ELSEVIER

Theoretical Computer Science 192 (1998) 31–54

Theoretical Computer Science

Bubbles in modularity

Massimo Marchiori*

Department of Pure and Applied Mathematics, University of Padova, Via Belzoni 7, 35131 Padova, Italy

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 responsibles of most of the problems occurring in modularity, and the concept of well-behaved (from the modularity point of view) reduction, called 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 elegantly 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.

Keywords: Term rewriting system; Left-linearity; Modularity; Verification

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 smaller safe submodules.

As far as the paradigm of term rewriting systems (TRSs) is concerned, the notion of modularity it 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. [15]).

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

* E-mail: max@math.unipd.it.

PII S0304-3975(97)00144-8

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

First, we 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 elegantly 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 paper 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 theorem, and apply it to all the basic properties of TRSs. Section 8 compares this technique with the original 'pile and delete' transformation introduced in [7]. Finally, Section 9 ends with some brief conclusive remarks.

2. Preliminaries

We assume knowledge of the basic notions regarding TRSs: the notation used is essentially the one in [4, 14]. Here we will just summarize some of the basic concepts that will be needed in the article.

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

We indicate with $\mathscr{T}(\Sigma,\mathscr{V})$ the set of terms built from a signature Σ and a (fixed) set of variables $\mathscr{V}$.

A *term rewriting system* (TRS) $\mathscr{R}$ consists of a signature $\Sigma_{\mathscr{R}}$ and a set of rewrite rules (sometimes called simply rules). A rewrite rule is an object of the form $l \to r$, where l and r are terms from $\mathscr{T}(\Sigma_{\mathscr{R}},\mathscr{V})$, such that l is not a variable and all the variables of r appear also in l. l and r are called, respectively, the left-hand side and the right-hand side of the rule. A rewrite rule is called *left-linear* if in the left-hand side every variable does not occur more than once (e.g. $f(g(X,g(Y,Z)) \to g(X,X))$. It is called *collapsing* if the right-hand side is a variable (e.g. $f(X) \to X$). It is called *duplicating* if there is a variable which occurs more times in the right-hand side than in the left-hand side (e.g. $f(X) \to g(X,X)$). It is called *erasing* if there is a variable in the left-hand side which is not present in the right-hand side (e.g. $g(X,Y) \to f(X)$). Also, we say a rule is *non-collapsing* (resp. *non-duplicating*, *non-erasing*) if it is not collapsing (resp. duplicating, erasing). Analogously, a term rewriting system is left-linear, non-collapsing, non-duplicating, non-erasing if each of its rewrite rules is, respectively, left-linear, non-collapsing, non-duplicating, non-erasing.

A *context* is a term built up using, besides function symbols and variables, the new special constants $\square_1, \square_2, \square_3, \ldots$ (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 $C[\square_1, \ldots, \square_n]$ and terms $t_1, \ldots, t_n$, $C[t_1, \ldots, t_n]$ stands for the term obtained from $C[\square_1, \ldots, \square_n]$ by replacing every occurrence of $\square_i$ with t_i $(1 \leqslant i \leqslant n)$.

A term rewriting system $\mathscr{R}$ determines a rewrite relation $\to_{\mathscr{R}}$ on $\mathscr{T}(\Sigma_{\mathscr{R}}, \mathscr{V})$, defined this way. Given two terms t and t', $t \to_R t'$ if $t = C[l\sigma]$ and $t' = C[r\sigma]$, for some context C, substitution σ, and rewrite rule $l \to r$ in $\mathscr{R}$. If $t_0 \to_{\mathscr{R}} t_1 \to_{\mathscr{R}} t_2 \ldots \to_{\mathscr{R}} t_n$ $(n > 0)$, then we say that t_0 *reduces to* t_n in $\mathscr{R}$; correspondingly, we call a *reduction* the sequence $t_0, t_1, \ldots, t_n$, together with the information on what rewrite rule $l_i \to r_i$ has been used to reduce t_i to t_{i+1} $(0 \leqslant i < n)$, and where it has been applied in t_i (i.e. what subterm of t_i the rule rewrites). Finally, $\twoheadrightarrow_{\mathscr{R}}$ denotes the transitive and reflexive closure of $\to_{\mathscr{R}}$. When $\mathscr{R}$ is clear from the context, we will simply write Σ, $\to$, and $\twoheadrightarrow$ in place of $\Sigma_{\mathscr{R}}$, $\to_{\mathscr{R}}$, and $\twoheadrightarrow_{\mathscr{R}}$.

Given a reduction $\rho : s \to s_1 \to s_2 \to \ldots$, 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 $s \underset{\zeta}{\twoheadrightarrow} t$, $\zeta \subseteq \rho$.

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

When two term rewriting systems $\mathscr{A}$ and $\mathscr{B}$ have disjoint signatures, we denote with $\mathscr{A} \oplus \mathscr{B}$ their *disjoint union*, that is to say the TRS having as signature the union of the signatures $\Sigma_{\mathscr{A}}$ and $\Sigma_{\mathscr{B}}$, and as rewrite rules both the rewrite rules of $\mathscr{A}$ and those of $\mathscr{B}$. A property $\mathscr{P}$ of term rewriting systems is then said to be *modular* if for every couple of TRSs $\mathscr{A}$ and $\mathscr{B}$ with disjoint signatures, $\mathscr{A} \in \mathscr{P}, \mathscr{B} \in \mathscr{P} \Leftrightarrow \mathscr{A} \oplus \mathscr{B} \in \mathscr{P}$. Throughout the paper we will indicate by $\mathscr{A}$ and $\mathscr{B}$ 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 union $\mathscr{A} \oplus \mathscr{B}$. For better readability, we will talk of function symbols belonging to $\mathscr{A}$ and $\mathscr{B}$ like *white* and *black* functions. Variables and holes, instead, have both the colours, and are thus also called *transparent* symbols. 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 $C[\Box_1,\dots,\Box_n]$ 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 $C[\Box_1,\dots,\Box_n]$.

