

From Conditional to Unconditional Rewriting

Grigore Roşu

Department of Computer Science,
University of Illinois at Urbana-Champaign.
`grosu@cs.uiuc.edu`

Abstract. An automated technique to translate conditional rewrite rules into unconditional ones is presented, which is suitable to implement, or compile, conditional rewriting on top of much simpler and easier to optimize unconditional rewrite systems. An experiment performed on world’s fastest conditional rewriting engines shows that speedups for conditional rewriting of an order of magnitude can already be obtained by applying the presented technique as a front-end transformation.

1 Introduction

Conditional term rewriting is a crucial paradigm in the algebraic specification of abstract data types, since it provides a natural means for executing equational specifications. Many specification languages today, including Maude [4], ELAN [3], OBJ [9], CafeOBJ [6], provide conditional rewriting engines to allow users to execute and reason about specifications. Conditional rewriting also plays a foundational role in functional logic programming [10]. Additionally, there are many researchers, including the author, considering rewriting a powerful programming paradigm by itself, who are often frustrated that conditional rewrite “programs” are significantly slower than unconditional ones doing the same thing.

Conditional rewriting is, however, rather inconvenient to implement directly. To reduce a term, a rewriting engine needs to maintain a *control context* for each conditional rule that is tried. Due to the potential nesting of rewrite rule applications, such a control context may grow arbitrarily. Our technique automatically translates conditional rewrite rules into unconditional rules, by *encoding the necessary control context into data context*. The obtained rules can be then executed on (almost) any unconditional rewriting engine, whose single task is to *match-and-apply* unconditional rules. Such a simplified engine can be seen as a *rewrite virtual machine*, which can be even implemented in hardware for increased efficiency, and our transformation technique can be seen as a compiler.

Experiments performed on two fast rewriting engines show that speedups of an order of magnitude can be obtained right now if one uses our transformation technique as a front-end. However, since these rewrite engines are optimized for conditional rewriting, we expect significant further increases in efficiency if one just focus on the much simpler problem of developing optimized unconditional rewrite engines and use our technique as a front-end. Even though presented as a translation of conditional rewrite systems into unconditional ones, our technique

can easily be adapted and used as a means to implement conditional rewriting also without applying an explicit transformation. We will discuss this elsewhere.

Related work. Stimulated by the benefits of transforming conditional term rewrite systems (CTRSs) into equivalent unconditional term rewrite systems (TRSs), there has been much research on this topic. Despite the apparent simplicity of most transformations, they typically work for restricted CTRSs and their correctness, when they are correct, is quite technical and tricky to prove. A large body of literature has been dedicated to transformations preserving only certain properties of CTRSs, e.g., termination and/or confluence. We do not discuss these here; the interested reader is referred, e.g., to Ohlebusch [14].

In this paper we focus on transformations that generate TRSs *computationally equivalent* to CTRSs, i.e., the TRSs can be *transparently* used to reduce terms in the original CTRSs. The first attempt in this category is due to Bergstra and Klop [2], for a restricted class of CTRSs (whose underlying unconditional TRS is left-linear and without superposition); unfortunately, this transformation was shown to be unsound by Dershowitz and Okada [5]. The transformation in Giovannetti and Moiso [8] works only under severe restrictions on the original CTRS: no superposition, simply terminating (enforced by the requirement of a simplification ordering), and non-overlapping of conditions with left-hand-side (lhs) terms. Hintermeier [11] proposes a technique where an “interpreter” for CTRS is defined as a TRS, providing explicit rewrite definitions for matching and applications of rewrite rules. Besides being technically very intricate and practically inefficient, this transformation is proven to be correct only when the original CTRS is confluent and strictly terminating (i.e., decreasing). Our work in this paper was motivated by efforts in *rewriting logic semantics* [12], where rewriting logic is used as a core mechanism to give operational semantics to concurrent programming languages. In this framework, as well as in many others, restrictions such as termination and/or confluence are unacceptably strong. Indeed, in any programming language there are programs which do not terminate, and concurrency leads quickly to non-confluence (e.g., data-races).

Our technique was presented at WADT’04 and was developed independently from that of Viry [16]. However, the two techniques have many similarities¹. They are both based on decorations of terms, obtained by adding as many auxiliary arguments to each operation f as conditional rules in the original CTRS having f at the top of their lhs. The procedure in [16] encodes the condition of each rule within a special data-structure that occurs as the corresponding auxiliary argument associated to the operation occurring at the top of its lhs. Two unconditional rules are added in the generated TRS for each conditional rule in the original CTRS, one for initializing the special data-structure and the other for continuing the rewriting process when the condition was evaluated. For example, the CTRS (taken from [16]) \mathcal{R} below is transformed into \mathcal{R}' :

¹ We thank Bernhard Gramlich for making us aware of Viry [16].

$$\begin{array}{l}
(\mathcal{R}) \left\{ \begin{array}{l} f(g(x)) \rightarrow p(x) \text{ if } c(x) \rightarrow^* \text{true} \\ f(h(x)) \rightarrow q(x) \text{ if } d(x) \rightarrow^* \text{true} \\ c(a) \rightarrow \text{true} \end{array} \right. \quad (\mathcal{R}') \left\{ \begin{array}{l} f(g(x) \mid \perp, z) \rightarrow f(g(x) \mid [c(x), (x)], z) \\ f(x \mid [\text{true}, (y)], z) \rightarrow p(y) \\ f(h(x) \mid z, \perp) \rightarrow f(h(x) \mid z, [d(x), (x)]) \\ f(x \mid z, [\text{true}, (y)]) \rightarrow q(y) \\ c(a) \rightarrow \text{true} \end{array} \right.
\end{array}$$

where “ \mid ” is syntactic sugar for “,”, separating the normal arguments from the auxiliary ones; “ \perp ” is a special constant whose occurrence states that the corresponding conditional rule has not been tried yet on the current position; a structure $[u, \vec{s}]$ occurring during a rewriting sequence as an auxiliary argument of an operation, means that u is the current reduction status of the corresponding condition that started to be evaluated at some point, and that \vec{s} was the substitution at that point that allowed the lhs of that rule to match. The substitution is needed by the second unconditional rule associated to a conditional rule, to correctly initiate the reduction of the rhs of the original conditional rule.

Despite being proved sound and complete by Viry [16], the procedure above, unfortunately, cannot be used *as is* to interpret any CTRS on top of a TRS. That is because it destroys the confluence of the original CTRS, thus leading to normal forms in the TRS which can be further reduced in the CTRS. Indeed, let us consider the following CTRS \mathcal{R} , from Antoy, Brassel and Hanus [1], together with Viry’s transformation \mathcal{R}' :

$$(\mathcal{R}) \left\{ \begin{array}{l} f(g(x)) \rightarrow x \text{ if } x \rightarrow^* 0 \\ g(g(x)) \rightarrow g(x) \end{array} \right. \quad (\mathcal{R}') \left\{ \begin{array}{l} f(g(x) \mid \perp) \rightarrow f(g(x) \mid [x, (x)]) \\ f(x \mid [0, (y)]) \rightarrow y \\ g(g(x)) \rightarrow g(x) \end{array} \right.$$

\mathcal{R} is confluent but \mathcal{R}' is not: $f(g(g(0)) \mid \perp)$ can be reduced to both 0 and $f(g(0) \mid [g(0), (g(0))])$; the latter occurs because the “conditional” rule is first tried and “failed”, then the “unconditional” one is applied successfully thus changing the context so that the “conditional” rule becomes conceptually applicable, but it fails to apply since it was already marked as “tried”. To solve this problem, Viry [16] proposes a reduction strategy within the generated TRS, called *conditional eagerness*, stating that t_1, \dots, t_n must be already in normal form before a “conditional” rule can be applied on a term $f(t_1, \dots, t_n \mid \perp, \dots, \perp)$. This way, in the example above, $g(g(0))$ is enforced to be first evaluated to $g(0)$ and only then $f(g(0) \mid \perp)$ is applied the “conditional” rule and eventually reduced to 0. However, conditional eagerness does not seem to be trivial to enforce in an unconditional rewriting engine, unless that is internally modified. One simple, but very restrictive, way to ensure conditional eagerness is to enforce innermost rewriting both in the original CTRS and in the resulting TRS.

A different fix to Viry’s technique was proposed by Antoy, Brassel and Hanus [1], namely to restrict the input CTRSs to *constructor-based* ones, i.e., ones in which the operations are split into *constructors* and *defined*, and the lhs of each rule is a term of the form $f(t_1, \dots, t_n)$, where f is defined and t_1, \dots, t_n are all

constructor terms. The problematic CTRS above is *not* constructor-based, so Viry's procedure is not guaranteed to work correctly on it. While constructor-baseness is an easy to check and automatic correctness criterion, we believe that it is an unnecessary strong restriction on the input CTRS, which may make the translation useless in many situations of practical interest.

An additional drawback of Viry's transformation is that it increases the number of rewrite rules having the same operator at the root of their lhs, which tends to be a source of matching overhead on many rewrite engines, especially in the context of very large CTRSs². Therefore, we are still left with no satisfactory translation of CTRSs into equivalent TRSs. In this paper we give a practical solution to this problem, which imposes *no restrictions* on the original CTRS, which adds exactly one unconditional rule for each conditional rule in the original CTRS, and which is shown to bring a significant speedup on current conditional rewrite engines if applied as a front-end transformation. Our translation is *almost* ideal, in that it still requires some special support from the underlying unconditional rewrite engine: to provide (1) a binary equality operation, denoted $equal?(t, t')$ in this paper, returning *true* iff the normal forms of t and t' are identical, and (2) a conditional $if(b, t, t')$ which is eager in b and lazy in t and t' . However, all rewriting engines that we know provide them [9, 3, 4, 6, 17]. They can also be easily defined if the rewriting engine provides support for simple contextual strategies, which all rewriting engines that we know do.

