

Bubbles in Modularity

Massimo Marchiori

Department of Pure and Applied Mathematics

University of Padova

Via Belzoni 7, 35131 Padova, Italy

`max@hilbert.math.unipd.it`

Abstract

We provide a global technique, called *neatening*, for the study of modularity of left-linear Term Rewriting Systems. Objects called *bubbles* are identified as the responsables of most of the problems occurring in modularity, and the concept of well-behaved (from the modularity point of view) reduction, said neat reduction, is introduced. Neatening consists of two steps: the first is proving a property is modular when only neat reductions are considered; the second is to ‘neaten’ a generic reduction so to obtain a neat one, thus showing that restricting to neat reductions is not limitative. This general technique is used to provide a unique, uniform method able to prove all the existing results on the modularity of every basic property of left-linear Term Rewriting Systems, and also to provide new results on the modularity of termination.

1 Introduction

Modularity is a field of computer science that has been receiving more and more interest along these years. Besides an interesting topic from a theoretical point of view, it is also of great practical importance: in program analysis, it allows to study a possibly big and complex program by decomposing it into smaller subparts; in program development, it allows to build a safe complex system by relying on small safe submodules.

As far as the paradigm of Term Rewriting Systems (TRSs for short) is concerned, the notion of modularity is that of disjoint union (i.e. the union of two TRSs having disjoint signatures): a property is said modular provided two TRSs enjoy it iff their disjoint union does.

This notion is somehow the basis from which to start for considering more and more complex combinations of TRSs (like composable or hierarchical, see e.g. [Ohl94a]).

In this paper we present a new technique, called *neatening*, as a global method to study modularity of left-linear TRSs. Neatening is able to cope with

all the basic properties of left-linear TRSs, proving all the so far known results on their modularity.

Here we will be rather informal in the presentation, only sketching the intuitive headlines.

We first focus on the intimate reasons that make modularity difficult to study: the major responsible is identified in the notion of *bubble*. A bubble, like the name suggests, is an object that has a potential unstability, since it could sooner or later ‘explode’ (collapse) with bad consequences on the global structure of the term. Therefore, we introduce the concept of *neat reduction*, where the ‘explosions’ of the bubbles are not dangerous (from a modularity viewpoint).

Then, to prove a property is modular, the method of neatening is introduced. Neatening, abstractly, consists of a two-step process.

First, prove that the property is *modularly neat*, that is to say it is modular when only neat reductions are considered.

Second, ‘neaten’ a generic reduction by translating it into a neat one, thus showing that restricting to neat reductions is not a limitation.

Neatening is an adequate global method for the study of modularity of TRSs under the left-linearity assumption: via this technique we obtain a meta-theorem from which all the known results on modularity, for every basic property of left-linear TRSs, are derived. Furthermore, it also provides a new sufficient criterion for the modularity of termination, and a new result on the structure of the counterexamples to the modularity of termination, for left-linear TRSs, that generalizes all the previous similar results.

The article is organized as follows.

In Section 2, some standard preliminary notions are introduced. In Section 3 the concept of bubble is presented, and in Section 4 that of neat reduction. Section 5 gives an abstract presentation of neatening, while Section 6 introduces the specific ‘neatening translation’ (\mathfrak{M}) that will be used in the practical application of neatening. In Section 7 we present the main (meta-)theorem, and apply it to all the basic properties of TRSs. Finally, Section 8 ends with some brief conclusive remarks.

2 Preliminaries

We assume knowledge of the basic notions regarding Term Rewriting Systems (TRSs for short): the notation used is essentially the one in [Klo92] and [Mid90].

For every property \mathcal{P} , $\neg\mathcal{P}$ denotes its complementary property (viz. a TRS enjoys $\neg\mathcal{P}$ iff it does not enjoy \mathcal{P}).

Given a reduction $\rho : s \rightarrow s_1 \rightarrow s_2 \dots$, the first term s is said the *start term*. Concatenation of two reductions ρ and ρ' will be indicated with $\rho \cdot \rho'$. We say a reduction ζ is contained in a reduction ρ (notation $\zeta \subseteq \rho$) if $\rho = \zeta \cdot \rho'$, for some ρ' . A term t belongs to ρ (notation $t \in \rho$) if $s \xrightarrow{\zeta} t$, $\zeta \subseteq \rho$.

Taken two reductions ρ and ρ' , we say that ρ' is *cofinal for* ρ (notation $\rho \twoheadrightarrow \rho'$) if $\forall s \in \rho \exists s' \in \rho'. s \twoheadrightarrow s'$.

A *context* is a term built up using, besides function symbols and variables, the new special constants $\square_1, \square_2, \square_3, \dots$ (said the *holes*). Contexts are as usual indicated with square brackets, e.g. $C[\square_1, \square_2]$ denotes a context with one occurrence of the hole \square_1 and one occurrence of the hole \square_2 . Given a context $C[\square_1, \dots, \square_n]$ and terms t_1, \dots, t_n , $C[t_1, \dots, t_n]$ stands for the term obtained from $C[\square_1, \dots, \square_n]$ by replacing every occurrence of \square_i with t_i ($1 \leq i \leq n$).

Throughout the paper we will indicate with \mathcal{A} and \mathcal{B} the two TRSs to operate on. When not otherwise specified, all symbols and notions not having a TRS label are to be intended operating on the disjoint sum $\mathcal{A} \oplus \mathcal{B}$. For better readability, we will talk of function symbols belonging to \mathcal{A} and \mathcal{B} like *white* and *black* functions. Variables and holes, instead, have both the colors. We also say a term/context is white (resp. black, transparent) if it is composed only by white (resp. black, transparent) symbols.

The *root* symbol of a term t is f provided $t = f(t_1, \dots, t_n)$, and t itself otherwise.

Let $t = C[t_1, \dots, t_n]$ and C not transparent; we write $t = C[[t_1, \dots, t_n]]$ if $C[\square_1, \dots, \square_n]$ is a white context and each of the t_i has a black and not transparent root, or vice versa (swapping the white and black attributes). The *topmost homogeneous part* (briefly *top*) of a term $C[[t_1, \dots, t_n]]$ is the context $C[\square_1, \dots, \square_n]$.

Definition 2.1 The *rank* of a term t ($rank(t)$) is 1 if t is black or white, and $\max_{i=1}^n \{rank(t_i)\} + 1$ if $t = C[[t_1, \dots, t_n]]$ ($n > 0$). \blacklozenge

The following well known lemma will be implicitly used in the sequel:

Lemma 2.2 ([Toy87b]) $s \rightarrow t \Rightarrow rank(s) \geq rank(t)$

Proof Clear. \blacklozenge

Definition 2.3 The multiset $S(t)$ of the *special subterms* of a term t is

1. $S(t) = \begin{cases} \{t\} & \text{if } t \text{ is black or white, and not transparent} \\ \emptyset & \text{if } t \text{ is transparent} \end{cases}$
2. $S(t) = \cup_{i=1}^n S(t_i) \cup \{t\}$ if $t = C[[t_1, \dots, t_n]]$ ($n > 0$)

The elements of $S(t)$ different from t are said the *proper special subterms* of t . \blacklozenge

Note that this definition is slightly different from the usual ones in the literature (for example in [Mid90]), since here variables are not considered special subterms.