Definition 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 $t=C[\![t_1,\dots,t_n]\!]$ $(n>0)$.

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

Lemma 2 (Toyama [20]). $s \twoheadrightarrow t \Rightarrow rank(s) \geqslant rank(t)$

Proof. Clear. $\Box$

Definition 3. The multiset $S(t)$ of the *special subterms* of a term t is

(i)

$$S(t)=\begin{cases}\{t\} & \text{if } t \text{ is black or white, and not transparent}\\ \emptyset & \text{if } t \text{ is transparent}\end{cases}$$

(ii) $S(t)=\bigcup_{i=1}^{n} S(t_i)\cup\{t\}$ if $t=C[\![t_1,\dots,t_n]\!]$ $(n>0)$.

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

Note that this definition is slightly different from the usual ones in the literature (for example in [14]), since here variables are not considered special subterms.

Given a term s, we indicate by $\|s\|$ the multiset of the ranks of the special subterms of s. Multisets of this kind are compared according to the usual multiset ordering (see e.g. [6]).

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 disjoint union $\mathscr{A}\oplus\mathscr{B}$ can be quite complicated w.r.t. the reductions in the components $\mathscr{A}$ and $\mathscr{B}$.

The disjointness requirement on $\mathscr{A}$ and $\mathscr{B}$ should ensure that symbols of one colour cannot interact with symbols of another colour. This is in a 'static' sense true, as we will see in Proposition 6.

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

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 $C[\Box_1,\dots,\Box_n]$ 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 $C[\Box_1,\dots,\Box_n]$.

Definition 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 $t=C[\![t_1,\dots,t_n]\!]$ $(n>0)$.

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

Lemma 2 (Toyama [20]). $s \twoheadrightarrow t \Rightarrow rank(s)\geqslant rank(t)$

Proof. Clear. □

Definition 3. The multiset $S(t)$ of the *special subterms* of a term t is

(i)

$$S(t)=\begin{cases}\{t\} & \text{if } t \text{ is black or white, and not transparent}\\ \emptyset & \text{if } t \text{ is transparent}\end{cases}$$

(ii) $S(t)=\bigcup_{i=1}^{n}S(t_i)\cup\{t\}$ if $t=C[\![t_1,\dots,t_n]\!]$ $(n>0)$.

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

Note that this definition is slightly different from the usual ones in the literature (for example in [14]), since here variables are not considered special subterms.

Given a term s, we indicate by $\|s\|$ the multiset of the ranks of the special subterms of s. Multisets of this kind are compared according to the usual multiset ordering (see e.g. [6]).

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 disjoint union $\mathscr{A}\oplus\mathscr{B}$ can be quite complicated w.r.t. the reductions in the components $\mathscr{A}$ and $\mathscr{B}$.

The disjointness requirement on $\mathscr{A}$ and $\mathscr{B}$ should ensure that symbols of one colour cannot interact with symbols of another colour. This is in a 'static' sense true, as we will see in Proposition 6.

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 that does not satisfy this property is also known in some literature as being *non-deterministically collapsing*, cf. [3, 16]).[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 non-trivial bubble at the top of a proper special subterm that collapses). In this case, we have the following situation:

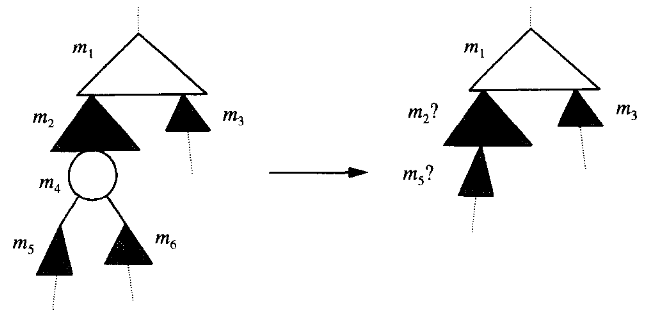

and we have a conflict between m_2 and m_5.

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.

Thus, the previous example would give (singletons like $\{m_1\}$ are written simply m_1):

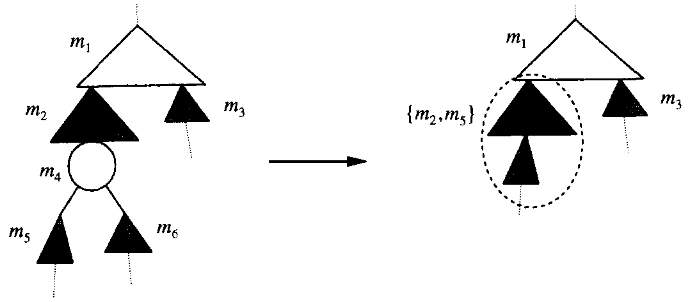

[1] Actually, the above definitions all concern the existence or not of a bubble of degree 2 (to be fussy, of a term reducing to two different variables), but, as just noticed, by Proposition 5 this is equivalent to talking about a non-deterministic bubble.

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 $t \twoheadrightarrow t'$ 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 and talk 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 disjoint union 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 every top as an independent term (cf. Proposition 6), 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 bubbles (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 non-trivial 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 (non-trivial) 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 $\mathscr{P}$ 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 $\mathscr{P}$ is modular is to

(i) show that $\mathscr{P}$ is modularly neat,

(ii) show that if $\mathscr{P}$ is modularly neat then $\mathscr{P}$ is modular.