2 Preliminaries

We recall some basic notions of conditional rewriting, referring the interested reader to [14] for more details. An (unsorted) *signature* Σ is a finite set of operational symbols, each having zero or more arguments. We let $\Sigma_n \subseteq \Sigma$ denote the set of operations of n arguments. The operations of zero arguments in Σ_0 are called *constants*. We assume an infinite set of *variables* \mathcal{X} . Given a signature Σ and a set of variables $X \subseteq \mathcal{X}$, we let $T_\Sigma(X)$ denote the algebra of Σ -terms over variables in X . A term without variables is called *ground*. A map $\theta : \mathcal{X} \rightarrow T_\Sigma(\mathcal{X})$ can be uniquely extended to a morphism of algebras $T_\Sigma(\mathcal{X}) \rightarrow T_\Sigma(\mathcal{X})$ replacing each variable in x by a term $\theta(x)$; to keep the notation simple, we let θ also denote this map. A conditional Σ -rewrite rule has the form

$$l \rightarrow r \text{ if } u_1 = v_1, \dots, u_m = v_m,$$

where $l, r, u_1, v_1, \dots, u_m, v_m$ are Σ -terms in $T_\Sigma(\mathcal{X})$. The term l is called the *left-hand-side (lhs)*, r is called the *right-hand-side (rhs)*, and $u_1 = v_1, \dots, u_m = v_m$ is called the *condition* of the rewriting rule above. As usual, we disallow rewriting rules whose lhs is a variable. Further, we assume that the lhs of a rewriting rule contains all the variables that occur in that rule, that is, following the terminology in [13] our rewrite systems are *of type 1*. If $m = 0$, the rewrite

² We have encountered CTRSs of thousands of rules in the context of rewriting logic semantics of programming languages [12]

rule is called *unconditional* and written $l \rightarrow r$. Unless specified differently, by conditional rule we mean a rule with $m \geq 1$. A *conditional (unconditional) Σ -term rewrite system* $\mathcal{R} = (\Sigma, R)$, abbreviated *CTRS (TRS)*, consists of a finite set R of conditional (unconditional) Σ -rewrite rules. Any Σ -rewrite system $\mathcal{R} = (\Sigma, R)$ generates a relation $\rightarrow_{\mathcal{R}}$ on $T_{\Sigma}(\mathcal{X})$, defined recursively as follows. For any $\theta : \mathcal{X} \rightarrow T_{\Sigma}(\mathcal{X})$, $t[\theta(l)] \rightarrow_{\mathcal{R}} t[\theta(r)]$ whenever there exists some s_i such that $\theta(u_i) \rightarrow_{\mathcal{R}}^* s_i$ and $\theta(v_i) \rightarrow_{\mathcal{R}}^* s_i$ for any $1 \leq i \leq m$, where t is a term having one occurrence of a special variable, say $*$, $t[\theta(l)]$ is the term obtained by substituting $*$ with $\theta(l)$ in t , and $\rightarrow_{\mathcal{R}}^*$ is the reflexive and transitive closure of $\rightarrow_{\mathcal{R}}$. Hence, $\alpha \rightarrow_{\mathcal{R}} \beta$ iff α has a subterm matching the lhs of a rule in R via some substitution, s.t. all the terms in each equality in the condition can be iteratively reduced to a common term. Such CTRSs are also called *join* or *standard* [14]. Alternative interpretations of equalities are also possible, and we will discuss transformations of those elsewhere soon. However, as their name suggests, standard conditional rewrite systems are the most common ones and major rewriting engines, e.g., Maude [4] and ELAN [3], support them. These systems perform millions of rewrites per second on standard PCs and are, at our knowledge, the fastest rewriting engines.

Terms which cannot be reduced any further in \mathcal{R} are called *normal forms* for \mathcal{R} . Rewriting of a given term may not terminate for two reasons: either the reduction of the condition of a rule does not terminate, or there are some rules that can be applied infinitely often on the given term. On systems like Maude or ELAN, the effect in both situations is the same: the system loops forever unless it crashes running out of memory. Because of this reason, we do not make any distinction between the two causes, and simply call a Σ -rewriting system *terminating* iff it always reduces any Σ -term to a normal form (we let this notion at an intuitive level here, but it can be formalized). Letting $;$, $-$ denote the composition of relations, a relation \rightarrow is *confluent* iff $\leftarrow^* ; \rightarrow^* \subseteq \rightarrow^* ; \leftarrow^*$.

3 Defining the Basic Infrastructure

We define several operators together with appropriate (unconditional) rules. Most rewriting engines have these basic operators built-in, but here we do not assume any existing operators and therefore define everything needed.

Let *true* and *false* be two constants which are assumed not defined within any given CTRS (otherwise change their name). Let us also assume a fresh binary operator \wedge , written in infix associative notation, together with the rules:

$$\begin{array}{ll} \text{true} \wedge \text{true} \rightarrow \text{true}, & \text{true} \wedge \text{false} \rightarrow \text{false}, \\ \text{false} \wedge \text{true} \rightarrow \text{false}, & \text{false} \wedge \text{false} \rightarrow \text{false}. \end{array}$$

These will be needed to evaluate conditions that will be translated into corresponding conjunctions of equalities; equalities will be defined shortly.

Let us now consider a special operator *if*($-, -, -$), together with the rules:

$$\text{if}(\text{true}, x, y) \rightarrow x, \quad \text{if}(\text{false}, x, y) \rightarrow y.$$

This operator is assumed *eager in its first argument and lazy in the others*. Most rewrite engines provide it as builtin, so the two rules above are not needed.

We need another special operator, $equal?(_,_)$, that reduces its arguments and returns *true* if they are identical and *false* otherwise. One obvious rule to add is $equal?(x, x) \rightarrow true$. Moreover, for all $\sigma \in \Sigma_n$ we add

$$equal?(\sigma(x_1, \dots, x_n), \sigma(y_1, \dots, y_n)) \rightarrow equal?(x_1, y_1) \wedge \dots \wedge equal?(x_n, y_n), \quad (1)$$

where $x_1, \dots, x_n, y_1, \dots, y_n$ are disjoint variables. These rules propagate the equality of two terms having the same operator as root to the equality of their corresponding sub-terms. Note that σ may be a constant in Σ_0 , in which case, by convention, $equal?(x_1, y_1) \wedge \dots \wedge equal?(x_n, y_n)$ is *true*, the unit of \wedge . The following rules, one for each pair $\sigma \in \Sigma_n, \tau \in \Sigma_m$ of different operations in Σ , state that terms having different operations at root are not equal:

$$equal?(\sigma(x_1, \dots, x_n), \tau(y_1, \dots, y_m)) \rightarrow false.$$

Note that $equal?$ needs to be *eager in both its arguments*. All rewrite engines we know have such an operator builtin, so these rules are not needed in practice.

For a given signature Σ , let Σ' denote the signature Σ extended with all the auxiliary operations above, and let $\mathcal{I}(\Sigma)$ be the Σ' -rewriting system containing all the rules above. We call $\mathcal{I}(\Sigma)$ the *infrastructure rewriting system of Σ* .

Proposition 1. *Let \mathcal{R} be a Σ -rewrite system, conditional or not. Then*

1. $\mathcal{I}(\Sigma)$ is a confluent and terminating unconditional Σ' -rewrite system;
2. If $u, v \in T_\Sigma(X)$ then $u \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* v$ iff $equal?(u, v) \rightarrow_{\mathcal{R} \cup \mathcal{I}(\Sigma)}^* true$;
3. If u, v are ground Σ -terms then a normal form of $equal?(u, v)$ in $\mathcal{R} \cup \mathcal{I}(\Sigma)$ is either *true* or *false*;
4. \mathcal{R} terminates if and only if $\mathcal{R} \cup \mathcal{I}(\Sigma)$ terminates;
5. If \mathcal{R} is confluent and terminates, i.e., it has unique normal forms, then $\mathcal{R} \cup \mathcal{I}(\Sigma)$ is also confluent and terminates.

Proof. 1. It can be proved independently, but we prefer to prove the more general result 5., and then to derive this one by taking \mathcal{R} to be the empty Σ -rewrite system in 5.

2. If $u \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* v$ then it is relatively clear that $equal?(u, v) \rightarrow_{\mathcal{R} \cup \mathcal{I}(\Sigma)}^* true$, because any two rewrite sequences proving the first translates into a possible rewriting sequence of the second. Indeed, let us assume the following rewriting sequences in \mathcal{R} , $u = u_0 \rightarrow u_1 \rightarrow \dots \rightarrow u_n$ and $v = v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_m$, where $u_n = v_m$ for some $n, m \geq 0$. Then one can generate the following rewriting sequence in $\mathcal{R} \cup \mathcal{I}(\Sigma)$: $equal?(u, v) = equal?(u_0, v_0) \rightarrow equal?(u_1, v_0) \rightarrow \dots \rightarrow equal?(u_n, v_0) \rightarrow equal?(u_n, v_1) \rightarrow \dots \rightarrow equal?(u_n, v_m)$. Further, since $u_n = v_m$, one can now apply the rule $equal?(x, x) \rightarrow true$.

Before we prove the other implication, we prove the following lemma:

Lemma 1. *If u, v, u', v' are all terms in $T_\Sigma(X)$ and $\text{equal?}(u, v) \rightarrow_{\mathcal{R} \cup \mathcal{I}(\Sigma)}^k \text{equal?}(u', v')$, i.e., there are k rewriting steps in $\mathcal{R} \cup \mathcal{I}(\Sigma)$, such that the rule (1) of equal? is not used in this rewriting sequence, then $u \rightarrow_{\mathcal{R}}^n u'$ and $v \rightarrow_{\mathcal{R}}^m v'$ with $k = n + m$.*

Proof. Since *true* and *false* are in normal form, which means that the rewriting sequence $\text{equal?}(u, v) \rightarrow_{\mathcal{R} \cup \mathcal{I}(\Sigma)}^k \text{equal?}(u', v')$ cannot use any of the rules of equal? , there is only one way to apply a rewriting rule in $\mathcal{R} \cup \mathcal{I}(X)$ to the term $\text{equal?}(u, v)$: to apply a rewriting rule in \mathcal{R} to either u or v . The term to rewrite will have the same form, an operator equal? with two Σ -terms as arguments. Therefore, k will be equal to the number of rules in \mathcal{R} applied to any of the arguments of equal? .

Conversely, we can prove by well-founded induction on the length of the rewriting sequence in $\mathcal{R} \cup \mathcal{I}(\Sigma)$ that if $\text{equal?}(u, v) \rightarrow_{\mathcal{R} \cup \mathcal{I}(\Sigma)}^* \text{true}$ then there is some Σ -term w such that $u \rightarrow_{\mathcal{R}}^* w$ and $v \rightarrow_{\mathcal{R}}^* w$. There are two situations to analyze. If the rewriting rule (1) is not used in the rewriting sequence $\text{equal?}(u, v) \rightarrow^* \text{true}$ in $\mathcal{R} \cup \mathcal{I}(\Sigma)$, then the only way for the rewriting sequence to exist is to reduce the *true* using the rule $\text{equal?}(x, x) \rightarrow \text{true}$, so $\text{equal?}(u, v) \rightarrow^* \text{equal?}(w, w) \rightarrow \text{true}$ for some Σ -term w . By Lemma 1, it follows that $u \rightarrow^* w$ and $v \rightarrow^* w$ in \mathcal{R} , so the claim holds in this case.