Given a term s , we indicate with $\|s\|$ the multiset of the ranks of the special subterms of s . Multisets of this kind are compared according to the usual multiset ordering.

If $t = C[[t_1, \dots, t_n]]$, the t_i are called the *principal* special subterms of t . Furthermore, a reduction step of a term t is called *outer* if the rewrite rule is not applied in the principal special subterms of t .

Given a term t , and taken two special subterms of it, t_1 and t_2 , we say that t_1 is *above* t_2 (or, equivalently, that t_2 is *below* t_1), if t_2 is a proper special subterm of t_1 .

3 Bubbles

When studying the modular behaviour of some property, the main difficulty one has to face is that the behaviour of the reductions in the direct sum $\mathcal{A} \oplus \mathcal{B}$ can be quite complicated w.r.t. the reductions in the components \mathcal{A} and \mathcal{B} .

The disjointness requirement on \mathcal{A} and \mathcal{B} should ensure that symbols of one color cannot interact with symbols of another color. This is in a ‘static’ sense true, as we will see in Proposition 6.1.

The problem, however, is that this static ‘modular structure’ given by the subdivision into (tops of) special subterms is not fixed and immutable, but changes dynamically with every reduction step: via some reduction some of these tops could disappear (collapsing), making one of the the tops below (formally, one of the proper special subterms) possibly merge, interacting, with the above top.

For example, consider the term $f(H(A, b), G(a, f(a, b)))$ (with f , a , and b white, G , H and A black symbols). If the ‘black’ rule $G(X, Y) \rightarrow X$ is applicated to it, we get the term $f(H(A, b), a)$: the top $G(\square_1, \square_2)$ has disappeared, and the tops $f(\square_1, \square_2)$ and a have merge into the single context $f(\square_1, a)$. So, for instance, the white rule $f(X, a) \rightarrow b$, that could not be applied to the original term, can now be applied, despite no other white rule was used (hence we cannot reason ‘locally’ on the component TRSs).

The ‘responsible’ of this bad behaviour can be given a formal definition:

Definition 3.1 A *bubble* is a black or white context that reduces to a transparent context. \blacklozenge

Hence, bubbles are objects that can have a ‘solid’ state (i.e. a well-defined color), but are also potentially unstable, since they can collapse into a transparent object, losing their color.

If the top of a special subterm is a bubble, we call it a *top-bubble*; analogously, if the special subterm is proper we talk about a *proper top-bubble*.

We say a bubble has *degree* k if it reduces to exactly k distinct holes, and write $B(\square_1, \dots, \square_k)$ to denote a bubble of degree k that reduces to $\square_1, \dots, \square_k$. If $B(\square_1, \dots, \square_k)$ and $B[\square_1, \dots, \square_k]$, we write $B(\{\square_1, \dots, \square_k\})$. Moreover, we call a bubble of degree 1 (resp. > 1) *deterministic* (resp. *nondeterministic*).

A bubble $B(\square_1, \dots, \square_k)$ will be graphically denoted using the symbol



where the lines in the lower part represent the ‘free slots’ $\square_1, \square_2, \dots, \square_{k-1}, \square_k$.

Note that every TRS has (*trivial*) bubbles of degree 1, namely the transparent contexts. As far as bubbles of higher degrees are concerned, the following result holds:

Proposition 3.2 *If a TRS has a nondeterministic bubble, it has bubbles of every degree.*

Proof Suppose the TRS has a nondeterministic bubble $B(\square_1, \dots, \square_k)$ (with $k > 1$). Then

$$B(X, \dots, X, \square_i, B(X, \dots, X, \square_{i-1}, B(\dots B(X, \dots, X, \square_1) \dots)))$$

is a bubble of degree i , for every $i \geq 1$. \blacklozenge


So, what really cares is whether or not a TRS has a nondeterministic bubble.

If a TRS does not have a nondeterministic bubble, it is said (cf. [Mar93, Mar95]) *consistent with respect to reduction* (briefly CON^\rightarrow).

If it has one, it is said (cf. [Gra94, Ohl94b]) *non-deterministically collapsing*.¹

4 Neat Reductions

To be able to describe the special subterms of a given term throughout a reduction, it is natural to develop a concept of (modular) marking. A first, naïve approach of modular marking for a term is to take an assignment from the multiset of its special subterms to a (fixed) set of markers. Then reductions steps, as usual, should preserve the markers. However, this simple definition presents a problem, since for one case there is ambiguity: when a collapsing rule makes a proper top-bubble vanish (that is, when there is a nontrivial bubble at the top of a proper special subterm that collapses). In this case, we have the following situation:




and we have a conflict between m_2 and m_5 .

This situation is dealt with by defining a *modular marking* for a term to be an assignment from the multiset of its special subterms to *sets* of markers, and taking in the ambiguous case just described the union of the marker sets of the two special subterms involved.

¹Actually, the above definitions all concern the existence or not of a bubble of degree 2 (to be fussy, of a term reducing to two different variables), but, as just noticed, by Proposition 3.2 this is equivalent to talking about a nondeterministic bubble.

Thus, the previous example would give (singletons like $\{m_3\}$ are written simply m_3):



When this situation occurs, we say that the special subterm m_5 has been *absorbed* by m_2 , and the special subterm m_4 has had a *modular collapsing* (briefly *m-collapsing*).

When dealing with reductions $t \rightarrow t'$ we will always assume, in order to distinguish all the special subterms, that the initial modular marking of t is injective and maps special subterms to singletons.

We call a reduction *neat* if it has no m-collapsings.

Inside a reduction a notion of descendant for every special subterm can be defined: in a reduction a special subterm is a *descendant* (resp. *pure descendant*) of another if the set of markers of the former contains (resp. is equal to) the set of markers of the latter. Note, en passant, that due to the presence of duplicating rules, there may be more than one descendant, or even none (due to erasing rules). Observe also that, since in a reduction without m-collapsings all the descendants are pure, the first special subterm to m-collapse in a generic reduction is a pure descendant. Hence it readily holds the following:

Fact 1: *A reduction has m-collapsings iff a pure descendant m-collapses.*

Since special subterms are in bijective correspondence with their tops, we will be often sloppy talking about the descendants of a top, meaning the descendants of the corresponding special subterm.

5 Neatening

As previously hinted, it is just the presence of m-collapsings that complicates a lot the behaviour of a reduction in a direct sum of TRSs, making possible the interaction of initially distinct tops. When these interactions are not possible (i.e. when reductions are neat), different tops remain different, and so one can separately reason of every top as an independent term (cf. Proposition 6.1), making the modularity analysis much easier.