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

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

Consider a modularity problem: to prove $\mathscr{P}$ is modular, one has to prove that for every couple of TRSs $\mathscr{A}$ and $\mathscr{B}$, $\mathscr{A} \in \mathscr{P} \ni \mathscr{B} \Leftrightarrow \mathscr{A} \oplus \mathscr{B} \in \mathscr{P}$. 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. So we get rid of it by directly considering only dense properties (this definition stems from [9], see also [11, 12]): a property $\mathscr{P}$ is said to be *dense* if whenever $\mathscr{A} \oplus \mathscr{B} \in \mathscr{P}$ then both $\mathscr{A}$ and $\mathscr{B}$ belong to $\mathscr{P}$. Therefore, what neatening has to prove is that $\mathscr{A} \in \mathscr{P} \ni \mathscr{B} \Rightarrow \mathscr{A} \oplus \mathscr{B} \in \mathscr{P}$.[2]

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

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

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

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

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 $\mathscr{P}$' (the property will be clear from the context).

We have seen that point (i) of 'bare-bones neatening' roughly corresponds to modularity under the non-collapsing assumption. In general, proving this is not a problem since this restriction is quite heavy: the problem lies in (ii).

---

[2] In fact, this latter implication is also often referred to as 'modularity of $\mathscr{P}$' (for a discussion, see e.g. [12, 11]).

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

$$\begin{gathered}\exists \text{ a provable counterexample to the modularity of } \mathscr{P}\\ \Downarrow\ (\star)\\ \exists \text{ a neat provable counterexample to the modularity of } \mathscr{P}\end{gathered}$$

If we had this, we could reason as follows:

$$\begin{gathered}\mathscr{P} \text{ is not modular}\\ \Downarrow\\ \exists \text{ a counterexample to the modularity of } \mathscr{P}\\ \Downarrow\\ \exists \text{ a provable counterexample to the modularity of } \mathscr{P}\\ \Downarrow\ (\star)\\ \exists \text{ a neat provable counterexample to the modularity of } \mathscr{P}\\ \Downarrow\\ \exists \text{ a neat counterexample to the modularity of } \mathscr{P}\\ \Downarrow\\ \mathscr{P} \text{ is not modularly neat}\end{gathered}$$

thus obtaining the contradiction to point (i).

The idea of neatening is to prove the missing implication ($\star$) 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 $\mathscr{P}$ is such that

(i) *it is modularly neat*;

(ii) *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 $\mathscr{P}$ is modular.

Observe that we have slightly stressed point (ii), since it would have sufficed to say that 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 (non-trivial) bubble, as seen, is a term that cannot properly have a fixed colour, since it can reduce to a transparent object: this way it assumes the colour of the objects it stays near. So, when a proper top-bubble is present, we have the unpleasant situation that two tops of one colour are separated by a potentially transparent object (the bubble) that has for the moment a different colour: a situation which is highly unstable.

The solution is to get rid of this bubble by attaching it to every top of its same colour which is above it (pile operation), and then change the bubble's colour (paint operation), so that the unstable situation disappears. 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 described above (that can lead to m-collapsings), and when a bubble is inside a non-bubble top of the same colour, even if it becomes transparent, the overall colour 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. *If a TRS is left-linear, then rewrite rules that have the possibility to act out on a special subterm t are exactly those that have the possibility to act on its top.*

Proof. Let $t = C[\![t_1, \ldots, t_n]\!]$: since $t_1, \ldots, t_n$ 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 of each other. □

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 7 (*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 $t = B(\!|(\!|t_1, \ldots, t_k|\!)|\!)$. 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 following two operations.

Pile: We 'pile' the bubble $B(\!|\square_1, t_2, \ldots, t_k|\!)$ 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 $r[\![r_1, \ldots, r_m]\!]$, with r_j above t, we pass from $r[\![r_1, \ldots, r_m]\!]$ to

$$r[\![r_1, \ldots, r_{j-1}, B(\!|r_j, t_2, \ldots, t_k|\!), r_{j+1}, \ldots, r_m]\!]$$

The operation is shown in Fig. 1.

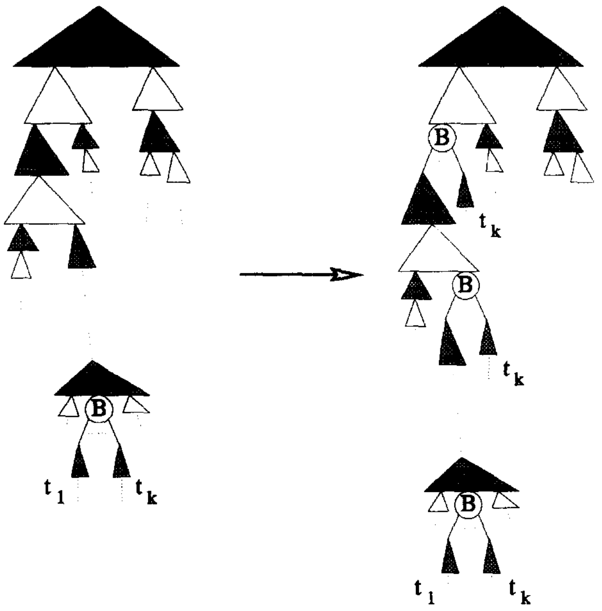

Fig. 1. The Pile operation.

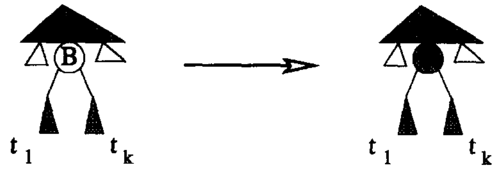

Fig. 2. The Paint operation.

Paint: We change the colour of the bubble B, replacing it with another black bubble $b(\!|\square_1,\ldots,\square_k|\!)$ of the same degree. So, t passes from $B(\!|t_1,\ldots,t_k|\!)$ to $b(\!|t_1,\ldots,t_k|\!)$. The operation is shown in Fig. 2.