If the rewriting rule (1) generating conjuncts of equalities is used one or more times in the rewriting sequence $\text{equal?}(u, v) \rightarrow^* \text{true}$, then let us consider its first use, that is, $\text{equal?}(u, v) \rightarrow^* \text{equal?}(\sigma(u_1, \dots, u_n), \sigma(v_1, \dots, v_n)) \rightarrow \text{equal?}(u_1, v_1) \wedge \dots \wedge \text{equal?}(u_n, v_n) \rightarrow^* \text{true}$ for some $\sigma \in \Sigma_n$ and some Σ -terms $u_1, v_1, \dots, u_n, v_n$. Since we considered the first use of rule (1), it follows by Lemma 1 that in fact $u \rightarrow^* \sigma(u_1, \dots, u_n)$ and $v \rightarrow^* \sigma(v_1, \dots, v_n)$ in \mathcal{R} . The only way for the conjunction of equalities to reduce to *true* is that each of the equalities reduces to *true*, that is, $\text{equal?}(u_i, v_i) \rightarrow^* \text{true}$ for each $1 \leq i \leq n$. By the induction hypothesis, there are some Σ -terms w_i such that $u_i \rightarrow^* w_i$ and $v_i \rightarrow^* w_i$ for each $1 \leq i \leq n$. Then one can combine all these rewriting sequences and construct the rewriting sequences $\sigma(u_1, \dots, u_n) \rightarrow^* \sigma(w_1, \dots, w_n)$ and $\sigma(v_1, \dots, v_n) \rightarrow^* \sigma(w_1, \dots, w_n)$. Letting now w denote the Σ -term $\sigma(w_1, \dots, w_n)$, we can construct the rewriting sequences $u \rightarrow^* w$ and $v \rightarrow^* w$ in \mathcal{R} , which completes the proof.

3. Suppose that u, v are ground Σ -terms and that w is a normal form of $\text{equal?}(u, v)$ in $\mathcal{R} \cup \mathcal{I}(\Sigma)$ different from *true* or *false*. We claim that in this case w must contain some subterm $\text{equal?}(u', v')$ for some ground Σ -terms u' and v' . Indeed, note that the operation equal? must occur in w because the rule (1) propagates it. Moreover, since rule (1) propagates as arguments of equal? subterms of the original argument terms of equal? , it follows that at any moment during the rewriting sequence of $\text{equal?}(u, v)$ any terms u'' and v'' that occur in any subterm $\text{equal?}(u'', v'')$ are Σ -terms; further, since rules do not contain new variables in their rhs, u'' and v'' are ground.

Therefore, the normal form w contains some subterm $\text{equal?}(u', v')$ for some ground Σ -terms u' and v' . However, this is a contradiction because equal? is completely defined on any terms having operations in Σ as roots, i.e., ground Σ -terms, so $\text{equal?}(u', v')$ cannot be in normal form.

4. If $\mathcal{R} \cup \mathcal{I}(\Sigma)$ terminates then \mathcal{R} also terminates, because any infinite rewrite sequence in \mathcal{R} translates into an infinite rewrite sequence in $\mathcal{R} \cup \mathcal{I}(\Sigma)$.

For the other implication, we first prove the following lemma:

Lemma 2. *Let $\mathcal{R} = (\Sigma, E)$ be any terminating finite term rewriting system, let σ be a new operation of n arguments, and let R_σ be a finite set of rewriting rules of the form*

$$\sigma(\alpha_1, \dots, \alpha_n) \rightarrow \gamma,$$

where $\alpha_1, \dots, \alpha_n$ are Σ -terms and γ is a $(\Sigma \cup \{\sigma\})$ -term whose subterms rooted in σ have the form $\sigma(\beta_1, \dots, \beta_n)$ with β_1, \dots, β_n subterms of $\alpha_1, \dots, \alpha_n$, respectively, at least one of them being a proper subterm. Then $\mathcal{R}' = (\Sigma \cup \{\sigma\}, R \cup R_\sigma)$ also terminates.

Proof. We show by well-founded induction on the number of operations σ that a $(\Sigma \cup \{\sigma\})$ -term t contains that any rewriting sequence starting with t must be finite. If t contains no σ then the result is immediate, because \mathcal{R} terminates. Let us assume that t contains at least one operation σ and that there is some infinite rewrite sequence in \mathcal{R}' starting with t . Since \mathcal{R} terminates, the only way to have an infinite rewriting of t is that t has some *maximal subterm rooted in σ* , i.e., one which is not a subterm of any other subterm of t rooted in σ , which generates an infinite rewriting sequence in \mathcal{R}' . Let $\sigma(t_1, \dots, t_n)$ be such a subterm of t which can be rewritten infinitely.

By the induction hypothesis, since t_1, \dots, t_n contain fewer operations σ than t , neither of t_1, \dots, t_n can be rewritten infinitely. Note that if a $(\Sigma \cup \{\sigma\})$ -term δ cannot be rewritten infinitely in \mathcal{R}' then there is only a *finite* number of terms that δ can be rewritten to: indeed, supposing that there are infinitely many terms that δ can be rewritten to then, since δ can be rewritten in one step to only a finite number of terms, there must be such a successor term δ' that can be rewritten to infinitely many terms. Iterating this argument on δ' , one can obtain an infinite rewriting sequence $\delta \rightarrow \delta' \rightarrow \dots$, violating the property that δ cannot be rewritten infinitely. For such a term δ , let $[\delta]$ be the finite set of all the subterms of all the terms that δ can be rewritten to in \mathcal{R}' .

Due to the form of the rewrite rules in R_σ , at any moment during the infinite rewriting sequence of $\sigma(t_1, \dots, t_n)$, the current term contains only *maximal* subterms rooted in σ of the form $\sigma(u_1, \dots, u_n)$, where (u_1, \dots, u_n) is a tuple in the finite set $[t_1] \times \dots \times [t_n]$. Since \mathcal{R} terminates and since t_1, \dots, t_n cannot be reduced infinitely, it follows that the rules in R_σ are applied infinitely often on *maximal* subterms rooted in σ of the current term during the infinite rewriting of $\sigma(t_1, \dots, t_n)$. Since there is only a finite number of tuples in $[t_1] \times \dots \times [t_n]$, there must be some tuple $(u_1, \dots, u_n) \in [t_1] \times \dots \times [t_n]$ such that $\sigma(u_1, \dots, u_n)$ rewrites in one or more steps to some term containing $\sigma(u_1, \dots, u_n)$ as a subterm, that is, $(u_1, \dots, u_n) \in [u_1] \times \dots \times [u_n]$. Since the rules in R_σ require that among

the arguments of any σ in the rhs of any rule there is least one which is a proper subterm of the corresponding subterm of the lhs of that rule, it follows that there is some $i \in \{1, \dots, n\}$ and some proper subterm u'_i of u_i such that $u_i \in [u'_i]$. This implies that there is some rewriting sequence $u'_i \rightarrow^* c[u_i]$ in \mathcal{R}' for some non-empty context c . But this leads to a contradiction, because $u_i \in [t_i]$ and t_i cannot be rewritten infinitely by the induction hypothesis.

We can now prove the termination of $\mathcal{R} \cup \mathcal{I}(\Sigma)$ by applying the lemma above incrementally for the infrastructure operators in this order: *true*, *false*, *if*(\neg , \neg), $\neg \wedge \neg$, *equal*?. Only the last operator is interesting, the others using degenerated variants of the lemma above, when the new operator does not occur in rhs of any newly added rule.

5. Let t be a Σ' -term and let $t \rightarrow t_1$ and $t \rightarrow t_2$ be one step rewrites in $\mathcal{R} \cup \mathcal{I}(\Sigma)$. We need to show that there is some t' such that $t_1 \rightarrow^* t'$ and $t_2 \rightarrow^* t'$ in $\mathcal{R} \cup \mathcal{I}(\Sigma)$. Therefore, there are some contexts c_1, c_2 and some rules $l_1 \rightarrow r_1$ if C_1 , $l_2 \rightarrow r_2$ if C_2 in $\mathcal{R} \cup \mathcal{I}(\Sigma)$ (the rules in $\mathcal{I}(\Sigma)$ are unconditional), such that $t = c_1[\theta_1(l_1)] = c_2[\theta_2(l_2)]$, $t_1 = c_1[\theta_1(r_1)]$, and $t_2 = c_2[\theta_2(r_2)]$ for some maps θ_1, θ_2 with $\theta_1(C_1)$ and $\theta_2(C_2)$ properly reducible in $\mathcal{R} \cup \mathcal{I}(\Sigma)$. If the two rewrites are applied at *disjoint* positions, i.e., if there is some two-hole context $c[\star_1, \star_2]$ such that $c_1 = c[\star_1, \theta_2(l_2)]$ and $c_2 = c[\theta_1(l_1), \star_2]$, then one can simply choose t' to be the term $c[\theta_1(r_1), \theta_2(r_2)]$ and note that $t_1 \rightarrow t'$ and $t_2 \rightarrow t'$ in one step in $\mathcal{R} \cup \mathcal{I}(\Sigma)$.

Next suppose that the two rules are *not* applied at disjoint positions, that is, there is some context δ such that, for example, $c_2 = c_1[\delta]$. There are several cases to analyze.

Suppose first that $l_1 \rightarrow r_1$ if C_1 is in \mathcal{R} . There are now two subcases to discuss. First, if δ contains any infrastructure operation on the path to its hole or if l_2 has an infrastructure operation at its top (in other words the second rule is in $\mathcal{I}(\Sigma)$), then, since \mathcal{R} does not involve any infrastructure operator in its rewriting rules, the term $\theta_2(l_2)$ occurs as a subterm of a subterm $\theta_1(x)$ of $\theta_1(l_1)$ for some variable x ; more precisely, $\theta_1(x) = \rho[\theta_2(l_2)]$ for some context ρ . Then note that $t_1 = c_1[\theta_1(r_1)]$ rewrites in as many steps as occurrences of x in r_1 to a term $t' = c_1[\theta'_1(r_1)]$, where θ'_1 is defined like θ_1 with the exception that $\theta'_1(x) = \rho[\theta_2(r_2)]$. Note that t_2 also rewrites to t' , because $t_2 = c_2[\theta_2(r_2)] = c_1[\delta[\theta_2(r_2)]]$ and $\delta[\theta_2(r_2)]$ rewrites in zero or more steps (as many as occurrences of x in l_1 minus one) to $\theta'_1(l_1)$. Second, if δ contains only operations in Σ on the path to its hole and so does l_2 at its top (so the second rule is also in \mathcal{R}), then we can “freeze” all the subterms of $\theta_1(l_1) = \delta[\theta_2(l_2)]$ rooted in an infrastructure operator, e.g., by temporarily replacing them with some fresh constants, and then use the confluence of \mathcal{R} to get a term u such that $\theta_1(r_1) \rightarrow^* u \leftarrow^* \delta[\theta_2(r_2)]$. Now we can take t' to be the term $c_1[u]$.