Historically, a first attempt to cope only with neat reductions was to syntactically limit the rewrite rules to ensure no bubble (but for the trivial ones of course) were present: if every rule is *non-collapsing* (viz. the right hand side

is not a transparent term), then readily no nontrivial bubble can exist, and so every reduction is automatically neat.

Indeed, every known property of interest is modular when (left-linear and non-collapsing TRSs are considered. Anyway, the restriction to non-collapsing TRSs is too heavy to be of great importance: it is the presence of collapsing rules that makes TRSs (and their combinations) so flexible.

So, avoiding the existence of (nontrivial) bubbles is effective for modularity but too restrictive. As a matter of fact, as seen, the real problem is not the presence of bubbles as such, but the presence of m-collapsings in reductions. So, the good ‘bottom’ notion of modularity is just that of modularity neatness: a property \mathcal{P} is said to be *modularly neat* if it is modular when only neat reductions are considered.

The general approach of ‘bare-bones neatening’ to prove a certain property \mathcal{P} is modular is so:

1. Show that \mathcal{P} is modularly neat
2. Show that if \mathcal{P} is modularly neat then \mathcal{P} is modular

In the paper we use, equivalently a reductio ad absurdum technique. We try to show that if \mathcal{P} is not modular, then it is not such even when only neat reductions are employed, hence contradicting Point 1.

After having sketched a ‘bare-bones’ version of neatening, we proceed on refining its definition.

Consider a modularity problem: to prove \mathcal{P} is modular, you have to prove that for every couple of TRSs \mathcal{A} and \mathcal{B} , $\mathcal{A} \in \mathcal{P} \ni \mathcal{B} \Leftrightarrow \mathcal{A} \oplus \mathcal{B} \in \mathcal{P}$. This means that in general two implications have to be considered. However, for all the properties of interest one of the two implications (\Leftarrow) is trivial. We so get rid of it by directly considering only dense properties: a property \mathcal{P} is said *dense* if whenever $\mathcal{A} \oplus \mathcal{B} \in \mathcal{P}$ then both \mathcal{A} and \mathcal{B} belong to \mathcal{P} .

Therefore, what neatening has to prove is that $\mathcal{A} \in \mathcal{P} \ni \mathcal{B} \Rightarrow \mathcal{A} \oplus \mathcal{B} \in \mathcal{P}$.

A *counterexample* (to the modularity of \mathcal{P}) is a pair of TRSs \mathcal{A} and \mathcal{B} such that $\mathcal{A} \in \mathcal{P} \ni \mathcal{B}$, $\mathcal{A} \oplus \mathcal{B} \notin \mathcal{P}$. A *provable counterexample* (to the modularity of \mathcal{P}) is a counterexample $(\mathcal{A}, \mathcal{B})$ to the modularity of \mathcal{P} together with a *proof* that $\mathcal{A} \oplus \mathcal{B} \notin \mathcal{P}$.

Readily,

$$\begin{array}{c}
 \exists \text{ a provable counterexample to the modularity of } \mathcal{P} \\
 \Downarrow \\
 \exists \text{ a counterexample to the modularity of } \mathcal{P} \\
 \Downarrow \\
 \mathcal{P} \text{ is not modular}
 \end{array}$$

Moreover, for all the dense properties it also holds the reverse implication

$$\begin{array}{c} \exists \text{ a counterexample to the modularity of } \mathcal{P} \\ \uparrow \\ \mathcal{P} \text{ is not modular} \end{array}$$

Hence, in the sequel we will tacitly assume that a property is not modular iff there is a provable counterexample to its modularity. Also, when talking about counterexamples we will often omit the appendix “to the modularity of \mathcal{P} ” (the property will be clear from the context).

We have seen that Point 1 of ‘bare-bones neatening’ roughly corresponds to modularity under the non-collapsing assumption. In general, proving this point is not a problem since this restriction is quite heavy: The problem lies in the second Point.

By the above implications, what we lack is only the implication

$$\begin{array}{c} \exists \text{ a provable counterexample to the modularity of } \mathcal{P} \\ \Downarrow (?) \\ \exists \text{ a neat provable counterexample to the modularity of } \mathcal{P} \end{array}$$

If we had this, we could reason as follows:

$$\begin{array}{c} \mathcal{P} \text{ is not modular} \\ \Downarrow \\ \exists \text{ a counterexample to the modularity of } \mathcal{P} \\ \Downarrow \\ \exists \text{ a provable counterexample to the modularity of } \mathcal{P} \\ \Downarrow (?) \\ \exists \text{ a neat provable counterexample to the modularity of } \mathcal{P} \\ \Downarrow \\ \exists \text{ a neat counterexample to the modularity of } \mathcal{P} \\ \Downarrow \\ \mathcal{P} \text{ is not modularly neat} \end{array}$$

thus obtaining the contradiction with Point 1.

The idea of neatening is to prove the missing implication (?) using a ‘neatening translation’ that transforms every generic reduction into a neat reduction. This way, it can be applied to the proof of the provable counterexample, yielding a neat provable counterexample.

Hence, the technique of (*abstract*) *neatening* is:

Suppose a dense property \mathcal{P} is such that

1. *it is modularly neat*
2. *if there is a counterexample, then it can be extended to a provable counterexample that is transformed via a ‘neatening translation’ into another neat provable counterexample*

Then \mathcal{P} is modular

Observe we have slightly stressed Point 2, since it would have sufficed to say: there is a provable counterexample that is transformed via a ‘neatening translation’ into another neat provable counterexample.

6 Pile and Paint

In this section we provide the formal definition of a ‘neatening translation’ that makes the neatening method work.

Visually, the intuition is that a (nontrivial) bubble, as seen, is a term that cannot properly have a color on itself, since it can reduce to a transparent object: this way it assumes the color of the objects it stays near. So, when a proper top-bubble is present, we have the unpleasant situation that two tops of one color are separated by a potentially transparent object (the bubble) that has for the moment a different color: a situation which is highly unstable.

The solution is: we get rid of this bubble by attaching it to every top of its same color which is above it (pile operation), and then change the bubble’s color (paint operation), so that the unstable situation disappear. Note that it is not dangerous to attach the bubble to other terms, as we do with the pile operation: presence of bubbles is in general unavoidable (recall the discussion on non-collapsing TRSs); what is dangerous is only the unstable situation above described (that can lead to m-collapsings), and when a bubble is inside a non-bubble top of the same color even if it gets transparent the overall color of the top does not change.