Remark 8. 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 so as 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 $\square_1$.

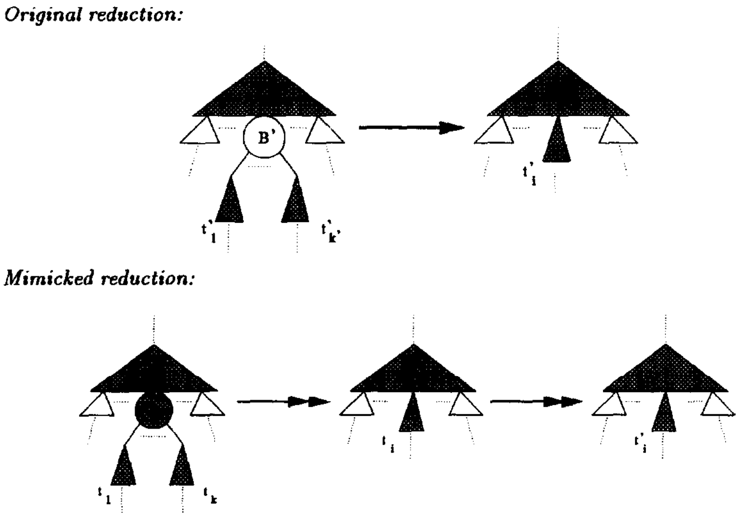

Fig. 3. Mimicking of the m-collapsing of the selected bubble.

When can π be applied? The only problematic step is the paint one, where we change the colour 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 degrees.

Equivalently, by Proposition 5, 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.

Definition 9 (*Mimicking* $\mathfrak{m}$). Given a reduction ρ of s, we define the corresponding *mimicking reduction* $\mathfrak{m}(\rho)$.

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 descendant of) t_i, then consider the corresponding term t':
    (i) Collapse t' into t_i
    (ii) Act with the corresponding reduction of ρ on (that descendant of) t_i.
    The situation is shown in Fig. 3.

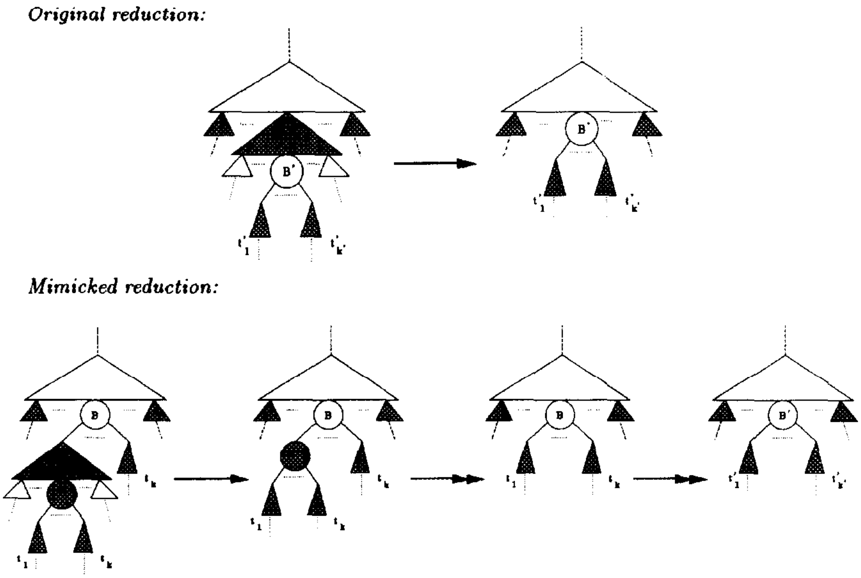

Fig. 4. Mimicking of the absorbtion of the selected bubble.

- 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':
    (i) Collapse t' (this recreates a fresh copy of t, see Fig. 4).
    (ii) Act with the corresponding reduction of ρ on (that descendant of) t on this newly created copy of t.
    The situation is illustrated in Fig. 4.
  - If a rule made a black special subterm m-collapse, collapse the bubble piled immediately above it (if it is not erased); see Fig. 5.
  - If a rule made a white special subterm m-collapse, collapse the bubble piled immediately below it (if it is not erased); see Fig. 6.

Lemma 10. *If in ρ eventually there are no pure descendants of the selected bubble B, then $\mathrm{m}(\rho)$ is cofinal for ρ.*

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

It is easy to see that $\rho' \twoheadrightarrow \mathrm{m}(\rho')$. 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 colour in place of B. When we reach a term having no pure descendants of B, it is immediate from the definition of mimicking that

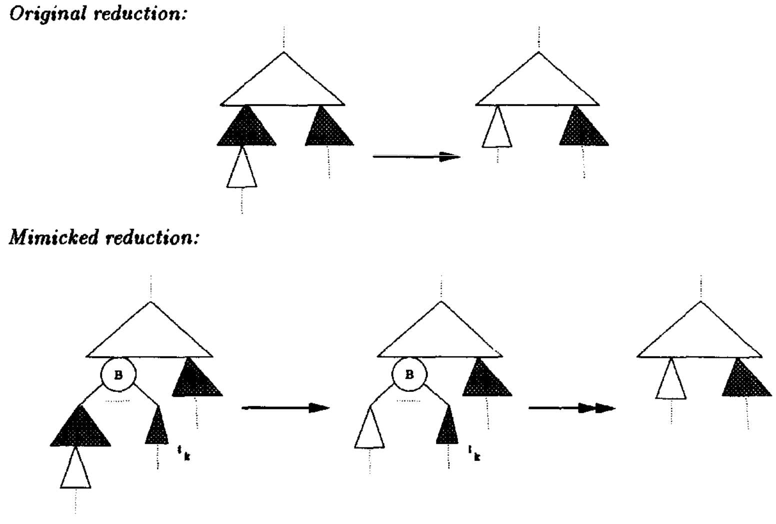

Fig. 5. Mimicking of the m-collapsing of a top black special subterm.