Suppose now that $l_1 \rightarrow r_1$ if C_1 is not in \mathcal{R} , that is, it is one of the unconditional infrastructure rules $l_1 \rightarrow r_1$ defined in subsection 3. Note that $l_1 \rightarrow r_1$ cannot be any of the four rules for boolean simplification, because otherwise the application of the other rule would not be possible. If it is one of the two rules

for *if* then note that the other rule may only apply to a subterm of an instance of x or y . If for example the first rule $l_1 \rightarrow r_1$ is $if(true, x, y) \rightarrow x$ and the second rule applies to a subterm of an instance of x , then the same rule for *if* can apply to t_2 to generate the same term as the one obtained from t_1 by appropriately applying the second rule. If the second rule applies to a subterm of an instance of y , then note that after applying the rule for *if* to t_2 one gets exactly t_1 . Finally, recall that the operation *equal?* is assumed to be reduced *eagerly*, so the concrete subterms of *equal?* are assumed in normal form whenever a rule for *equal?* is applied. This means that $l_1 \rightarrow r_1$ *cannot* be any of the rules of *equal?*.

By 2., one can replace any equality $u = v$ in the condition of a rule in \mathcal{R} by $equal?(u, v) = true$. Note that the restriction on u and v to be ground is crucial in 3. Suppose, e.g., that u is a variable, say x . Then there is no rule to reduce the term $equal?(x, v)$ to *true* or *false*. Moreover, one does *not* want to add rules of the form $equal?(x, \tau(y_1, \dots, y_m)) \rightarrow false$ to $\mathcal{I}(\Sigma)$ because one would destroy the confluence of $\mathcal{I}(\Sigma)$ and thus the correctness of the definition of *equal?*: indeed, $equal?(\tau(y_1, \dots, y_m), \tau(y_1, \dots, y_m))$ would reduce to both *true* and *false* in $\mathcal{I}(\Sigma)$.

4 The Main Transformation

The major reason for which conditional rules are inconvenient to implement in a rewriting engine is because, in order to reduce a term, the rewriting engine needs to maintain a *control context* for each conditional rule that is tried to be applied. By control context we here mean the status of the evaluation of the condition (note that a condition is a set of equalities) plus the right hand term that needs to replace the left hand one in case the condition evaluates to *true*. Due to the potential nesting of rewrite rule applications, such a control context may grow arbitrarily, meaning that the rewriting engine needs to pay special care to choosing appropriate data-structures to maintain it and to recover the computation in case the evaluation of a condition fails.

Example 1. Let us consider natural numbers built with 0 and successor s , together with the following, on purpose inefficient, conditional rules defining *odd* and *even* operators on natural numbers:

$$\begin{array}{ll} odd(0) \rightarrow 0, & even(0) \rightarrow s(0), \\ odd(s(x)) \rightarrow 0 \text{ if } even(x) = 0, & even(s(x)) \rightarrow 0 \text{ if } odd(x) = 0, \\ odd(s(x)) \rightarrow s(0) \text{ if } even(x) = s(0), & even(s(x)) \rightarrow s(0) \text{ if } odd(x) = s(0). \end{array}$$

In order to check whether a natural number n , i.e., a term consisting of n successor operations applied to 0, is odd, a rewriting engine may need $\mathcal{O}(2^n)$ rewrites in the worst case. Indeed, if $n > 0$ then either the second or the third rule of *odd* can be applied at the first step; however, in order to apply any of those rules one needs to reduce the even of the predecessor of n , twice. Iteratively, the evaluation of each even involves the reduction of two odds, and so on. Moreover, the rewriting engine needs to maintain a control context data-structure, storing

the status of the application of each (nested) rule that is being tried in a reduction. It is the information stored in this control context that allows the rewriting engine to backtrack and find an appropriate rewriting sequence. \square

A challenging question motivating the present work is the following: would it be possible to automatically replace conditional rules like the above by unconditional ones, so that a rewriting engine's single job would be to *match-and-apply* rules, without worrying about any control context aspects? A positive answer to this question could potentially lead to a new generation of efficient rewriting engines, which would take advantage of today's increasingly highly parallel computing architectures and would potentially allow optimizations that were not possible for conditional rewriting. In this section we show how a conditional rewrite system \mathcal{R} can be automatically transformed into an unconditional one $\overline{\mathcal{R}}$, which practically preserves all the properties of \mathcal{R} . The major idea is, like in the use of continuations (see [15] for a discussion on several independent discoveries of continuations, and [7] for a pragmatic presentation of continuation), to convert the control context into data context. This way, the term to be rewritten is enriched at appropriate positions to contain *all* the information needed to continue its reduction. The rewriting engine does not need to maintain any auxiliary information about the status of the rewriting process: it only needs to find a redex in the term to rewrite and apply a corresponding unconditional rewrite rule, a simple process amenable to high parallelization and optimization.

4.1 An Unsatisfactory Transformation

Once one generates the infrastructure (unconditional) Σ' -rewrite system $\mathcal{I}(\Sigma)$, a simple-minded way to transform a conditional Σ -rewrite system R into an unconditional one is to translate each conditional rule

$$l \rightarrow r \text{ if } u_1 = v_1, \dots, u_m = v_m$$

into an unconditional rewrite rule

$$l \rightarrow \text{if}(\text{equal?}(u_1, v_1) \wedge \dots \wedge \text{equal?}(u_m, v_m), r, l).$$

Such a transformation has the desirable property that both the conditional rewrite system and its unconditional variant can “reach”, by reduction in zero or more steps starting with a given Σ -term, the same set of Σ -terms. In other words, if a and b are Σ -terms then $a \rightarrow^* b$ in the conditional Σ -rewrite system if and only if $a \rightarrow^* b$ in the unconditional Σ' -rewrite system. Therefore, if reachability analysis is what one is interested in then this simple translation provides an effective method to reduce the problem to unconditional rewrite systems. This rewrite system transformation can be useful in systems like Maude, providing commands of the form “**search** $a \Rightarrow^* b$ ” searching for a sequence of applications of rewrite rules transforming a into b .

However, this translation cannot be used to execute conditional rewriting on top of an unconditional rewriting engine. Indeed, if the conjunction of equalities reduces to *false* then the unconditional rewrite system leads to an infinite

rewriting sequence, by keeping applying the rule above. Would it be possible to properly *mark* the term to rewrite whenever a rule is tried and its condition reduces to *false*, so that that rule will not be applied anymore on that position?.

4.2 Adding Control Context Arguments

Like in Viry [16], the idea is to add a few auxiliary arguments to some operators to keep the necessary control context information. This way, terms to rewrite will store information about the conditional rules that can be potentially applied on each of their subterms. Let $\mathcal{R} = (\Sigma, E)$ be any Σ rewriting system. For each n and each $\sigma \in \Sigma_n$, let us associate a unique number between 1 and k_σ to each conditional rewrite rule in R whose lhs is rooted in σ , that is, a rule of the form

$$\sigma(t_1, \dots, t_n) \rightarrow r \text{ if } u_1 = v_1, \dots, u_m = v_m,$$

with $t_1, \dots, t_n, r, u_1, v_1, \dots, u_m, v_m$ terms and $m \geq 1$, where k_σ is the total number of such rules. Note that k_σ is 0 if there is no rule having σ as a root of its lhs, or if all such rules are unconditional.

Let us next define a signature $\overline{\Sigma}$, replacing each $\sigma \in \Sigma_n$ by an operator of $n + k_\sigma$ arguments, $\overline{\sigma} \in \overline{\Sigma}_{n+k_\sigma}$. The additional k_σ arguments are written at the right of the other n arguments, and they can take only two possible values (or constant terms): *true* or *false*. An important step in our transformation technique is to replace all the operations in Σ by corresponding operations in $\overline{\Sigma}$. The intuition for the additional arguments comes from the overall idea of passing the control context (due to conditional rules) into data context: the additional i -th argument of an operation $\overline{\sigma}$ staying at some position in a term to rewrite, tells whether the i -th rule having σ at the root of its lhs is enabled or not at that position; if *true* then it means that the rule can potentially be applied, and if *false* then it means that the rule has been already tried at that position but its condition failed to evaluate to *true*, so there is no need to try it anymore. Let us extend this to Σ -terms, by letting the variables unchanged and replacing each operator σ by $\overline{\sigma}$ with the k_σ additional arguments all *true*. Formally, let $\tau : T_\Sigma(\mathcal{X}) \rightarrow T_{\overline{\Sigma}}(\mathcal{X})$ be a map from Σ -terms to $\overline{\Sigma}$ -terms defined inductively as

- $\overline{x} = x$ for any variable $x \in \mathcal{X}$, and
- $\overline{\sigma(t_1, \dots, t_n)} = \overline{\sigma}(\overline{t_1}, \dots, \overline{t_n}, \text{true}, \dots, \text{true})$ for any $\sigma \in \Sigma_n$ and any terms $t_1, \dots, t_n \in T_\Sigma(\mathcal{X})$.

Let's define another useful map from Σ -terms to $\overline{\Sigma}$ -terms, $\tilde{\cdot}^X : T_\Sigma(X) \rightarrow T_{\overline{\Sigma}}(\mathcal{X})$, but this time indexed by a finite set of variable $X \subseteq \mathcal{X}$, as follows:

- $\tilde{x}^X = x$ for any variable $x \in X$, and
- $\overline{\sigma(\widetilde{t_1, \dots, t_n})^X} = \overline{\sigma}(\tilde{t_1}^X, \dots, \tilde{t_n}^X, b_1, \dots, b_{k_\sigma})$ for any $\sigma \in \Sigma_n$ and any terms $t_1, \dots, t_n \in T_\Sigma(X)$, where $b_1, \dots, b_{k_\sigma} \in \mathcal{X} - X$ are some arbitrary but fixed different fresh variables that do not occur neither in X nor in $\tilde{t_1}^X, \dots, \tilde{t_n}^X$.