The following simple proposition (that will be often considered understood) is nevertheless fundamental, explaining why left-linearity is so important:

Proposition 6.1 *If a TRS is left-linear, then rewrite rules that have the possibility to act outer on a special subterm t are exactly those that have the possibility to act on its top.*

Proof Let $t = C[[t_1, \dots, t_n]]$: since t_1, \dots, t_n have a root belonging to the other TRS (with respect to C), they are matched by variables from any rewrite rule applicable to C , and for the left-linearity assumption these variables are independent each other. \blacklozenge

Roughly speaking, the proposition says that when left-linearity is present, rewrite rules that are applied to the top of a special subterm do not ‘look below’, i.e. they do not care at all about the special subterms that are below. This means that we can modify all these special subterms, without preventing the application of such rewrite rules (that act, so to say, ‘locally’).

Assumption: *From now on, every TRS, unless otherwise specified, is understood to be left-linear.*

Definition 6.2 (Pile and Paint)

The Pile and Paint transformation of a term s (notation $\pi(s)$) is obtained as follows.

Select the leftmost (in writing order) proper special subterm of s that has rank minimal amongst the ones with a bubble as top: say $t = B(\langle t_1, \dots, t_k \rangle)$. Without loss of generality, we suppose it is top white.

If no such t is present, we leave the term unchanged (i.e. $\pi(s) = s$).

Otherwise, we define $\pi(s)$ as the term obtained from s after the two following operations.

Pile:

We ‘pile’ the bubble $B(\langle \square_1, t_2, \dots, t_k \rangle)$ just below the tops of all the above white special subterms. That is, if a top black special subterm of s is of the form $r[r_1, \dots, r_m]$, with r_j above t , we pass from r to

$$r[r_1, \dots, r_{j-1}, B(\langle r_j, t_2, \dots, t_k \rangle), r_{j+1}, \dots, r_m]$$

The operation is shown in Figure 1.

Paint:

We change the color of the bubble B , replacing it with another black bubble $b(\langle \square_1, \dots, \square_k \rangle)$ of the same degree. So, t passes from $B(\langle t_1, \dots, t_k \rangle)$ to $b(\langle t_1, \dots, t_k \rangle)$. The operation is shown in Figure 2. \blacklozenge

Remark 6.3 The transformation π chooses at the beginning the leftmost proper special subterm of s that has rank minimal amongst the ones with a bubble as top (roughly speaking, it selects the leftmost and uppermost proper top-bubble). The requirement of being the leftmost, however, is completely arbitrary for our purposes, since it can be dropped. However, we use it not to heaven the transformation using an additional parameter indicating which top-bubble has been *selected*.

Analogously, in the pile operation we inserted r_j in the first slot of B : this is not necessary, since every slot could be used, but for commodity we fix one (the first). Hence in the sequel, when saying that B *collapses* without specifying to what, we will mean to its first slot \square_1 . \blacklozenge

When can π be applied? The only problematic step is the paint one, where we change the color of a bubble replacing it with another from the other TRS having the same degree. Hence, a sufficient condition for the applicability of π is:

Fact 2: π can be applied if the two TRSs have bubbles of the same degree.

Equivalently, for Proposition 3.2, the above fact can be restated as: *the two TRSs must be either both CON^{\rightarrow} or both $\neg CON^{\rightarrow}$.*

We now show how from $\Pi(s)$ we can still mimic the old reduction. The intuition is that the bubbles that we piled can be needed if during the original reduction, via other bubbles’ collapsings, the selected bubble was absorbed. The ‘painted’ bubble, instead, is needed when the original selected bubble collapsed: we make this new bubble collapse to the same ‘slot’. Also, when all these bubbles (piled and painted) are not needed any more, they can be deleted by simply making them collapse.




Figure 1: The Pile operation.




Figure 2: The Paint operation.

Definition 6.4 (Mimicking m)

Given a reduction ρ of s , we define the corresponding *mimicking reduction* $\mathbf{m}(\rho)$.

The start term is $\pi(\rho)$.

Then, we simply use the rules of ρ with the following modifications:

- If the rule was applied to a pure descendant of t :
 - If the rule made a pure descendant of t m-collapse to a (a descendant of) t_i , then consider the corresponding term t' :
 1. Collapse t' into t_i
 2. Act with the corresponding reduction of ρ on (that descendant of) t_i .

The situation is shown in Figure 3.

- Otherwise, skip that rule.
- If the rule was not applied to a pure descendant of t , the rule is applied, and moreover:
 - If the rule made a pure descendant of t be absorbed, then consider the corresponding term t' :
 1. Collapse t' (this recreates a fresh copy of t , see Figure 4).
 2. Act with the corresponding reduction of ρ on (that descendant of) t on this newly created copy of t .

The situation is illustrated in Figure 4.

- If a rule made a black special subterm m-collapse, collapse the bubble piled immediately above it (if it is not erased). See Figure 5.
- If a rule made a white special subterm m-collapse, collapse the bubble piled immediately below it (if it is not erased). See Figure 6.

◆


Lemma 6.5 *If in ρ eventually there are no pure descendants of the selected bubble B , then $\mathbf{m}(\rho)$ is cofinal for ρ .*

Proof The assumption says that we have $s \xrightarrow[\rho']{\rightarrow} s'$, $\rho' \subseteq \rho$, and in s' there are no pure descendants of the bubble B .

It is easy to see that $\rho' \rightarrow \mathbf{m}(\rho')$. Indeed, the only differences between the original reduction and its mimicked counterpart are the extra presence in the mimicking of the piled bubbles, and a bubble of different color in place of B . When we reach a term having no pure descendants of B , it is immediate from the definition of mimicking that the sole difference now present is the piled bubbles. So, it suffices to collapse all of them to get back the original term.

Now we can prove that $\rho \rightarrow \mathbf{m}(\rho)$: take $t_1 \in \rho$ ($s \xrightarrow[\zeta]{\rightarrow} t_1$, $\zeta \subseteq \rho$). We choose a $\rho' \subseteq \rho$ sufficiently big such that $\zeta \subseteq \rho'$ and reducing s by ρ' yields no pure

Original reduction:



Mimicked reduction:





Figure 3: Mimicking of the m-collapsing of the selected bubble.

Original reduction:



Mimicked reduction:





Figure 4: Mimicking of the absorption of the selected bubble.

Original reduction:



Mimicked reduction:





Figure 5: Mimicking of the m-collapsing of a top black special subterm.