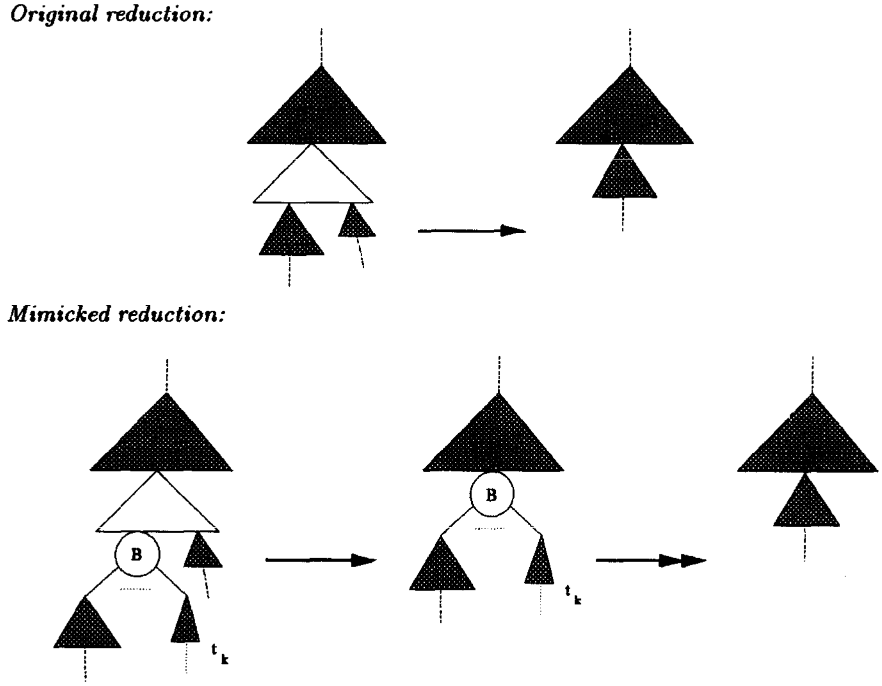

Fig. 6. Mimicking of the m-collapsing of a top white special subterm.

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 $\rho \twoheadrightarrow \mathfrak{m}(\rho)$: take $t_1 \in \rho$ ($s \underset{\zeta}{\twoheadrightarrow} t_1$, $\zeta \subseteq \rho$). We choose a $\rho' \subseteq \rho$ sufficiently big such that $\zeta \subseteq \rho'$ and reducing s by ρ' yields no pure 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 $\zeta(\subseteq \rho)$. □

6.1. Multiple pile and paint

The transformation π (together with $\mathfrak{m}$) makes the structure of a term (of a reduction) stabler 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 11. *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. □

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 12. *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 its every reduction is neat. □

We call $\mathfrak{M}(\rho)$ the mimicking reduction associated with Π, obtained from ρ by repeatedly applicating $\mathfrak{m}$ 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 12 not only it gives neat reductions, but even reductions without proper top-bubbles.

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

Lemma 13. *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 10, once the transitivity of the cofinality relation is noticed. □

7. Modularity

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

Definition 14. A property $\mathscr{P}$ is called *pseudo-deterministic* (respectively *pseudo-non-deterministic*) if $\exists T \forall T'. T' \in \mathscr{P} \Rightarrow T \oplus T' \in \mathscr{P} \wedge \mathrm{CON}^{\rightarrow}$ (respectively $\in \mathscr{P} \wedge \neg \mathrm{CON}^{\rightarrow}$).

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

Lemma 15. *A dense property $\mathscr{P}$ is pseudo-deterministic iff $\mathscr{P}$ implies $CON^{\rightarrow}$.*

Proof. The if direction is always satisfied, since taking T as the empty TRS we have $T' \in \mathscr{P} \Rightarrow T \oplus T' = T \in \mathscr{P}$. For the only if direction, observe that $\mathscr{P}$ and $\mathrm{CON}^{\rightarrow}$ being dense, then also $\mathscr{P} \wedge \mathrm{CON}^{\rightarrow}$ is such, and so from the pseudo-determinism of $\mathscr{P}$ we get $T' \in \mathscr{P} \Rightarrow T' \in \mathscr{P} \wedge \mathrm{CON}^{\rightarrow}$, which implies $\mathscr{P} \Rightarrow \mathrm{CON}^{\rightarrow}$. □

Theorem 16 (Main). *Suppose that a dense property $\mathscr{P}$ is either pseudo-deterministic or pseudo-non-deterministic, and*

(i) *$\mathscr{P}$ is modularly neat.*

(ii) *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 $\mathscr{P}$ is modular.

Proof. Suppose $\mathscr{P}$ is pseudo-deterministic. If $\mathscr{P}$ is not modular, there is a counterexample and so by point (ii) there is also a provable counterexample obtained translated via $\mathfrak{M}$ ($\mathfrak{M}$ can be applied by Fact 2 and Lemma 15): but this provable counterexample must be neat by Lemma 12, hence contradicting point (i).

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

By point (ii), 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 $\mathscr{A} \oplus T$ and $\mathscr{B} \oplus T$ are both $\neg \mathrm{CON}^{\rightarrow}$, and so using Fact 2).

But this new provable counterexample is neat by Lemma 12, hence contradicting point (i). □

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

7.1. Termination

A TRS is terminating if all its reductions are finite. Termination is, in general, not a modular property (see e.g. [19]). Via Theorem 16 we will prove the state-of-the-art results on its modularity for left-linear TRSs, and also provide two new results (actually, we will even manage to prove the best possible results, as we will see later).