Therefore, \tilde{t}^X transforms the Σ -term t into a $\overline{\Sigma}$ -term, replacing each operation $\sigma \in \Sigma$ by $\overline{\sigma} \in \overline{\Sigma}$ and adding distinct variables for the additional arguments, following some arbitrary but deterministic conventions. Given a Σ -term t in $T_\Sigma(X)$ of the form $\sigma(t_1, \dots, t_n)$ for some operation $\sigma \in \Sigma_n$, and given a natural number i between 1 and k_σ , then we let $\tilde{t}_{i/true}^X$ denote the $\overline{\Sigma}$ -term $\overline{\sigma}(\tilde{t}_1^X, \dots, \tilde{t}_n^X, b_1, \dots, b_{i-1}, true, b_{i+1}, \dots, b_{k_\sigma})$, that replaces b_i in \tilde{t}^X by *true*. Similarly, $\tilde{t}_{i/false}^X$ denotes $\overline{\sigma}(\tilde{t}_1^X, \dots, \tilde{t}_n^X, b_1, \dots, b_{i-1}, false, b_{i+1}, \dots, b_{k_\sigma})$, that replaces b_i in \tilde{t}^X by *false*. Thus, $\tilde{t}_{i/true}^X$ (resp. $\tilde{t}_{i/false}^X$) contains the additional control context information whether the i -th conditional rule of σ is enabled.

4.3 An Almost Correct Transformation

For a given conditional Σ -rewrite system \mathcal{R} , we can now define an unconditional $\overline{\Sigma}'$ -rewrite system³ \mathcal{R}' by adding to $\mathcal{I}(\overline{\Sigma})$ the following unconditional $\overline{\Sigma}'$ -rewrite rules. For each conditional ($m \geq 1$) rule $l \rightarrow r$ **if** $u_1 = v_1, \dots, u_m = v_m$ over variables X in \mathcal{R} , say the i -th among the conditional rewrite rules in \mathcal{R} having the root operation of l as a root of their lhs, add to \mathcal{R}' the unconditional rule

$$\tilde{l}_{i/true}^X \rightarrow \text{if}(\text{equal}?(u_1, v_1) \wedge \dots \wedge \text{equal}?(u_m, v_m), \tilde{r}, \tilde{l}_{i/false}^X).$$

For each unconditional rewrite rule $l \rightarrow r$ in \mathcal{R} over variables X , add to \mathcal{R}' an unconditional rewriting rule $\tilde{l}^X \rightarrow \tilde{r}$.

Therefore, for each conditional rule in \mathcal{R} we add an unconditional one in \mathcal{R}' , whose corresponding additional argument of its transformed lhs is *true*. By throwing the control context's ball into matching's court, this intuitively says that such a rule can be applied on a (sub)term only if it is “enabled” in that (sub)term. Its rhs term has a conditional operation at its root, which first evaluates the conjunction of all the equalities of pairs of terms occurring in the condition of the conditional rule; note that these terms are properly transformed into $\overline{\Sigma}$ -terms enabling all possible rules at any of their positions. If the conjunction evaluates to *true* then the rhs of the conditional rule is returned, also modified to enable all possible rules on it. If the condition reduces to *false* then the only thing to do is to “disable” this current rule. Due to the change of the corresponding argument from *true* to *false*, note that matching will disallow this rule to be applied anymore on that (sub)term. Since unconditional rules are always enabled, they are transformed into unconditional rules ignoring the control context arguments in the lhs and enabling all possible rules on its rhs. Note that \mathcal{R}' does not modify any rule in \mathcal{R} if \mathcal{R} already contains only unconditional rules.

Example 2. Let us apply the translation technique above on the conditional rewriting system for odd/even in Example 1. Since there are two conditional rules whose root of lhs is *odd* and two whose root of lhs is *even*, each of these operators will be enriched with two additional arguments. The new, unconditional rewriting system is then:

³ Note that $(\overline{\Sigma}') = (\overline{\Sigma})'$, so we take the liberty to denote this signature $\overline{\Sigma}'$.

$$\begin{aligned}
\overline{odd}(0, b_1, b_2) &\rightarrow 0, \\
\overline{odd}(s(x), true, b_2) &\rightarrow if(equal?(\overline{even}(x, true, true), 0), 0, \overline{odd}(s(x), false, b_2)), \\
\overline{odd}(s(x), b_1, true) &\rightarrow if(equal?(\overline{even}(x, true, true), s(0)), s(0), \overline{odd}(s(x), b_1, false)), \\
\overline{even}(0, b_1, b_2) &\rightarrow s(0), \\
\overline{even}(s(x), true, b_2) &\rightarrow if(equal?(\overline{odd}(x, true, true), 0), 0, \overline{even}(s(x), false, b_2)), \\
\overline{even}(s(x), b_1, true) &\rightarrow if(equal?(\overline{odd}(x, true, true), s(0)), s(0), \overline{even}(s(x), b_1, false)).
\end{aligned}$$

The unconditional rule for \overline{odd} says that 0 is not an odd number, regardless of the control context. The first conditional rule for \overline{odd} has the constant *true* as the first auxiliary argument of its lhs, telling the matching procedure that this rule can be applied only if it was not previously disabled. If the condition of *if* evaluates to *true* then 0 is returned, otherwise the same term as the lhs, except that *true* is replaced by *false*, thus disabling the current conditional rule to avoid getting into non-terminating rewriting. The variable argument b_2 says that it does not matter whether the second conditional rule is enabled or not (but this information will be preserved in case the first conditional rule is disabled). The other conditional equations are similar. If one wants to test whether a number n , i.e., n consecutive applications of successor on 0, is odd, one should reduce the term $\overline{odd}(n)$, i.e., $\overline{odd}(n, true, true)$, under the unconditional rewrite system. Note that the operations 0 and s are not added auxiliary arguments because they do not occur as a root of a lhs of any conditional rule in the original conditional rewriting system. \square

Unfortunately, the translation above suffers from the same problem as that of Viry [16]: for some CTRSs, the generated TRSs have additional normal forms corresponding to terms which could be further reduced in the original CTRS.

Example 3. Consider the problematic CTRS from Section 1:

$$\begin{aligned}
f(g(x)) &\rightarrow x \text{ if } x = 0, \\
g(g(x)) &\rightarrow g(x),
\end{aligned}$$

whose corresponding TRS, according to the transformation above, is:

$$\begin{aligned}
f(g(x), true) &\rightarrow if(equal?(x, 0), x, f(g(x), false)), \\
g(g(x)) &\rightarrow g(x).
\end{aligned}$$

Then note that even if $f(g(g(0)))$ admits a unique normal form in the original CTRS, $f(g(g(0)), true)$ admits two normal forms in its corresponding TRS:

$$\begin{aligned}
f(g(g(0)), true) &\rightarrow f(g(0), true) \rightarrow if(equal?(0, 0), 0, f(g(0), false)) \rightarrow^* 0, \\
f(g(g(0)), true) &\rightarrow if(equal?(g(0), 0), g(0), f(g(g(0)), false)) \rightarrow^* \\
&\rightarrow^* f(g(g(0)), false) \rightarrow f(g(0), false).
\end{aligned}$$

The latter cannot be further reduced with the rules in the TRS. \square

The problem here, like in Viry’s transformation [16], is that a successful application of a rewrite rule may enable some application of a conditional rule that has already been tried before, but at that time failed to apply. One unsatisfactory way to fix this problem is, like in [16], to enforce conditional eagerness on the generated TRS; another, even more unsatisfactory, is to reduce the applicability of the transformation to only innermost, or eager, CTRSs. We next show how to fix this problem in general.

4.4 The Correct Transformation

To fix the problem in the previous subsection, we need a mechanism to “inform” the term to reduce, after each successful application of a rewrite rule, that some “conditional” rules that have been tried before and failed may succeed now. More precisely, we need to traverse the term along the path from the current position (where the successful rule was applied) to its root, and make all the auxiliary arguments of the operations on this path *true*. This can be accomplished, for example, by considering a new (unary) operator, say $\{-\}$, stating that the enclosed term has just been modified, together with appropriate rewrite rules to propagate this information upwards, updating the “applicability bits”: for each $\sigma \in \Sigma_n$ and each $1 \leq i \leq n$, consider a rule

$$\begin{aligned} \sigma(x_1, \dots, x_{i-1}, \{x_i\}, x_{i+1}, \dots, x_n, b_1, \dots, b_{k_\sigma}) &\rightarrow \\ &\rightarrow \{\sigma(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n, \text{true}, \dots, \text{true})\}. \end{aligned}$$

The applicability information of an operation can be updated from several of its subterms; to keep this operation idempotent, we also consider the rule

$$\{\{x\}\} \rightarrow \{x\}.$$

Formally, for a given conditional Σ -rewrite system \mathcal{R} , we let $\overline{\Sigma}'_{\{\cdot\}}$ define the signature $\overline{\Sigma}'$ in the previous subsection extended with the unary operator $\{-\}$ above, and we let $\overline{\mathcal{R}}$ be the unconditional $\overline{\Sigma}'_{\{\cdot\}}$ -rewrite system extending $\mathcal{I}(\overline{\Sigma})$ with the operator $\{-\}$ together with its unconditional rewrite rules above, as well as with the following rules. For each conditional ($m \geq 1$) rule $l \rightarrow r$ if $u_1 = v_1, \dots, u_m = v_m$ over variables X in \mathcal{R} , say the i -th among the conditional rewrite rules in \mathcal{R} having the root operation of l as a root of their lhs, add to $\overline{\mathcal{R}}$:

$$\tilde{l}_{i/\text{true}}^X \rightarrow \text{if}(\text{equal?}(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge \text{equal?}(\{\overline{u_m}\}, \{\overline{v_m}\}), \{\overline{r}\}, \tilde{l}_{i/\text{false}}^X).$$

For each unconditional rewrite rule $l \rightarrow r$ in \mathcal{R} over variables X , add to $\overline{\mathcal{R}}$ an unconditional rewriting rule $\tilde{l}^X \rightarrow \{\overline{r}\}$.

Before we formalize the exact relationship between CTRSs and their unconditional variants, let us define another map of terms, this time from $\overline{\Sigma}$ -terms to Σ -terms. Let $\hat{\cdot}: T_{\overline{\Sigma}}(\mathcal{X}) \rightarrow T_{\Sigma}(\mathcal{X})$ be the map defined inductively as

- $\widehat{x} = x$ for any variable $x \in \mathcal{X}$, and
- $\widehat{\sigma}(t'_1, \dots, t'_n, s_1, \dots, s_{k_\sigma}) = \sigma(\widehat{t'_1}, \dots, \widehat{t'_n})$ for any operator $\sigma \in \Sigma_n$ and any terms $t'_1, \dots, t'_n, s_1, \dots, s_{k_\sigma} \in T_{\Sigma'}(\mathcal{X})$.