Original reduction:



Mimicked reduction:




Figure 6: Mimicking of the m-collapsing of a top white special subterm.

descendants of the selected bubble B . We have seen that $\rho' \twoheadrightarrow \mathfrak{m}(\rho')$. Moreover, $\mathfrak{m}(\rho') \subseteq \mathfrak{m}(\rho)$ (this stems from the general fact that $\xi_1 \subseteq \xi_2 \Rightarrow \mathfrak{m}(\xi_1) \subseteq \mathfrak{m}(\xi_2)$), and so ($\zeta \subseteq \rho \twoheadrightarrow \mathfrak{m}(\rho') \subseteq \mathfrak{m}(\rho)$) we get $\zeta \twoheadrightarrow \mathfrak{m}(\rho)$, and hence just $\rho \twoheadrightarrow \mathfrak{m}(\rho)$ by the arbitrariness of $\zeta (\subseteq \rho)$. \blacklozenge

6.1 Multiple Pile and Paint

The transformation π (together with \mathfrak{m}) makes the structure of a term (of a reduction) more stable in the sense that it gets rid of a bubble. It is therefore natural to try to repeat this simplification process as far as possible. The following lemma shows that the iteration of this process is indeed terminating:

Lemma 6.6 *If $\pi(s) \neq s$, $\|\pi(s)\| < \|s\|$.*

Proof Immediate, since the paint operation drops a special subterm, whereas the pile operation possibly adds only special subterms of strictly inferior rank. \blacklozenge

We can so repeat the application of π until we obtain a term having no proper top-bubble: this happens in a finite number of steps because of the above lemma.

We indicate with $\Pi(s)$ the output of this process.

Readily, the applicability conditions for π (Fact 2) still hold for Π .

Note that the measure $\|\cdot\|$ shows also that the termination process of Π is a basic ‘syntactical’ property, not depending on ‘semantical’ arguments (the bubble).

Π enjoys the following property:

Lemma 6.7 *Every reduction of $\Pi(s)$ is neat.*

Proof $\Pi(s)$ has, by definition, no top-bubble. Thus, no pure descendant can \mathfrak{m} -collapse, and this implies by Fact 1 that every its reduction is neat. \blacklozenge

We call $\mathfrak{M}(\rho)$ the mimicking reduction associated with Π , obtained from ρ by repeatedly applying \mathfrak{m} until the start term is $\Pi(s)$ (where s is the start term of ρ): this is the ‘neatening reduction’ that we will use.

Incidentally, observe that \mathfrak{M} is even more powerful than required by neatening, since by Lemma 6.7 not only it gives neat reductions, but even reductions without proper top-bubbles.

\mathfrak{M} inherits from \mathfrak{m} the following result:

Lemma 6.8 *If in ρ eventually there are no pure descendants of all the proper tops, then $\mathfrak{M}(\rho)$ is cofinal for ρ .*

Proof Immediate from Lemma 6.5, once noticed the transitivity of the cofinality relation. \blacklozenge

7 Modularity

It has arrived the moment to apply the machinery we have developed, stating the main theorem (recall the all the TRSs are assumed to be left-linear). First we need a definition:

Definition 7.1 A property \mathcal{P} is said *pseudo-deterministic* (respectively *pseudo-nondeterministic*) if $\exists T \forall T'. T' \in \mathcal{P} \Rightarrow T \oplus T' \in \mathcal{P} \wedge \text{CON}^\rightarrow$ (respectively $\in \mathcal{P} \wedge \neg \text{CON}^\rightarrow$). \blacklozenge

Pseudo-determinism is in tight relationship with consistency w.r.t. reduction:

Lemma 7.2 A dense property \mathcal{P} is pseudo-deterministic iff \mathcal{P} implies CON^\rightarrow .

Proof The if direction is always satisfied, since the empty TRS is CON^\rightarrow . For the only if direction, observe that being \mathcal{P} and CON^\rightarrow dense, than also $\mathcal{P} \wedge \text{CON}^\rightarrow$ is such, and so from the pseudo-determinism of \mathcal{P} we get $T' \in \mathcal{P} \Rightarrow T' \in \mathcal{P} \wedge \text{CON}^\rightarrow$, which implies $\mathcal{P} \Rightarrow \text{CON}^\rightarrow$. \blacklozenge

Theorem 7.3 (Main) Suppose that a dense property \mathcal{P} is either pseudo-deterministic or pseudo-nondeterministic, and

1. \mathcal{P} is modularly neat
2. If there is a counterexample, then it can be extended to a provable counterexample that is translated by \mathfrak{M} into another provable counterexample

Then \mathcal{P} is modular.

Proof Suppose \mathcal{P} is pseudo-deterministic. If \mathcal{P} is not modular, there is a counterexample and so by Point 2 there is also a provable counterexample obtained translated via \mathfrak{M} (\mathfrak{M} can be applied by Fact 2 and Lemma 7.2): but this provable counterexample must be neat by Lemma 6.7, hence contradicting Point 1.

On the other hand, suppose \mathcal{P} is pseudo-nondeterministic. If \mathcal{P} is not modular, there is a counterexample (to the modularity of \mathcal{P}), viz. $\mathcal{A} \in \mathcal{P} \ni \mathcal{B}$, $\mathcal{A} \oplus \mathcal{B} \notin \mathcal{P}$. Being \mathcal{P} pseudo-nondeterministic, there is a TRS T such that $\mathcal{A} \oplus T \in \mathcal{P} \wedge \text{CON}^\rightarrow \ni \mathcal{B} \oplus T$. Also, by the density of \mathcal{P} it follows $\mathcal{A} \oplus \mathcal{B} \notin \mathcal{P} \Rightarrow (\mathcal{A} \oplus T) \oplus (\mathcal{B} \oplus T) \notin \mathcal{P}$. Hence, $\mathcal{A} \oplus T$ and $\mathcal{B} \oplus T$ give again a counterexample.

By Point 2, we can extend it to a provable counterexample, that is translated by \mathfrak{M} into another provable counterexample (note \mathfrak{M} can be applied since $\mathcal{A} \oplus T$ and $\mathcal{B} \oplus T$ are readily both $\neg \text{CON}^\rightarrow$, and so using Fact 2).

But this new provable counterexample is neat by Lemma 6.7, hence contradicting Point 1. \blacklozenge

We now apply this general theorem to several properties. For the sake of clarity, we repeat in all the results the so far understood assumption of left-linearity.

Termination

Termination is in general not a modular property (see e.g. [Toy87a]). Via Theorem 7.3 we will prove the state-of-the-art results on its modularity for left-linear TRSs, and also provide two new results.

Lemma 7.4 *Termination is modularly neat for left-linear TRSs.*

Proof Suppose a term s has an infinite reduction. Then at least one of its special subterms has an infinite number of rewrite rules applied to it (s descendants). Take one with minimal rank, say t . If the infinite reduction is neat, an infinite number of rewrite rules applies also to the top of t , thus obtaining an infinite reduction of a term with only one color. \blacklozenge

Note that the above result also holds in the non left-linear case, using the same proof with slight modifications (in place of the top of t , say $C[\square_1, \dots, \square_n]$, the context $C[\square_1, \dots, \square_1]$ must be used, for the possible presence of non left-linear rewrite rules). Since a non-collapsing TRS is also modularly neat, this generalizes the result of Rusinowitch ([Rus87]) stating the modularity of termination for non-collapsing TRSs.