Lemma 17. *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 colour. □

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 $C[\square_1,\ldots,\square_n]$, the context $C[\square_1,\ldots,\square_1]$ 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 [17] stating the modularity of termination for non-collapsing TRSs.

Corollary 18. *Termination is modular for left-linear and pseudo-deterministic TRSs and for left-linear and pseudo-non-deterministic TRSs.*

Proof. The above lemma shows point (i) of Theorem 16. For point (ii), 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. □

This result entails the main results of [8, 18]:

Corollary 19. *Termination is modular for left-linear and consistent w.r.t. reduction TRSs.*

Proof. By the above Corollary 18 and Lemma 15. □

Now we consider the other 'dual' result that Corollary 18 offers.

Call OR the TRS $\{or(X,Y)\to X, or(X,Y)\to Y\}$; a TRS T is said *termination preserving under non-deterministically collapses* (briefly $\mathscr{C}_{\mathscr{E}}$-terminating) if $T \oplus OR$ is terminating. Gramlich in [3] proved that $\mathscr{C}_{\mathscr{E}}$-termination is modular for finitely branching TRSs. Later, Ohlebusch (cf. [16]) 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 20. *$\mathscr{C}_{\mathscr{E}}$-termination is modular for left-linear TRSs.*

Proof. It follows from Corollary 18 once it is observed that $\mathscr{C}_{\mathscr{E}}$-termination is pseudo-non-deterministic and implies termination. □

Another criterion for the modularity of termination was proven by Middeldorp in [13]: he showed that whenever one of two terminating TRSs is both non-collapsing and non-duplicating, then their disjoint union is terminating. Using the two above corollaries, we can not only entail this result in the left-linear case, but also even properly generalize it with the following new result:

Corollary 21. *Suppose two left-linear TRSs are terminating. Then if one of them is both* $CON^{\rightarrow}$ *and* $\mathscr{C}_{\mathscr{E}}$*-terminating, their disjoint union is terminating.*

Proof. Every TRS is either $\mathrm{CON}^{\rightarrow}$ or $\neg\mathrm{CON}^{\rightarrow}$. So, take two terminating TRSs, with one of the two $\mathrm{CON}^{\rightarrow}$ and $\mathscr{C}_{\mathscr{E}}$-terminating: if the other is $\mathrm{CON}^{\rightarrow}$, their disjoint union is terminating by Corollary 19, otherwise if it is $\neg\mathrm{CON}^{\rightarrow}$ then it is also $\mathscr{C}_{\mathscr{E}}$-terminating, and so their disjoint union is terminating by Corollary 20. □

Corollary 22. *Suppose two left-linear TRSs are terminating. Then if one of them is both non-collapsing and non-duplicating, their disjoint union is terminating.*

Proof. By the above corollary, since non-collapsing $\Rightarrow$ $\mathrm{CON}^{\rightarrow}$, and a TRS which is both terminating and non-duplicating is $\mathscr{C}_{\mathscr{E}}$-terminating (as it is easy to show, see e.g. [3, 12, 11]). □

We now turn our attention to the structure of counterexamples to the modularity of termination. So far, two main results are known. Ohlebusch in [16] (again, extending a result of Gramlich in [3] for finitely branching TRSs), showed that in every counterexample one of the TRSs is not $\mathscr{C}_{\mathscr{E}}$-terminating and the other is collapsing. Schmidt-Schauß, Marchiori and Panitz showed in [18] that, in the left-linear case, in every counterexample one of the TRSs is $\mathrm{CON}^{\rightarrow}$ and the other is $\neg\mathrm{CON}^{\rightarrow}$. Both of these results require a non-trivial 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 [18]:

Corollary 23. *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRS is* $CON^{\rightarrow}$ *and the other is* $\neg CON^{\rightarrow}$.

Proof. Since every TRS is either $\mathrm{CON}^{\rightarrow}$ or $\neg\mathrm{CON}^{\rightarrow}$, only three cases are possible: (1) both are $\mathrm{CON}^{\rightarrow}$, (2) both are $\neg\mathrm{CON}^{\rightarrow}$, (3) like in the statement of this corollary. But (1) is not possible by Corollary 19, whereas (2) is not possible by Corollary 20 (since every terminating and $\neg\mathrm{CON}^{\rightarrow}$ TRS is trivially $\mathscr{C}_{\mathscr{E}}$-terminating). □

Next we show a somewhat dual result:

Corollary 24. *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRSs is* $\mathscr{C}_{\mathscr{E}}$*-terminating and the other is* $\neg\ \mathscr{C}_{\mathscr{E}}$*-terminating.*

Proof. Completely analogous to the proof of the above corollary. □

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

Corollary 25. *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRSs is* $\neg\,\mathrm{CON}^{\rightarrow}$ *and the other is* $\neg\ \mathscr{C}_{\mathscr{E}}$*-terminating.*

Proof. From Corollaries 23 and 24, in every counterexample only two cases are possible: (1) one of the TRSs is $\mathscr{C}_{\mathscr{E}}$-terminating$\wedge\,\mathrm{CON}^{\rightarrow}$ and the other is $\neg\ \mathscr{C}_{\mathscr{E}}$-terminating $\wedge\neg\,\mathrm{CON}^{\rightarrow}$, or (2) one of the TRSs is $\mathscr{C}_{\mathscr{E}}$-terminating$\wedge\neg\,\mathrm{CON}^{\rightarrow}$ and the other is $\neg\ \mathscr{C}_{\mathscr{E}}$-terminating$\wedge\mathrm{CON}^{\rightarrow}$. But the first case is ruled out by Corollary 21. On the other hand, by the fact that for terminating TRSs $\neg\,\mathrm{CON}^{\rightarrow} \Rightarrow \mathscr{C}_{\mathscr{E}}$-termination, it follows right away that case (2) is just the statement of this corollary, since for terminating TRSs $\mathscr{C}_{\mathscr{E}}$-termination$\wedge\neg\,\mathrm{CON}^{\rightarrow} \Leftrightarrow \neg\,\mathrm{CON}^{\rightarrow}$ and $\neg\ \mathscr{C}_{\mathscr{E}}$-termination$\wedge\mathrm{CON}^{\rightarrow} \Leftrightarrow \neg\,\mathscr{C}_{\mathscr{E}}$-termination. □