Therefore, $\widehat{t'}$ forgets all the auxiliary arguments of each operation occurring in t' . Note in particular that $\widehat{\widehat{t}} = \widehat{t} = t$ for any $t \in T_\Sigma(\mathcal{X})$.

Example 4. Let us consider the problematic CTRS in example 3. Its corresponding TRS generated as above contains the following rules:

$$\begin{aligned} \{\{x\}\} &\rightarrow \{x\} \\ f(\{x\}, b) &\rightarrow \{f(x, \text{true})\} \\ g(\{x\}) &\rightarrow \{g(x)\} \\ f(g(x), \text{true}) &\rightarrow \text{if}(\text{equal?}(\{x\}, \{0\}), \{x\}, f(g(x), \text{false})), \\ g(g(x)) &\rightarrow \{g(x)\}. \end{aligned}$$

Then the term $f(g(g(0)), \text{true})$ admits the normal form $\{0\}$:

$$\begin{aligned} f(g(g(0)), \text{true}) &\rightarrow \text{if}(\text{equal?}(\{g(0)\}, \{0\}), \{g(0)\}, f(g(g(0)), \text{false})) \rightarrow^* \\ &\rightarrow^* f(g(g(0)), \text{false}) \rightarrow f(\{g(0)\}, \text{false}) \rightarrow \\ &\rightarrow \{f(g(0), \text{true})\} \rightarrow \{0\}. \end{aligned}$$

Note that the normal form $\{0\}$ is possible exactly because the information that a subterm has been rewritten is transmitted upwards via the operator $\{-\}$ and its associated rules. The obtained TRS is not confluent, because the term above also admits the normal form 0 , but this time the (at most two) normal forms that a term can have (t and/or $\{t\}$) are very closely related and one can easily infer the desired normal form in the original CTRS. In order to have a unique normal form in the TRS, we will actually enclose the original term into curly brackets before we reduce it, as the theorem below suggests.

Theorem 1. *If \mathcal{R} is a conditional Σ -rewriting system then*

1. *For any ground Σ -terms α and β , $\alpha \rightarrow_{\mathcal{R}}^* \beta$ if and only if there is some ground $\overline{\Sigma}$ -term γ such that $\widehat{\gamma} = \beta$ and $\{\overline{\alpha}\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma\}$;*
2. *\mathcal{R} terminates on a ground Σ -term α if and only if $\overline{\mathcal{R}}$ terminates on $\{\overline{\alpha}\}$;*
3. *$\overline{\mathcal{R}}$ terminates if and only if it terminates on all terms $\{\overline{\alpha}\}$ with α a Σ -term;*
4. *If γ in 1. is a normal form (in $\overline{\mathcal{R}}$) then β is also a normal form (in \mathcal{R});*
5. *If \mathcal{R} terminates then \mathcal{R} is ground confluent iff $\overline{\mathcal{R}}$ is ground confluent.*

Proof. 1. Let us first prove that if $\alpha \rightarrow_{\mathcal{R}}^* \beta$ then there is some ground $\overline{\Sigma}$ -term γ such that $\{\overline{\alpha}\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma\}$ and $\widehat{\gamma} = \beta$. We can show it by well-founded induction on the total number of applications of rewrite rules, including those to derive the equalities in conditions. If $\alpha = \beta$ then the property is true because $\widehat{\overline{\alpha}} = \alpha$. Let us isolate the first rewrite step in the derivation of α : let τ be a ground σ -term such that $\alpha \rightarrow_{\mathcal{R}} \tau \rightarrow_{\mathcal{R}}^* \beta$. Let $l \rightarrow r$ if $u_1 = v_1, \dots, u_m = v_m$ be a rule in

\mathcal{R} , let θ be an appropriate substitution such that $\theta(u_i) \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* \theta(v_i)$ for each $1 \leq i \leq n$, and let c be a context such that $\alpha = c[\theta(l)]$ and $\tau = c[\theta(r)]$. By the induction hypothesis, there are some ground $\overline{\Sigma}$ -terms ρ_{u_i} and ρ_{v_i} such that $\{\overline{\theta(u_i)}\} \rightarrow_{\mathcal{R}}^* \{\rho_{u_i}\}$, $\{\overline{\theta(v_i)}\} \rightarrow_{\mathcal{R}}^* \{\rho_{v_i}\}$, and $\widehat{\rho_{u_i}} = \widehat{\rho_{v_i}}$. Let us prove the following lemma:

Lemma 3. *If α_1, α_2 are Σ -terms and γ_1, γ_2 are $\overline{\Sigma}$ -terms such that $\{\overline{\alpha_1}\} \rightarrow_{\mathcal{R}}^* \{\gamma_1\}$, $\{\overline{\alpha_2}\} \rightarrow_{\mathcal{R}}^* \{\gamma_2\}$, and $\widehat{\gamma_1} = \widehat{\gamma_2}$, then $\{\overline{\alpha_1}\} \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* \{\overline{\alpha_2}\}$.*

Proof. Since γ_1 and γ_2 may differ only in the corresponding *true/false* constant subterms, where the *false* ones reflect the fact that corresponding “conditional” rules have been tried and failed, one can generate another ground $\overline{\Sigma}$ -term, written $\wedge(\gamma_1, \gamma_2)$, with $\widehat{\wedge(\gamma_1, \gamma_2)} = \widehat{\gamma_1} = \widehat{\gamma_2}$, as follows:

- If $\gamma_1 = \gamma_2$ then $\wedge(\gamma_1, \gamma_2)$ is the same $\overline{\Sigma}$ -term;
- If

$$\begin{aligned} \gamma_1 &= \overline{\sigma}(t_1^1, \dots, t_n^1, b_1^1, \dots, b_{k_\sigma}^1), \text{ and} \\ \gamma_2 &= \overline{\sigma}(t_1^2, \dots, t_n^2, b_1^2, \dots, b_{k_\sigma}^2), \end{aligned}$$

then

$$\wedge(\gamma_1, \gamma_2) = \overline{\sigma}(\wedge(t_1^1, t_1^2), \dots, \wedge(t_n^1, t_n^2), b_1, \dots, b_{k_\sigma}),$$

where b_i is *true* if and only if both b_i^1 and b_i^2 are *true*.

The definition of $\wedge(\gamma_1, \gamma_2)$ above is correct because $\widehat{\gamma_1} = \widehat{\gamma_2}$. We next show by structural induction (note that γ_1, γ_2 , and $\wedge(\gamma_1, \gamma_2)$ have the same structure, except for the constants *true* and *false*) that $\{\gamma_1\} \rightarrow_{\mathcal{R}}^* \{\wedge(\gamma_1, \gamma_2)\} \leftarrow_{\mathcal{R}}^* \{\gamma_2\}$. The case when $\gamma_1 = \gamma_2$ is trivial. Otherwise, if γ_1 and γ_2 are as above, then by the induction hypothesis it follows that $\{t_i^1\} \rightarrow_{\mathcal{R}}^* \{\wedge(t_i^1, t_i^2)\} \leftarrow_{\mathcal{R}}^* \{t_i^2\}$ for any $1 \leq i \leq n$, so

$$\begin{aligned} \{\gamma_1\} &\rightarrow_{\mathcal{R}}^* \{\overline{\sigma}(\wedge(t_1^1, t_1^2), \dots, \wedge(t_n^1, t_n^2), b_1^1, \dots, b_{k_\sigma}^1)\}, \text{ and} \\ \{\gamma_2\} &\rightarrow_{\mathcal{R}}^* \{\overline{\sigma}(\wedge(t_1^1, t_1^2), \dots, \wedge(t_n^1, t_n^2), b_1^2, \dots, b_{k_\sigma}^2)\}. \end{aligned}$$

If $(b_1^1, \dots, b_{k_\sigma}^1)$ and $(b_1^2, \dots, b_{k_\sigma}^2)$ are not already equal, then let i be some index with $1 \leq i \leq n$ such that $b_j^1 = b_j^2$ for all $1 \leq j < i$ and $b_i^1 \neq b_i^2$. Suppose that $b_i^1 = \text{true}$ and $b_i^2 = \text{false}$ (the other case is treated similarly). Since there is only one way for the i -th auxiliary argument of the operation $\overline{\sigma}$ to become *false*, namely when the i -th rule having $\overline{\sigma}$ at the top of its left-hand-side term, say

$$\widetilde{l}_{i/\text{true}}^X \rightarrow \text{if}(\text{equal?}(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge \text{equal?}(\{\overline{u_m}\}, \{\overline{v_m}\}), \{\overline{r}\}, \widetilde{l}_{i/\text{false}}^X),$$

is applied and the conjunction of equalities reduces to *false*, one can replay the exact rewriting sequence that led to $b_i^2 = \text{false}$ in γ_2 on γ_1 and thus obtain

$$\{\gamma_1\} \rightarrow_{\mathcal{R}}^* \{\overline{\sigma}(\wedge(t_1^1, t_1^2), \dots, \wedge(t_n^1, t_n^2), b_1^1, \dots, b_{i-1}^1, \text{false}, b_{i+1}^1, \dots, b_{k_\sigma}^1)\}.$$

Iterating this process for the remaining pairs of unequal auxiliary subterms *true* or *false*, we eventually obtain $\{\gamma_1\} \rightarrow_{\mathcal{R}}^* \{\wedge(\gamma_1, \gamma_2)\} \leftarrow_{\mathcal{R}}^* \{\gamma_2\}$. \square

Let us now return to the proof of the main theorem. By Lemma 3, it follows that $\{\overline{\theta(u_i)}\} \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* \{\overline{\theta(v_i)}\}$, so by 2. in Prop. 1, $equal?(\{\overline{\theta(u_i)}\}, \{\overline{\theta(v_i)}\}) \rightarrow_{\mathcal{R}}^* true$ for each $1 \leq i \leq n$. Therefore, since $\overline{\alpha} = \overline{c[\theta(l)]}$ it follows that $\{\overline{\alpha}\} \rightarrow_{\mathcal{R}}^* \{\overline{c[\theta(r)]}\}$. Further, since $c[\theta(r)] \rightarrow_{\mathcal{R}}^* \beta$, by the induction hypothesis there is some ground $\overline{\Sigma}$ -term γ such that $\{c[\theta(r)]\} \rightarrow_{\mathcal{R}}^* \{\gamma\}$ and $\widehat{\gamma} = \beta$. The rest follows from the fact that $\overline{c[\theta(r)]} = \overline{c[\theta(r)]}$.