Corollary 7.5 *If the TRSs are left-linear and either both pseudo-deterministic or both pseudo-nondeterministic, then termination is modular.*

Proof The above lemma shows Point 1 of Theorem 7.3. For Point 2, take an infinite reduction with the minimum rank of the start term. This reduction must have an infinite number of rewrite rules applied on the top of the start term (since all the proper special subterms are terminating by hypothesis). But \mathfrak{M} does not modify these rules, and hence the obtained reduction is still infinite. \blacklozenge

This result entails the main results of [Mar95, SSMP95]:

Corollary 7.6 *Termination is modular for left-linear and consistent w.r.t. reduction TRSs*

Proof By the above Corollary 7.5 and Lemma 7.2. \blacklozenge

Now we consider the other ‘dual’ result that Corollary 7.5 offers.

Call OR the TRS $\{or(X, Y) \rightarrow X, or(X, Y) \rightarrow Y\}$; a TRS T is said *termination preserving under non-deterministically collapses* (briefly $\mathcal{C}_\mathcal{E}$ -terminating) if $T \oplus OR$ is terminating. Gramlich in [Gra94] proved that $\mathcal{C}_\mathcal{E}$ -termination is modular for finitely branching TRSs. Later, Ohlebusch (cf. [Ohl94b]) extended this result to arbitrary TRSs dropping the finitely branching condition. We can entail Gramlich’s and Ohlebusch’s result in the left-linear case:

Corollary 7.7 *$\mathcal{C}_\mathcal{E}$ -termination is modular for left-linear TRSs.*

Proof It follows from Corollary 7.5 once observed that $\mathcal{C}_\mathcal{E}$ -termination is pseudo-nondeterministic and implies termination. \blacklozenge

Another criterion for the modularity of termination was proven by Middeldorp in [Mid89]: he showed that whenever one of two terminating TRSs is both non-collapsing and non-duplicating, then their disjoint sum is modular. Using the two above corollaries, we can not only entail this result in the left-linear case, but even properly generalize it with the following new result:

Corollary 7.8 *Suppose two left-linear TRSs are terminating. Then if one of them is both CON^\rightarrow and $\mathcal{C}_\mathcal{E}$ -terminating, their disjoint sum is terminating.*

Proof Every TRS is either CON^\rightarrow or $\neg\text{CON}^\rightarrow$. So, take two terminating TRSs, with one of the two CON^\rightarrow and $\mathcal{C}_\mathcal{E}$ -terminating: if the other is CON^\rightarrow , their sum is terminating by Corollary 7.6, otherwise if it is $\neg\text{CON}^\rightarrow$ then it is also $\mathcal{C}_\mathcal{E}$ -terminating, and so their sum is terminating by Corollary 7.7. \blacklozenge

Corollary 7.9 *Suppose two left-linear TRSs are terminating. Then if one of them is both non-collapsing and non-duplicating, their disjoint sum is terminating.*

Proof By the above corollary, since non-collapsing $\Rightarrow \text{CON}^\rightarrow$, and a TRS which is both terminating and non-duplicating is $\mathcal{C}_\mathcal{E}$ -terminating (as it is easy to show, see e.g. [Gra94]). \blacklozenge

We now turn our attention to the structure of counterexamples to the modularity of termination. So far, two main results are known. Ohlebusch in [Ohl94b] (again, extending a result of Gramlich in [Gra94] for finitely branching TRSs), showed that in every counterexample one of the TRSs is not $\mathcal{C}_\mathcal{E}$ -terminating and the other is collapsing. Schmidt-Schauß, Marchiori and Panitz showed in [SSMP95] that, in the left-linear case, in every counterexample one of the TRSs is CON^\rightarrow and the other is $\neg\text{CON}^\rightarrow$. Both of these results require a nontrivial proof. Here we show how we can easily obtain not only the previous two results (in the left-linear case of course), but even a single result that properly generalizes both of them.

First, we prove the result of [SSMP95]:

Corollary 7.10 *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRS is CON^\rightarrow and the other is $\neg\text{CON}^\rightarrow$.*

Proof Since every TRS is either CON^\rightarrow or $\neg\text{CON}^\rightarrow$, only three cases are possible: 1) both are CON^\rightarrow , 2) both are $\neg\text{CON}^\rightarrow$, 3) like in the statement of this corollary. But 1) is not possible by Corollary 7.6, whereas 2) is not possible by Corollary 7.7 (since every terminating and $\neg\text{CON}^\rightarrow$ TRS is trivially $\mathcal{C}_\mathcal{E}$ -terminating). \blacklozenge

Next we show a somehow dual result:

Corollary 7.11 *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRS is $\mathcal{C}_\mathcal{E}$ -terminating and the other is $\neg\mathcal{C}_\mathcal{E}$ -terminating.*

Proof Completely analogous to the proof of the above corollary. \blacklozenge

We can now prove the following result that generalizes all the previous ones:

Corollary 7.12 *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRS is $\neg\text{CON}^\rightarrow$ and the other is $\neg\mathcal{C}_\mathcal{E}$ -terminating.*

Proof From Corollaries 7.10 and 7.11, in every counterexample only two cases are possible: 1) one of the TRSs is $\mathcal{C}_\mathcal{E}$ -terminating $\wedge\text{CON}^\rightarrow$ and the other is $\neg\mathcal{C}_\mathcal{E}$ -terminating $\wedge\neg\text{CON}^\rightarrow$, or 2) one of the TRSs is $\mathcal{C}_\mathcal{E}$ -terminating $\wedge\neg\text{CON}^\rightarrow$ and the other is $\neg\mathcal{C}_\mathcal{E}$ -terminating $\wedge\text{CON}^\rightarrow$. But the first case is ruled out by Corollary 7.8. On the other hand, by the fact that for terminating TRSs $\neg\text{CON}^\rightarrow \Rightarrow \mathcal{C}_\mathcal{E}$ -termination, it follows right away that case 2) is just the statement of this corollary, since for terminating TRSs $\mathcal{C}_\mathcal{E}$ -termination $\wedge\neg\text{CON}^\rightarrow \Leftrightarrow \neg\text{CON}^\rightarrow$ and $\neg\mathcal{C}_\mathcal{E}$ -termination $\wedge\text{CON}^\rightarrow \Leftrightarrow \mathcal{C}_\mathcal{E}$ -termination. \blacklozenge

As previously claimed, this result properly generalizes (besides Corollary 7.11) the result of [SSMP95] (Corollary 7.10) and the result of (Gramlich and Ohlebusch ([Gra94, Ohl94b]):

Corollary 7.13 *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRS is $\neg\mathcal{C}_\mathcal{E}$ -terminating and the other is collapsing.*

Proof Trivial by the above corollary, since $\neg\text{CON}^\rightarrow \Rightarrow \text{collapsing}$. \blacklozenge

Uniqueness of Normal Forms w.r.t. Reduction

A TRS is said to have the unique normal forms w.r.t. reduction (briefly, to be UN^\rightarrow), if every term has at most one normal form. The UN^\rightarrow property is not modular in general (cf. [Mid90]); whether it is modular or not for left-linear TRSs was the last open problem in the modularity of the basic properties of TRSs (cf. [DJK91]); this problem was finally shown to have a positive solution in [Mar93] (see also the discussion in Subsection 7.1). We now show how also this result can be obtained.

