\text{top}\}$. For every $f^\lambda \in \text{lab}(\Sigma_{\text{top}})$ with $\lambda = \langle a_1, \dots, a_{\#f} \rangle$ we define f^λ to have input sorts $\langle a_1, \dots, a_n \rangle$ and output sort $[f](a_1, \dots, a_n)$, except for top for which we fix output sort top . An adaptation of [7, Proposition 5.5.24] for μ TRSs together with non-collapsingness of $\Delta_{\mathcal{A}}^\pi(R)$ yields the existence of a well-sorted infinite $\Delta_R^\pi(\mathcal{A})$ rewrite sequence τ . By Lemma 4.3 we have a ground term for every sort in A . Thus by applying a ground substitution to τ we get a well-sorted infinite ground term rewrite sequence τ' . Well-sortedness implies correct labeling: for every well-sorted term $t \in \text{Ter}(\text{lab}(\Sigma_{\text{top}}), \emptyset)$ there exists a term $t' \in \text{Ter}(\Sigma_{\text{top}}, \emptyset)$ such that $t = \text{lab}(t')$. Moreover, by well-sortedness the symbol top can only occur at the top of a term; without loss of generality we assume that every term in τ' has top as root. Hence to arrive at a contradiction it suffices to show that for all terms $s, t \in \text{Ter}(\Sigma, \emptyset)$ with $\text{lab}(\text{top}(s)) \rightarrow_{\Delta_R^\pi(\mathcal{A})} \text{lab}(\text{top}(t))$ we have $s \xrightarrow{\text{out}}_R t$. By construction, every rule in $\Delta_R^\pi(\mathcal{A})$ is a context extension plus labeling of a rule in R . Let $\rho : s \rightarrow_R t$ be the corresponding step. What remains to be shown is that ρ is an outermost step. Assume there would be a redex u above the rewrite position. Then by Lemma 4.7 we have $u \in \mathcal{L}(\mathcal{A}_R)$ where \mathcal{A}_R is the redex-algebra for R (Definition 4.4). From Lemma 5.3 it follows that $u \in \mathcal{L}(\mathcal{A})$ as \mathcal{A} is the minimization of \mathcal{A}_R . By definition of maximal labeling we get $\text{root}(\text{lab}(u)) \in \Sigma^{\text{red}}$. But then this symbol must be in $\text{lab}(\text{top}(s))$, either above the applied $\Delta_R^\pi(\mathcal{A})$ rule or within the prepended context. Both cases yield a contradiction: the former since $\mu(\text{root}(\text{lab}(u))) = \emptyset$ would prohibit the μ -step, and the latter since we do not prepend symbols from Σ^{red} . \square

Example 6.6. We revisit Example 6.3, but this time we use maximal labeling:

$$\begin{aligned} \text{inf}^\perp(x) &\rightarrow \text{cons}^{\perp, \perp}(x, \text{inf}^s(s^\perp(x))), & \text{inf}^s(x) &\rightarrow \text{cons}^{s, \perp}(x, \text{inf}^s(s^s(x))), \\ \text{cons}^{s, \perp}(s^\perp(x), y) &\rightarrow \text{nil}, & \text{cons}^{s, \perp}(s^s(x), y) &\rightarrow \text{nil}, \\ \text{cons}^{s, s}(s^\perp(x), y) &\rightarrow \text{nil}, & \text{cons}^{s, s}(s^s(x), y) &\rightarrow \text{nil}, \end{aligned}$$

with $\mu(\text{inf}^\perp) = \mu(\text{inf}^s) = \mu(\text{cons}^{s, \perp}) = \mu(\text{cons}^{s, s}) = \emptyset$. This μ TRS is indeed terminating as opposed to the μ TRS constructed in Example 6.3.

7 Discussion

For arbitrary TRSs our transformation (including the construction of C -models) is sound, and for quasi left-linear TRSs it is complete (see Theorem 6.5). The redex-algebra we construct recognizes redexes with respect to left-linear rules. As a consequence, in the μ TRS $\Delta_{\mathcal{A}}^{\pi}(R)$ rewriting is forbidden only inside such redex positions. This corresponds to a weakening of the outermost rewriting strategy: contraction of redexes is allowed as long as they are not strictly contained in a redex occurrence with respect to a left-linear rule. We stress that our transformation with maximal labeling is complete for termination with respect to this rewrite strategy for all TRSs.

An open question is whether there are interesting labelings between minimal and maximal. In particular, are there more efficient complete labelings? Here efficiency is measured in the size of the signature and the number of rules of the transformed system. In Example 6.6 it would have been sufficient to label *cons* with the interpretation of the left argument, saving two symbols and two rules of the transformed system.

References

1. B. Alarcón, F. Emmes, C. Fuhs, J. Giesl, R. Gutiérrez, S. Lucas, P. Schneider-Kamp, and R. Thiemann. Improving Context-Sensitive Dependency Pairs. In *LPAR*, volume 5330 of *LNCS*, pages 636–651. Springer, 2008.
2. T. Arts and J. Giesl. Termination of Term Rewriting Using Dependency Pairs. *Theoretical Computer Science*, 236:133–178, 2000.
3. H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree Automata Techniques and Applications. Available at <http://www.grappa.univ-lille3.fr/tata>, 2007.
4. J. Endrullis, J. Waldmann, and H. Zantema. Matrix Interpretations for Proving Termination of Term Rewriting. In U. Furbach and N. Shankar, editors, *IJCAR 2006*, volume 4130 of *LNAI*, pages 574–588. Springer, 2006.
5. O. Fissore, I. Gnaedig, and H. Kirchner. Cariboo, a Termination Proof Tool for Rewriting-Based Programming Languages with Strategies, Version 1.0, 2004.
6. S. Lucas. Context-Sensitive Computations in Functional and Functional Logic Programs. *Journal of Functional and Logic Programming*, 1998(1), 1998.
7. E. Ohlebusch. *Advanced Topics in Term Rewriting*. Springer, New York, 2002.
8. M. Raffelsieper and H. Zantema. A Transformational Approach to Prove Outermost Termination Automatically. *ENTCS*, 237:3–21, 2009.
9. Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.
10. Termination Competition. <http://www.termination-portal.org/>.
11. R. Thiemann. From Outermost Termination to Innermost Termination. In *SOFSEM '09*, pages 533–545, 2009.
12. H. Zantema. Termination of Term Rewriting by Semantic Labelling. *Fundamenta Informaticae*, 24:89–105, 1995.