For the other implication, we can prove a more general result stating that if δ and γ are ground $\overline{\Sigma}$ -terms such that $\{\delta\} \rightarrow_{\mathcal{R}}^* \{\gamma\}$, then $\widehat{\delta} \rightarrow_{\mathcal{R}}^* \widehat{\gamma}$; the desired result follows then by taking δ to be $\overline{\alpha}$. We show this more general result by well-founded induction on the length of the derivation $\{\delta\} \rightarrow_{\mathcal{R}}^* \{\gamma\}$. First note that the property trivially holds when $\delta = \gamma$. Also, if there is some intermediate ground $\overline{\Sigma}$ -term τ such that $\{\delta\} \rightarrow_{\mathcal{R}}^+ \{\tau\} \rightarrow_{\mathcal{R}}^+ \{\gamma\}$, then by induction it follows that $\widehat{\delta} \rightarrow_{\mathcal{R}}^* \widehat{\tau} \rightarrow_{\mathcal{R}}^* \widehat{\gamma}$. Let us next assume that there is no other $\overline{\Sigma}$ -term like τ above, and let us consider the first step in the derivation of δ : let

$$\widehat{l}_{i/true}^X \rightarrow if(equal?(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge equal?(\{\overline{u_m}\}, \{\overline{v_m}\}), \{\overline{r}\}, \widehat{l}_{i/false}^X),$$

be the rule in $\overline{\mathcal{R}}$ (corresponding to a rule $l \rightarrow r$ if $u_1 = v_1, \dots, u_m = v_m$ in \mathcal{R}) which is first applied on δ , and let c be the context where this rule is applied, that is, $\delta = c[\theta(\widehat{l}_{i/true}^X)]$ for some appropriate substitution θ . Let τ be the term

$$c[\theta(if(equal?(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge equal?(\{\overline{u_m}\}, \{\overline{v_m}\}), \{\overline{r}\}, \widehat{l}_{i/false}^X))].$$

Since there is no rule in $\overline{\mathcal{R}}$ containing the $if(-, -, -)$ operation in its left-hand-side term and since γ contains no such operators either, we deduce that there are two possibilities in order for $\{\tau\} \rightarrow_{\mathcal{R}}^* \{\gamma\}$ to take place:

(a) The subterm of τ rooted in $if(-, -, -)$ is not needed in the rewriting sequence and therefore discarded via a rule which contains fewer variables in its right-hand-side than in its left-hand-side. If this is the case, then it is also the case that $c \rightarrow_{\mathcal{R}}^* \gamma$ using the same number of rewrite steps. Regarding c as a ground term (replacing its “hole” by some constant), by the induction hypothesis we get that $\widehat{c} \rightarrow_{\mathcal{R}}^* \widehat{\gamma}$. However, this leads to $\widehat{c[l]} \rightarrow_{\mathcal{R}}^* \widehat{\gamma}$ by using the exact rewrite steps (but with a different substitution); the rest follows because $\widehat{c[l]} = \widehat{\gamma}$.

(b) The subterm of τ rooted in $if(-, -, -)$ is needed in the rewriting sequence. That means that that subterm reduces to either $\{\overline{r}\}$ or $\widehat{l}_{i/false}^X$, and, by potentially reordering the applications of rewrite rules, that either $\{c[\{\overline{r}\}]\}$ or $\{c[\widehat{l}_{i/false}^X]\}$ reduces to $\{\gamma\}$ in $\overline{\mathcal{R}}$, respectively. In the latter case, the result follows immediately because $\widehat{\delta} = \widehat{c[l]}$. In the former case, which can happen only if

$$\{\theta(if(equal?(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge equal?(\{\overline{u_m}\}, \{\overline{v_m}\}), \{\overline{r}\}, \widehat{l}_{i/false}^X))\}$$

reduces to $\{\theta(\overline{r})\}$, since $\widehat{\tau} \rightarrow_{\mathcal{R}}^* \widehat{c[r]}$ it suffices to show that $\widehat{\delta} \rightarrow_{\mathcal{R}}^* \widehat{\tau}$. Note that each of $equal?(\{\overline{u_i}\}, \{\overline{v_i}\})$ for $1 \leq i \leq m$ reduces to $true$ in $\overline{\mathcal{R}}$, so by 1. in Proposition 1 we get that $\{\overline{u_i}\} \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* \{\overline{v_i}\}$ for each $1 \leq i \leq m$. Moreover, since the

number of rewrite steps are preserved in the proof of 1. in Proposition 1, we can use again the induction hypothesis and infer that $\widehat{u_i} \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* \widehat{v_i}$, that is, that $u_i \rightarrow_{\mathcal{R}}^* \leftarrow_{\mathcal{R}}^* v_i$ for each $1 \leq i \leq m$. That means that the original conditional rule $l \rightarrow r$ if $u_1 = v_1, \dots, u_m = v_m$ in \mathcal{R} applies, so $l = \widehat{\widehat{\theta}(l)} \rightarrow_{\mathcal{R}} \widehat{\widehat{\theta}(r)}$.

2. In the proof of **1.** it has been shown that if a rewrite step in \mathcal{R} can be applied on a Σ -term α , then a corresponding non-empty sequence of rewrite steps can be applied in $\overline{\mathcal{R}}$ on $\{\overline{\alpha}\}$. That means that if $\overline{\mathcal{R}}$ terminates on $\{\overline{\alpha}\}$ then \mathcal{R} terminates on α . For the other implication, suppose that there is some infinite rewriting sequence in $\overline{\mathcal{R}}$ starting with $\{\overline{\alpha}\}$, while \mathcal{R} terminates on α . We recall the reader that, in this paper, a conditional rewrite system terminates on a term if and only if there is no infinite rewriting sequence starting with that term *and* the condition of any rule that can be potentially applied along such a rewriting sequence terminates (and, of course, so on recursively). Since there is no rule in $\overline{\mathcal{R}}$ containing operators $if(-, -, -)$ at non-top positions in their left-hand-side terms, there are only two possibilities in order for an infinite rewriting sequence of $\overline{\alpha}$ to exist in $\overline{\mathcal{R}}$:

(a) The number of (chain-)nested operators $if(-, -, -)$ grows unbounded in the rewriting sequence of $\overline{\alpha}$, that is, for any natural number p there is some $\overline{\Sigma'}$ -term γ such that $\{\overline{\alpha}\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma\}$ and γ contains at least p nested $if(-, -, -)$ operators. Since these operators are introduced by the unconditional rules in $\overline{\mathcal{R}}$ in order to evaluate the condition of the corresponding conditional rule in \mathcal{R} , we can deduce that it is the case that during the rewriting of α in \mathcal{R} , a conditional rule (e.g., the one corresponding to the topmost $if(-, -, -)$ operator) is attempted to be applied but its condition does not terminate to evaluate. Therefore, in this case \mathcal{R} does not terminate on α .

(b) There is some bound for the number of nested $if(-, -, -)$ operators that can occur in the rewriting sequence of $\{\overline{\alpha}\}$ in $\overline{\mathcal{R}}$. Then we claim that there is some $\overline{\Sigma}$ -term γ_0 of a term occurring in the infinite rewriting sequence of $\{\overline{\alpha}\}$, which has the property that there is an infinite number of $\overline{\Sigma}$ -terms $\gamma_1, \gamma_2, \dots$, such that $\{\gamma_0\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma_1\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma_2\} \dots$. Indeed, we can search for such a γ_0 as follows. First, we attempt to pick $\overline{\alpha}$ as γ_0 . This is possible only if there is an infinite number of $\overline{\Sigma}$ -terms in the rewriting sequence of $\{\overline{\alpha}\}$. If this is not the case, then we can attempt to search γ_0 inductively among the terms $\overline{u_1}, \overline{v_1}, \dots, \overline{u_m}, \overline{v_m}$ that occur in the subterm of the form

$$if(equal?(\{\overline{u_1}\}, \{\overline{v_1}\}) \wedge \dots \wedge equal?(\{\overline{u_m}\}, \{\overline{v_m}\}), t, t')$$

of the term following *the last* $\overline{\Sigma}$ -term in the derivation of $\{\overline{\alpha}\}$. Since there is a bound for the number of nested $if(-, -, -)$ operators, we will eventually find such a $\overline{\Sigma}$ -term γ_0 . There are now again two possibilities. If γ_0 occurs as a subterm in a subterm whose root is an $if(-, -, -)$ operator, then it means that the corresponding conditional rule in \mathcal{R} does not evaluate its condition in a finite number of steps, so \mathcal{R} does not terminate. If γ_0 does not occur under an $if(-, -, -)$ operator, then one can generate an infinite rewriting sequence of α in \mathcal{R} : indeed, as shown in the proof of **1.**, each rewrite sequence $\{\gamma_i\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma_{i+1}\}$ corresponds to a non-empty rewrite sequence $\widehat{\gamma_i} \rightarrow_{\mathcal{R}} \widehat{\gamma_{i+1}}$ in \mathcal{R} .

3. The only thing to show is that $\overline{\mathcal{R}}$ terminates when it terminates on all terms $\{\overline{\alpha}\}$ with α a Σ -term. We can show by induction on the number of $if(-, -, -)$ and \neg -operators that it contains, that if there is any $\overline{\Sigma}'$ -term γ generating an infinite rewriting sequence in $\overline{\mathcal{R}}$, then there is some $\overline{\Sigma}$ -term δ which also generates an infinite rewriting sequence in $\overline{\mathcal{R}}$. If γ has no such operators then we can simply take δ to be γ . If γ contains such an operator, then we can either apply the induction hypothesis and take one of the subterms of the subterm rooted in that operator in case it admits an infinite rewriting sequence, or otherwise first reduce the subterm rooted in that operator to its normal form and then distinguish two cases: (a) if the operator “dissolves” due to a rule of the form

$$\begin{aligned} true \wedge true &\rightarrow true, \\ true \wedge false &\rightarrow false, \\ false \wedge true &\rightarrow false, \\ false \wedge false &\rightarrow false, \\ if(true, x, y) &\rightarrow x, \\ if(false, x, y) &\rightarrow y. \end{aligned}$$

then apply the induction hypothesis on the new term obtained from γ after dissolving that operator and thus obtain a desired δ , and (b) if the operator does not dissolve during the rewriting of γ then, since there is no rule in $\overline{\mathcal{R}}$ whose left-hand-side terms contain the special operators except the six ones above, then we can safely replace the subterm rooted in the special operator by any other term containing no special operator, and still have an infinite rewriting of the new term obtained this way from γ . This new term has one less special operator, so the induction hypothesis can apply. Therefore, if $\overline{\mathcal{R}}$ does not terminate then there is some $\overline{\Sigma}$ -term δ on which $\overline{\mathcal{R}}$ does not terminate either. Notice now that if $\overline{\mathcal{R}}$ does not terminate on such a δ , then it will not terminate on the term δ either, which proves the result.