Lemma 7.14 *UN^\rightarrow is modularly neat for left-linear TRSs.*

Proof By rank induction: using neat reductions, every term can be reduced to normal form by separately reducing its top (being of rank 1, it has an unique normal form), and its principal special subterms (they have unique normal forms by rank induction). \blacklozenge

Corollary 7.15 *UN^\rightarrow is modular for left-linear TRSs.*

Proof The above lemma shows Point 1 of Theorem 7.3. For Point 2, take a counterexample to the modularity of UN^\rightarrow , i.e. a term s reducing to two distinct normal forms n_1 (via ρ_1) and n_2 (via ρ_2). Since n_1 and n_2 are normal forms, no bubble can be present, and hence by Lemma 6.8 $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$. But, again, being n_1 and n_2 normal forms implies that $\mathfrak{M}(\rho_1)$ reduces $\Pi(s)$ to n_1 and $\mathfrak{M}(\rho_2)$ reduces $\Pi(s)$ to n_2 , hence giving a counterexample (which is neat by Lemma 6.7). \blacklozenge

Consistency w.r.t. Reduction

Although not modular in general ([Mar93]), CON^\rightarrow is modular for left-linear TRSs, as shown for the first time in [Mar93]. We now prove this result.

Lemma 7.16 *CON^\rightarrow is modularly neat for left-linear TRSs.*

Proof If a term reduces to a variable via a neat reduction, then its top does also. \blacklozenge

Corollary 7.17 *CON^\rightarrow is modular for left-linear TRSs.*

Proof The above lemma shows Point 1 of Theorem 7.3. For Point 2, take a counterexample to the modularity of CON^\rightarrow , viz. a term s reducing to two distinct variables X (via ρ_1) and Y (via ρ_2). No bubbles are readily present in X and Y , and hence by Lemma 6.8 $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$. But being X and Y variables implies that $\mathfrak{M}(\rho_1)$ reduces $\Pi(s)$ to X and $\mathfrak{M}(\rho_2)$ reduces $\Pi(s)$ to Y , thus giving a counterexample (neat by Lemma 6.7). \blacklozenge

The importance of this result, besides theoretical, lies in the fact that it allows to use the result on the modularity of termination obtained in Corollary 7.5 for more than two TRSs, since the disjoint sum of two left-linear, terminating and either both CON^\rightarrow or both $\neg\text{CON}^\rightarrow$ TRSs is still left-linear, terminating and either CON^\rightarrow or $\neg\text{CON}^\rightarrow$.

Confluence

Toyama in his famous paper [Toy87b] (see also [KMTdV94]) proved that confluence is a modular property: we can entail this result in the left-linear case.

Lemma 7.18 *Confluence is modularly neat for left-linear TRSs.*

Proof By rank induction. Suppose a term $s = C[[t_1, \dots, t_n]]$ reduces to t_1 (via ρ_1) and to t_2 (via ρ_2). The rewrite steps of ρ_1 (and of ρ_2) that act outer on a descendant of s can be applied to its top as well (Proposition 6.1), and every pure descendant of s in ρ_1 has its top joinable with the top of every pure descendant of s in ρ_2 . On the other hand, the descendants of the special subterms t_1, \dots, t_n in ρ_1 are joinable to the corresponding descendants in ρ_2 by the induction hypothesis: hence every term in ρ_1 is joinable with every term in ρ_2 . \blacklozenge

Corollary 7.19 *Confluence is modular for left-linear TRSs.*

Proof The above lemma shows Point 1 of Theorem 7.3. For Point 2, suppose a term s reduces to t_1 (via ρ_1) and to t_2 (via ρ_2), and that t_1 and t_2 are not joinable. Without loss of generality, we can suppose t_1 and t_2 have no proper top-bubbles: if it is not the case, finitely extend every reduction by repeatedly selecting a proper top-bubble of maximal rank and m-collapsing it. By Lemma 6.8, $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$, and so $\Pi(s)$ reduces via the neat (cf. Lemma 6.7) reductions $\mathfrak{M}(\rho_1)$ and $\mathfrak{M}(\rho_2)$ to two terms that are still not joinable, and hence *a fortiori* not joinable using neat reductions. \blacklozenge

Weak Normalization

Weak Normalization (briefly WN) was contemporary proven to be modular by several authors (see [Mid90] for some references): we can entail this result in the left-linear case.

Lemma 7.20 *WN is modularly neat for left-linear TRSs.*

Proof By rank induction. Taken a term s , its top reduces to a normal form (being of an unique color); by Proposition 6.1 we can apply these rules to s as well (not this reduction is neat). The obtained term has its top in normal form, and so we can reduce to normal form its proper special subterms (by the induction hypothesis), obtaining a normal form. \blacklozenge

Corollary 7.21 *WN is modular for left-linear TRSs.*

Proof We first prove that weak normalization is pseudo-nondeterministic. Consider the TRS $OR = \{or(X, Y) \rightarrow X, or(X, Y) \rightarrow Y\}$. Take a TRS T' which is WN: $T' \oplus OR \in WN \wedge \neg CON^{\rightarrow}$. Indeed, $T' \oplus OR \in \neg CON^{\rightarrow}$ is trivial; on the other hand, $T' \oplus OR \in WN$: taken a term s , we can normalize it w.r.t. OR , so obtaining a term in T' that is normalizable by hypothesis.

We can so apply Theorem 7.3: the above Lemma 7.20 shows Point 1; for Point 2, take a term s not having a normal form: if a top-bubble is present in it, repeatedly collapse it (no matter to what ‘slot’), till you obtain a term s' still having no normal forms. Thus, $\Pi(s') = s'$ has not normal forms, and so *a fortiori* has not a normal form reachable by neat reductions. \blacklozenge

Completeness

Completeness, as well known, is the conjunction of confluence and termination. Despite it is not modular in general, it was proven to be modular for left-linear TRSs by Toyama, Klop and Barendregt in their ingenious paper [TKB89] (see also [TKB94]); the proof of such result, however, is ‘rather intricate and not easily digested’ (citing the same authors). This result can instead be obtained as a simple corollary:

Corollary 7.22 *Completeness is modular for left-linear TRSs.*

Proof Since completeness equals to termination and uniqueness of normal forms w.r.t. reduction, the result follows from Corollaries 7.6 and 7.19. \blacklozenge

Note that a direct proof of the above result via Theorem 7.3 is also easy to obtain.