As previously claimed, this result properly generalizes (besides Corollary 24) the result of [18] (Corollary 23) and the result of (Gramlich and) Ohlebusch [3, 16] in the left-linear case:

Corollary 26. *In every counterexample to the modularity of termination for left-linear TRSs, one of the TRSs is* $\neg\ \mathscr{C}_{\mathscr{E}}$*-terminating and the other is collapsing.*

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

To comment on the results on the modularity of termination that we have obtained, we said at the beginning that we were going to prove the state-of-the-art results for left-linear modular termination; in fact, via the theory of *vaccines* (cf. [9, 12, 11]) it has been proved much more: the above main results (Corollaries 18 and 25) are *the best* we can obtain for left-linear TRSs (see the above references for precise statements of this claim). Hence, neatening allows to prove the strongest possible results for left-linear TRSs.

7.2. Uniqueness of normal forms w.r.t. reduction

A TRS is said to have the unique normal forms w.r.t. reduction, $\mathrm{UN}^{\rightarrow}$ for short, if every term has at most one normal form (recall that a term t is in normal form for a TRS if there is no other term t' such that $t \rightarrow t'$, i.e., t cannot be reduced). The $\mathrm{UN}^{\rightarrow}$ property is not modular in general (cf. [14]); 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. [2]); this problem was finally shown to have a positive solution in [7] (see also the discussion in Section 8). We now show how also this result can be obtained.

Lemma 27. *$UN^{\rightarrow}$ 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). □

Corollary 28. *$UN^{\rightarrow}$ is modular for left-linear TRSs.*

Proof. The above lemma shows Point 1 of Theorem 16. 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 13 $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$. But, again, n_1 and n_2 being 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 12). □

7.3. Consistency w.r.t. reduction

Although not modular in general [7, 10], $CON^{\rightarrow}$ is modular for left-linear TRSs, as shown for the first time in [7] (see also [10, 18]). We now prove this result.

Lemma 29. *$CON^{\rightarrow}$ is modularly neat for left-linear TRSs.*

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

Corollary 30. *$CON^{\rightarrow}$ is modular for left-linear TRSs.*

Proof. The above lemma shows point (i) of Theorem 16. For point (ii), 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 13, $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$. But X and Y being 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 12). □

The importance of this result, besides theoretical, also lies in the fact that it allows to use the result on the modularity of termination obtained in Corollary 18 for more than two TRSs, since the disjoint union of two left-linear, terminating and either both $CON^{\rightarrow}$ or both $\neg CON^{\rightarrow}$ TRSs is still left-linear, terminating and either $CON^{\rightarrow}$ or $\neg CON^{\rightarrow}$.

7.4. Confluence

As well known, a TRS is confluent if for every term t reducing to two terms t_1 and t_2, there is a term s such that both t_1 and t_2 reduce to s. Toyama in his famous paper [20] (see also [5]) proved that confluence is a modular property: we can entail this result in the left-linear case.

Lemma 31. *Confluence is modularly neat for left-linear TRSs.*

Proof. By rank induction. Suppose a term $s = C[\![t_1, \ldots, t_n]\!]$ 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), 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 $t_1, \ldots, t_n$ 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. □

Corollary 32. *Confluence is modular for left-linear TRSs.*

Proof. The above lemma shows point (i) of Theorem 16. For point (ii), 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 13, $\rho_1 \twoheadrightarrow \mathfrak{M}(\rho_1)$ and $\rho_2 \twoheadrightarrow \mathfrak{M}(\rho_2)$, and so $\Pi(s)$ reduces via the neat (cf. Lemma 12) 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. □

7.5. *Weak normalization*

A TRS is said weakly normalizing if every term has at least one normal form. Weak normalization (WN) was at the same time proven to be modular by several authors (see [14] for some references): we can entail this result in the left-linear case.

Lemma 33. *Weak normalization is modularly neat for left-linear TRSs.*

Proof. By rank induction. Considering a term s, its top reduces to a normal form (being of an unique colour); by Proposition 6 we can apply these rules to s as well (note 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. □

Corollary 34. *Weak normalization is modular for left-linear TRSs.*

Proof. We first prove that weak normalization is pseudo-non-deterministic. Consider the TRS $OR = \{or(X,Y) \to X, or(X,Y) \to Y\}$. Take a TRS T' which is WN: $T' \oplus OR \in WN \wedge \neg\mathrm{CON}^{\to}$. Indeed, $T' \oplus OR \in \neg\mathrm{CON}^{\to}$ 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 16: the above Lemma 33 shows point (i); for point (ii), 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 one obtains 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. □

7.6. *Completeness*

Completeness, as well known, is the conjunction of confluence and termination. Despite not being modular in general, it was proven to be modular for left-linear TRSs by Toyama, Klop and Barendregt in their ingenious paper [21] (see also [22]); the proof of such a result, however, is 'rather intricate and not easily digested' (citing the same authors). This result can instead be obtained as a simple corollary:

Corollary 35. *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 19 and 28. □

Note that a *direct* proof of the above result via Theorem 16 is also easy to obtain.

7.7. *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 36. *Semi-completeness is modular for left-linear TRSs.*

Proof. From Corollaries 32 and 34. □

Again, note that it is easy to obtain a *direct* proof of the above result via Theorem 16.