4. It follows immediately from the proof of 1..

5. It follows by 1. and 2..

5 Putting Them All Together

We can now present the main result of this paper, namely a technique providing a rewriting engine that accepts conditional rewrite rules, obtained by appropriately wrapping a simpler rewriting engine that only accepts unconditional rules.

Input: a conditional Σ -rewrite system \mathcal{R} and a Σ -term α over variables X to be reduced with \mathcal{R} .

Step 1: Add the variables in X as fresh constants into Σ , so that α becomes a ground Σ -term;

Step 2: Generate $\overline{\Sigma}$ like in Subsection 4.2, by adding to each operator $\sigma \in \Sigma$ as many auxiliary arguments as conditional rules of lhs rooted in σ are in \mathcal{R} ;

Step 3: Generate the infrastructure unconditional $\overline{\Sigma}'$ -rewrite system $\mathcal{I}(\overline{\Sigma})$ by adding to $\overline{\Sigma}$ the operators *equal?* and *if* ($_$, $_$, $_$) as well as their corresponding rules described in Section 3;

Step 4: Generate the unconditional $\overline{\Sigma}'$ -rewrite system $\overline{\mathcal{R}}$ by adding to $\mathcal{I}(\overline{\Sigma})$ the operation $\{_ \}$ and its rules, as well as the unconditional $\overline{\Sigma}'$ -rewrite rules associated to the rules in \mathcal{R} as shown in Subsection 4.4;

Step 5: Reduce the term $\{\overline{\alpha}\}$ to a normal form in $\overline{\mathcal{R}}$, say $\{\gamma\}$, using any engine for unconditional rewriting;

Step 6: Return the Σ -term $\hat{\gamma}$.

We claim that the steps above, by applying a series of simple and totally automatic syntactic transformations to the input conditional rewrite system and term to rewrite, yield a rewriting engine that accepts conditional rewrite systems.

Step 1 shows a usual way to reduce terms with variables: interpret the variables as constants. However, in our framework it is quite important to add these variables explicitly as constants early in the reduction process. This is because the equality operator will need to consider these constants as distinct operations, so that it can add appropriate rules, including ones of the form “*equal?*($a, \sigma(x_1, \dots, x_n)$) \rightarrow *false*” for any such constant a and operation $\sigma \in \Sigma_n$.

Step 2 modifies the signature Σ into $\overline{\Sigma}$, by analyzing the rules in \mathcal{R} and adding an appropriate number of arguments to operations in Σ . Since the constants added to Σ at Step 1 are fresh, no rule in \mathcal{R} has them as lhs terms, so these constants will not be changed in $\overline{\Sigma}$. Note, however, that other constants in Σ may get translated into operations with several arguments.

Step 3 adds the auxiliary equality and conditional operators that are needed to translate the conditional rules into unconditional ones. Note, again, that the constants added at Step 1 will increase the number of rules for *equal?*.

Step 4 generates the unconditional rewrite system $\overline{\mathcal{R}}$, by adding exactly one unconditional rule per conditional or unconditional rule in \mathcal{R} . Once $\overline{\mathcal{R}}$ is available, any engine for unconditional rewriting can be used to reduce $\overline{\alpha}$ under $\overline{\mathcal{R}}$, as done in Step 5. Note that the normal form γ of $\overline{\alpha}$ in $\overline{\mathcal{R}}$, if it exists, is a $\overline{\Sigma}$ -term. Therefore, since the hat function is defined as $\hat{\cdot} : T_{\overline{\Sigma}}(\mathcal{X}) \rightarrow T_{\Sigma}(\mathcal{X})$, the term $\hat{\gamma}$ returned at Step 6 is indeed a Σ -term, as the result of reducing α under \mathcal{R} is expected to be.

Theorem 2. *The algorithm above, taking a conditional Σ -rewrite system \mathcal{R} and a Σ -term α as input, terminates iff \mathcal{R} terminates on α . If the algorithm terminates and outputs a Σ -term β , then β is a normal form of α in \mathcal{R} .*

Proof. The algorithm terminates if and only if its Step 5 terminates, because all the other steps are nothing but simple syntactic translations over a finite signature and number of rewrite rules. Step 5 terminates if and only if $\overline{\mathcal{R}}$ terminates on $\{\overline{\alpha}\}$. By 2 in Theorem 1, this is equivalent to saying that \mathcal{R} terminates on α . Now suppose that the algorithm terminates and that it outputs a Σ -term β . By Steps 5 and 6 and the discussion preceding this theorem, this happens if and only if there is some $\overline{\Sigma}$ -term γ such that $\{\overline{\alpha}\} \rightarrow_{\overline{\mathcal{R}}}^* \{\gamma\}$ and $\hat{\gamma} = \beta$, which, by 1 in

Theorem 1, is equivalent to $\alpha \rightarrow_{\mathcal{R}}^* \beta$. Since γ is a normal form in $\overline{\mathcal{R}}$, it follows by 4 in Theorem 1 that β is a normal form in \mathcal{R} .

6 Preliminary Experiments

As mentioned previously in the paper, our original purpose for translating conditional rewrite systems into unconditional ones was to ease the process of implementing rewriting engines, at the same time aiming at highly efficient implementations of rewriting based on a fast and simple *rewriting virtual machine*. To test the computational equivalence between conditional rewriting systems and their corresponding unconditional variants, we have applied the translation presented in Section 5 manually for the conditional rewrite system in Example 1, and performed some experiments on a 2.4GHz PC machine using two major rewriting engines: Maude [4] and Elan [3]. The results of these experiments are listed in Figure 1, in seconds; the “-” should be read “core dump”.

n	Maude		Elan	
	Conditional	Unconditional	Conditional	Unconditional
11	0.007	0.003	0.604	0.061
13	0.027	0.009	2.309	0.244
15	0.142	0.037	8.922	0.959
17	0.459	0.154	34.2	3.8
21	7.9	2.6	548.1	62.9
23	29.2	10.4	-	-
25	117.2	49.3	-	-
31	7431	2489	-	-

Fig. 1. Times in seconds to reduce $odd(n)$ using Maude and ELAN, using both the conditional rewrite system in Example 1 and its unconditional variant in Example 2.

It was an unexpected surprise to see that the unconditional variant was much faster than its corresponding conditional rewrite system: almost 3 times faster for Maude and 10 times for ELAN. Since the implementation details of these rewrite engines are not well documented, we do not know the exact reasons for which conditional rewriting is so slow in these systems in comparison to unconditional rewriting. However, the preliminary results in Figure 1 tell us that maintaining the control context required to backtrack through conditional rewrite rules is a non-trivial matter, and that the translation technique proposed in this paper can be perhaps implemented in these systems to bring immediate benefits to conditional rewriting.

7 Conclusion and Future Work

An automatic technique to transform a conditional term rewriting system into an unconditional one was presented, which preserves all the major properties. The technique consists of adding some key auxiliary arguments to certain operations, which are used to maintain the control context as data context.

Only first steps towards a generic transformation procedure for general conditional rewrite systems have been made here. We have not considered rewriting modulo axioms yet, such as associativity, commutativity and identity, but these will be considered soon. Also, the current technique can be improved: not all the auxiliary arguments added to operations seem to be necessary. Can we statically reduce their number?

Future work will also investigate how well analysis techniques for unconditional rewriting systems translate into corresponding analysis techniques for conditional ones via our transformation. One complication here is that some of the operations that we introduce require to be evaluated eagerly in some arguments. This is also a drawback of our technique if one wants to use it for theorem proving purposes. Can one find a transformation which imposes no evaluation strategies?

References

1. S. Antoy, B. Brassel, and M. Hanus. Conditional narrowing without conditions. In *5th ACM SIGPLAN international conference on Principles and practice of declarative programming (PPDP'03)*, pages 20–31. ACM Press, 2003.
2. J. Bergstra and J. Klop. Conditional rewrite rules: Confluence and termination. *Journal of Computer and System Sciences*, 32(3):323–362, 1986.
3. P. Borovansky, H. Cirstea, H. Dubois, C. Kirchner, H. Kirchner, P. Moreau, C. Ringeissen, and M. Vittek. *ELAN: User Manual*, 2000. Loria, Nancy, France.
4. M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. *Maude 2.0 Manual*, 2003. <http://maude.cs.uiuc.edu/manual>.
5. N. Dershowitz and M. Okada. A rationale for conditional equational programming. *Theoretical Computer Science*, 75:111–138, 1990.
6. R. Diaconescu and K. Futatsugi. *CafeOBJ Report: The Language, Proof Techniques, and Methodologies for Object-Oriented Algebraic Specification*. World Scientific, 1998. AMAST Series in Computing, volume 6.
7. D. P. Friedman, C. T. Haynes, and M. Wand. *Essentials of programming languages*. MIT Press, 1992.
8. E. Giovannetti and C. Moiso. Notes on the elimination of conditions. In *1st International Workshop on Conditional Term Rewriting Systems (CTRS'87)*, volume 308 of *LNCS*, pages 91–97. Springer, 1987.
9. J. Goguen, T. Winkler, J. Meseguer, K. Futatsugi, and J.-P. Jouannaud. Introducing OBJ. In *Software Engineering with OBJ: algebraic specification in action*, pages 3–167. Kluwer, 2000.
10. M. Hanus. The integration of functions into logic programming: From theory to practice. *The Journal of Logic Programming*, 19 & 20:583–628, 1994.

11. C. Hintermeier. How to transform canonical decreasing ctrss into equivalent canonical trss. In *4th International Workshop on Conditional and Typed Rewriting Systems (CTRS'94)*, volume 968 of *LNCS*, pages 186–205, 1994.
12. J. Meseguer and G. Roşu. Rewriting logic semantics: From language specifications to formal analysis tools. In *2nd International Joint Conference on Automated Reasoning (IJCAR'04)*, *LNCS*, to appear, 2004.
13. A. Middeldorp and E. Hamoen. Completeness results for basic narrowing. *Journal of Applicable Algebra in Eng., Communication and Computing*, 5:313–353, 1994.
14. E. Ohlebusch. *Advanced Topics in Term Rewriting*. Springer, 2002.
15. J. C. Reynolds. The discoveries of continuations. *LISP and Symbolic Computation*, 6(3–4):233–247, 1993.
16. P. Viry. Elimination of conditions. *Journal of Symbolic Computation*, 28:381–401, Sept. 1999.
17. E. Visser. Stratego: A language for program transformation based on rewriting strategies. System description of Stratego 0.5. In *Rewriting Techniques and Applications (RTA'01)*, volume 2051 of *LNCS*, pages 357–361. Springer, May 2001.