Semi-Completeness

Semi-completeness is the property obtained by the conjunction of confluence and weak normalization. It is immediate to prove its modularity for left-linear TRSs:

Corollary 7.23 *Semi-completeness is modular for left-linear TRSs.*

Proof From Corollaries 7.19 and 7.21. \blacklozenge

Again, note that a direct proof of the above result via Theorem 7.3 is easy to obtain.

The Other Properties

So far, we mentioned all the main properties of TRSs, but for these last four: local confluence (WCR), consistency (CON), uniqueness of normal forms (UN) and the normal form property (NF) (for their definition, see e.g. [DJ90, Klo92]). It is not difficult to see that even these remaining properties can be proven to be modular for left-linear TRSs using Theorem 7.3. The only point worth mentioning is that all these properties are pseudo-deterministic but for local confluence, which can be proven to be pseudo-nondeterministic using the TRS $\{f(X, Y) \rightarrow X, f(X, Y) \rightarrow g(X, Y), g(X, Y) \rightarrow f(X, Y), g(X, Y) \rightarrow Y\}$.

7.1 Paint vs. Delete

The reader may have noticed a kind of duality inside Theorem 7.3, since the property is required to be *either* pseudo-deterministic *or* pseudo-nondeterministic.

As we have seen, requiring pseudo-determinism essentially equals to requiring consistency w.r.t. reduction (Lemma 7.2). So in this case every bubble is by definition of degree one. But, as noticed in Section 3, every TRS has trivial bubbles of degree one, namely the transparent contexts. Hence, when in the Paint operation we change color to the bubble, we can do it by always using a trivial bubble (viz. a hole). This corresponds, in practice, to *delete* the selected top-bubble. This is just what was done in the ‘pile and delete’ technique that was introduced in [Mar93] for the study of the modularity of UN^\rightarrow (and later used in [Mar95] and [SSMP95]), of which this transformation is a refinement and a generalization.

So, when coping only with pseudo-deterministic properties we can use the method presented in this paper slightly simplified using the ‘delete’ operation in place of the more general paint one, and dropping the concepts of pseudo-determinism and pseudo-nondeterminism (by Lemma 7.2 we can modify Theorem 7.3 by directly requiring that the property \mathcal{P} implies CON^\rightarrow). This allows to treat the great majority of the considered properties. What we lose is: treatment of the properties that essentially require pseudo-nondeterminism ($\mathcal{C}_\mathcal{E}$ -termination, weak normalization and local confluence), the criterion for the modularity of termination given by Corollary 7.8, and all the results on the structure of counterexamples (Corollaries 7.10, 7.11, 7.12, 7.13).

8 Conclusions

We have introduced a uniform technique which is able to successfully deal with the modularity of all the basic properties of TRSs in the left-linear case, and also to provide some new results on the modularity of termination. Moreover, the technique is intuitively appealing, since it relies on visual arguments, making the involved reasonments more intuitive and easier to grasp.

This can be seen as a first step towards the ambitious task of providing a global technique to cope with modularity (i.e., dropping the left-linearity requirement). In our opinion, such a technique can be developed on the basis of the ideas underlying the method. Indeed, note that left-linearity is only explicitly required in the construction of the specific ‘neatening translation’ \mathfrak{M} , not by abstract neatening. So, a promising line of research would be trying to develop a suitable neatening translation such that abstract neatening can work even in the presence of non left-linear rewrite rules.

References

- [DJ90] N. Dershowitz and J. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 6, pages 243–320. Elsevier – MIT Press, 1990.
- [DJK91] N. Dershowitz, J. Jouannaud, and J.W. Klop. Open problems in rewriting. In R.V. Book, editor, *Proceedings of the Fourth International Conference on Rewriting Techniques and Applications*, volume 488 of *LNCS*, pages 445–456. Springer–Verlag, 1991.
- [Gra94] B. Gramlich. Generalized sufficient conditions for modular termination of rewriting. *Applicable Algebra in Engineering, Communication and Computing*, 5:131–158, 1994.
- [Klo92] J.W. Klop. Term rewriting systems. In S. Abramsky, Dov M. Gabbay, and T.S.E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 2, chapter 1, pages 1–116. Clarendon Press, Oxford, 1992.
- [KMTdV94] J.W. Klop, A. Middeldorp, Y. Toyama, and R. de Vrijer. Modularity of confluence: A simplified proof. *Information Processing Letters*, 49(2):101–109, 1994.
- [Mar93] M. Marchiori. Modularity of UN^{\rightarrow} for left-linear term rewriting systems. Draft, March 1993. Extended and revised version: Technical Report CS-R9433, CWI, Amsterdam.
- [Mar95] M. Marchiori. Modularity of completeness revisited. In J. Hsiang, editor, *Proceedings of the Sixth International Conference on Rewriting Techniques and Applications*, volume 914 of *LNCS*, pages 2–10. Springer–Verlag, April 1995.
- [Mid89] A. Middeldorp. A sufficient condition for the termination of the direct sum of term rewriting systems. In *Proceedings of the Fourth IEEE Symposium on Logic in Computer Science*, pages 396–401, 1989.

- [Mid90] A. Middeldorp. *Modular Properties of Term Rewriting Systems*. PhD thesis, Vrije Universiteit, Amsterdam, November 1990.
- [Ohl94a] E. Ohlebusch. *Modular Properties of Composable Term Rewriting Systems*. PhD thesis, Universität Bielefeld, 33501 Bielefeld, FRG, May 1994. Appeared as Technical Report 94-01.
- [Ohl94b] E. Ohlebusch. On the modularity of termination of term rewriting systems. *Theoretical Computer Science*, 136(2):333–360, December 1994.
- [Rus87] M. Rusinowitch. On termination of the direct sum of term rewriting systems. *Information Processing Letters*, 26:65–70, 1987.
- [SSMP95] M. Schmidt-Schauß, M. Marchiori, and S. Panitz. Modular termination of r-consistent and left-linear term rewriting systems. *TCS*, 1995. To appear.
- [TKB89] Y. Toyama, J.W. Klop, and H.P. Barendregt. Termination for the direct sum of left-linear term rewriting systems. In N. Derzhovitz, editor, *Proceedings of the Third International Conference on Rewriting Techniques and Applications*, volume 355 of *LNCS*, pages 477–491, Chapel Hill, April 1989. Springer-Verlag. Extended version: Report CS-R8923, CWI, Amsterdam.
- [TKB94] Y. Toyama, J.W. Klop, and H.P. Barendregt. Termination for direct sums of left-linear complete term rewriting systems. *Journal of the ACM*, 1994. To appear.
- [Toy87a] Y. Toyama. Counterexamples to termination for the direct sum of term rewriting systems. *Information Processing Letters*, 25:141–143, 1987.
- [Toy87b] Y. Toyama. On the Church-Rosser property for the direct sum of term rewriting systems. *Journal of the ACM*, 1(34):128–143, 1987.