7.8. *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. [1,4]). It is not difficult to see that even these remaining properties can be proven to be modular for left-linear TRSs using Theorem 16. The only point worth mentioning is that all these properties are pseudo-deterministic but for local confluence, which can be proven to be pseudo-non-deterministic using the TRS $\{f(X,Y)\to X, f(X,Y)\to g(X,Y), g(X,Y)\to f(X,Y), g(X,Y)\to Y\}$.

8. Paint vs. delete

The reader may have noticed a kind of duality inside Theorem 16, since the property is required to be *either* pseudo-deterministic *or* pseudo-non-deterministic.

As we have seen, requiring pseudo-determinism essentially equals to requiring consistency w.r.t. reduction (Lemma 15). 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 colour 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 [7] for the study of the modularity of $UN^{\rightarrow}$ (and later used in [8] and with some modifications in [18]), 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-non-determinism (by Lemma 15 we can modify Theorem 16 by directly requiring that the property $\mathscr{P}$ 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-non-determinism ($\mathscr{C}_{\mathscr{E}}$-termination, weak normalization and local confluence), the criterion for the modularity of termination given by Corollary 21, and all the results on the structure of counterexamples (Corollaries 23–26).

9. 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' $\mathfrak{N}$, 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

[1] N. Dershowitz, J.-P. Jouannaud, Rewrite systems, in: J. van Leeuwen (Ed.), Handbook of Theoretical Computer Science, vol. B, Ch. 6, Elsevier, Amsterdam, 1990, pp. 243–320.

[2] N. Dershowitz, J.-P. Jouannaud, J.W. Klop, Open problems in rewriting, in: R.V. Book (Ed.), Proc. 4th Internat. Conf. on Rewriting Techniques and Applications, Lecture Notes in Computer Science, vol. 488, Springer, Berlin, 1991, pp. 445–456.

[3] B. Gramlich, Generalized sufficient conditions for modular termination of rewriting, Appl. Algebra Eng. Comm. Comput. 5 (1994) 131–158.

[4] J.W. Klop, Term rewriting systems, in: S. Abramsky, D.M. Gabbay, T.S.E. Maibaum (Eds.), Handbook of Logic in Computer Science, vol. 2, Ch. 1, Clarendon Press, Oxford, 1992, pp. 1–116.
[5] J.W. Klop, A. Middeldorp, Y. Toyama, R. de Vrijer, Modularity of confluence: a simplified proof, Inform. Process. Lett. 49 (2) (1994) 101–109.
[6] Z. Manna, R. Waldinger, The Logical Basis for Computer Programming, vol. 1: Deductive Reasoning, Addison-Welsey, Reading, MA, 1985.
[7] M. Marchiori, Modularity of $UN^{\rightarrow}$ for left-linear term rewriting systems, Draft, March 1993. Extended and revised version, Tech. Report CS-R9433, CWI, Amsterdam.
[8] M. Marchiori, Modularity of completeness revisited, in: J. Hsiang (Ed.), Proc. 6th Internat. Conf. on Rewriting Techniques and Applications, Lecture Notes in Computer Science, vol. 914, Springer, Berlin, 1995, pp. 2–10.
[9] M. Marchiori, The theory of vaccines, Tech. Report 27, Department of Pure and Applied Mathematics, University of Padova, 1995.
[10] M. Marchiori, On the modularity of normal forms in rewriting, J. Symbolic Comput. 22 (2) (1996) 143–154.
[11] M. Marchiori, Local analysis and localizations, Ph.D. Thesis, Department of Pure and Applied Mathematics, University of Padova, 1997 (in Italian).
[12] M. Marchiori, The theory of vaccines, in: Proc. 24th Internat. Coll. on Automata, Languages, and Programming (ICALP), Lecture Notes in Computer Science, vol. 1256, Springer, Berlin, 1997, pp. 660–670.
[13] A. Middeldorp, A sufficient condition for the termination of the direct sum of term rewriting systems, in: Proc. 4th IEEE Symp. on Logic in Computer Science, 1989, pp. 396–401.
[14] A. Middeldorp, Modular properties of term rewriting systems, Ph.D. Thesis, Vrije Universiteit, Amsterdam, 1990.
[15] E. Ohlebusch, Modular properties of composable term rewriting systems, Ph.D. Thesis, Universität Bielefeld, 33501 Bielefeld, FRG, 1994; Appeared as Tech. Report 94-01.
[16] E. Ohlebusch, On the modularity of termination of term rewriting systems, Theoret. Comput. Sci. 136 (2) (1994) 333–360.
[17] M. Rusinowitch, On termination of the direct sum of term rewriting systems, Inform. Process. Lett. 26 (1987) 65–70.
[18] M. Schmidt-Schauß, M. Marchiori, S.E. Panitz, Modular termination of r-consistent and left-linear term rewriting systems, Theoret. Comput. Sci. 149 (2) (1995) 361–374.
[19] Y. Toyama, Counterexamples to termination for the direct sum of term rewriting systems, Inform. Process. Lett. 25 (1987) 141–143.
[20] Y. Toyama, On the Church-Rosser property for the direct sum of term rewriting systems, J. ACM 1 (34) (1987) 128–143.
[21] Y. Toyama, J.W. Klop, H.P. Barendregt, Termination for the direct sum of left-linear term rewriting systems, in: N. Dershowitz (Ed.), Proc. 3rd Internat. Conf. on Rewriting Techniques and Applications, Lecture Notes in Computer Science, vol. 355, Chapel-Hill, 1989, pp. 477–491; Springer, Berlin, 1989. Extended version: Report CS-R8923, CWI, Amsterdam.
[22] Y. Toyama, J.W. Klop, H.P. Barendregt, Termination for direct sums of left-linear complete term rewriting systems, J. ACM 42 (6) (1995) 1275–1304.